

Principal-Agent Models

Inés Macho-Stadler^a and David Pérez-Castrillo^a

^a Universitat Autònoma de Barcelona, Spain.

Article Outline

Glossary

- I. Definition
- II. Introduction
- III. The Base Game
- IV. Moral Hazard
- V. Adverse Selection
- VI. Future Directions
- VII. Bibliography

Glossary

Information Economics

Information economics studies how information and its distributions among the players affect economic decisions.

Asymmetric Information

In a relationship or a transaction, there is asymmetric information when one party has more or better information than other party concerning relevant characteristics of the relationship or the transaction. There are two types of asymmetric information problems: Moral Hazard and Adverse Selection.

Principal – Agent

The principal-agent model identifies the difficulties that arise in situations where there is asymmetric information between two parties and finds the best contract in such environments. The “principal” is the name used for the contractor while the “agent” corresponds to the contractee. Both principal and agent could be individuals, institutions, organizations, or decision centres. The optimal solutions propose mechanisms that try to align the interests of the agent

with those of the principal, such as piece rates or profit sharing; or that induce the agent to reveal the information, such as self-reporting contracts.

Moral Hazard (hidden action)

The term moral hazard initially referred to the possibility that the redistribution of risk (such as insurance which transfers risk from the insured to the insurer) changes people's behaviour. This term, which has been used in the insurance industry for many years, was studied first by Kenneth Arrow.

In principal-agent models, the term moral hazard is used to refer to all environments where the ignorant party lacks information about the behaviour of the other party *once* the agreement has been signed, in such a way that the asymmetry arises *after* the contract is settled.

Adverse Selection (hidden information)

The term adverse selection was originally used in insurance. It describes a situation where, as a result of private information, the insured are more likely to suffer a loss than the uninsured (such as offering a life insurance contract at a given premium may imply that only the people with a risk of dying over the average take it).

In principal-agent models, we say that there is an adverse selection problem when the ignorant party lacks information while negotiating a contract, in such a way that the asymmetry is *previous* to the relationship.

I. Definition of the Subject and Its Importance

Principal-Agent models provide the theory of *contracts under asymmetric information*. Such a theory analyzes the characteristics of optimal contracts and the variables that influence these characteristics, according to the behaviour and information of the parties to the contract. This approach has a close relation to *game theory* and *mechanism design*: it analyzes the strategic behaviour by agents who hold private information and proposes mechanisms that minimize the inefficiencies due to such strategic behaviour. The costs incurred by the principal (the contractor) to ensure that the agents (the contractees) will act in her interest are some type of *transaction cost*. These costs include the tasks of investigating and selecting appropriate agents, gaining information to set performance standards, monitoring agents, bonding payments by the agents, and residual losses.

Principal-agent theory (and *Information Economics* in general) is possibly the area of economics that has evolved the most over the past twenty-five years. It was initially developed in parallel with the new economics of *Industrial Organization* although its applications include now almost all areas in economics, from finance and political economy to growth theory.

Some early papers centred on incomplete information in insurance contracts, and more particularly on moral hazard problems, are Spence and Zeckhauser (1971) and Ross (1973). The theory soon generalized to dilemmas associated with contracts in other contexts (Jensen and Meckling, 1976; Harris and Raviv, 1978). It was further developed in the mid-seventies by authors such as Pauly (1968, 1974), Mirrlees (1975), Harris and Raviv (1979), and Holmström (1979). Arrow (1985) worked on the analysis of the optimal incentive contract when the agent's effort is not verifiable.

A particular case of adverse selection is the one where the type of the agent relates to his valuation of a good. Asymmetric information about buyers' valuation of the objects sold is the fundamental reason behind the use of auctions. Vickrey (1961) provides the first formal analysis of the first and second-prize auctions. Akerlof (1970) highlighted the issue of adverse selection in his analysis of the market for second-hand goods. Further analyses include the early work of Mirrlees (1971), Spence (1974), Rothschild and Stiglitz (1976), Mussa and Rosen (1978), as well as Baron and Myerson (1982), and Guesnerie and Laffont (1984).

The importance of the topic has also been recognized by the Nobel Foundation. James A. Mirrlees and William Vickrey were awarded with the Nobel Prize in Economics in 1996 "for their fundamental contributions to the economic theory of incentives under asymmetric information". Five years later, in 2001, George A. Akerlof, A. Michael Spence and Joseph E. Stiglitz also obtained the Nobel Prize in Economics "for their analyses of markets with asymmetric information".

II. Introduction

The objective of the Principal-Agent literature is to analyze situations in which a contract is signed under asymmetric information, that is, when one party knows certain relevant things of which the other party is ignorant. The simplest situation concerns a bilateral relationship: the contract between one principal and one agent. The objective of the contract is for the agent to

carry out actions on behalf of the principal; and to specify the payments that the principal will pass on to the agent for such actions.

In the literature, it is always assumed that the principal is in charge of designing the contract. The agent receives an offer and decides whether or not to sign the contract. He will accept it whenever the utility obtained from it is greater than the utility that the agent would get from not signing. This utility level that represents the agent's outside opportunities is his *reservation utility*. In order to simplify the analysis, it is assumed that the agent cannot make a counter offer to the principal. This way of modelling implicitly assumes that the principal has all the bargaining power, except for the fact that the reservation utility can be high in those cases where the agent has excellent outside opportunities.

If the agent decides not to sign the contract, the relationship does not take place. If he does accept the offer, then the contract is implemented. It is crucial to notice that the contract is a reliable promise by both parties, stating the principal and agent's obligations for all (contractual) contingencies. It can only be based on *verifiable variables*, that is, those for which it is possible for a third party (a court) to verify whether the contract has been fulfilled. When some players know more than others about relevant variables, we have a situation with asymmetric information. In this case, *incentives* play an important role.

Given the description of the game played between principal and agent, we can summarize its timing in the following steps:

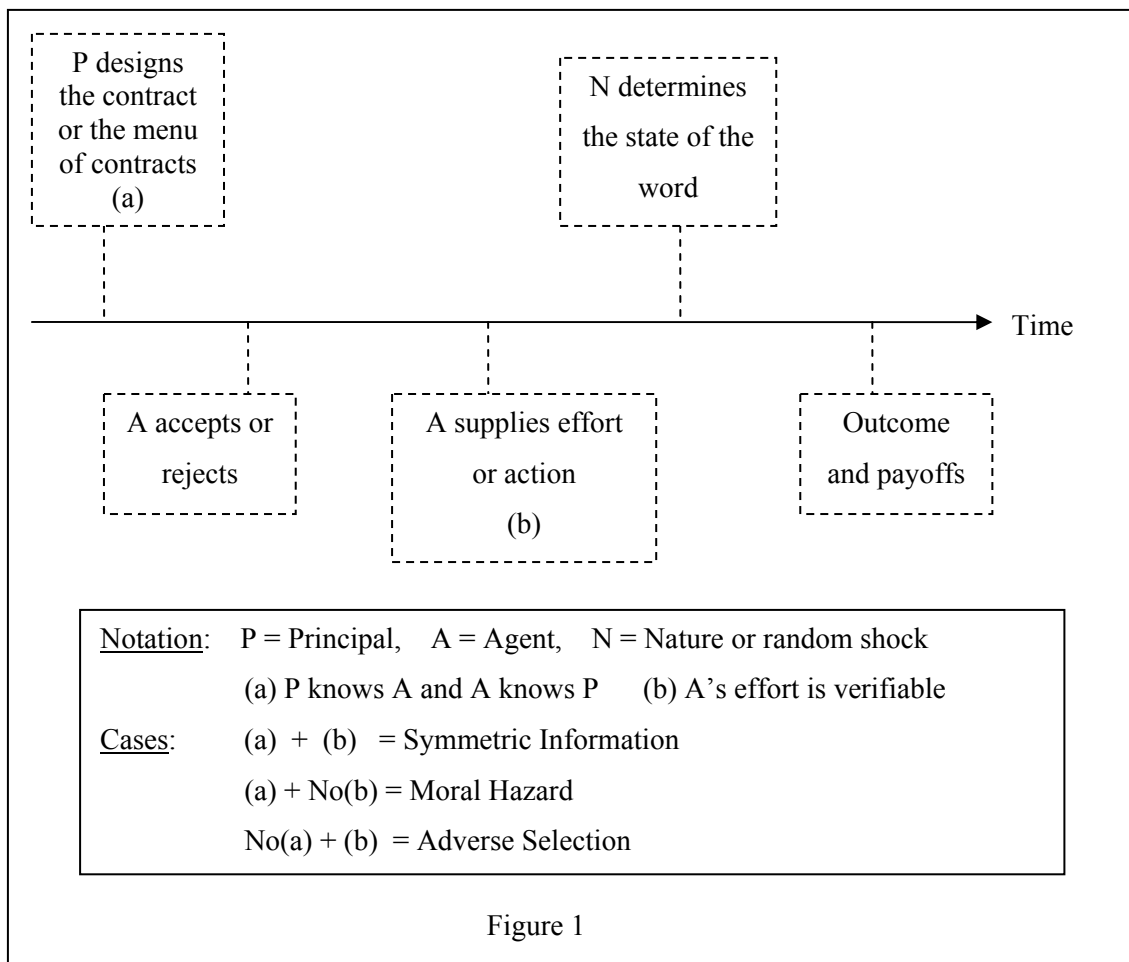
- i) The principal designs the contract (or set of contracts) that she will offer to the agent, the terms of which are not subject to bargaining.
- ii) The alternatives opened to the agent are to accept or to reject the contract. The agent accepts it if he desires so, that is, if the contract guarantees him greater expected utility than any other (outside) opportunities available to him.
- iii) The agent carries out an action or effort on behalf of the principal.
- iv) The outcome is observed and the payments are done.

From these elements, it can be seen that *the agent's objectives may be in conflict with those of the principal*. When the information is asymmetric, the informed party tries to take advantage, while the uninformed player tries to control this behaviour via the contract. Since a Principal-Agent problem is a sequential game, the solution concept to use is *Subgame (Bayesian) Perfect Equilibrium*.

The set-up gives rise to three possible scenarios:

- 1) The *Symmetric Information* case, where the two players share the same information, even if they both may ignore some important elements (some elements may be uncertain).
- 2) The *Moral Hazard* case, where the asymmetry of information arises once the contract has been signed: the decision or the effort of agent is not verifiable and hence it cannot be included in the contract.
- 3) The *Adverse Selection* case, where the asymmetry of information is previous to the signature of the contract: a relevant characteristic of the agent is not verifiable and hence the principal cannot include it in the contract.

Figure 1 summarizes the timing of the relationship and the three cases as a function of the information available to the participants.



To see an example of moral hazard, consider a laboratory or research centre (the principal) that contracts a researcher (the agent) to work on a certain project. It is difficult for the principal to distinguish between a researcher who is thinking about how to push the project through, and a

researcher who is thinking about how to organize his evening. It is precisely this difficulty in controlling effort inputs, together with the inherent uncertainty in any research project, what generates a moral hazard problem, which is a non-standard labour market problem.

For an example of adverse selection, consider a regulator who wants to set the price of the service provided by a public monopoly equal to the average costs in the firm (to avoid subsidies). This policy (as many others) is subject to important informational requirements. It is not enough that the regulator asks the firm to reveal the required information in order to set the adequate price, since the firm would attempt to take advantage of the information. Therefore, the regulator should take this problem into account.

III. The Base Game

Consider a contractual relationship between a principal and an agent, who is contracted to carry out a task. The relationship allows a certain *result* to be obtained, whose monetary value will be referred to as x . For the sake of exposition, the set of possible results X is assumed to be finite, $X = \{x_1, \dots, x_n\}$. The final result depends on the *effort* that the agent devotes to the task, which will be denoted by e , and the value of a random variable for which both participants have the same prior distribution. The probability of result x_i conditional on effort e can be written as:

$$\text{Prob}[x = x_i | e] = p_i(e), \text{ for } i \in \{1, 2, \dots, n\},$$

with $\sum_{i=1}^n p_i(e) = 1$. Let us assume that $p_i(e) > 0$ for all e, i , which implies that no result can be ruled out for any given effort level.

The Base Game is the reference situation, where principal and agent have the same information (even the one concerning the random component that affects the result). Since uncertainty exists, participants react to risk. Risk preferences are expressed by the shape of their *utility functions* (of the von Neumann-Morgenstern type). The principal, who owns the result and must pay the agent, has preferences represented by the utility function

$$B(x - w),$$

where w represents the payoff made to the agent. $B(\cdot)$ is assumed to be increasing and concave: $B' > 0$, $B'' \leq 0$ (where the primes represent, respectively, the first and second derivatives). The concavity of the function $B(\cdot)$ indicates that the principal is either risk-neutral or risk averse.

The agent receives a monetary pay-off for his participation in the relationship, and he supplies an effort which implies some cost to him. For the sake of simplicity we represent his utility function as:

$$U(w, e) = u(w) - v(e),$$

additively separable in the components w and e . This assumption implies that the agent's risk-aversion does not vary with the effort he supplies (many results can be generalized for more general utility functions). The utility derived from the wage, $u(w)$, is increasing and concave:

$$u'(w) > 0, \quad u''(w) \leq 0,$$

thus the agent may be either risk-neutral, $u''(w) = 0$, or risk-averse $u''(w) < 0$. In addition, greater effort means greater disutility. We also assume that the marginal disutility of effort is not decreasing: $v'(e) > 0, \quad v''(e) \geq 0$.

A *contract* can only be based in verifiable information. In the base game, it can depend on the characteristics of the principal and the agent and includes both the effort e that the principal demands from the agent, and the wages $\{w(x_i)\}_{i=1, \dots, n}$.

If the agent rejects the contract, he will have to fall back on the outside opportunities that the market offers him. These other opportunities, that by comparison determine the limit for participation in the contract, are summarized in the agent's *reservation utility*, denoted by \underline{U} . So the agent will accept the contract as long as he earns an expected utility equal or higher than his reservation utility. Since, the principal's problem is to design a contract that the agent will accept, (by backward induction) the optimal contract must satisfy the *participation constraint* and it is the solution to the following maximization problem:

$$\begin{aligned} & \text{Max}_{[e, \{w(x_i)\}_{i=1, \dots, n}]} \sum_{i=1}^n p_i(e) B(x_i - w(x_i)) \\ & \text{s.t.} \quad \sum_{i=1}^n p_i(e) u(w(x_i)) - v(e) \geq \underline{U}. \end{aligned}$$

The above problem corresponds to a *Pareto Optimum* in the usual sense of the term. The solution to this problem is conditional on the value of the parameter \underline{U} , so that even those cases where the agent can keep a large share of the surplus are taken into account.

The principal's program is well behaved with respect to payoffs given the assumptions on $u(w)$. Hence the Kuhn-Tucker conditions will be both necessary and sufficient for the global solution of the problem. However, we cannot ascertain the concavity (or quasi-concavity) of the functions with respect to effort given the assumptions on $v(e)$, because these functions also

depend on all the $p_i(e)$. Hence it is more difficult to obtain global conclusions with respect to this variable.

Let us denote by e° the efficient effort level. From the first-order Kuhn-Tucker conditions with respect to the wages in the different contingencies, we can analyze the associated pay-offs $\{w^\circ(x_i)_{i=1, \dots, n}\}$. We obtain the following condition:

$$\lambda^\circ = \frac{B'(x_i - w^\circ(x_i))}{u'(w^\circ(x_i))}, \text{ for all } i \in \{1, 2, \dots, n\},$$

where λ° is the multiplier associated with the participation constraint. When the agent's utility is additively separable, the participation constraint binds (λ° is positive). The previous condition equates marginal rates of substitution and indicates that the optimal distribution of risk requires that the ratio of marginal utilities of the principal and the agent to be constant irrespective of the final result.

If the principal is risk-neutral ($B''(\cdot) = 0$), then the optimal contract has to be such that $u'(w^\circ(x_i)) = \text{constant}$ for all i . In addition, if the agent is risk-averse ($u''(\cdot) < 0$), he receives the same wage, say w° , in all contingencies. This wage only depends on the effort demanded and is determined by the participation constraint. If the agent is risk-neutral ($u''(\cdot) = 0$) and the principal is risk-averse ($B''(\cdot) < 0$), then we are in the opposite situation. In this scenario, the optimal contract requires the principal's profit to be independent of the result. Consequently, the agent bears all the risk, insuring the principal against variations in the result. When both the principal and the agent are risk-averse, each of them needs to accept a part of the variability of the result. The precise amount of risk that each of them supports depends on their degrees of risk-aversion. Using the Arrow-Pratt measure of absolute risk-aversion $r_p = -B''/B'$ and $r_a = -u''/u'$ for the principal and the agent respectively, we can show that:

$$\frac{dw^\circ}{dx_i} = \frac{r_p}{r_p + r_a},$$

which indicates how the agent's wage changes given an increase in the result x_i . Since $r_p/(r_p + r_a) \in (0,1)$, when both participants are risk-averse, the agent only receives a part of the increased result via a wage increase. The more risk-averse is the agent, that is to say, the greater is r_a , the less the result influences his wage. On the other hand, as the risk-aversion of the principal increases, greater r_p , changes in the result correspond to more important changes in the wage.

IV. Moral Hazard

IV.1. Basic Moral Hazard Model

Here we concentrate on *moral hazard*, which is the case in which the informational asymmetry relates to the agent's behaviour during the relationship. We analyze the optimal contract when the agent's effort is not verifiable. This implies that effort cannot be contracted upon, because in case of the breach of contract, no court of law could know if the contract had really been breached or not. There are many examples of this type of situation. A traditional example is accident insurance, where it is very difficult for the insurance company to observe how careful a client has been to avoid accidents.

The principal will state a contract based on any signals that reveal information on the agent's effort. We will assume that only the result of the effort is verifiable at the end of the period and, consequently, it will be included in the contract. However, if possible, the contract should be contingent on many other things. Any information related to the state of nature is useful, since it allows better estimations of the agent's effort thus reducing the risk inherent in the relationship. This is known as the *sufficient statistic result*, and it is perhaps the most important conclusion in the moral hazard literature (Holmström, 1979). The empirical content of the sufficient statistic argument is that a contract should exploit all available information in order to filter out risk optimally.

The timing of a moral hazard game is the following. In the first place, the principal decides what contract to offer the agent. Then the agent decides whether or not to accept the relationship, according to the terms of the contract. Finally, if the contract has been accepted, the agent chooses the effort level that he most desires, given the contract that he has signed. This is a free decision by the agent since effort is not a contracted variable. Hence, the principal must bear this in mind when she designs the contract that defines the relationship.

To better understand the nature of the problem faced by the principal, consider the case of a risk-neutral principal and a risk-averse agent, which implies that, under the symmetric information, the optimal contract is to completely insure the agent. However, if the principal proposes this contract when the agent's effort is not a contracted variable, once he has signed the contract the agent will exert the effort level that is most beneficial for him. Since the agent's wage does not depend on his effort, he will use the lowest possible effort.

The idea underlying an incentive contract is that the principal can make the agent interested in the consequences of his behaviour by making his pay-off dependent on the result obtained. Note

that this has to be done at the cost of distorting the optimal risk sharing among both participants. The trade-off between *efficiency*, in the sense of the optimal distribution of risk, and *incentives* determines the optimal contract.

Formally, since the game has to be solved by backwards induction, the *optimal contract* under moral hazard is the solution to the maximization problem:

$$\begin{aligned} & \text{Max}_{[e, \{w(x_i)\}_{i=1, \dots, n}]} \sum_{i=1}^n p_i(e) B(x_i - w(x_i)) \\ \text{s.t.} \quad & \sum_{i=1}^n p_i(e) u(w(x_i)) - v(e) \geq \underline{U} \\ & e \in \text{Arg Max}_{\hat{e}} \left\{ \sum_{i=1}^n p_i(\hat{e}) u(w(x_i)) - v(\hat{e}) \right\} \end{aligned}$$

The second restriction is the *incentive compatibility constraint* and the first restriction is the participation constraint. The incentive compatibility constraint, and not the principal as under symmetric information, determines the effort of the agent.

The first difficulty in solving this program is related to the fact that the incentive compatibility constraint is a maximization problem. The second difficulty is that the expected utility may fail to be concave in effort. Hence, to use the first order condition of the incentive compatibility constraint may be incorrect. In spite of this, there are several ways to proceed when facing to this problem. (a) Grossman and Hart (1983) propose to solve it in steps, identifying first the optimal payment mechanism for any effort and then, if possible, the optimal effort. This can be done since the problem is concave in payoffs. (b) The other possibility is to consider situations where the agent's maximization problem is well defined. One possible scenario is when the set of possible efforts is finite, in which case the incentive compatibility constraint takes the form of a finite set of inequalities. Another scenario is to write the incentive compatibility as the first-order condition of the maximization problem, and introduce assumptions that allow doing it. The last solution is known as the *first-order approach*.

Let us assume that the first-order approach is adequate, and substitute the incentive compatibility constraint in the previous program by

$$\sum_{i=1}^n p'_i(\hat{e}) u(w(x_i)) - v'(\hat{e}) = 0.$$

Solving the principal's program with respect to the payoff scheme, and denoting by λ (resp., μ) the Lagrangean multiplier of the participation constraint (resp., the incentive compatibility constraint), we obtain that for all i :

$$\frac{1}{u'(w(x_i))} = \lambda + \mu \frac{p'_i(e)}{p_i(e)} .$$

This condition shows that the wage should not depend at all on the value that the principal places on the result. It depends on the results as a measure of how informative they are as to effort, in order to serve as an incentive for the agent. The wage will be increasing in the result as long as the result is increasing in effort. Hence, it is optimal that the wage will be increasing in the result only in particular cases. The necessary condition for a wage to be increasing with results is $\frac{p'_i(e)}{p_i(e)}$ to be decreasing in i . In statistics, this is called the *monotonous likelihood quotient property*. It is a strong condition; for example, first-order stochastic dominance does not guarantee the monotonous likelihood property.

IV.2. Extensions of Moral Hazard Models

The basic moral hazard setup, with a principal hiring and an agent performing effort, has been extended in several directions to take into account more complex relationships.

Repeated Moral Hazard

Certain relationships in which a moral hazard problem occurs do not take place only once, but they are repeated over time (for example, work relationships, insurance, etc.). The duration aspect (the repetition) of the relationship gives rise to new elements that are absent in static models.

Radner (1981) and Rubinstein and Yaari (1983) consider infinitely repeated relationships and show that frequent repetition of the relationship allows us to converge towards the efficient solution. Incentives are not determined by the payoff scheme contingent on the result of each period, but rather on average effort, and the information available is very precise when the number of periods is large. A sufficiently threatening punishment, applied when the principal believes that the agent on average does not fulfil his task, may be sufficient to dissuade him from shirking.

When the relationship is repeated a finite number of times, the analysis of the optimal contract concentrates on different issues relating *long-term* agreements and *short-term* contracts. Note

that in a repeated set up, the agent's wage and the agent's consumption in a period need not be equal. Lambert (1983), Rogerson (1985), and Chiappori and Macho-Stadler (1990) show that long-term contracts have memory (i.e., the pay-offs in any single period will depend on the results of all previous periods) since they internalize agent's consumption over time, which depends on the sequence of payments received (as a function of the past contingencies). Malcomson and Spinnewyn (1988), Fudenberg, Holmström, and Milgrom (1990), and Rey and Salanié (1990) study when the optimal long-term contract can be implemented through the sequence of optimal short-term contracts. Chiappori, Macho-Stadler, Rey, and Salanié (1994) show that, in order for the sequence of optimal short-term contracts to admit the same solution as the long-term contract, two conditions must be met. First, the optimal sequence of single-period contracts should have memory. That is why, when the reservation utility is invariant (is not history dependent), the optimal sequence of short-term contracts will not replicate the long-term optimum unless there exist means of smoothing consumption, that is, the agent has access to credit markets. Second, the long-term contract must be renegotiation-proof. A contract is said to be renegotiation-proof if at the beginning of any intermediate period, no new contract or renegotiation that would be preferred by all participants is possible. When the long term contract is not renegotiation-proof (i.e., if it is not possible for participants to change the clauses of the contract at a certain moment of time even if they agree), it cannot coincide with the sequence of short term contracts.

One Principal and Several Agents

When a principal contracts with more than one agent, the stage where agents exert their effort, which is translated into the incentive compatibility, depends on the *game among the agents*. If the agents behave as a coordinated and cooperating group, then the problem is similar to the previous one where the principal hires a team. A more interesting case appears when agents play a non-cooperative game and their strategies form a Nash equilibrium.

Holmström (1979) and Mookherjee (1984), in models where there is *personalized information* about the output of each agent, show that the principal is interested in paying each agent according to his own production and that of the other agents if these other results can inform on the actions of the agent at hand. Only if the results of the other agents do not add information or, in other words, if an agent's result is a *sufficient statistic* for his effort, then he will be paid according to his own result.

When the only verifiable outcome is the final result of teamwork (*joint production models*), the optimal contract can only depend on this information and the conclusions are similar to those

obtained in models with only one agent. Alchian and Demsetz (1972) and Holmström (1982) show that joint production cannot lead to efficiency when all the income is distributed amongst the agents, i.e., if the budget constraint always binds. Another player should be contracted to control the productive agents and act as the residual claimant of the relationship.

Tirole (1986) and Laffont (1990) have studied the effect of coalitions among the agents in an organization on their payment scheme. If collusion is bad for the organization, it adds another dimension of moral hazard (the colluding behaviour). The principal may be obliged to apply rules that are collusion-proof, which implies more constraints and simpler contracts (more bureaucratic). When coordination can improve the input of a group of agents, the optimal contract has to find payment methods that strengthen group work (see Itoh, 1990, and Macho-Stadler and Pérez-Castrillo, 1993).

Another principal's decision when she hires several agents is the organization with which she will relate. This includes such fundamental decisions as how many agents to contract and how should they be structured. These issues have been studied by Demski and Sappington (1986), Melumad and Reichelstein (1987), and Macho-Stadler and Pérez-Castrillo (1998).

Holmström and Milgrom (1991) analyze a situation in which the agent carries out *several tasks*, each one of which gives rise to a different result. They study the optimal contract when tasks are complementary (in the sense that exerting effort in one reduces the costs of the other) or substitutes. Their model allows to build a theory of job design and to explain the relationship among responsibility and authority.

Several Principals and One Agent

When one agent works for (or signs his contracts with) several principals simultaneously (*common agency* situation), in general, the principals are better off if they cooperate. When the principals are not able to achieve the coordination and commitment necessary to act as a single individual and they do not value the results in the same way, they each demand different efforts or actions from the agent. Bernheim and Whinston (1986) show that the effort that principals obtain when they do not cooperate is less than the effort that would maximize their collective profits. However, the final contract that is offered to the agent minimizes the cost of getting the agent to choose the contractual effort.

V. Adverse Selection

V.1. Basic Adverse Selection Model

Adverse selection is the term used to refer to problems of asymmetric information that appear before the contract is signed. The classic example of Akerlof (1970) illustrates very well the issue: the buyer of a used car has much less information about the state of the vehicle than the seller. Similarly, the buyer of a product knows how much he appreciates the quality, while the seller only has statistical information about a typical buyer's taste (Mussa and Rosen, 1978); or the regulated firm has more accurate information about the marginal cost of production than the regulator.

A convenient way to model adverse selection problems is to consider that the agent can be of different *types*, and that the agent knows his type before any contract is signed while the principal does not know it. In the previous examples, the agent's type is the particular quality of the used car, the level of appreciation of quality, or the firm's marginal cost. How can the principal deal with this informational problem? Instead of offering just one contract for every (or several) types of agents, she can propose several contracts so that each type of agent chooses the one that is best for him. A useful result in this literature is the *revelation principle* (Gibbard, 1973, Green and Laffont, 1977, Myerson, 1979) that states that any mechanism that the principal can design is equivalent to a *direct revelation mechanism* by which the agent is asked to reveal his type and a contract is offered according to his declaration. That is, a direct revelation mechanism offers a *menu of contracts* to the agent (one contract for each possible type), and the agent can choose any of the proposed contracts. Clearly, the mechanism must give the agent the right incentives to choose the appropriate contract, that is, it must be a *self-selection mechanism*. Menus of contracts are not unusual. For instance, insurance companies offer several possible insurance contracts between which clients may freely choose their most preferred. For example, car insurance contracts can be with or without deductible clauses. The second goes to more risk averse or more frequent drivers while deductibles attract less risk averse or less frequent drivers.

Therefore, the timing of an adverse selection game is the following. In the first place, the agent's characteristics (his "type") are realized, and only the agent learns them. Then, the principal decides the menu of contracts to offer to the agent. Having received the proposal, the agent decides which one of the contracts (if any) to accept. Finally, if one contract has been accepted, the agent chooses the predetermined effort and receives the corresponding payment.

A simple model of adverse selection is the following. A risk-neutral principal contracts an agent (who could be risk-neutral or risk-averse) to carry out some verifiable effort on her behalf. Effort e provides an expected payment to the principal of $\Pi(e)$, with $\Pi'(e) > 0$ y $\Pi''(e) < 0$. The agent could be either of *two types* that differ with respect to the disutility of effort, which is $v(e)$ for type G (good), and $k v(e)$, with $k > 1$ for type B (bad). Hence, the agent's utility function is either $U^G(w, e) = u(w) - v(e)$ or $U^B(w, e) = u(w) - k v(e)$. The principal considers that the probability for an agent to be type- G is q , where $0 < q < 1$.

The principal designs a menu of contracts $\{(e^G, w^G), (e^B, w^B)\}$, where (e^G, w^G) is directed towards the most efficient type of agent, while (e^B, w^B) is intended for the least efficient type. For the menu of contracts to be a sensible proposal, the agent must be better off by truthfully revealing his type than by deceiving the principal. The principal's problem, is therefore to maximize her expected profits subject to the restrictions that (a) after considering the contracts offered, the agent decides to sign with the principal (*participation constraints*), and (b) each agent chooses the contract designed for his particular type (*incentive compatibility constraints*):

$$\begin{aligned} & \text{Max}_{[(e^G, w^G), (e^B, w^B)]} && q [\Pi(e^G) - w^G] + (1 - q) [\Pi(e^B) - w^B] \\ \text{s.t.} &&& u(w^G) - v(e^G) \geq \underline{U} \\ &&& u(w^B) - k v(e^B) \geq \underline{U} \\ &&& u(w^G) - v(e^G) \geq u(w^B) - v(e^B) \\ &&& u(w^B) - k v(e^B) \geq u(w^G) - k v(e^G) \end{aligned}$$

The main characteristics of the optimal contract menu $\{(e^G, w^G), (e^B, w^B)\}$ are the following:

- (i) The contract offered to the good agent (e^G, w^G) is efficient ('*non distortion at the top*'). The optimal salary w^G however is higher than under symmetric information: this type of agent receives an *informational rent*. That is, the most efficient agent profits from his private information and in order to reveal this information he has to receive a utility greater than his reservation level.
- (ii) The participation condition binds for the agent when he has the highest costs (he just receives his reservation utility). Moreover, a distortion is introduced into the efficiency condition for this type of agent. By distorting, the principal loses efficiency with respect to type- B agents, but she pays less informational rent to the G -types.

V.2 Principals Competing for Agents in Adverse Selection Frameworks

Starting with the pioneer work by Rothschild and Stiglitz (1976) on insurance markets, there have been many studies on markets with adverse selection problems where there is competition among principals to attract agents. We move from a model where one principal maximizes her profits subject to the above constraints, to a game theory environment where each principal has to take into account the actions by others when deciding which contract to offer. In this case, the adverse selection problem may be so severe that we may find ourselves in situations in which no equilibrium exists.

To highlight the main results in this type of models, consider a simple case in which there are two possible risk-averse agent types: good (G) and bad (B) with G being more productive than B . In particular, we assume that G is more careful than B , in the sense that he commits fewer errors. When the agent exerts effort, the result could be either a success (S) or a failure (F). The probability that it is successful is p^G when the agent is type- G and p^B when he is type B , where $p^G > p^B$. The principal values a successful result more than a failure. The result is observable, so that the principal can pay the agent according to the result she desires.

There are several risk-neutral principals. Therefore, we look for the set of equilibrium contracts in the game played by principals competing to attract agents. Equilibrium contracts must satisfy that there does not exist a principal who can offer a different contract that would be preferred by all or some of the agents and that gives that principal greater expected profits. This is why, if information was symmetric, the equilibrium contracts would be characterized by the following properties: (i) principals' expected profits are zero; and (ii) each contract must be efficient. Hence the agent receives a fixed contract insuring him against random events. In particular, the equilibrium salary that the agent receives under symmetric information is higher when he is of type G than when he is of type B .

When the principals cannot observe the type of the agent, the previous contracts cannot be longer an equilibrium: all the agents would claim to be a good type. An equilibrium contract pair $\{C^G, C^B\}$ must satisfy the condition that no principal can add a new contract that would give positive expected profits to the agents that prefer this new contract to C^G and C^B . If the equilibrium contracts for the two agent types turn out to be the same, that is, there is only one contract that is accepted by both agent types, then the equilibrium is said to be *pooling*. On the other hand, when there is a different equilibrium contract for each agent type, then we have a *separating* equilibrium. In fact, pooling equilibria never exist, since pooling contracts always give room for a principal to propose a profitable contract that would only be accepted by the G -

types (the best agents). If an equilibrium does exist, it must be such that each type of agent is offered a different contract.

If the probability that the agent is “good” is large enough, then a separating equilibrium does not exist either. That is, an adverse selection problem in a market may provoke the absence of any equilibrium in that market. Nevertheless, a separating equilibrium does exist. In the separating equilibrium, the results are similar to the ones under moral hazard in spite of the differences in the type of asymmetric information and in the method of solving. That is, contingent pay-offs are offered to the best agent to allow the principal to separate them from the less efficient ones. In this equilibrium, the least efficient agents obtain the same expected utility (and even sign the same contract) as under symmetric information, while the best agents lose expected utility due to the asymmetric information.

V.3. Extensions of Adverse Selection Models

Repeated Adverse Selection

In this extension, we consider whether the repetition of the relationship during several periods helps the principal and how it influences the form of the optimal contract. Note first that if the agent’s private information is different in each period and the information is not correlated among periods, then any current information revealed does not affect the future and hence the repeated problem is equivalent to simple repetition of the initial relationship. The optimal intertemporal contract will be the sequence of optimal single-period contracts.

Consider the opposite situation where the agent's type is constant over time. If the agent decides to reveal his type truthfully in the first period, then the principal is able to design efficient contracts that extract all surpluses from the agent. Hence, the agent will have very strong incentives to misrepresent his information in the early periods of the relationships. In fact, Baron and Besanko (1984) show that if the principal can commit herself with a contract that covers all the periods, then the optimal contract is the repetition of the optimal static contract. This implies that the contract is not *sequentially rational*, and it is also non *robust to renegotiation*: once the first period is finished, the principal “knows” the agent’s type and a better contract for both parties is possible.

It is often the case that the principal cannot commit not to renegotiate a long-term contract. Laffont and Tirole (1988) show that, in this case, it may be impossible to propose perfect revelation (separating) contracts in the first periods. This is known as the *ratchet effect*. Also,

Freixas, Guesnerie, and Tirole (1985), and Laffont and Tirole (1987) have proven that, even when separating contracts exist, they may be so costly that they are often non optimal and we should expect that information be revealed progressively over time. Baron and Besanko (1987) and Laffont and Tirole (1990) also introduce frameworks in which it is possible to propose perfect revelation contracts but they are not optimal.

Relationships with Several Agents: Auctions

One particularly interesting case of relationship among one principal and several agents is that of a seller who intends to sell one or several items to several interested buyers, where buyers have private information about their valuation for the item(s). A very popular selling mechanism in such a case is an *auction*. As Klemperer (2004) would put it, auction theory is one of economics' success stories in both practical and theoretical terms. Art galleries generally use *English auctions*: the agents bid "upwards"; while fish markets are generally examples of *Dutch auctions*: the seller reduces the price of the good until someone stops the auction by buying. Public contracts are generally awarded through (*first price* or *second price*) *sealed-bid auctions* where buyers introduce their bid in a closed envelope, the good is sold to the highest bidder and the prize is either the winner's own bid or the second highest bid.

Vickrey (1961, 1962) was the first to establish the key result in auction theory, the *Revenue Equivalence Theorem* which, subject to some reasonable conditions, says that the seller can expect equal profits on average from all the above (and many other) types of auctions, and that buyers are also indifferent among them all. Auctions are efficient, since the buyer who ends up with the object is the one with the highest valuation. Hence, the existence of private information does not generate any distortions with respect to who ends up getting the good, but the revenue of the seller is less than under symmetric information.

Myerson (1981) solves the general mechanism design problem of a seller who wants to maximize her expected revenue, when the bidders have independent types and all agents are risk-neutral. In general, the optimal auction is more complex than the traditional English (second price) or Dutch (first price) auction. His work has been extended by many other authors. When the buyers' types are affiliated (i.e., they are not negatively correlated in any subset of their domain), Milgrom and Weber (1982) show that the revenue equivalence theorem breaks down. In fact, in this situation, McAfee, McMillan, and Reny (1989) show that the seller may extract the entire surplus from the bidders as if there was no asymmetric information. Starting with Maskin and Riley (1989), several authors have also analyzed auctions of multiple units.

Finally, Clarke (1971) and Groves (1973) initiated another group of models in which the principal contracts with several agents simultaneously, but does not attempt to maximize her own profits. This is the case of the provision of a public good through a mechanism provided by a benevolent regulator.

Relationships with Several Agents: Other Models and Organizational Design

Adverse selection models have attempted to analyze the optimal task assignment, the advantages of delegation, or the optimal structure of contractual relationships, when the principal contracts with several agents. Riordan and Sappington (1987) analyze a situation where two tasks have to be fulfilled and show that if the person in charge of each task has private information about the costs associated with the task, then the assignment of tasks within the organization is an important decision. For example, when the costs are positively correlated, then the principal will prefer to take charge of one of the phases herself while she will prefer to delegate the task when the costs are negatively correlated.

In a very general framework, Myerson (1982) shows a powerful result: in adverse selection situations, centralization cannot be worse than decentralization, since it is always possible to replicate a decentralized contract with a centralized one. This result is really a generalization of the revelation principle. Baron and Besanko (1992) and Melumad, Mookherjee, and Reichelstein (1995) show that, if the principal can offer complex contracts in a decentralized organization, then a decentralized structure can replicate a centralized organization. When there are problems of communication between principal and agents, the equivalence result does not hold: Melumad and Reichelstein (1987) show that *delegation* of authority can be preferable if communication between the principal and the agents is difficult. Still concerning the optimal design of the organization, Dana (1993) analyzes the optimal hierarchical structure in industries with several productive phases, when firms have private information related to their costs. They show that structures that concentrate all tasks to a single agent are superior since, the incentives to dishonestly reveal the costs of each of the phases are weaker. Da-Rocha-Alvarez and De-Frutos (1999) argue that the absolute advantage of the centralized hierarchy is not maintained if the differences in costs between the different phases are sufficiently important.

Several Principals

Stole (1990) and Martimort (1996) point out the difficulty of extending the revelation principle to situations where an agent with private information is contracted by several principals who act

separately. Given that not only one contract (or menu of contracts) is offered to the agent, but several contracts coming from different principals, it is not longer necessarily true that the best a principal can do is to offer a “truth-telling mechanism”.

Consider a situation with two principals that are hiring a single agent. If we accept that agent's messages are restricted to the set of possible types that the agent may have, we can obtain some conclusions. If the activities or efforts that the agent carries out for the two principals are substitutes (for instance, a firm producing for two different customers), then the usual result on the distortion of the decision holds: the most efficient type of agent supplies the efficient level of effort while the effort demanded from the least efficient type is distorted. However, due to the lack of cooperation between principals, the distortion induced in the effort demanded from the less efficient type of agent is lower than the one maximizing the principals' aggregate profits. On the other hand, if the activities that the agent carries out for the principals are complementary (for example, the firm produces a final good that requires two complementary intermediate goods in the production process), then the comparison of the results under cooperation and under no cooperation between the principals reveals that: if a principal reduces the effort demanded from the agent, in the second case, this would imply that it is also profitable for the other principal to do the same. Therefore, the distortion in decisions is greater to that produced in the case in which principals cooperate.

Models of Moral Hazard and Adverse Selection

The analysis of principal-agent models where there are simultaneously elements of moral hazard and adverse selection is a complex extension of classic agency theory. Conclusions can be obtained only in particular scenarios. One common class of models considers situations where the principal cannot distinguish the part corresponding to effort from the part corresponding to the agent's efficiency characteristic because both variables determine the production level. Picard (1987) and Guesnerie, Picard, and Rey (1989) propose a model with risk-neutral participants and show that, if the effort demanded from the different agents is not decreasing in the characteristic (if a higher value of this parameter implies greater efficiency), then the optimal contract is a menu of distortionary deductibles designed to separate the agents. The menu of contracts includes one where the principal sells the firm to the agent (aiming at the most efficient type), and another contract where she sells only a part of the production at a lower prize (aiming at the least efficient type). However, there are also cases where fines are needed to induce the agents to honestly reveal their characteristic.

In fact, the principle message of the previous literature is that the optimal solution for problems that mix adverse selection and moral hazard does not imply efficiency losses with respect to the pure adverse selection solution when the agent's effort is observable. However, in other frameworks (see Laffont and Tirole, 1986), a true problem of asymmetric information appears only when both problems are mixed when, and efficiency losses are evident. Therefore, the same solution as when only the agent's characteristic is private information cannot be achieved.

VI. Future directions

VI.1 Empirical Studies of Principal-Agent Models

The growing interest on empirical issues related to asymmetric information started in the mid nineties (see the survey by Chiappori and Salanie, 2003). A very large part of the literature is devoted to test the predictions of the canonical models of moral hazard and adverse selection, where there is only one dimension in which information is asymmetric. A great deal of effort is devoted to try to ascertain whether it is moral hazard, or adverse selection, or both prevalent in the market. This is a difficult task because both adverse selection and moral hazard generate the same predictions in a cross section. For instance, a positive correlation between insurance coverage and probability of accident can be due to either the intrinsically riskier drivers selecting into contracts with better coverage (as the Rothschild and Stiglitz, 1976 model of adverse selection will predict) or to drivers with better coverage exerting less effort to drive carefully (as the canonical moral hazard model will predict). Chiappori, Jullien, Salanié and Salanié (forthcoming) have shown that the positive correlation between coverage and risk holds more generally than in the canonical models as long as the competitive assumption is maintained.

Future empirical approaches are likely to incorporate market power (as in Chiappori *et al*, forthcoming), multiple dimensions of asymmetric information (as in Finkelstein and McGarry, 2006), as well as different measures of asymmetric information (as in Vera-Hernandez, 2003). These advances will be partly possible thanks to richer surveys which collect subjective information regarding agents' attributes usually unobserved by principals or agent's subjective probability distributions. The wider availability of panel data will mean that it will become easier to disentangle moral hazard from adverse selection (as in Abbring, Chiappori, Heckman and Pinquet, 2003). Much is to be learnt by using field experiments that allow randomly varying

contract characteristics offered to individuals and hence disentangling moral hazard from adverse selection (as in Karlan and Zinman, 2006).

VI.2 Contracts and Social Preferences

Although principal-agent theory has proved fundamental in expanding our understanding of contract situations, real-life contracts frequently do not exactly match its predictions. Many contracts are linear and simpler, incentives are often stronger and wage gaps more compressed than expected. One possible explanation is that the theory has mainly focused on economic agents exclusively motivated by their own monetary incentives. However, this assumption leaves aside issues such as social ties, team spirit or work morale, which the human resources literature highlights. A recent strand of economic literature, known as “behavioural contract theory”, has recently tried to incorporate social aspects into the economic modelling of contracts.

Such theory has been motivated by two types of empirical support. On the one hand, extensive interview studies with firm managers and employees (Bewley, 1999) has shown not only that agents care about social comparisons such as internal pay equity or effort diversity, but that their incentives to work hard are affected by them and that principals are aware of it and design their contracts accordingly. On the other hand, one of the most influential contributions of the experimental literature has been to show that, assuming that economic agents are not completely selfish (but exhibit some form of social preferences), helps organizing many laboratory data. Experiments replicating labour markets (starting with Fehr, Kirchsteiger and Riedl, 1993) confirm Akerlof’s (1982) insight that contracts may be understood as a form of gift exchange in which principals may offer a “generous” wage and agents may respond with more than the minimum effort required.

Incorporating social and psychological aspects in a systematic manner into agents’ motivations has given rise to several forms of utility functions reflecting inequality aversion (Fehr and Schmidt, 1999, Bolton and Ockenfels, 2000, Desiraju and Sappington, 2007), fairness (Rabin, 1993) and reciprocity (Dufwenberg and Kirchsteiger, 2004). More recently, such utility functions have been included into standard contract theory models and have helped in shortening the gap between theory predictions and real-life contracts. In particular, issues such as employees’ feelings of envy or guilt towards their bosses (Itoh, 2004), utility comparisons among employees (Grund and Sliwka, 2005, Rey-Biel, 2002) or peer-pressure motivating effort decisions (Lazear, 1995, Huck and Rey-Biel, 2006) have proved important in widening the scope of issues principal-agent theory can help to understand.

VI.3 Principal-Agent Markets

The literature has been treating each principal-agent relation as an isolated entity. Thus, it normally takes a given relationship between a principal and an agent (or among several principals and/or several agents), and analyzes the optimal contract. In particular, the principal assumes all the bargaining power as she has the right to offer the contract she likes the most, and agent's payoff is determined by his exogenously given reservation utility. However, in markets there is typically not a single partnership but there are several. It is then interesting to consider the simultaneous determination of both the identity of the pairs that meet (i.e., the *matching* between principals and agents) and the contracts these partnerships sign. The payoffs to each principal and agent will then depend on the other principal-agent relationships being formed in the market. This analysis requires a general equilibrium-like model.

Game theory provides a very useful tool to deal with the study of markets where heterogeneous players from one side can form partnerships with heterogeneous players from the other side: the two-sided matching models. Examples of classic situations studied in two-sided matching models (see Shapley and Shubik, 1972, and Roth and Sotomayor, 1990) are the marriage market, the college admissions model, or the assignment market (where buyers and sellers transact). Several papers extend this game theory models to situations where each partnership involves contracts and show that the simultaneous consideration of matching and contracts has important implications. Dam and Pérez-Castrillo (2006) show that, in an economy where landowners contract with tenants, a government willing to improve the situation of the tenants can be interested in creating wealth asymmetries among them. Otherwise, the landowners would appropriate all the incremental money that the government is willing to provide to the agents. Serfes (forthcoming) shows that higher-risk projects do not necessarily lead to lower incentives, which is the prediction in the standard principal-agent theory, and Alonso-Paulí and Pérez-Castrillo (2007) apply the theory to markets where contracts (between shareholders and managers) can include Codes of Best Practice. On the empirical side, Akerberg and Botticini (2002) find strong evidence for endogenous matching between landlords and tenants and that risk sharing is an important determinant of contract choice.

Future research will extend the general equilibrium analysis of principal-agent contracts to other markets. In addition, the literature has only studied one-to-one matching models. This should be extended to situations where each principal can hire several agents, or where each agent deals with several principals. The interplay between (external) market competition and

(internal) collaboration between agents or principals can provide useful insights about the characteristics of optimal contracts in complex environments.

VII.- Bibliography

Primary Literature

- Abbring J, Chiappori PA, Heckman JJ, Pinquet J (2003) Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish? *J European Econ Ass* 1: 512-521.
- Akerberg DA, Botticini M (2002) Endogenous Matching and the Empirical Determinants of Contract Form. *J Political Econ* 110: 564-592.
- Akerlof G (1970) The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism. *Quarterly J Econ* 89: 488-500.
- Akerlof G (1982) Labor Contracts as a Partial Gift exchange. *Quarterly J Econ* 97: 543-569.
- Alchian A, Demsetz H (1972) Production, Information Costs, and Economic Organization. *Am Econ Rev* 62: 777-795.
- Alonso-Pauli E, Pérez-Castrillo D (2007) Codes of Best Practice in Competitive Markets. Mimeo.
- Arrow K (1985). The Economics of Agency. In J. Pratt and R. Zeckhauser, eds., *Principals and agents: The Structure of business*. Boston: Harvard University Press.
- Baron D, Besanko D (1984) Regulation and Information in a Continuing Relationship. *Information Econ Policy* 1: 267-302.
- Baron D, Besanko D (1987) Commitment and Fairness in a Dynamic Regulatory Relationship. *Rev Econ Studies* 54: 413-436.
- Baron D, Myerson R (1982) Regulating a Monopoly with Unknown Costs. *Econometrica* 50: 911-930.
- Baron D, Besanko D (1992) Information, Control and Organizational Structure. *J Econ Management Strategy* 1 (2): 237-275.
- Bernheim BD, Whinston MD (1986) Common Agency, *Econometrica* 54: 923-942.
- Bewley T (1999) *Why Rewards Don't Fall During a Recession*. Harvard University Press.
- Bolton G, Ockenfels A (2000) ERC: A Theory of Equity, Reciprocity and Competition. *Am Econ Rev* 90: 166-193.

- Chiappori PA, Macho-Stadler I (1990) Contrats de Travail Répétés: Le Rôle de la Mémoire. *Annales Economie Statistique* 17: 4770.
- Chiappori PA, Macho-Stadler I, Rey P, Salanié B (1994) Repeated Moral Hazard: The Role of Memory, Commitment and Access to Credit Markets. *European Econ Rev* 38: 1527-1553.
- Chiappori PA, Jullien B, Salanié B, Salanié F (forth) Asymmetric Information In Insurance: General Testable Implications. *Rand J Econ*.
- Chiappori PA, Salanié B (2003) Testing Contract Theory: a Survey of Some Recent Work. In *Advances in Economics and Econometrics 1*, Dewatripont, Hansen and Turnovsky eds, Cambridge University Press.
- Clarke E. (1971) Multipart Pricing of Public Goods. *Public Choice* 11:17-33.
- Da-Rocha-Alvarez JM, De-Frutos MA (1999) A Note on the Optimal Structure of Production. *J Econ Theory*: 234-246.
- Dam K, Pérez-Castrillo D (2006) The Principal-Agent Matching Market. *Frontiers Econ Theory, Berkeley Electronics* 2 (1).
- Dana Jr. JD (1993) The Organization and Scope of Agents: Regulating Multiproduct Industries. *J Econ Theory* 59: 288-310.
- Demski JS, Sappington D (1986) Line-item Reporting, Factor Acquisition and Subcontracting. *J Accounting Research* 24: 250-269.
- Desiraju R, Sappington D (2007) Equity and Adverse Selection. *J Econ Management Strategy* 16: 285-318.
- Dufwenberg M, Kirchsteiger G (2004) A Theory of Sequential Reciprocity. *Games Econ Behavior* 47: 268-298.
- Fehr E, Kirchsteiger G, Riedl A (1993) Does Fairness Prevent Market Clearing?. *Quarterly J Econ* 108: 437-460.
- Fehr E, Schmidt K (1999) A Theory of Fairness, Competition and Cooperation. *Quarterly J Econ* 114: 817-868.
- Finkelstein A, McGarry K (2006) Multiple Dimensions of Private Information: Evidence from the Long-term Care Insurance Market. *Am Econ Rev* 96: 938-958.
- Freixas X, Guesnerie R, Tirole J (1985) Planning under Information and the Ratchet Effect. *Rev Econ Studies* 52: 173-192.
- Fudenberg D, Holmström B, Milgrom B (1990) Short-term Contracts and Long-term Agency Relationships. *J Econ Theory* 51: 1-31.
- Gibbard A (1973) Manipulation for Voting Schemes. *Econometrica* 41: 587-601.
- Green JR, Laffont JJ (1977) Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods. *Econometrica* 45: 427-438.

- Grossman SJ, Hart OD (1983) An Analysis of the Principal-Agent Problem. *Econometrica* 51: 7-45.
- Groves T (1973) Incentives in Teams. *Econometrica* 41: 617-631.
- Grund C, Sliwka D (2005) Envy and Compassion in Tournaments. *J Econ Management Strategy* 14: 187-207.
- Guesnerie R, Laffont JJ (1984) A Complete Solution to a Class of Principal-Agent Problems with Application to the Control of a Self-Managed Firm. *J Public Econ* 25: 329-369.
- Guesnerie R, Picard P, Rey P (1989) Adverse Selection and Moral Hazard with Risk Neutral Agents. *European Econ Rev* 33: 807-823.
- Harris M, Raviv A (1978) Some Results on Incentive Contracts with Applications to Education and Employment, Health Insurance and Law Enforcement. *Am Econ Rev* 68: 20-30.
- Harris M, Raviv A (1979) Optimal Incentive Contracts with Imperfect Information. *J Econ Theory* 2: 231-259.
- Holmström B (1979) Moral Hazard and Observability. *Bell J Econ* 10: 74-91.
- Holmström B (1982) Moral Hazard in Teams. *Bell J Econ* 13: 324-340.
- Holmström B, Milgrom P (1991) Multitask Principal-Agent Analysis: Incentive Contracts, Assets Ownership, and Job Design. *J Law Econ Organization* 7 (Suppl), 24-52.
- Huck S, Rey-Biel P (2006) Endogenous Leadership in Teams. *J Inst Theoretical Econ* 162: 1-9.
- Itoh H (1990) Incentives to Help in Multi-Agent Situations. *Econometrica* 59: 611-636.
- Itoh H (2004) Moral Hazard and Other-Regarding Preferences. *Japanese Econ Rev* 55: 18-45.
- Jensen M, Meckling W (1976) The Theory of the Firm, Managerial Behavior, Agency Costs and Ownership Structure. *J Finan Econ* 3: 305-360.
- Karlan D, Zinman J (2006) Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment. Mimeo.
- Laffont JJ (1990) Analysis of Hidden Gaming in a Three Levels Hierarchy. *J Law Econ Organization* 6: 301-324.
- Laffont JJ, Tirole J (1986) Using Cost Observation to Regulate Firms. *J Polit Econ* 94: 614-641.
- Laffont JJ, Tirole J (1987) Comparative Statics of the Optimal Dynamic Incentive Contract. *European Econ Rev* 31: 901-926.

- Laffont JJ, Tirole J (1988) The Dynamics of Incentive Contracts. *Econometrica* 56: 1153-1176.
- Laffont JJ, Tirole J (1990) Adverse Selection and Renegotiation in Procurement. *Rev Econ Studies* 57: 597-626.
- Lambert R (1983) Long Term Contracts and Moral Hazard. *Bell J Econ* 14: 441-452.
- Lazear E (1995) *Personnel Economics*. MIT Press.
- Macho-Stadler I, Pérez-Castrillo D (1993) Moral Hazard with Several Agents: The Gains from Cooperation. *Int J Industrial Organization* 11: 73-100.
- Macho-Stadler I, Pérez-Castrillo D (1998) Centralized and Decentralized Contracts in a Moral Hazard Environment. *J Industrial Econ* 46: 489-510.
- Malcomson JM, Spinnewyn F (1988) The Multiperiod Principal-Agent Problem. *Rev Econ Studies* 55: 391-408.
- Martimort D (1996) Exclusive Dealing, Common Agency and Multiprincipals Incentive Theory. *Rand J Econ* 27: 1-31.
- Maskin E, Riley J (1989) Optimal Multi-Unit Auctions. In F. Hahn (ed.), *The Economics of Missing Markets, Information, and Games*, 312-335, Oxford University Press, Oxford.
- McAfee P, McMillan J, Reny P (1989) Extracting the Surplus in the Common Value Auction. *Econometrica* 57: 1451-1459.
- Melumad N, Mookherjee D, Reichstein S (1995) Hierarchical Decentralization of Incentive Contracts. *Rand J Econ* 26: 654-672.
- Melumad N, Reichstein S (1987) Centralization vs Delegation and the Value of Communication. *J Accounting Research* 25: 1-18.
- Milgrom P, Weber RJ (1982) A Theory of Auctions and Competitive Bidding. *Econometrica* 50: 1089-1122.
- Mirrlees J (1971) An Exploration in the Theory of Optimum Income Taxation. *Rev Econ Studies* 38: 175-208.
- Mirrlees J (1975) *The Theory of Moral Hazard and Unobservable Behavior, Part I*, WP Nuffield College, Oxford.
- Mookherjee D (1984) Optimal Incentive Schemes with Many Agents. *Rev Econ Studies* 51: 433-446.
- Mussa M, Rosen S (1978) Monopoly and Product Quality. *J Econ Theory* 18: 301-317.
- Myerson R (1979) Incentive Compatibility and the Bargaining Problem. *Econometrica* 47: 61-73.
- Myerson R (1981) Optimal Auction Design. *Math Op Research* 6: 58-73.

- Myerson R (1982) Optimal Coordination Mechanisms in Generalized Principal-Agent Models. *J Math Econ* 10: 67-81.
- Pauly MV (1968) The Economics of Moral Hazard. *Am Econ Rev* 58: 531-537.
- Pauly MV (1974) Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection. *Quarterly J Econ* 88: 44-62.
- Picard P (1987) On the Design of Incentive Schemes under Moral Hazard and Adverse Selection. *J Public Econ* 33: 305-331.
- Rabin M (1993) Incorporating Fairness into Game Theory and Economics. *Am Econ Rev* 83: 1281-1302.
- Radner R (1981) Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship. *Econometrica* 49: 1127-1148.
- Rey P, Salanié B (1990) Long term, Short term and Renegotiation. *Econometrica* 58: 597-619.
- Rey-Biel P (2002) Inequity Aversion and Team Incentives. ELSE WP.
- Riordan MH, Sappington DE (1987) Information, Incentives, and the Organizational Mode. *Quarterly J Econ* 102: 243-263.
- Rogerson W (1985) Repeated Moral Hazard. *Econometrica* 53: 69-76.
- Ross SA (1973) The Economic Theory of Agency: The Principal's Problem. *Am Econ Rev* 63: 134-139.
- Roth AE, Sotomayor M (1990) Two-sided matching: A study in game-theoretic modeling and analysis. Cambridge University Press, New York and Melbourne
- Rothschild M, Stiglitz J (1976) Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Imperfect Information. *Quarterly J Econ* 90: 629-650.
- Rubinstein A, Yaari ME (1983) Repeated Insurance Contracts and Moral Hazard. *J Econ Theory* 30: 74-97.
- Serfes K (forthcoming) Endogenous Matching in a Market with Heterogeneous Principals and agents. *International J Game Theory*
- Shapley LS, Shubik M (1972) The Assignment Game I: The Core. *International J Game Theory* 1: 111-130.
- Spence M (1974) *Market Signaling*. Harvard University Press, Cambridge, Mass.
- Spence M, Zeckhauser R (1971) Insurance Information, and Individual Action. *Am Econ Rev* 61: 380-387.
- Stole L (1990) Mechanism Design under Common Agency. WP, MIT.
- Tirole J (1986) Hierarchies and Bureaucracies: On the Role of Collusion in Organizations. *J Law Econ Organization* 2: 181-214.

- Vera-Hernandez M (2003) Structural Estimation of a Principal-Agent Model: Moral Hazard in Medical Insurance. *Rand J Econ* 34: 670-693.
- Vickrey W (1961) Counterspeculation, Auctions and Competitive Sealed Tenders. *J Finance* 16: 8-37.
- Vickrey W (1962) Auction and Bidding Games. In *Recent Advances in Game Theory*. Princeton, NJ, The Princeton University Conference, 15-27.

Books and Reviews

- Hart O, Holmström B (1987) *The Theory of Contracts*. In Bewley T (ed) *Advances in Economic Theory, Fifth World Congress*. Cambridge University Press, Cambridge.
- Hirshleifer J, Riley JG (1992) *The Analytics of Uncertainty and Information*. Cambridge University Press, Cambridge.
- Klemperer P (2004) *Auctions: Theory and Practice*. Princeton University Press, Princeton.
- Laffont JJ, Martimort D (2002) *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, Princeton.
- Laffont JJ, Tirole J (1993) *A Theory of Incentives in Procurement and Regulation*. MIT Press, Cambridge, Mass.
- Macho-Stadler I, Pérez-Castrillo D (1997) *An Introduction to the Economics of Information: Incentives and Contracts*. Oxford University Press, Oxford.
- Milgrom P, Roberts J (1992) *Economics, Organization and Management*. Prentice-Hall, Englewood Cliffs, NJ.