

Bertrand and Walras equilibria under moral hazard.*

A. Bennardo[†]

P.A. Chiappori[‡]

Revised version

November 2001

Abstract

We consider a simple model of competition under moral hazard with constant return technologies. We consider preferences that are not separable in effort: marginal utility of income is assumed to increase with leisure, especially for high income levels. We show that, in this context, Bertrand competition may result in positive equilibrium profit. This result holds for purely idiosyncratic shocks when only deterministic contracts are considered, and extends to unrestricted contract spaces in the presence of aggregate uncertainty. Finally, these findings have important consequences upon the definition of an equilibrium. We show that, in this context, a Walrasian general equilibrium a la Prescott-Townsend may fail to exist: any 'equilibrium' must involve rationing.

1. Introduction

Consider an economy where a large number of insurers compete by offering insurance contracts to a continuum of agents owing a risky endowment. These endowments are affected by two types of shocks. One is idiosyncratic; the other is an aggregate shock affecting the economy as a whole. The probability of occurrence of the individual shock is affected by some unobservable prevention effort supplied by the agent, while the aggregate shock is exogenous. Finally, agents have general VNM preferences exhibiting risk aversion, and the insurance technology is linear and involves neither frictions nor fixed costs. What will the outcome of competition be in this economy?

Two alternative representations of competition under asymmetric information have been adopted in the literature. Following the 'strategic' approach, insurance companies simultaneously offer contracts that are then purchased by the agents, and the competitive outcome is modelled as a Nash equilibrium of this two-stage game¹. From this perspective, our economy

*Paper presented at seminars in Chicago, Los Angeles, Philadelphia, Milan, Paris and Toulouse and at the NBER Conference on General Equilibrium (Berkeley, 1998). We thank the participants and Fernando Alvarez, Richard Arnott, Alberto Bisin, Piero Gottardi, Bruno Jullien, Joseph Ostroy, Jose Scheinkman, Lars Stole, Rob Townsend, Bill Zame and especially the editor (Lars Hansen) and three anonymous referees for helpful comments. Errors are ours.

[†] University of California at Los Angeles. Email : abennardo@econ.ucla.edu

[‡] University of Chicago. Email : pchiappo@midway.uchicago.edu

¹The introduction of further stages, as discussed by Hellwig (1983) in an adverse selection setting, would not affect the results in our moral hazard framework.

can be viewed as a moral hazard counterpart of Rothschild and Stiglitz's model of competition under adverse selection. Alternatively, one can, following the seminal contribution of Prescott and Townsend (1984a), analyze the economy from a Walrasian viewpoint. Then an equilibrium is defined as a set of market-clearing 'prices', where the corresponding 'commodities' are contracts in a general sense, i.e. contingent lotteries on consumption and effort. In both cases, it has been recognized that the outcome of competition depends on the set of available contracts. For instance, if individual are not able to sign (and enforce) exclusive contracts, then the strategic equilibrium need not coincide with a Pareto efficient allocation². Also, it is well known that efficiency may in this context require randomized contracts. If lotteries are not enforceable (i.e., only 'deterministic' contracts can be implemented), then again the competitive outcome may fail to be (second best) efficient.³

Albeit the results presented below hold both when lottery contracts are enforceable and when only deterministic contracts can be implemented, our main interest, in the present paper, is the *pure moral hazard (PMH)* context, where unobservability of effort is the *only* restriction on agents' trades (i.e., consumptions and trades are observable and contractible, lotteries are enforceable, etc.). The realism of this framework can be questioned, and raises several difficult issues (for instance, is it possible to monitor individual consumptions or to enforce lottery contracts?). Still, we believe that the PMH case is the natural benchmark for a study of this kind, if only because it allows to disentangle the impact of moral hazard per se from that of other restrictions on trades. Once the logic of PMH situations has been understood, it becomes easier to predict the outcomes of alternative models where additional constraints (say, non enforceability of random contracts) are introduced.

In the PMH case, existing results in the literature, as well as conventional wisdom, suggest that three main conclusions should hold in our framework: (i) competition drives profits to zero, (ii) a Walrasian equilibrium exists and is (second best) efficient, and (iii) the strategic equilibrium coincides with the Walrasian equilibrium⁴. The main claim of the present paper is that existing results, which have been developed in particular frameworks, are not robust, and that conventional wisdom is actually wrong. In fact, none of the three properties just stated holds in general. We exhibit a robust example in which the equilibrium of the strategic market game entails positive profits for the insurers, while the Walrasian equilibrium fails to exist. This conclusion does not rely on specific assumptions such as arbitrary limitations on the contract space or exogenous unobservability of particular transactions. Rather, it stems from the interaction of two features of our model, namely non separable preferences and the presence of an aggregate shock. These features are, if anything, more general than what can be

²The intuition is that the second-best optimal contract between an agent and an insurer typically entails partial coverage in order to provide the agents with the right incentives to prevention. However, in the absence of exclusivity, the agent could then purchase a complementary coverage from some competitor, which would unravel the initial contract. See Helpman and Laffont (1975), Arnott and Stiglitz (1987) or Bisin and Guaitoli (1999) for precise investigations.

³Actually, Bennardo (2001) shows that in situations where second best efficiency requires randomization, the deterministic equilibrium contract may fail to be efficient *even within the set of deterministic contracts*.

⁴The last two conclusions are obviously specific to the moral hazard case. Under adverse selection, Rothschild and Stiglitz have showed that a strategic equilibrium may fail to exist and/or to generate efficient outcomes. This feature is due to the externality between contracts that characterizes adverse selection (whether a contract makes profit or losses depends on the contracts proposed by the competitors). No externality of this type occurs in a moral hazard context. The same remark applies to the existence of a Walrasian equilibrium in a la Prescott-Townsend under adverse selection (see Bisin and Gottardi (2002) for a thorough investigation).

found in most of the literature; it is precisely the generality of our perspective that explains the discrepancies with previous works.

The intuition of the main result can be summarized as follows. We assume that leisure and consumption are gross complements, i.e., that the marginal utility (resp. disutility) of leisure (resp. effort) increases with wealth, and that this effect becomes especially important as the agent gets wealthier. It should be noted that this assumption, if anything, has better empirical support than the alternative hypothesis that is adopted by most on the literature - namely, that preferences are additively separable in leisure and consumption. Under moral hazard, this form of complementarity gives rise to specific income effects that prevent "usual" undercutting strategies from eliminating profits. Specifically, we show that there exists some (finite) consumption bundle that maximizes the agent's utility on the set of contracts inducing the higher prevention effort level. From this bundle, it is impossible to increase the agent's utility (say, by increasing his consumption when the 'good' outcome is realized) without violating the incentive constraint, the idea being that additional expected consumption increases the marginal cost of effort, and that this effect, for large enough consumption levels, overcompensates the standard incentive effect. If the technology is such that expected wealth in the economy exceeds the expected consumption corresponding to the optimal bundle, then the principal makes a positive profit at equilibrium.

For the sake of readability, we stick throughout the model to the insurance story just described. However, the scope of our conclusions is much broader. They can be applied to labor contracts, executive compensation, sharecropping or credit relationships, just to name a few. They have several surprising implications. One is that the zero profit assumption, often viewed as a shortcut for competition, should be handled with care. Competition generally implies that agents receive the best contract available under several constraints, including a non negative profit condition. Our result suggests that there is no reason to expect the non negative profit restriction to always bind. A second consequence is that the Walrasian mechanism, powerful as it is (at least in the Prescott-Townsend version), may fail to work in a moral hazard context. The interpretation of our non existence result is that rationing may be needed to achieve efficiency. To the best of our knowledge, this conclusion is new, at least when fully non linear (exclusive) contracts are assumed to be enforceable. Finally, our conclusions have direct consequences for the literature on incentives. That increasing an agent's reward when the outcome is 'good', while keeping it unchanged in 'bad' cases, may in fact *reduce* the agent's overall incentives (especially for high levels of wealth or income) may shed a new light on such issues as managerial incentives and executive compensation. These applications are considered in a companion paper (Bennardo and Chiappori 2002).

The paper is organized as follows. In section 2, we describe the basic model and stress its main properties. Section 3 is devoted to the study of equilibria a la Bertrand. The general equilibrium perspective is developed in Section 4. Conclusions are discussed in the last section.

2. The model

2.1. Uncertainty

The economy consists of two class of individuals: a continuum of ex ante identical risk averse agents and a 'large number' of risk neutral principals competing by offering contracts to the

agents. Agents have a risky endowment that is affected by two types of shocks. One is an exogenous aggregate shock influencing the economy as a whole; i.e., there exist two different, ex post verifiable aggregate (or 'collective') states of the world, $s = 1$ and $s = 2$, that occur with respective probabilities λ and $(1 - \lambda)$. In addition, the endowment is subject to an idiosyncratic (or 'individual', following the distinction emphasized by Cass and al. 1997) shock. Individual shocks are identically and independently distributed (iid) throughout the population, and the shock affecting an agent depends on an effort e supplied by this agent in an accident avoidance technology. The level of effort is not observable by the insurer; we assume for simplicity that it can take only two values, e_l and e_h , with $e_h > e_l$. In each aggregate state s , individual wealth can take two values $y_s^a > y_s^b$; the 'good' outcome a obtains with probability $P(e)$, where $P(e_h) = P$ and $P(e_l) = p < P$. We assume furthermore that 2 is a 'bad' aggregate state, i.e. $y_1^a > y_2^a$ and $y_1^b > y_2^b$. An important assumption is that each agent has to exert the effort *before* the aggregate state is revealed (i.e., before knowing whether the endowment space is $Y_1 = \{y_1^a, y_1^b\}$ or $Y_2 = \{y_2^a, y_2^b\}$).

A natural interpretation of this setting is that agents are households, while principals are insurance companies mutualizing individual risks. Each household may incur an accidental damage $\xi_s = y_s^a - y_s^b$ to its property (the occurrence of an accident then defines the 'bad' outcome, b). The occurrence probability of a damage can be reduced if the agent provides some labor e in a prevention or maintenance activity that is not observed by the insurer. Finally, the two aggregate states can be interpreted as weather conditions: state 1 is 'normal weather', whereas state 2 corresponds to a 'hurricane'. The presence of aggregate uncertainty is one important difference between our framework and most of the moral hazard literature. Of special interest is the case where y_2^b is very small; i.e., failing to provide the maintenance effort required is always bad, but can reveal disastrous in case of a hurricane.

Alternative interpretations can of course be considered (e.g., agents and principals could be respectively thought of as workers and managers, tenants and landlords, etc.). Following most of the literature on competition under asymmetric information, we will stick to the insurance interpretation. Throughout the paper, $\bar{Y}_s = P y_s^a + (1 - P) y_s^b$ denotes the expected production in state s when effort is e_h , while $\bar{y}_s = p y_s^a + (1 - p) y_s^b$ is the expected production in state s when effort is e_l . From the previous assumptions, we have that

$$\bar{Y}_s > \bar{y}_s, \bar{y}_1 > \bar{y}_2 \text{ and } \bar{Y}_1 > \bar{Y}_2$$

2.2. Preferences

Preferences of both the principals and the agents are state independent. Principals are risk neutral profit maximizers. For any non negative consumption c and any effort $e \in \{e_0, e_1\}$, agent's preferences are represented by the VNM utility function $u(c, e) = v(c, 1 - e)$, where $1 - e$ is the agent's leisure, and v is twice continuously differentiable, strictly concave, strictly increasing, and satisfies $\lim_{c \rightarrow 0} \partial v / \partial c = +\infty$. Let c_s^x , where $x = a, b$ and $s = 1, 2$, denote the agent's consumption, contingent on each aggregate state and on the realization of the output. A key point is that, in contrast with most of the literature on moral hazard, *we assume that u is not separable in effort and consumption*. This assumption fits the particular interpretation we just suggested, where effort is interpreted as working time. Non separability of leisure (with

respect to consumption) is a standard finding of the empirical literature on labor supply (see for instance Browning and Meghir (1997)); it reflects the very natural intuition that, in general, the marginal utility of leisure increases with wealth, if only because number of consumable goods (travel, services,...) are obvious complements of leisure.⁵ Surprisingly enough, many standard results of the moral hazard literature turn out to be very sensitive to the separability assumption, although most of the time the latter is only made for commodity and cannot be considered as particularly realistic in any sense, as already recognized in the seminal contribution of Grossman-Hart (1983).

2.3. Contracts

In our insurance setting with exclusive relationships, a deterministic contract specifies a required effort level and a transfer d_s^x from the insurer to the agent, contingent on each realization of the output (x) and of the aggregate state (s); typically, one expects that $d_1^a < 0$ (i.e., the agent pays an insurance premium when no damage has occurred, at least in the good aggregate state). For each transfer, the agent's consumption is $c_s^x = y_s^x + d_s^x$. It is of course equivalent, and often more convenient, to write the contract in consumption terms (instead of transfers). That is, a deterministic contract is defined by a 5-uple

$$\gamma = (e; c_1^a, c_1^b; c_2^a, c_2^b)$$

that prescribes the required effort level e and the contingent consumptions.

Such a contract is called deterministic albeit it is contingent on the state of world. 'Deterministic', here, means that the contract does not entail additional randomness. It is well known that efficiency may require some explicit randomization, which in full generality can be of two types: ex ante (whereby the contract prescribes some probability distribution on the recommended level of effort) and ex post (whereby the agent's contingent payment, conditional on the realization of the shocks, is stochastic). Hence the fully general form of a contract is:

$$\Gamma = (\alpha, \mu_{h,1}^a, \mu_{h,1}^b, \mu_{l,1}^a, \mu_{l,1}^b, \mu_{h,2}^a, \mu_{h,2}^b, \mu_{l,2}^a, \mu_{l,2}^b) \quad (2.1)$$

where :

- α is the ex ante probability of effort e_h (then e_l is chosen with probability $1 - \alpha$)
- $\mu_{h,s}^a$ (resp. $\mu_{h,s}^b$) is the ex post probability distribution of consumption in state s , conditional on output y_s^a (resp. y_s^b) and the contractually required effort e_h

⁵A Beckerian justification relies upon the existence of a domestic production function that produce some agent-specific commodity, using time and the consumption good as complementary inputs. Note that, in this case, the marginal utility of consumption is likely to increase with leisure - i.e., $\frac{\partial^2 u}{\partial c \partial e} < 0$. For instance, if well-being is proportional to the consumption of a single household good ξ , produced from some constant return to scale technology :

$$\xi = f(c, 1 - e) = (1 - e) \phi\left(\frac{c}{1 - e}\right)$$

where ϕ is increasing concave, then $\frac{\partial^2 u}{\partial c \partial e} = \frac{\partial^2 f}{\partial c \partial e}$ is always negative.

- similarly, $\mu_{l,s}^a$ (resp. $\mu_{l,s}^b$) is the ex post probability distribution of consumption in state s , conditional on output y_s^a (resp. y_s^b) and effort e_l .

The above construction is equivalent to the one proposed by Prescott and Townsend (1984a,b). Technically, they define random contracts as joint probability distributions on effort, output and consumption; the equivalence with our definition comes from the fact that lotteries on lotteries are lotteries. In the remainder, we stick to our notation, where ex ante and ex post randomization are explicitly distinguished, for two reasons. One is clarity; we believe that the basic intuitions (and, specifically, the role of randomization) can be better grasped with this less compact but may be more pedagogical presentation. Also, the distinction is particularly useful for the characterization of equilibrium contracts; indeed, ex ante and ex post randomization play very different roles in the equilibrium outcomes, as it will become clear later on.

The general timing of actions is summarized in Figure 1 below.

Insert Figure 1 here

2.4. Incentive and feasibility constraints

Feasibility Allocations are constrained by a feasibility and an incentive compatibility condition. The feasibility constraint reflects the fact that the insurance company must break even *in each aggregate state*. For deterministic contracts, this implies that

$$P(e) c_s^a + (1 - P(e)) c_s^a \leq P(e) y_s^a + (1 - P(e)) y_s^a, \quad s = 1, 2 \quad (2.2)$$

For instance, if the contract induces the choice of the high effort level, the constraint becomes

$$P c_s^a + (1 - P) c_s^b \leq \bar{Y}_s, \quad s = 1, 2 \quad (2.3)$$

while with a low level it becomes

$$p c_s^a + (1 - p) c_s^b \leq \bar{y}_s, \quad s = 1, 2 \quad (2.4)$$

In the general case (allowing for random contracts), the constraint is

$$\alpha (P E_{h,s}^a + (1 - P) E_{h,s}^b) + (1 - \alpha) (p E_{l,s}^a + (1 - p) E_{l,s}^b) \leq \alpha \bar{Y}_s + (1 - \alpha) \bar{y}_s, \quad s = 1, 2$$

where $E_{h,s}^a = \int c d\mu_{h,s}^a(c)$ is the mean consumption corresponding to the distribution $\mu_{h,s}^a$.

Incentive compatibility Incentive compatibility imposes that the prescribed effort coincides with the agent's best choice. With random contracts, the constraint takes the form:

$$\sum_{s=1}^2 [(P E U_{h,s}^a + (1 - P) E U_{h,s}^b) - (p E U_{l,s}^a + (1 - p) E U_{l,s}^b)] \geq 0 \quad (2.5)$$

where $EU_{x,s}^a = \int u(c, e_x) \mu_{x,s}^a(c)$ is the expected utility corresponding to effort x and the distribution of consumption $\mu_{x,s}^a$ with $x = h, l$. In particular, a *deterministic* allocation $(e_h, c_1^a, c_1^b; c_2^a, c_2^b)$ is incentive-compatible if it satisfies:

$$\lambda F(c_1^a, c_1^b) + (1 - \lambda) F(c_2^a, c_2^b) \geq 0 \quad (2.6)$$

where

$$F(c_s^a, c_s^b) = Pu(c_s^a, e_h) + (1 - P)u(c_s^b, e_h) - pu(c_s^a, e_l) - (1 - p)u(c_s^b, e_l) \quad (2.7)$$

denotes the difference (in some aggregate state s) between expected utilities under the high and low efforts respectively. It is natural to assume that the incentive problem is non degenerate, in the sense that the high effort can be implemented by *some* contract- i.e., that there exists some strictly positive consumption level g^a such that

$$F(g^a, 0) > 0$$

We now state our main assumption on the agents' preferences:

Assumption D : *Marginal utility of consumption uniformly increases with leisure; i.e., there exists some $k > 0$ such that, for all (c, e)*

$$\frac{\partial^2 u}{\partial c \partial e}(c, e) \leq -k < 0 \quad (D)$$

In addition,

$$\lim_{c \rightarrow +\infty} \frac{\partial u(c, e_h) / \partial c}{\partial u(c, e_l) / \partial c} = 0 \quad (D_\infty)$$

$$\lim_{c \rightarrow 0} \frac{\partial u(c, e_h) / \partial c}{\partial u(c, e_l) / \partial c} = 1 \quad (D_0)$$

In words, we assume that marginal utility of consumption always increases with leisure, and that this effect is particularly strong for wealthy agents, whereas it is almost negligible for very low consumption levels.

The geometry of the incentive frontier It is convenient to fix the consumption plan in state s to some arbitrary value $(\bar{c}_s^a, \bar{c}_s^b)$ and to consider, in the alternative state $s' \neq s$, the set of consumption plans $(c_{s'}^a, c_{s'}^b)$ that induce the high effort level. These must satisfy an equation of the form:

$$F(c_{s'}^a, c_{s'}^b) \geq H \quad (2.8)$$

where H is a constant. In the (c^a, c^b) plane, the equation $F(c_{s'}^a, c_{s'}^b) = H$ defines a family of curves κ_H indexed by the constant H . Any of these curves splits the plane into two areas. Within one of them, the incentive-compatible set K_H , all contracts induce the high effort level e_h , whereas e_l is optimal in the complement \bar{K}_H .

As it will become clear later on, the shape of the κ_H curves just defined plays a key role in the derivation of our result, because it characterizes of the set of incentive compatible allocations for both deterministic and randomized contracts. In the standard case where utility is separable in effort and consumption, κ_H is increasing. We now show that this is no longer the case under Assumption D: in the (c^a, c^b) plane, κ_H reaches a global maximum and decreases afterwards. Furthermore, for any given constant H the agent's expected utility is bounded over the the incentive-compatible set K_H - a fact that will be directly linked to the existence of contracts generating positive profit. Formally:

Lemma 2.1. *Assume Assumption D holds true. Then :*

- κ_H is a continuously differentiable curve in the (c^a, c^b) plane
- For (c^a, c^b) small enough, κ_H is increasing, with a slope smaller than 1.
- There exist a point $(\bar{c}^a(H), \bar{c}^b(H))$ on κ_H in the (c^a, c^b) plane where c^b is locally maximum, and such that, for any $c^a > \bar{c}^a(H)$, κ_H is decreasing in the (c^a, c^b) plane. Moreover, the values of c^a at which the maxima are reached are the same for all curves κ_H .
- For $H' > H$, $\kappa_{H'}$ is below κ_H in the (c^a, c^b) plane; i.e., if $(\bar{c}^a, c^b) \in \kappa_H$ and $(\bar{c}^a, c'^b) \in \kappa_{H'}$ then $c'^b < c^b$

Proof: See Appendix

Insert here Figure 2

An illustration of this result is provided in Figure 2 for the simple case of a unique global maximum: in the (c^a, c^b) plane, κ_H first increases, reaches a maximum, then declines. The interpretation of the maximum level for c^b is straightforward : a contract providing a consumption larger than this maximum when the low production level is realized cannot induce the maximum effort level, *whatever* consumption may be in the other state. In other words, a 'too generous' allocation of consumption in the bad state cannot be offset by an even more generous provision in the good state. The intuition is that increasing c^a has two effects. On the one hand, it broadens the difference between consumption in the good and in the bad states. This has good incentives properties, since choosing the high effort increases the probability of receiving this difference. Note, however, that this effect is related to marginal utility of income, so that its magnitude decreases with wealth. On the other hand, a higher c^a increases total (expected) consumption, which, from Assumption D, raises marginal disutility of effort. This clearly reduces incentives. For c^a large enough, the latter effect may well dominate; it actually does under Assumption D.

A crucial consequence of Lemma 2 is that the agent's utility is bounded over the incentive compatible set K_H , and reaches its maximal, as stated in the following Proposition:

Proposition 2.2. *There exists a finite bundle $c^*(H) = (c^{a*}(H), c^{b*}(H))$ that maximizes the agent's expected utility under high effort on the set K_H .*

Proof: See Appendix

As it can be seen on Figure 2, increasing the agent's utility beyond its level at $c^*(H)$ would require a larger c^a, c^b or both, but this cannot be compatible with the incentive constraint. An important consequence of this result is that the agent's preferences on the contract space exhibit a *satiation* property: no increase in consumption can improve the agent's welfare.

3. Competitive equilibria: the strategic point of view

We now consider our model from a strategic point of view. Competition is represented as a two-stage game. At stage one, each principal offers one or several contracts. At stage two, each agent selects a contract and chooses an effort; then the state of the world is revealed and consumption takes place. An equilibrium is then defined as a subgame perfect equilibrium of this game. We will call it a Bertrand-Rothschild-Stiglitz (BRS) equilibrium because a feasible, incentive-compatible contract is a Nash equilibrium outcome of our game if and only if it satisfies Rothschild and Stiglitz's (1976) condition of robustness to the introduction of additional profitable contracts.

3.1. Preliminary characterization

A first step toward the characterization of BRS equilibria is provided by the following result:

Lemma 3.1. *At any BRS equilibrium, agents' ex ante utility is maximized under incentive compatibility and non negative profit (feasibility) constraints.*

Proof. Assume not. Then there exists some contract that is incentive compatible, makes non negative profit, and provides the agent with strictly higher expected utility than the equilibrium one. Decrease the agent's consumption in the state $(1, b)$ by some 'small' amount $\epsilon > 0$. The contract thus obtained is still incentive compatible and is still preferred by the agent to the equilibrium one, but makes strictly positive profits - a contradiction with the equilibrium condition. ■

It should be noted that this Lemma applies not only to general contract spaces, but also to particular subclasses of contracts (e.g., deterministic ones); in both cases, the BRS equilibrium generates the maximum level of utility attainable over the set of contracts at stake. A first and simple consequence is the following:

Corollary 3.2. *A BRS equilibrium always exists. Moreover, every competitive equilibrium is constrained Pareto optimal.*

The proof of this result is immediate since the agent's program coincide with the definition of the (constrained) efficient outcome preferred by the agents; again, the Corollary holds for any contract space. Note, however, that the constrained Pareto efficiency of BRS equilibria is not a

very robust result. For instance, it is well known that it does not hold under adverse selection.⁶ Even in the moral hazard context, the conclusion strongly depends on the assumption that only one good is consumed in the economy, as proved in Bennardo (1997).

We shall now derive our first result, namely that the BRS equilibrium may entail positive profits. For the sake of clarity, it is useful to convey the basic intuition in the simpler case where only *deterministic* contracts can be offered, and then to analyze how the argument is affected by the introduction of randomization.

3.2. The basic intuition: competition with deterministic contracts

Assume that the equilibrium deterministic contract involves the high effort level e_h ; and let $(\hat{c}_s^a, \hat{c}_s^b)$ denote the corresponding equilibrium consumptions in state s . Defining \hat{H}_s by $F(\hat{c}_s^a, \hat{c}_s^b) = \hat{H}_s$, incentive compatibility requires that $\lambda \hat{H}_1 + (1 - \lambda) \hat{H}_2 = 0$;⁷ hence for each state s the equilibrium allocation $(\hat{c}_s^a, \hat{c}_s^b)$ must, in the (c_s^a, c_s^b) plane, be located on the curve $\kappa_{\hat{H}_s}$. On this curve, from Lemma 3.1, the agent's expected utility is maximum at $(c^{a*}(\hat{H}_s), c^{b*}(\hat{H}_s))$. The question, now, is whether this optimal bundle is feasible given the economy's production in that state, \bar{Y}_s . If it is not, then at equilibrium the feasibility constraint is binding, i.e.

$$P\hat{c}_s^a + (1 - P)\hat{c}_s^b = \bar{Y}_s$$

In that case, the equilibrium profit is zero, and the optimal consumption is located at an intersection of the feasibility constraint above and the curve $\kappa_{\hat{H}_s}$. This case is illustrated in Figure 3a. In the alternative situation, illustrated in Figure 3b, the expected production (conditional on the high effort being exerted) exceeds the expected consumption corresponding to the optimal bundle $c^*(\hat{H}_s)$ - formally:

$$Pc^{a*}(\hat{H}_s) + (1 - P)c^{b*}(\hat{H}_s) < \bar{Y}_s \quad (Z_s)$$

Then from Lemma 3.1 the equilibrium allocation $(\hat{c}_s^a, \hat{c}_s^b)$ must coincide with the optimal one $(c^{a*}(\hat{H}_s), c^{b*}(\hat{H}_s))$, and (Z_s) implies that the agents receive less than the total amount produced; the difference is the principal's profit *in the aggregate state s* . This argument suggests that when aggregate productions, \bar{Y}_1 and \bar{Y}_2 , are both 'large enough', then profit should be positive in at least one state. This intuition turns out to be correct, although its proof is somewhat tricky because the optimal incentive level \hat{H}_s is endogenous and depends in particular on aggregate productions. Formally, the following result holds:

Proposition 3.3. *Under Assumption D, there exists an open set of parameters for which the BRS equilibrium $(e_h; \hat{c}_1^a, \hat{c}_1^b; \hat{c}_2^a, \hat{c}_2^b)$ requires the high effort level and entails positive profits in both states.*

Proof: see Appendix

⁶See Chassagnon and Chiappori (2002) for a recent discussion.

⁷As it is well-known, the incentive compatibility constraint must be binding at equilibrium.

Insert here Figure 3

The economic intuition for this result is a clear consequence of what has been said above. In the standard case (say, with separable preferences), equilibrium profits are zero because of undercutting. Should profits be positive, a new entrant could attract all consumers and make a positive profit. The trick - a direct generalization of undercutting - is to propose a contract corresponding to an infinitesimal move along the Pareto frontier, in the direction of increased consumer's welfare (and decreased profits).

This argument, however, requires the possibility of a trade-off between welfare and profit. In our case, such a trade-off may not exist, since consumer's welfare cannot be increased beyond the level reached at $(c^{a*}(\hat{H}_s), c^{b*}(\hat{H}_s))$. Then undercutting is just not feasible. In other words, because of competitive pressure, principals would be willing to give away part of their profits to the agents, in order to attract a larger number of them. However, this is not possible here, because the incentive constraint results in an upper bound on the agents' (contingent) wealth: making them richer would kill incentives to work, which would in turn result in much lower welfare for all.

Finally, the positive profit situation has an interesting interpretation in terms of the shape of the Pareto frontier. Assume for a moment that $\lambda = 0$ (i.e., there is only one aggregate state), and that parameters are such that inducing the high effort is optimal; the incentive constraint is now $F(c_{s'}^a, c_{s'}^b) = 0$. Also, assume that the consumption bundle $(c^{a*}(0), c^{b*}(0))$ is feasible, and let $u^*(0) = Pu(c^{a*}(0), e_h) + (1 - P)u(c^{b*}(0), e_h)$ denote the corresponding expected utility. While the firm gets a profit equal to

$$\pi^*(0) = \bar{Y} - Pc^{a*}(0) - (1 - P)c^{b*}(0) > 0$$

it is impossible for the agent to reach an expected utility level greater than $u^*(0)$. Starting from $(u^*(0), \pi^*(0))$, it is impossible to change the contract so as to decrease the principal's profit *and* increase the agent's expected utility: no matter the contract, the agent's welfare cannot go beyond $u^*(0)$. In a standard representation of the Pareto frontier where the horizontal (resp. vertical) axis represents the agent's expected utility (resp. the principal's profit), the frontier has a *vertical* segment⁸. This fact is reminiscent of another, more well-known result already mentioned by Grossmann and Hart (1983): in a standard principal-agent model with moral hazard and non separable utility, the agent's participation constraint may not be binding - which would correspond to an *horizontal* segment on the Pareto frontier. Note, however, that the two results are unrelated. In particular, one can exhibit utility functions that satisfy one property but not the other.

3.3. Competition with unrestricted contract space

Although the previous argument provides the basic intuition underlying our results, it still relies on the assumption that only deterministic contracts are implementable. It is however well-known in contract theory that random contracts may be mutually profitable for principals and

⁸To be precise, only the upper limit of the segment, A, belongs to the Pareto frontier; all other points are Pareto-dominated by A. Also, note that reaching the points below A requires free disposal, since it must be possible to withdraw resources from the principal without increasing the agent's wealth.

agents in the presence of moral hazard (see Prescott and Townsend (1984), Gjesdal (1985), and Arnott-Stiglitz (1988)). We now characterize competitive equilibria in the space of randomized contracts.

First, we can use Lemma 3.1 to characterize the equilibrium as an efficient allocation. Formally, if the contract $\Gamma = (\alpha, \mu_{h,1}^a, \mu_{h,1}^b, \mu_{l,1}^a, \mu_{l,1}^b, \mu_{h,2}^a, \mu_{h,2}^b, \mu_{l,2}^a, \mu_{l,2}^b)$ is an equilibrium, it must solve the following program:

$$\max_{\Gamma} \sum_{s=1}^2 (PEU_{h,s}^a + (1-P)EU_{h,s}^b)$$

under the constraints

$$\alpha (PE_{h,s}^a + (1-P)E_{h,s}^b) + (1-\alpha) (pE_{l,s}^a + (1-p)E_{l,s}^b) \leq \alpha \bar{Y}_s + (1-\alpha) \bar{y}_s, \quad s = 1, 2 \quad (3.1)$$

and

$$\sum_{s=1}^2 [(PEU_{h,s}^a + (1-P)EU_{h,s}^b) - (pEU_{l,s}^a + (1-p)EU_{l,s}^b)] \geq 0 \quad (3.2)$$

This program, however, is not particularly simple or attractive. We shall now discuss its main properties.

Ex post randomization We indicated above that, in its most general form, the contract Γ entails both ex ante and ex post randomization. As it turns out, these two types of randomization play completely different roles in our context. Specifically, ex post randomization plays no specific role in our setting, and is actually useless under a mild strengthening of Assumption D, as stated by the following result:

Proposition 3.4. *At any BRS equilibrium, whenever the low effort is exerted the contract entails no ex post randomization. Moreover, if the concavity of the utility function decreases with effort, i.e.*

$$\frac{\partial^3 u}{\partial c^2 \partial e} < 0 \quad (\text{DC})$$

then the equilibrium contract never entails ex post randomization

The first part of the Proposition is obvious: whenever the agent has chosen the low effort level, no incentive problem arises, and increasing the risk born by the (risk averse) agent can only reduce efficiency. The second part is less obvious, and was first demonstrated by Arnott and Stiglitz (1988). The intuition for ex post randomization is that in some cases it can help relaxing the incentive constraints. Assume, for instance, that $u(c, e)$ is such that the agent is risk neutral when $e = e_h$, whereas he is risk averse for $e = e_l$. Then ex post randomization can be used as an effective incentive device : it increases the relative cost of the low effort without changing either expected profits or the agent's expected utility conditional on e_h . In fact, in this extreme case, (sufficient) ex post randomization can implement the first best. This, however, requires particular properties of the third cross derivative of the utility function (i.e., the way $\partial^2 u / \partial c^2$ changes with effort) that are ruled out by assumption (DC). Technically, (DC) implies in particular that risk aversion increases with effort (decreases with leisure)⁹, which rules out

⁹This effect is somewhat natural. It has been argued, for instance, that more leisure provides more flexibility to cope with adverse income shocks (say, by increasing labor supply). See for instance Gollier (1999)

the 'incentive' role of ex post randomization.

Also, it should be noted that (DC) is fully consistent with Assumption D. Indeed, the latter states that the increase in marginal utility of income due to an increase in leisure is negligible when consumption is small, and dominant when it is large, while DC requires that it increases monotonically in-between.

Ex ante randomization On the contrary, ex ante randomization turns out to play a key role in our context. To see why, start from the situation described above, where the equilibrium with deterministic contracts entails some positive profit π_s in both states $s = 1, 2$. Now, a principal can successfully propose the following, ex ante randomized contract: with probability α (where α is smaller than but very close to 1), the agent receives the deterministic equilibrium contract; with probability $1 - \alpha$, he receives an alternative contract, entailing low effort and some consumption level C_s independent of the individual shock (but contingent on the aggregate state), with

$$C_s = \bar{y}_s + \frac{\alpha}{1 - \alpha} \pi_s - \frac{\epsilon}{1 - \alpha}$$

This contract is feasible, and entails a profit equal to ϵ . For $1 - \alpha$ and ϵ small enough, C_s is arbitrarily large, and $u(C_s, e_l)$ is larger than the utility generated at the deterministic equilibrium. It follows that the new contract is strictly preferred to the deterministic one, hence will attract all agents, while still generating positive profit - a contradiction with the definition of an equilibrium. This shows that whenever the deterministic contract entails positive profit in *both* aggregate states, the equilibrium contract cannot be deterministic, but *must* entail randomization.

In the randomized contract just described, each agent can either be 'lucky', in the sense that he receives a no effort, high consumption contract; or he can get the standard, deterministic contract, in which case he faces the same incentives to work as before. The intuition is simple. Because of competitive pressures, profits should be dissipated in wage increases. This, however, cannot be done in a deterministic way, because it would destroy incentives to work (that's the previous result). Ex ante randomization solves the problem by exploiting the separability of expected utility across states; the idea being to distribute the profit in a contingent way, i.e., in some states of the world (as defined by the ex ante random device). Technically, randomization creates a (low probability) contingency where agents are prescribed to exert low effort, so that the incentive constraint is not binding and profits can be freely distributed. From the agent's ex ante viewpoint, this randomized contract is better, because of the positive (even if small) probability of being lucky.

It must be stressed that, at equilibrium, some agents will work ($e = e_h$) whereas other will not ($e = e_l$). This has some cost in term of total production, since the latter will be equal to $(1 - \alpha) \bar{y}_s + \alpha \bar{Y}_s$, which is strictly less than \bar{Y}_s . However, the loss may be very small (since α can be arbitrarily close to one). On a more theoretical side, it is important to note that the coexistence of different equilibrium effort levels *does not* require the corresponding utilities to be equal. On the contrary, ex ante randomization is useful precisely because it allows the coexistence of different and non equivalent (in terms of the agent's utility) effort levels. This intuition reveals crucial in more general contexts, and especially when different commodities are produced (see Benmardo 1997).

The main result Finally, the intuition given above for the role of randomization assumes that (deterministic) equilibrium profit is positive in *both* states. We now concentrate on the more interesting case where, in the absence of randomization, the equilibrium effort is e_h and the optimal contract entails positive profit in the good state (state 1) but zero profits in the alternative state 2. This will typically be the case when \bar{Y}_1 is large while \bar{Y}_2 is small (see Appendix D for a formal proof). We also assume that \bar{y}_2 is much smaller than \bar{y}_1 - i.e., although prevention is always good, it becomes absolutely crucial in case of a hurricane.

In such a context, ex ante randomization can still be used to redistribute the profit that is generated in the good state of the world, but this can no longer be done at no cost. The crucial point is that, given the prevention technology, effort must be chosen (possibly in a random way) *before* the aggregate state of the world is known. If a fraction $1 - \alpha$ of the agents receive the low effort contract, then with probability $(1 - \lambda)$ the average production of these agents will be at the very low level \bar{y}_2 . Clearly, this risk is non insurable, since it directly reflects aggregate uncertainty (technically, the resource constraint must be satisfied in each aggregate state). As a result, redistribution of profit in state 1 is realized at a cost - namely, less total resources are available in the bad state 2, where they are most needed. The question, now, is whether the benefit of randomization - the ability to redistribute profit in the good state - is worth the cost - namely, reducing everybody's consumption in the bad state, although all agents would be willing to transfer resources *to* that state.

The answer ultimately depends on the properties of the utility function, as stated the next proposition.

Proposition 3.5. *Assume that Assumptions D and DC are satisfied. If utility with the low effort is unbounded above :*

$$\lim_{x \rightarrow +\infty} u(x, e_l) = +\infty$$

then the BRS equilibrium contract always entails ex ante randomization and zero profit. Conversely, if utility is bounded above :

$$\lim_{x \rightarrow +\infty} u(x, e_l) < +\infty$$

then there exist an open range of parameters for which any BRS equilibrium contract is deterministic and equilibrium profits are positive in state 1.

Proof: see Appendix

The intuition is that a very small $1 - \alpha$ leads to practically infinite consumption in the most lucky case (low effort contract *and* good state of the world), and this perspective, however unlikely, is sufficient to decide in favor of randomization whenever utility is *unbounded*. Conversely, if utility is bounded, randomization may, depending on the parameters, become too costly. Then the equilibrium contract is the deterministic one, and entails positive profit. Note that in our context, a bounded utility has the following interpretation : there exists some consumption level \bar{C} , some (very small) probability $1 - \bar{\alpha}$ and some (very low) consumption level \bar{A} such that a lottery paying \bar{A} with probability $\bar{\alpha}$ and some arbitrarily large amount B with probability $1 - \bar{\alpha}$

will never be preferred to \bar{C} , *whatever* B . Or, to put it differently: paying an infinite amount with a low probability is not sufficient to compensate for a large (but finite) and very likely loss.

From now on, our analysis will mainly focus on the situation just described where the equilibrium contract is deterministic and generates positive profit in the aggregate state 1 and zero profit in the aggregate state 2; let $\hat{\gamma} = (e_h; \hat{c}_1^a, \hat{c}_1^b; \hat{c}_2^a, \hat{c}_2^b)$ denote the corresponding equilibrium contract.

4. A general equilibrium viewpoint

4.1. The framework

In the previous sections, we adopted a strategic point of view, where a finite number of principals simultaneously propose contracts. This setting explicitly takes into account the strategic nature of competition, and relies upon an equilibrium concept borrowed from game theory (non cooperative Nash in our case). We now consider an alternative, 'general equilibrium' perspective where agents are price takers in a general commodity space. The approach is borrowed from the seminal work of Prescott and Townsend (1984a, b). We assume that the economy consists of two types of agents, households and intermediaries - the latter corresponding to the principals in the BRS context.¹⁰ We use the 'Walrasian' concept of equilibrium under asymmetric information defined by Prescott and Townsend. In particular, we use the modeling technology provided in their papers. No restriction is placed on the space of enforceable contracts; i.e., we consider the pure moral hazard case defined above, where the set of available (random) contracts allows to exploit all potential gains from trade.

In this setting, a contract is a distribution on the space of allowable consumptions, productions and effort. Following Prescott Townsend, we make the simplifying assumption that the set of possible consumption levels and, consequently, the set of allowable contractual contingent payoffs are finite. Technically, a contract is a vector x of the simplex:

$$x = (\dots, P(c, y_1^a, e_i), P(c, y_1^b, e_i), \dots, P(c, y_2^a, e_i), P(c, y_2^b, e_i), \dots)$$

where c varies within the finite set of feasible consumption levels, and $i \in \{l, h\}$ indexes the effort level. Specific constraints have to be satisfied, that reflect the nature of the available technology (e.g., the marginal probability of producing y^x in the aggregate state s conditional on supplying effort e_i is given by nature) and the probability nature of the commodity (e.g., the sum of conditional probabilities must equal one)¹¹. Also, the following notation will reveal convenient:

$$x = (x_1, x_2)$$

where

$$x_s = (\dots, P(c, y_s^a, e_l), P(c, y_s^b, e_l), P(c, y_s^a, e_h), P(c, y_s^b, e_h), \dots)$$

¹⁰An alternative and equivalent interpretation assumes away intermediaries. The agents in the economy can create insurance firms in the usual way, i.e. by allocating property rights (shares) among them, and profits (if any) are distributed among shareholders.

¹¹These constraints are exactly analogous to those in Prescott-Townsend; for brevity, we do not restate them here.

denotes the restriction of x to the aggregate state s (i.e., the distribution of consumption, production and effort when state s is realized).

Each commodity (contract) x has a price. Prices of commodities are linear with respect to the probabilities $P(c, y_s^x, e_i)$; i.e., if π_{c, y_1^x, e_i} is the price of a contract paying c contingent on e_i and on the realization of the endowment y_s^x , and zero otherwise, the value of the commodity x is then given by the scalar product

$$\Pi \cdot x$$

where

$$\Pi = (\dots, \pi_{c, y_1^a, e_i}, \pi_{c, y_1^b, e_i}, \dots, \pi_{c, y_2^a, e_i}, \pi_{c, y_2^b, e_i}, \dots)$$

Again, it is convenient to use the notation

$$\Pi = (\Pi_1, \Pi_2)$$

so that

$$\Pi \cdot x = \Pi_1 \cdot x_1 + \Pi_2 \cdot x_2$$

Contracts are offered by risk-neutral intermediaries that are price takers and supply contracts in a profit-maximizing way, subject to feasibility constraints of the form¹²

$$r_s \cdot x \leq 0$$

where

$$r_s = (\dots, c - y_s^a, c - y_s^b, c - y_s^a, c - y_s^b, \dots)$$

In words, this constraint requires that, in each aggregate state of the world, intermediaries cannot deliver more of the consumption good than what is actually produced.

Finally, households buy contracts at market price so as to maximize utility under budget constraint, and under the additional constraint that any purchased contract must belong to the consumption set. A specific feature of Prescott and Townsend's approach is the treatment of the incentive compatibility constraint. Prescott and Townsend suggest to include the constraint within the definition of the agent's consumption set (in other words, the agent cannot 'consume' a contract that is not incentive compatible).

A first consequence of this setting is the following:

Lemma 4.1. *At any equilibrium, profits must be zero. In particular, at any equilibrium with positive trade, the price vector Π_s must be either identically null or proportional to r_s .*

¹²If intermediaries are allowed to trade among them commodities across states of the world using Arrow-Debreu securities (say, selling one unit of the commodity in state 1 in exchange for k units of the commodity in state 2), then only one constraint is needed. However, since such trades cannot take place at equilibrium, the equilibrium will actually satisfy the two constraints $r_s x_s = 0$. Our results hold under both interpretations.

Proof. Assume that, at equilibrium, a contract generates positive profits. With a linear technology, its supply would then be infinite, and market clearing cannot obtain, a contradiction.

Assume, now, that Π_s is not identically null, and consider the two half spaces $E_{\Pi_s} = \{x / \Pi_s x > 0\}$ and $E_{r_s} = \{x / r_s x \leq 0\}$. If Π_s is not proportional to r_s , then the two spaces have a non empty intersection, which is a convex cone. If this cone has an empty intersection with the positive orthant, no feasible contract generates a non-negative profit; then no trade takes place at equilibrium. If, on the other hand, the intersection of the cone with the positive orthant is not empty, there exists a contract that is feasible and generates a strictly positive profit, a contradiction. ■

The intuition is that if profits are positive, then any producer is willing to supply an infinite amount of the corresponding contracts. Since demand is necessarily finite, there must be excess supply and markets do not clear. It follows that the Prescott-Townsend approach explicitly requires either zero profit or zero prices as an equilibrium condition.

4.2. Existence of an equilibrium

We now concentrate on the existence of an equilibrium. A first remark is that, in our setting, a zero-price equilibrium cannot exist. Indeed, prices are used to decentralize the feasibility constraints at the agents' level. If they are zero, there is a positive demand for contracts that are not feasible and an equilibrium cannot exist. Formally:

Lemma 4.2. *At any equilibrium, neither Π_1 nor Π_2 can equal the null vector.*

Proof. Assume that Π_s is identically null. Consider a contract $x = (x_1, x_2)$ such that x_s requires the low effort level and provides some very high level of consumption C with probability one, whatever the idiosyncratic outcome. This contract is incentive-compatible, hence belongs to the agents's consumption set. Also, buying this contract cannot be incompatible with the budget constraint, since its price is zero whatever C . For C large enough, it is preferred to any incentive-compatible contract inducing high effort (remember the utility provided by such contracts is bounded). At zero prices, there will be a positive demand for some contracts of this kind. However, they cannot be supplied, since they violate the feasibility constraint. ■

We can now show our main result.

Proposition 4.3. *Consider the set \mathcal{E} of economies in which all Nash equilibrium contract entail positive profit in the aggregate state 1 and zero profit in the aggregate state 2. No Walrasian equilibrium exists in any economy belonging to the set \mathcal{E} .*

Proof. From Lemmas 4.1 and 4.2, at any equilibrium with non zero trade, the price vector must be proportional to (r_1, r_2) . Equilibrium prices are thus defined up to a multiplicative constant. After normalization, the price of any contract (including $\hat{\gamma}$) must be such that the producer makes zero profit. It follows that, whenever the conditions for Proposition 4.3 are satisfied, the

price of $\hat{\gamma}$ must be negative (i.e., the agent receives money when purchasing the contract) and equal to minus the profit made on $\hat{\gamma}$.

Consider, now, the agent's perspective. If he chooses $\hat{\gamma}$, his budget constraint is not binding. So he can afford a contract that is slightly more expensive. For instance, he can buy some contract that entails, in state 2, a consumption bundle (c_2^a, c_2^b) located on the incentive frontier beyond $(\hat{c}_2^a, \hat{c}_2^b)$ (i.e., such that $c_2^a > \hat{c}_2^a, c_2^b > \hat{c}_2^b$ - remember such a contract exists by construction). Since Π_2 is proportional to r_2 , the difference in price between the two contracts is proportional to $c_2^a - \hat{c}_2^a + c_2^b - \hat{c}_2^b$, hence can be made smaller than the profit made on $\hat{\gamma}$. This means that a contract of the form $(e_h; \hat{c}_1^a, \hat{c}_1^b; c_2^a, c_2^b)$ belongs to the agent's budget and consumption sets, and is strictly preferred to $\hat{\gamma}$. At any equilibrium, the demand of a contract of this sort will be positive, whereas its supply must be zero since it violates the feasibility constraint in state 2 - a contradiction. ■

Two comments are in order at this point. Firstly, the previous Proposition provides an example of an economy where no Walrasian equilibrium a la Prescott-Townsend exists. This conclusion may seem surprising, because of the close relationship between the Prescott-Townsend approach and standard general equilibrium (GE) theory. A Prescott-Townsend economy is but a particular case of a production economy with constant return to scale technologies; the existence of an equilibrium obtains as a consequence of the standard existence theorem in GE. However, the analogy between moral hazard economies and standard GE models requires a reinterpretation of such basic concepts as prices or commodities. Through the reinterpretation process, some assumptions that seem totally natural in the standard framework may become much more restrictive, and may actually fail to be satisfied in many contexts. For instance, monotonicity of preferences, a natural hypothesis in standard consumer theory, does not obtain in the new setting. As it is well known, the existence theorem in GE (see Debreu 1959 or Arrow and Hahn 1971) holds without monotonicity of preferences, but it then requires local non satiation and resource-relatedness. With deterministic contracts only, local non satiation is not satisfied in our context (utility is locally satiated at $\hat{\gamma}$). When lottery contracts are enforceable, local non satiation is satisfied, but resource relatedness is not.

Secondly, the result has an important implication for the links between Walrasian and Bertrand equilibria¹³. It can be easily showed that the two sets coincide whenever, at each BRS equilibrium, firms make zero profits. For this class of economies, the Walrasian equilibrium concept can be interpreted as a reduced equilibrium form of a Bertrand-Nash game. In the alternative case where BRS equilibria entail positive profits, a completely different result obtains. BRS and Walrasian equilibria do not coincide; Walrasian equilibria do not exist in that case, whereas BRS equilibria always exist and are constrained efficient. The key intuition is that positive profits contradicts market clearing. Assume that the BRS equilibrium contract entails positive profit (which implies in particular that the supply of these contracts is rationed at equilibrium). By definition, this contract is preferred by agents to any alternative, feasible one - including those entailing zero profit. As argued above, market clearing requires that only contracts yielding zero profit can be exchanged. But any principal could then profitably deviate by proposing a positive profit contract that is preferred by agents as well. In other words, the Walrasian equilibrium concept does not satisfy the basic condition of absence of profitable devi-

¹³See Bennardo (1997) for a detailed analysis of the links between Walrasian and Bertrand equilibria within a general, multi-commodity setting.

ation defining non cooperative Nash equilibria. What market clearing conditions do is imposing ad hoc constraints on the set of mutually convenient trades that agents can undertake (in our case, an unfortunate consequence is that no equilibrium satisfying these constraints exist). In this sense, the Walrasian-like equilibrium notion can be considered as a quite unsatisfactory equilibrium concept.

Finally, a different modeling choice can be made; namely, the incentive constraint can be imposed on the set of contracts offered by intermediaries instead of being introduced through the definition of the consumption set. Then Walrasian equilibria always exists, even when Bertrand equilibrium contracts generate positive profit; however, Walrasian equilibria are typically constrained inefficient (because they entail low effort). In this case, imposing market clearing typically results in welfare losses.

5. Related literature

Our findings are related to three branches of the literature on asymmetric information: general equilibrium approach, strategic competition with decentralized trading and mechanism design.

Helpman and Laffont (1975), Prescott and Townsend (1984), Gale (1996) and Bisin and Gottardi (1998) are among the main contributors to the literature on general equilibrium with asymmetric information. The relationships between our model and Prescott-Townsend have been discussed in the previous section. In another seminal contribution, Helpman and Laffont (1975) characterize Walrasian equilibria with moral hazard and anonymous trades and show that Walrasian equilibria may not exist because of the discontinuity of the assets' demand correspondence generated by moral hazard under linear prices; when they exist, they are generically constrained inefficient. Bisin and Gottardi (1997) generalize these results to a larger class of economies and also demonstrate that existence can be restored by imposing (exogenous) bounds on agents' trades. Such bounds on trades are suggestive of rationing phenomena. In a Walrasian setting, however, trades' restrictions ensuring existence have to be exogenously imposed (whether they can *endogenously* emerge from market interactions remain unclear), furthermore, we rule out these problems by assuming non-anonymous, exclusive trades. Finally, Gale (1996) studies competition with adverse selection in a non strategic-setting where markets are assumed to clear through rationing. Although our approach and that of Gale are to some extent complementary, the hypotheses on the trading mechanisms are different. In the equilibrium notion used by Gale, rationing is the only market clearing mechanism available (markets clear by rationing instead of prices), while in our context rationing emerge only if it allows to exploit gains from trade that cannot be achieved at market clearing prices.

Regarding the game-theoretic literatures on decentralized trading and mechanism design, our findings are related to three types of results in these literature. One is the existence of Nash equilibria with positive profit in insurance economies with anonymous trades which has been studied in several papers by Arnott and Stiglitz (1987, 1988, 1993), and recently by Bisin and Guaitoli (1999). A second related result is the existence of equilibria with rationing, which is a standard result of the efficiency wage/credit rationing literature. Finally, the optimality of budget breaking mechanisms for environments with asymmetric information and externalities has been studied in the mechanism design literature and in the literature on moral hazard in teams. The first two classes of models (moral hazard with anonymous trades and efficiency wage) depart explicitly from the pure moral hazard case we consider in the present paper, in

the sense that they restrict their attention to a particular subset of the contract space (i.e., they assume away either exclusive or randomized contracts - or both), while the literature on budget breaking with moral hazard focus on situation where budget breaking only occurs out of the equilibrium path.

Specifically, in the single-good insurance setting with anonymous trades studied by Arnott-Stiglitz and Bisin-Guaitoli, positive profits are a consequence of a discontinuity in the demand for insurance. Such a discontinuity is in turn caused by the externality that appears in the absence of exclusive relationship - a point first emphasized by Helpman and Laffont (1974) (see also Hellwig (1983) and Segal (1999) on this point). The idea, here, is that when agents can freely buy insurance from different insurers, the profitability of the contract offered by each insurer depends (through incentives) on the total amount of insurance that the agent receives. In our setting, positive profits equilibria exist even though exclusive contracts rule out this externality. In addition, the positive profit result in Arnott-Stiglitz does not extend to economies where agents can trade randomized contracts.

The efficiency wages literature (see Shapiro-Stiglitz (1984) and Malcomson-MacLeod (1989) among others) is based on another type of externality - namely, the profitability of a labor contract depends on the unemployment rate of the economy (i.e. on the hiring decisions of all the other firms in the economy). The initial Shapiro-Stiglitz (1984) model also imposes exogenous bounds upon the 'punishment' that can be inflicted to an agent (typically, he can only be laid off at worse). Even under these restrictions, however, the rationing result is not robust to the introduction of non linear pricing devices such as bonding, as showed by Carmichael (1985). Malcomson and Mac-Leod, on the other hand, assume 'double moral hazard', in the sense that neither the agent's effort, nor the principal's measurement of output are verifiable. Under this assumption, the emergence of efficiency wages does not require limited liability; it stems from the restrictions on the set of enforceable contracts due to the principals' incentive constraints. However, rationing, again, is not robust to the introduction of random contracts, and specifically of tournaments, as argued by Malcomson (1984). Similar considerations also apply to the credit market model of Stiglitz and Weiss (1981). Finally, an interesting difference with our results is that in both efficiency wage and credit rationing models, the informed party (the agent) is rationed, while in our setting it is the non informed party (the principal) who is rationed and gets a rent in equilibrium.

Equilibrium positive profits are finally reminiscent of a general result in the mechanism design literature, namely that budget breaking mechanisms are generally welfare enhancing in the presence of externalities among privately informed agents. The fundamental contribution of Holmstrom (1982) showed that budget breaking mechanism generally improve incentives in team production with moral hazard. In this context, the non measurability of the individual contributions to the collective output (the non-separability of the production function) implies that feasible payoffs to one agents depend on the effort exerted by all others. In the Holmstrom model, however, the first best (or second best) allocation is achievable through a mechanism inducing principals' profit or losses out of equilibrium (i.e. if agents choose out of equilibrium actions-efforts) but such that principals get zero profit given the equilibrium actions. Under this mechanism, agents consume exactly what they produce. In our setting, in contrast, the technology is fully separable, and principals make positive profit in equilibrium.

6. Conclusion

Our paper studies an economy in which identical agents interact with linear principals in a moral hazard context with observable trades. This economy is successively considered from two viewpoints. We first adopt a 'strategic' perspective. The main result, here, is the existence of BRS equilibria in which the producers, albeit using a constant return to scale technology, make positive profits and have a rationed effective supply at equilibrium. When only deterministic contracts are enforceable, this conclusion obtains in the absence of aggregate uncertainty, under simple and fairly general assumptions on preferences and technologies; namely, the marginal utility of consumption must increase with leisure, especially when income is high. To the best of our knowledge, this result is new, and may shed a new light on the properties of strategic competition under moral hazard in situations where exogenous limitations prevent the use of randomized contracts.

A second result is that when randomized contracts are allowed for, positive profit equilibria may still appear, although not in the absence of aggregate uncertainty. Hence the general intuition that competition drives profits to zero may fail to hold in an asymmetric information context, even with totally unrestricted contract spaces.

These results have a counterpart in the alternative, general equilibrium perspective. Whenever BRS equilibria entail positive profits, the standard Walrasian concept, which imposes market clearing, cannot be interpreted as a reduced form equilibrium of a game in which agents set prices in order to exploit potential gains from trade. That is, a producer and a worker can profitably deviate from the Walrasian equilibrium by exchanging out-of-equilibrium contracts at relative prices that are not compatible with Walrasian market clearing.

One important related point is that, once rationing emerges, contracts must be allocated among the agents in the rationed side of the market (the supply side in our model) in a way that is not entirely determined by market prices. The precise form of the equilibrium will typically depend on the particular rationing scheme adopted - a choice that necessarily involves some degree of arbitrariness. However, while different rationing schemes may decentralize different allocations, the mere *existence* of rationing is purely endogenous. In situations similar to the one we just described, rationing is a necessary consequence of the assumed absence of profitable deviations that defines Nash equilibria. In that sense, it emerges as a by-product of individual rationality - and specifically of the agents' ability to implement any type of trades that turn out to be mutually beneficial. This feature is to be contrasted, in particular, with the so-called 'disequilibrium' approach initiated by Benassy (1975), Dreze (1982) and others, where rationing was a consequence of exogenously imposed restrictions on price movements (typically, prices were assumed rigid (at least in the short run), albeit prices changes would have improved welfare).

Two questions remain open at this stage. One concerns the robustness of our results. Although our setting is highly specific (two outcomes, two effort levels, identical agents) our results can be extended to more general environments (see Bennardo-Chiappori 1999). Similarly, the exclusivity assumption can be viewed as imposing an extreme form on non-linear pricing; however, it can be replaced with the milder hypothesis that contractual terms of trades are verifiable at sufficiently low costs and non linear (but not necessarily exclusive) contracts are enforceable. Finally, we do believe that our conclusions are by no means pathological, and that on the contrary they are likely to hold in many (sufficiently complex) second best contexts. A general property is that competition among principals results in equilibrium allocations which maxi-

mize agent's welfare subject to several constraints, including non negative profits and incentive compatibility conditions. However, there is simply no general argument suggesting that the non negative profit constraint should be binding in all cases.

The second fundamental issue concern the modeling of strategic interactions between different types of agents (possibly including large intermediaries) in a fully developed general equilibrium model where agents interact in several markets and several commodities are possibly produced and exchanged. A natural question, in this context, is whether a *reduced form* equilibrium concept can be defined and used. By reduced form, we mean a model that, while abstracting from the detailed game theoretic specification of the trading mechanisms, can select a set of allocations that coincides with the Nash equilibrium outcomes of a class of games in which maximizing agents and firms exploit arbitrage possibilities by choosing their trading plans. These questions are addressed in Bennardo (1997, 2001) where the Bayesian perfect equilibria of a multicommodity general equilibrium economy with moral hazard are characterized, and a (reduced form) equilibrium concept is also proposed, that replaces Walrasian market clearing with the Nash-type requirement of absence of profitable deviations.

Appendix

A. Proof of lemma 2.1

Note first that, from Assumption D,

$$\frac{\partial F}{\partial c^b} = (1 - P) \frac{\partial u(c^b, e_h)}{\partial c} - (1 - p) \frac{\partial u(c^b, e_l)}{\partial c} < 0 \quad (\text{A.1})$$

Consider the equation

$$F(c^a, c^b) = H \quad (\text{A.2})$$

From the implicit function theorem, (A.2) defines (when it admits solutions) c^b as a continuously differentiable function of c^a and H . A first consequence of (A.1) is that, for any given c^a , c^b is a decreasing function of H . Consider, now, c^b as a function of c^a for some given H ; this is the definition of the curve κ_H . Its slope is given by:

$$\left(\frac{\partial c^b}{\partial c^a} \right)_{\kappa} = - \frac{\partial F / \partial c^a}{\partial F / \partial c^b} = \frac{P \frac{\partial u(c^a, e_h)}{\partial c} - p \frac{\partial u(c^a, e_l)}{\partial c}}{- \left[(1 - P) \frac{\partial u(c^b, e_h)}{\partial c} - (1 - p) \frac{\partial u(c^b, e_l)}{\partial c} \right]}$$

In this fraction, the denominator is strictly positive. Consider, now, the numerator, say N . A first remark is that, when c^a is small, N is positive (this is a direct consequence of Assumption D). Then

$$\left(\frac{\partial c^b}{\partial c^a} \right)_{\kappa} = \frac{P \frac{\partial u(c^a, e_h)}{\partial c} - p \frac{\partial u(c^a, e_l)}{\partial c}}{P \frac{\partial u(c^b, e_h)}{\partial c} - p \frac{\partial u(c^b, e_l)}{\partial c} + \left[\frac{\partial u(c^b, e_l)}{\partial c} - \frac{\partial u(c^b, e_h)}{\partial c} \right]} < 1$$

Also, N can be written as

$$N = \frac{\partial u(c^a, e_l)}{\partial c} \left(P \frac{\partial u(c^a, e_h) / \partial c}{\partial u(c^a, e_l) / \partial c} - p \right)$$

From Assumption D, this implies that N is negative for c^a large enough; then κ_H is decreasing.

Thus κ_H is increasing for c^a small enough, decreasing for c^a large enough, and continuously differentiable in between. Hence it has one (global) maximum value (and possibly several local maxima); $\bar{c}^a(H)$ corresponds to the greatest local maximum.

Note, finally, that any local maximum must satisfy the first order condition .

$$P \frac{\partial u(c^a, e_h)}{\partial c} = p \frac{\partial u(c^a, e_l)}{\partial c}$$

This equation does not depend on H . Hence the values of c^a at which the maxima are reached are the same for all curves κ_H . Note that if the ratio $\frac{\partial u(c, e_h) / \partial c}{\partial u(c, e_l) / \partial c}$ is monotonic, there exists one local (and global) maximum only. ■

B. Proof of proposition 2.2

We first establish the following Lemma:

Lemma B.1. *For any given \bar{H} , there exist a value $v^a(\bar{H})$ such that for any $H \geq \bar{H}$, at any point (c^a, c^b) on the incentive frontier κ_H with $c^a > v^a(\bar{H})$, the slope of the agent's indifference curve through (c^a, c^b) , assuming effort e_h , is negative and smaller (in absolute value) than that of κ_H .*

Proof of the Lemma

The slope of the indifference curve is given by

$$\left(\frac{\partial c^b}{\partial c^a}\right)_u = \frac{-P \frac{\partial u(c^a, e_h)}{\partial c}}{(1-P) \frac{\partial u(c^b, e_h)}{\partial c}}$$

We thus want to show that, for c^a large enough,

$$\frac{-P \frac{\partial u(c^a, e_h)}{\partial c}}{(1-P) \frac{\partial u(c^b, e_h)}{\partial c}} \geq \frac{P \frac{\partial u(c^a, e_h)}{\partial c} - p \frac{\partial u(c^a, e_l)}{\partial c}}{-\left[(1-P) \frac{\partial u(c^b, e_h)}{\partial c} - (1-p) \frac{\partial u(c^b, e_l)}{\partial c}\right]}$$

Note, first, that in the above expression both denominators are positive. Hence it is equivalent to:

$$\frac{P \frac{\partial u(c^a, e_h)}{\partial c}}{1-P \frac{\partial u(c^a, e_l)}{\partial c}} \leq \frac{p \frac{\partial u(c^b, e_h)}{\partial c}}{1-p \frac{\partial u(c^b, e_l)}{\partial c}}$$

For c^a large enough, the left hand side term goes to zero, while the right hand side term is bounded away from zero. Indeed, from Lemma 2.1, $\kappa_{\bar{H}}$ is bounded above in the (c^a, c^b) plane, and this upper bound is also an upper bound for all κ_H with $H \geq \bar{H}$. Hence the conclusion. ■

Lemma B.1 essentially states that, from any point on κ_H where c^a is large enough (so that κ_H is decreasing in the (c^a, c^b) plane), an agent who is constrained to move along κ_H always prefers the direction in which c^a is reduced (and c^b is increased).

We can now conclude the proof of the Proposition. If K_H is bounded, the result is obvious. Assume it is not, and let $v^b(H)$ be such that $(v^a(H), v^b(H))$ belongs to κ_H . A consequence of Lemma B.1 is that the agent prefers $(v^a(H), v^b(H))$ to any (c^a, c^b) on κ_H such that $c^a > v^a(H)$. Then the contract that maximizes expected utility on K must belong to the set

$$K_H \cap [0, v^a(H)] \times [0, v^b(H)]$$

This set is compact, hence the result. A characterization of $c^*(H)$ is that the slope of the incentive constraint and that of the indifference curve are equal:

$$\frac{-P \frac{\partial u(c_H^*, e_h)}{\partial c}}{(1-P) \frac{\partial u(c_H^*, e_h)}{\partial c}} = \frac{P \frac{\partial u(c_H^*, e_h)}{\partial c} - p \frac{\partial u(c_H^*, e_l)}{\partial c}}{-\left[(1-P) \frac{\partial u(c_H^*, e_h)}{\partial c} - (1-p) \frac{\partial u(c_H^*, e_l)}{\partial c}\right]}$$

■

C. Proof of Proposition 3.3

Pick up two values H_1 and H_2 such that $\lambda H_1 + (1 - \lambda) H_2 = 0$. Define $v^a(H_s)$ as in Lemma B.1 above. Pick up two values Y_1 and Y_2 such that

$$Y_s > \max_{\substack{(c^a, c^b) \in \kappa_{H_s} \\ c^a \leq v^a(H_s)}} P c^a + (1 - P) c^b$$

Let $(\hat{c}_s^a, \hat{c}_s^b)$ denote the equilibrium consumptions in state s corresponding to (Y_1, Y_2) , and define \bar{H}_s by $F(\hat{c}_s^a, \hat{c}_s^b) = \bar{H}_s$, $s = 1, 2$. Since $\lambda H_1 + (1 - \lambda) H_2 = \lambda \bar{H}_1 + (1 - \lambda) \bar{H}_2 = 0$, we have that $H_s \leq \bar{H}_s$ for at least one s . Then $\kappa_{\bar{H}_s}$ is below κ_{H_s} , and since from Lemma B.1 $\hat{c}_s^a < v^a(H_s)$, it must be the case that

$$P \hat{c}_s^a + (1 - P) \hat{c}_s^b < \bar{Y}_s$$

■

D. Proof of Proposition 3.5

Take some given, 'large' Y_1 . We first show that for Y_2 and $(P - p)$ small enough, the deterministic equilibrium is such that profit is zero in the bad state 2. Assume not, and define \hat{H}_s as above. We show that by reducing \hat{H}_2 by dH , hence relaxing the incentive constraint in state 2, while increasing \hat{H}_1 by $\frac{1-\lambda}{\lambda} dH$, one can increase welfare, which contradicts the efficiency of the equilibrium. The change dH generates a gain dW_2 in state 2 and a loss $-dW_1$ in state 1. Define

$$W(H) = \max_{F(c^a, c^b) = H} P u(c^a, e_h) + (1 - P) u(c^b, e_h)$$

Then

$$W'(H) = \frac{(1 - P) \frac{\partial u(c^b, e_h)}{\partial c}}{(1 - P) \frac{\partial u(c^b, e_h)}{\partial c} - (1 - p) \frac{\partial u(c^b, e_l)}{\partial c}}$$

If Y_2 is small, so is c_2^b , and from Assumption D, the gain in state 2 is

$$dW_2 = \frac{(1 - P) \frac{\partial u(c^b, e_h)/\partial c}{\partial u(c^b, e_l)/\partial c}}{(1 - P) \frac{\partial u(c^b, e_h)/\partial c}{\partial u(c^b, e_l)/\partial c} - (1 - p)} \cdot dH \simeq \frac{1 - P}{P - p} \cdot dH$$

If Y_1 is large, on the other hand, c_1^b is bounded away from zero, so

$$\frac{\partial u(c^b, e_h)/\partial c}{\partial u(c^b, e_l)/\partial c} \leq \lambda$$

for some $\lambda < 1$ over a neighborhood of c^b . It follows that

$$-dW_1 = \frac{(1 - P) \frac{\partial u(c^b, e_h)/\partial c}{\partial u(c^b, e_l)/\partial c}}{(1 - p) - (1 - P) \frac{\partial u(c^b, e_h)/\partial c}{\partial u(c^b, e_l)/\partial c}} \cdot dH \leq \frac{1 - P}{(1 - p) - (1 - P) \lambda} \cdot dH$$

When $P - p$ is small enough, $dW_2 > -dW_1$, QED.

Starting from the deterministic equilibrium, we now consider first the consequences of the introduction of some infinitesimal ex ante randomization, as described in subsection 3.2. Switching from the deterministic equilibrium to a randomized contract of this kind has a benefit dB_1 in state 1 and a cost $-dB_2$ in state 2; for $1 - \alpha$ small enough, both the benefit and the cost are infinitesimal, and proportional to $1 - \alpha$ at the first order. We now compare the first order benefit and cost. In state 1, the benefit comes from the fact that, with probability $1 - \alpha$, the agent gets a large consumption C with no effort instead of the deterministic equilibrium contract $(e_h, \hat{c}_1^a, \hat{c}_1^b)$; hence

$$dB_1 = (1 - \alpha) \cdot \Delta u$$

where

$$\Delta u = u(C, e_l) - (Pu(\hat{c}_1^a, e_h) + (1 - P)u(\hat{c}_1^b, e_h))$$

In state 2, the total amount of resources available is decreased by $(1 - \alpha)(\bar{Y}_2 - \bar{y}_2)$; $-dB_2$ can be approximated by a product of the form $(1 - \alpha)(\bar{Y}_2 - \bar{y}_2)K$, where K is related to the expected marginal utility of income in that state (technically, K is larger than the lower bound of marginal utility of income in some neighborhood of the contract in state 2). Note that if \bar{Y}_2 is small enough, K is arbitrarily large.

When α goes to 1, the consumption level C goes to infinity. If u is unbounded, Δu goes to infinity, and so does dB_1 , whereas dB_2 remains bounded; the first order effect of the randomization is positive. If, on the other hand, u is bounded by some \bar{u} , then Δu is bounded by $\bar{u} - (Pu(\hat{c}_1^a, e_h) + (1 - P)u(\hat{c}_1^b, e_h))$. For \bar{Y}_2 small enough, the marginal utility of income term K is arbitrarily large, and the first order effect of the randomization is negative.

Finally, consider the case of a non infinitesimal ex ante randomization. By the same argument used above, the gain it generates is bounded above since Δu is bounded. On the other hand the cost of a non-infinitesimal randomization must be strictly larger than that of an infinitesimal one, as it is immediate to verify, and becomes arbitrarily large for \bar{Y}_2 small enough. Hence we can conclude that no ex ante randomization can be profitable. ■

References

Arnott, R., and Stiglitz J.E., "Equilibrium in Competitive Markets with Moral Hazard", discussion paper 4, John M. Olin Program for the Study of Economic Organization and Public Policy. Princeton University, 1987.

Arnott R., and Stiglitz, J.E., "The Basic Analytics of Moral Hazard", *Scandinavian Journal of Economics*, 90, 1988, pp. 383-413.

Arnott, R., and Stiglitz, J.E., "Randomization with Asymmetric Information", *Rand Journal of Economics*, V. 19, 1988, pp. 1-24.

Arnott, R. and Stiglitz, J.E., "Price Equilibrium Efficiency, and Decentralizability in Insurance Markets with Moral Hazard", *Chaire d'économie et d'économétrie de l'assurance, Working Paper DELTA 05, 1993*.

Arrow, K.J., "The Organization of Economic Activity: Issues Pertinent to the Choice of Market Versus Nonmarket Allocation", in *The Evaluation and Analysis of Public Expenditures: The PBB System*, Joint Economic Committee Compendium, 91 Congress Session, Vol. 1, 1969.

Arrow, K.J., and Hahn F., "General Competitive Analysis", Holden Day, 1971.

Benassy J.P., "The Economics of Market Disequilibrium" *Academic Press*, New York, 1982.

Bennardo, A., "Competitive Equilibria and Pareto Optima in Multicommodity Economies with Moral Hazard", Mimeo, 1997.

Bennardo, A., and P.A. Chiappori, "General equilibrium under asymmetric information: competition, positive profits and welfare-decreasing market clearing", Working Paper, University of Chicago, 1999.

Bester, H., "Screening versus Rationing in Credit Markets with Imperfect Information", *American Economic Review*, 75, 850-5, 1985.

Bisin, A., and Guaitoli, D., "Moral Hazard and Non-Exclusive Contracts ", CEPR W.P., 1998.

Bisin, A. and Gottardi, P., "Competitive Equilibria with Asymmetric Information", *Journal of Economic Theory* v87, n1 (July 1999): 1-48

Bisin, A. and Gottardi, P., "Efficient Competitive Equilibria With Adverse Selection", mimeo, NYU, 2000

Browning, Martin and Costas Meghir, "The Effects of Male and Female Labor Supply on Commodity Demands", *Econometrica*, 59(4),(1991), 925-51.

Carmichael, L., "Can Unemployment Be Involuntary? Comment" *American Economic Review* 75, 5, 1985, pp.1213-1214.

Chassagnon, A., and P.A. Chiappori, "Insurance under Moral Hazard and Adverse Selection: The Competitive Case", mimeo, DELTA, 1997.

Debreu, G., *Theory of Value*, Yale University press, 1959.

Dreze J., "Existence of an Exchange Equilibrium under Price Rigidities" *International Economic Review*, 16, 1975, pp.301-320.

Gale, D., "Equilibria and Pareto Optima of Markets with Adverse Selection" *Economic Theory* 7, 1996, pp. 207-36.

Gjesdal, F., "Information and Incentives: The Agency Information Problem", *Review of Economic Studies* 49, 1982, pp. 373-390.

Gollier, C., *The Economics of Risk and Time*, Mimeo, 1999.

Greenwald B., and Stiglitz J., Externalities in Economies with Imperfect Information and Price Rigidities" *Quarterly Journal of Economics*, 101 (2) 1986, pp.229-264.

Grossman, S. and O. Hart, "An Analysis of the Principal-Agent Problem", *Econometrica*, V. 57, 1983, pp. 1-45.

Guesnerie, R., "The Arrow-Debreu Paradigm Faced with Modern Theories of Contracting: A Discussion of Selected Issues Involving Information and Time", *The Nobel symposia*, 1992, Basil Blackwell.

Hahn, F. "On Equilibrium with Market Dependent Information" *Quantitative Wirtschaftschun*, E Helmstadler and R. Henn Ed., 1978.

Hart, O. and Holmstrom, B., "The Theory of Contracts", in T. Bewley (ed.) *Advances in Economic Theory:Fifth World Congress*, Cambridge University Press, 1987.

Hellwig, M.F., "On Moral Hazard and Non-price Equilibrium in Competitive Insurance Markets, 1983b, University of Boston discussion paper #109.

Helpman, E., and Laffont J. J., "On Moral Hazard in General Equilibrium" *Journal of Economic Theory*" 10, 1975, pp.8-23.

Jewitt, I., "Justifying the First Order Approach for Agency Problems" *Econometrica*, 56 (5) 1988, pp. 1177-90.

Malcomson J.M., "Work Incentives, Hierarchy, and Internal Labor Markets", *Journal of Political Economy*, 92 (1984), 486-507.

McLeod, B., Malcomson J., "Implicit Contracts, Incentive Compatibility and Involuntary Unemployment" *Econometrica*, 57 (2), 1989, pp. 447-480.

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

Prescott, E.C. and Townsend, R.M., "General Competitive Analysis in an Economy with Private Information." *International Economic Review* 1984b pp. 1-20.

Prescott, E.S. and Townsend, R.M., "The Theory of the Firm: Applied Mechanism Design" 1996, mimeo.

Rothschild, M., and J.E. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information", *Quarterly Journal of Economics* 90, 1976, pp. 629-49.

Shapiro C. and Stiglitz J., "Equilibrium Unemployment as a Worker Discipline Device" *American Economic Review* 74 (3), 1984, pp.433-444.

Segal E. "Contracting with externalities" *Quarterly Journal of Economics*, 114 (2) 1999, pp. 337-388

Stiglitz, J.E., "Risk, Incentives, and Insurance: The Pure Theory of Moral Hazard", *The Geneva Papers on Risk and Insurance* 8, 1983, pp.4-33.

Stiglitz, J.E. and L. Weiss "Credit Rationing in Markets with Asymmetric Information", *American Economic Review*, 1981, 71 (3), pp.393-410