The B.E. Journal of Theoretical Economics

Topics

Volume 8, *Issue* 1 2008 *Article* 13

On Competitive Equilibria with Asymmetric Information

Jérome Pouyet* Bernard Salanié†
Francois Salanié‡

*Ecole Polytechnique, jerome.pouyet@polytechnique.edu
†Columbia University, bs2237@columbia.edu
‡INRA Toulouse, salanie@toulouse.inra.fr

Recommended Citation
Jérome Pouyet, Bernard Salanié, and Francois Salanié (2008) "On Competitive Equilibria with Asymmetric Information," *The B.E. Journal of Theoretical Economics*: Vol. 8: Iss. 1 (Topics), Article 13.

On Competitive Equilibria with Asymmetric Information*

Jérome Pouyet, Bernard Salanié, and Francois Salanié

Abstract

Asymmetric information is widely supposed to impair the functioning of markets. We show that the presence of competition may invalidate this intuition. Consider a market in which principals compete for attracting heterogeneous agents by offering contracts. Suppose contracts are exclusive, and there are constant returns to trade. When the agents' types are publicly observed under mild conditions, competitive equilibria are efficient. Efficiency is also obtained when types are privately observed, provided that principals do not directly care about the agents' private information (the private value case). Thus hidden information only matters in competitive markets if it affects common values.

KEYWORDS: competition, information, adverse selection, moral hazard

*We thank Pierre-André Chiappori, Marie-Cécile Fagart, Paolo Siconolfi, Robert Townsend, the editor and three referees for helpful comments and discussions. The usual caveat applies.

Introduction

On most markets, there is at least one element of asymmetric information: a seller almost never knows the buyer's preferences perfectly, and a buyer usually knows little about production costs, and sometimes product quality. According to the economic literature, this represents a serious impediment to the working of the economy. Indeed, when studied in bilateral relationships hidden information has been shown to provide a rationale for rebates, discounts, and more generally for non-linear tariffs. It thereby creates distortions that drive the resulting allocation away from the complete information allocation. Its impact on competitive markets may even be dramatic; in Akerlof (1970) competitive equilibria are strongly inefficient due to adverse selection, while pure-strategy equilibria may fail to exist in the insurance model of Rothschild-Stiglitz (1976).

This paper proposes to check the impact of competition on equilibrium allocations by relying on a well-known distinguo. A piece of hidden information is classified as a *private value* if it does not impact directly the payoff of other agents, for given trades between these agents. For example, the preferences of a buyer are usually private values, because the seller only cares about production costs and transfers, and not directly about the willingness-to-pay of the buyer. The case of *common values* occurs when for example a buyer cares about the quality of the good sold to him (as in Akerlof's model, or in financial trading), or when a bank cares on the reliability of a borrower, or an insurer cares about the riskiness of an insuree: here profits depend not only on the contract traded, but also on the information detained by the informed party. It is important to note that the same piece of information may be classified in either category, depending on the situation. For example, whether a client is a reckless driver matters directly to the car insurer, but not to the seller of the car.

To illustrate further, consider procurement contracts. If the buyer is a private firm, then for given contract terms, the unobserved production costs of a supplier are a matter of indifference to the buyer. On the other hand, this need not be the case for public procurement, as the public agency may give a positive weight to the supplier's profit (as in Baron-Myerson (1982) or Laffont-Tirole (1986)); then we are back to a case of common values.

In a little-noticed paper, Fagart (1996, Propositions 6 and 7) proved under some simplifying assumptions that *when contracts are exclusive and values are private, competitive equilibria do not depend on whether information is verifiable or hidden, and competitive outcomes are efficient*. Thus hidden information does not change the set of competitive equilibria and does not harm efficiency, unless it bears at least partly on common values. This note generalizes her result and proves that it holds in a very general model of exclusive contracting, with arbitrary preferences[1]. We only have to assume that for a given agent's type, in the principal-agent utility space the Pareto frontier is continuous and strictly decreasing. The following discussion illustrates the scope of this efficiency result.

First consider the case in which the type of the agent is verifiable, so that there is no hidden information. Though apparently simple, this case includes situations with moral hazard, in which the agent may choose a hidden action. As an application, consider the case when employers compete by offering incentive contracts, and employees choose a contract, and then exert an unobservable effort. As announced above, all competitive equilibria are efficient[2]. This confirms the result in for example Dam and Perez-Castrillo (2006), which we extend to (almost) arbitrary models with pure moral hazard.

We then show that efficiency still obtains when hidden information is introduced, and bears on private values. To illustrate, consider the market for new cars, in which buyers differ along many dimensions: income, taste for engine power, car size, or safety. Even though producers are allowed to use arbitrary non-linear tariffs, competition is strong enough to align tariffs with costs, so that the efficient allocation obtains. In other words, under private values any observed inefficiency in trade can be attributed to imperfect competition, and hidden information alone cannot create any inefficiency.

Finally, under common values we show that competitive equilibria are efficient only in special cases: for example, in a standard model with smooth preferences and continuously distributed types, types must be "locally private", a condition expressing that an uninformed party trading an efficient contract is indifferent to small changes in the type of the agent. This condition is rather stringent. Consider for example the moral hazard model introduced above. Introducing hidden information on, say, the agent's cost of effort would give rise to a model with common value: indeed, for a given incentive contract the principal still cares on the agent's type, since this type determines the choice of effort. According to our results, efficiency only obtains in special cases.

Hence the distinction between private and common values seems to be the relevant one when one wants to characterize the impact of hidden information on the efficiency of competitive equilibria. Still the model imposes restrictions that we acknowledge. Firstly, we only study the "exclusivity" case when the informed party can only contract with one uninformed party: thus the model is

[1] We discuss the differences with Fagart's setting at the end of Section 3.

[2] Here efficiency refers to second-best efficiency—taking into account the employee's incentive constraints.

best-fitted for markets with 0-1 demand, or when exclusivity is legally enforced – car insurance but not life insurance, for instance. The "common agency" case in which each agent can contract with more than one principal is also important, but we leave it for future research. Secondly, we assume that competition is perfect. Interestingly, Inderst (2001) has shown in a model of imperfect competition with matching that the set of equilibria does not depend on whether information is public or hidden, provided it bears on private values. In his model, when two agents are matched one of them is drawn at random and makes a take-it-or-leave-it offer to the other. By contrast, we maintain the usual model of competition, assuming that uninformed parties propose contracts among which informed parties freely select.

We also assume that the payoff of the uninformed parties is the (expected) sum of payoffs realized with each informed party, an assumption that we call "constant returns to trade": loosely said, the marginal payoff from signing an additional contract is constant. This clearly makes competition tougher.

We concentrate here on a partial equilibrium model. General equilibrium models were first studied in Prescott-Townsend (1984), giving rise to an important literature (see Bisin-Gottardi (2006), Prescott-Townsend (2006) and Rustichini-Siconolfi (2006) for recent contributions). These papers aim at extending the welfare theorems to economies with hidden action and hidden information. They thus study the existence of a set of markets and prices that implement (constrained) efficient allocations, when agents take these linear prices as given. By contrast, our paper deals with price-setters that can use arbitrary (non-linear) contracts, and studies whether competition among these agents is efficient. Walrasian-like models with incomplete information have also been studied by Gul-Postlewaite (1992) and McLean-Postlewaite (2002). These papers do not study competitive equilibria; but they exhibit sufficient conditions for the existence of a mechanism that can approximate competitive outcomes arbitrarily closely when the economy is replicated. The essential condition is that each agent must have an "informational size" that goes to zero as the economy grows. As they point out, this condition holds under private values. Thus our results and theirs can be seen as complementary.

1 The Model

We consider a market in which $N \geq 2$ identical principals face heterogenous agents, and compete to attract them by offering contracts. Agents can be of type $i \in \mathcal{I}$ — we put no restrictions on the set $\mathcal{I}$. Feasible exchanges, or contracts, between a principal and an agent must belong to a set $\mathcal{C}$, resulting from technological, informational or institutional constraints. If a principal trades a contract $C \in \mathcal{C}$ with a type-i agent, then the principal gets a profit $b_i(C)$, and the type-i agent gets a payoff $v_i(C)$. Each agent can only contract with a single principal (exclusivity). The total payoff of a principal is simply the expectation of $b_i(C)$, using the distribution of all trades (sell C to i) agreed with the agents: there are thus constant returns to trade. This assumption renders identical the case of a distribution of agents, and the case of a single agent whose type is randomly drawn. For simplicity, we adopt the former convention. We model "no trade" by introducing in $\mathcal{C}$ a contract $C^{\emptyset}$, such that for any type i $b_i(C^{\emptyset}) = 0 = v_i(C^{\emptyset})$.

We distinguish two situations. In the "Verifiable Types" case, types i are publicly observed by both agents and principals and discrimination is allowed, so that a principal may decide to make a contract available only for a subset of types. Under constant returns to trade, the situation is thus equivalent to one in which there is only one type of agent. On the other hand, some actions may not be verifiable; thus this case includes what the literature calls "hidden action" or moral hazard.

In the "Hidden Information" case, types are privately observed by agents (or discrimination is not allowed), so that principals cannot exclude some types of agents from choosing a given available contract. If $b_i(C) = b(C)$ for all $i \in \mathcal{I}$ and all $C \in \mathcal{C}$, we say that there are *private values*, and otherwise that there are *common values*.

Such an abstract model allows us to cover a variety of situations. In models of trade, a contract is a pair (q, t), meaning that the buyer must pay a monetary transfer t for delivery of a bundle of goods described by q; q may gather any relevant informations on quantities or qualities. The buyer may be informed on his preferences, or the seller may be informed on his costs, as in procurement models; both cases are instances of private values. Common values occur for example when the seller of a good has a private information on its quality.

In pure moral hazard models, a contract writes $(w(.),(e_i)_{i\in I})$, where the incentive scheme $w(.)$ maps a set of public outcomes into a transfer from the principal to the agent, and each e_i is the effort chosen by a type-i agent when faced with $w(.)$; thus each e_i must be incentive-compatible for type i. The restriction to incentive-compatible contracts is without loss of generality. Notice that when there are multiple types, $b_i(C)$ will typically depend on i through the agent's choice of effort, so that we are in the common value case.

Finally note that by modifying the definition of $\mathcal{C}$, one can allow for lotteries in a contract (i.e. the price paid for a given quantity may be stochastic). Lotteries on contracts ca also be allowed by considering the set of distributions on the set of feasible contracts.

We can now study the Pareto frontier in the utility space, that we parameterize by the principal's profit $B \geq 0$:

$$\mathcal{V}^i(B) = \sup_{C \in \mathcal{C}} \{v_i(C) \text{ s.t. } b_i(C) \geq B\} \tag{1}$$

At this stage we only know that $\mathcal{V}^i$ is non-increasing. Our first assumption is mainly technical, and is likely to be verified in most economic situations:

Assumption 1. *There exists $B < +\infty$ such that for every types $\mathcal{V}^i(B)$ is negative (or $-\infty$). Moreover, for any type $i \in \mathcal{I}$, if $0 < \mathcal{V}^i(0) < +\infty$, then $\mathcal{V}^i$ is continuous at the right of zero.*

As we shall see later relaxing the second part would allow for some zero-profit equilibria that are nevertheless inefficient. Our main assumption is in fact the following:

Assumption 2. *For any type i, any contract C with $b_i(C) > 0$ and $v_i(C) > 0$, and any positive number ε, there exists a contract C' such that $v_i(C') > v_i(C)$ and $b_i(C') \geq b_i(C) - \varepsilon$.*

Technically, Assumption 2 is easily shown to be equivalent to: $\mathcal{V}^i(B)$ is decreasing for $B > 0$. In words, the assumption says that a principal can undercut any profitable contract that attracts a client. Principals must thus be able to redistribute profits to the agent. This is verified in models of trade, since the transfer specified in a transaction can be reduced. Similarly, in moral hazard models one could think of increasing the agent's wage. This can be done without modifying incentives when the agents' utility function is separable in effort and income. However, Bennardo-Chiappori (2003) consider a pure moral hazard model where the agent's utility function is non-separable in effort and income. If a richer agent is less prone to effort, they show that the Pareto frontier may exhibit a horizontal segment. This invalidates assumption 2, at least with contracts that do not involve randomization.

On the other hand, if one allows for randomization on contracts then assumption 2 holds quite generally. For example, if there exists a contract C_0 such that $v_i(C_0) > \mathcal{V}^i(0)$, then from any contract C one can build a new contract that is a lottery defined on $\{C, C_0\}$, with a small enough probability on C_0 so that it both increases the agent's payoff and hardly reduces the principal's. Assumption 2 then follows[3].

[3] A similar trick can also make $\mathcal{V}^i$ continuous as required in the second part of assumption 1, thus avoiding the difficulties investigated in Arnott-Stiglitz (1988).

Finally the game between principals and agents is the following. First, principals simultaneously offer contracts to agents. Each principal may offer any subset of $\mathcal{C}$ that includes $C^{\emptyset}$. Second, each agent chooses one contract (possibly randomly) among the contracts offered by principals.

Notice that in our framework all agents always end up trading exactly one contract, that can be the empty contract $C^{\emptyset}$. As usual, we label competitive equilibria the subgame-perfect equilibria of this game[4]. Note that we do not assume free entry (this would only reinforce our results). Our aim is to evaluate the power of competition, under both verifiable and hidden information. A natural benchmark is:

Definition 1. *A contract C is efficient for type i if $b_i(C) = 0$ and $v_i(C) = \mathcal{V}^i(0)$. A family $(C_i)_{i\in\mathcal{I}}$ of contracts is efficient if for any $i \in \mathcal{I}$ contract C_i is efficient for type i. A competitive equilibrium is efficient if each type only trades contracts that are efficient for this type.*

2 Verifiable Types

Consider the situation in which the type i of each agent is perfectly observable by all principals. (Again, this includes models of pure moral hazard.) We can then reason as if there was a single type of agent. Our first result characterizes the set of contracts that can be traded in equilibrium:

Proposition 1. *Under verifiable types, for any $i \in \mathcal{I}$ we have:*

i) If C is efficient for type i, then there exists a competitive equilibrium in which C is sold with probability one.

ii) Under Assumptions 1 and 2, all competitive equilibria are efficient.

Proof: i) is immediate: if all principals propose contract C, then all type-i agents agree to buy this contract from a (possibly randomly chosen) principal. Moreover no principal can profitably deviate, since by definition of $\mathcal{V}^i(0)$ to attract any type i this principal would have to incur losses.

Let us now prove ii). At an equilibrium, note that a type-i agent must get the same payoff V_i with each contract that he buys with positive probability. Let $\overline{B}$ be the supremum of profits b_i on this subset of contracts (this subset cannot be empty, since no-trade means trading $C^{\emptyset}$). Notice that $\mathcal{V}^i(\overline{B}) \geq V_i \geq 0$, so that $\overline{B}$ is finite from the first part of assumption 1.

Suppose that $\overline{B} > 0$. Aggregate profits on type i are at most equal to $\overline{B}$, so that one of the N principals – say, principal 1 – gets a profit at most equal

[4] We only consider pure strategies for principals.

to $\overline{B}/N$. Choose B' such that $\overline{B}/N < B' < \overline{B}$. Assumption 2 then implies the existence of a contract C such that $b_i(C) \geq B'$ (and thus the profit on C is strictly higher than the initial profit of Principal 1) and $v_i(C) > \mathcal{V}^i(\overline{B})$ (and thus C attracts the agent since we know that $\mathcal{V}^i(\overline{B}) \geq V_i$). Hence principal 1 could profitably deviate by offering C at the first stage of the game, a contradiction.

Therefore $\overline{B} = 0$, and consequently all contracts sold with positive probability must yield zero profit. Moreover, if $V_i < \mathcal{V}^i(0)$, then by the continuity part in assumption 1 we have $V_i < \mathcal{V}^i(\varepsilon)$ for ε positive and small, and thus a principal could deviate to a profit-making contract that would attract the agent[5]. ■

Thus competition achieves efficiency when types are verifiable. This result is a straightforward illustration of the power of competition when principals compete à la Bertrand. Still it is general enough to be applied to quite sophisticated models, such as moral hazard models. To fix ideas, consider the model in Dam and Perez-Castrillo (2006). Feasible contracts specify an observable investment, an incentive wage, and an incentive-compatible effort. It is easily seen that assumptions 1 and 2 hold in their model. Proposition 1 thus confirms the result in Dam and Perez-Castrillo (2006) that equilibrium contracts must be second-best efficient (taking incentive-compatibility into account), and that principals get no rent[6].

This seems contradictory with Bennardo-Chiappori (2003), who exhibit competitive equilibria with positive profits for principals. As explained above, agents in their model have nonseparable utilities; if only non-random contracts are allowed, this invalidates our assumptions. With random contracts, their example requires that the economy be subject to aggregate shocks – a feature that we ruled out from the start.

---

[5]Relaxing the continuity part in assumption 1 would not alter the zero-profit result, but would allow for inefficient equilibria with lower payoffs for the agent. Indeed let $\mathcal{V}^i(0^+)$ be the supremum of $\mathcal{V}^i(B)$ on $\{B \ s.t. \ B > 0\}$. If there exists a contract C such that $b_i(C) = 0$ and $\mathcal{V}^i(0^+) \leq v_i(C) < \mathcal{V}^i(0)$, then one can build an equilibrium in which all principals propose C, and the agent trades it: by construction no other profitable contract would attract the agent.

[6]In the case when the number of agents is greater than the number of principals, one can derive similar efficiency results; but now (1) maximizes the profits of a principal, under a participation constraint for the agent.

3 Hidden Information

We now tackle the case when the types of the agents are unknown to the principals. The following result corresponds to Proposition 1.i:

Proposition 2. *Suppose that there exists an efficient family of contracts* $(C_i)_{i\in I}$ *that is moreover incentive-compatible:*

$$\forall\ i, j \in \mathcal{I} \qquad v_i(C_j) \leq v_i(C_i).$$

Then whether information is verifiable or hidden, there exists a competitive equilibrium in which all principals propose $(C_i)_{i\in\mathcal{I}}$*, and each* C_i *is chosen by type* i*.*

Proof: Under verifiable information, the result follows from Proposition 1.i. Under hidden information, suppose that all principals propose $(C_i)_{i\in\mathcal{I}}$. Because this family is incentive-compatible, the type-i agent maximizes his payoff by choosing C_i. Moreover no principal can profitably deviate, since by definition of $\mathcal{V}^i(0)$ to attract any type i this principal would have to incur losses. ■

Now consider the private values case, in which the profit of a principal does not depend on the agent's type: $b_i(C) = b(C)$ for all types i and all contracts C.

Proposition 3. *Under private values, any efficient family of contracts is incentive-compatible.*

Proof: if the family $(C_i)_{i\in\mathcal{I}}$ is efficient, then for any j $b_j(C_j) \geq 0$; under private values, this implies for all i $b_i(C_j) = b_j(C_j) \geq 0$. Hence C_j verifies the constraint in the program $\mathcal{V}^i(0)$, whose solution is C_i. Therefore $v_i(C_i) \geq v_i(C_j)$, so incentive-compatibility holds. ■

Proposition 2 and 3 together imply that under private values all efficient families of contracts can be obtained as equilibrium outcomes, both under verifiable and under hidden information. In other words, efficient equilibria under verifiable information remain equilibria under hidden information. We now tackle the converse result, namely whether hidden information generates new equilibrium outcomes that did not exist under verifiable information. There is a technical difficulty here: an agent indifferent between two contracts may change the probability with which he chooses each if a principal proposes an additional contract, even if the new contract does not attract this agent. In the literature, this difficulty is typically dealt with by assuming that a principal can break ties in his favor. In fact, we can manage with a much weaker assumption. We define robust competitive equilibria as follows:

Definition 2. *A competitive equilibrium is robust if a principal cannot profitably deviate by adding contracts to his offer, assuming that agents that do not strictly prefer one of these additional contracts choose the same contracts as before the deviation.*

Loosely speaking, this assumes that each agent's behavior satisfies a form of independence of irrelevant alternatives. Then we obtain the following result, which extends Proposition 1.ii (the proof is similar and is relegated to the Appendix):

Proposition 4. *Assume private values, and Assumptions 1 and 2. Then all robust competitive equilibria are efficient.*

As mentioned in the introduction of this paper, Fagart (1996) obtains similar results, using stronger assumptions: two types of agents, with concavity/convexity assumptions on preferences and feasible contracts. Fagart also assumes that principals can undercut an incentive-compatible family of contracts by proposing another incentive-compatible family that attracts all types (ties are broken in favor of the principal) and only slightly reduces the principal's payoff on each type. In contrast, in our framework ties are dealt with by imposing robustness; and we only assume that introducing new contracts to undercut some existing contracts is possible, without requiring that the new set of contracts be incentive-compatible. These differences are mainly technical, but they allow us to greatly extend the scope of the results, as we now discuss.

4 Some Results on the Common Value Case

Our results show that under private values hidden information does not matter: it does not change the set of competitive equilibria, and each contract traded is efficient for the type that chooses it. When for instance buyers are privately informed on their preferences, values are private since sellers' payoffs only depend on transfers and costs. Therefore any inefficiency in trade must be linked to the existence of market power.

Things are more difficult in the common value case, since an efficient family of contracts need not be incentive-compatible. As it turns out, we can say a bit more. Assume that types are one-dimensional and indexed by θ, that a contract specifies a continuous quantity q and a monetary transfer t, that principals have utility function $b(q,t,\theta)$ and agents have utility function $v(q,t,\theta)$. Assume smooth preferences, with opposite signs for v'_t and b'_t; assume also that the efficient contract $(q^*(\theta), t^*(\theta))$ is interior, and differentiable in θ. Then the first best is given by:

$$\frac{b'_t}{v'_t} = \frac{b'_q}{v'_q} \qquad \text{and} \qquad b(q^*(\theta), t^*(\theta), \theta) = 0.$$

Now the gain for type θ from deviating from θ to θ' is

$$G(\theta,\theta') = v(q^*(\theta'), t^*(\theta'), \theta) - v(q^*(\theta), t^*(\theta), \theta).$$

Denote $H(\theta,\theta')$ the partial derivative of G with respect to θ'. By definition:

$$G(\theta,\theta') = \int_{\theta}^{\theta'} H(\theta,x)dx,$$

We want the first-best to be incentive compatible, so $G(\theta,\theta') \leq 0$ for all θ,θ', or:

$$\int_{\theta}^{\theta'} H(\theta,x)dx \leq 0 \;\text{ for all } \theta,\theta'. \qquad (H)$$

When types are continuously distributed, we can take the limit as θ' goes to θ, either from below or from above. This implies that $H(\theta,\theta) = 0$ for any interior θ, otherwise the agent would deviate in one or the other direction. Standard calculations show that

$$H(\theta,\theta) = -\left(\frac{v'_t}{b'_t} b'_\theta\right)(q^*(\theta), t^*(\theta), \theta).$$

Since v'_t/b'_t is negative, $H(\theta,\theta)$ can only be zero if $b'_\theta(q^*(\theta), t^*(\theta), \theta) = 0$. Thus competitive equilibria with common values (b depends on θ) can only be efficient if values are "locally private" at the first best contract: given the efficient contract, principals must be indifferent to small changes in θ.

Locally private common values are of course a very special case. The insurance sector provides an illustration. Consider the following straightforward generalization of the Rothschild-Stiglitz model. The agent of type θ faces the risk of a stochastic loss $\tilde{L}(\theta)$, and may buy one insurance contract in the set $\mathcal{C}$. A contract specifies a premium P and an indemnity $I(L)$ conditioned by the observed loss; the insurer's payoff then is $P - EI(\tilde{L}(\theta))$. The agent's payoff is $u(P, I(.), \theta)$, which for example can be written $EU(I(\tilde{L}(\theta)) - P - \tilde{L}(\theta))$ in the expected utility case.

Note that hidden information here bears both on the agent's preferences and on his riskiness. The former is private values, but the latter is common values as it changes the insurer's profits for any given contract. Now assume that the agent is risk-averse, and the principal is risk-neutral, so that under verifiable information the competitive contract is the full insurance contract $I(L) = L$, with the actuarially fair premium $P(\theta) = E\tilde{L}(\theta)$. Then incentive-compatibility holds if and only if the expected loss $E\tilde{L}(\theta)$ is the same for all types; this corresponds to our definition of locally private values. This of course does not hold in the seminal Rothschild-Stiglitz model, in which agents differ in the probability that they incur a given loss. Still recall that insurers usually gather very detailed information on their insurees; the variance of the expected loss conditional on this information may in practice be small, which would make competitive equilibria close to efficient even though the distribution of losses may differ across types.

Introducing moral hazard in insurance models makes the picture more complex because both the agent's preferences and riskiness matter for the choice of effort, so that both characteristics become common values. Then the property of incentive-compatibility is more difficult to check, and typically does not hold. In particular, Jullien-Salanié-Salanié (2007) solves an insurance model with adverse selection on the absolute risk-aversion index θ of the insuree, and moral hazard on the prevention effort. Then the first best contract usually does not fully insure the agent; for this contract the principal's payoff will depend on the agent's risk-aversion since the latter underlies the choice of effort. Therefore values are not locally private and competitive equilibria (if any) will be inefficient if types are continuously distributed.

When the distribution of types is discrete, it is possible to construct reasonable examples of models with common values in which first-best contracts are incentive-compatible. If there are m possible types then condition (H) yields $m(m-1)$ inequalities, and it may be possible to satisfy all of them. On the other hand, when m increases this requires stronger and stronger restrictions on preferences; in the limit they impose that values be locally private. As announced in the introduction, the distinction between private and common values seems to be the relevant one when one wants to characterize the impact of hidden information on the efficiency of competitive equilibria.

Appendix

Proof of Proposition 4: The proof begins much as that of Proposition 1. At an equilibrium each type i must get the same payoff V_i on the subset of contracts that he buys with positive probability. Denote $\overline{B}_i$ the supremum of $b(.)$ on this set. Define:

$$\mathcal{J}(B) = \{i \in \mathcal{I}, \overline{B}_i > B\},$$

and let $\overline{B}$ be the supremum of those B such that $\mathcal{J}(B)$ has positive probability (this construction allows us to ignore types with possibly high profits but zero-probability). Notice that $\mathcal{V}^i(\overline{B}_i) \geq V_i \geq 0$, so that $\overline{B}$ is finite from the first part of assumption 1.

Suppose that $\overline{B} > 0$. Choose B such that $\overline{B}/N < B < \overline{B}$. Aggregate profits on each type in $\mathcal{J}(B)$ are at most equal to $\overline{B}$, so one of the principals – say, principal 1 – earns at most $\overline{B}/N$ on average on each type in $\mathcal{J}(B)$. We now show how principal 1 can profitably attract all such types without losing anything on other types.

For each $j \in \mathcal{J}(B)$, assumption 2 ensures the existence of a contract C'_j such that $v_j(C'_j) > \mathcal{V}^j(\overline{B}_j) \geq V_j$ and $b(C'_j) \geq B$. Now let principal 1 deviate by adding to his equilibrium offer the contracts $\{C'_j, j \in \mathcal{J}(B)\}$. Types $j \in \mathcal{J}(B)$ strictly benefit from this deviation and buy one of the additional contracts with probability one, and thus principal 1 gets a profit at least equal to B on each of them. Types $j \notin \mathcal{J}(B)$ that do not strictly benefit from the deviation can be assumed not to change their behavior, since by assumption the equilibrium is robust. Types $j \notin \mathcal{J}(B)$ that strictly benefit from the deviation must choose one of the additional contracts; but given the definition of $\mathcal{J}(B)$, any such type generated a profit no larger than B in our equilibrium and so this cannot hurt principal 1.

Overall principal 1 can obtain $B > \overline{B}/N$ on each type in $\mathcal{J}(B)$, without losing anything on other agents. Since $\mathcal{J}(B)$ has positive probability, the deviation yields strictly higher profits, and we have obtained a contradiction.

Therefore $\overline{B} = 0$, and all contracts traded yield zero-profit to the principals. Finally, if a non-negligible subset of types get strictly less than $\mathcal{V}^i(0)$, then thanks to assumptions 1 and 2 a principal could deviate by offering profitable contracts that would attract these types (and maybe others, which is not a problem under private values), once more a contradiction. ■

References

Akerlof, G.A. (1970), "The market for 'lemons': Quality uncertainty and the market mechanism", *The Quarterly Journal of Economics*, 89, 488–500.

Arnott, R.J. and J.E. Stiglitz (1988), "The Basic Analytics of

Moral Hazard", *Scandinavian Journal of Economics*, 90(3): 383-413.

Baron, D. and R. Myerson (1982), "Regulating a Monopolist with Unknown Costs", *Econometrica*, 50(4): 911-930.

Bennardo, A. and P.A. Chiappori (2003), "Bertrand and Walras Equilibria under Moral Hazard", *Journal of Political Economy*, 111: 785-817.

Bisin, A. and P. Gottardi (2006), "Efficient Competitive Equilibria with Adverse Selection", *Journal of Political Economy*, 114: 485-516.

Dam, K. and D. Perez-Castrillo (2006), "The Principal-Agent Matching Market", *The B.E. Journal of Theoretical Economics*, Vol. 2: Iss. 1 (Advances), Article 1.

Fagart, M.C. (1996), "Concurrence en contrats, anti-sélection et structures d'information", *Annales d'Economie et de Statistique*, 43: 1-27.

Gul, F. and A. Postlewaite (1992), "Asymptotic Efficiency in Large Exchange Economies With Asymmetric Information", *Econometrica*, 60: 1273-1292.

Inderst, R. (2001), "Screening in a Matching Market", *Review of Economic Studies*, 68: 849-868.

Jullien, B., B. Salanié and F. Salanié (2007), "Screening Risk-averse Agents Under Moral Hazard: Single-crossing and the CARA Case", *Economic Theory*, 1: 151-169.

Laffont, J.-J. and J. Tirole (1986), "Using Cost Observation to Regulate Firms", *Journal of Political Economy*, 94(3): 614-41.

McLean, R. and A. Postlewaite (2002), "Informational Size and Incentive Compatibility", *Econometrica*, 70: 2421-2453.

Prescott, E.C. and R.M. Townsend (1984), "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard", *Econometrica* 52: 21-45.

Prescott, E.S. and R.M. Townsend (2006), "Firms as Clubs in Walrasian Markets with Private Information", *Journal of Political Economy*, 114(4): 644-671.

Rothschild, M. and J. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect

Information", *The Quarterly Journal of Economics*, 90: 629-649.

Rustichini, A. and P. Siconolfi (2006), "General Equilibrium in Economies with Adverse Selection", Working paper, University of Minnesota and Columbia University.