Information Retention and Inefficiency in Competitive Markets for Services *

Bertrand Villeneuve†
Institut D'Economie Industrielle
Université de Toulouse 1
Place Anatole France
31042 Toulouse cedex
France
E-mail: villeneu@cict.fr

This Version: May 2000.

Abstract

We study competitive markets for services where the providers are also the experts in the assessement of needs. In terms of information economics, we deal with markets for credence goods with common value and multi-principals. In spite of the absence of technological or spatial asymmetries, and though the market is competitive, we show that neither revelation of information nor efficiency given the information transmitted are warranted, and that the surplus may be captured by some (not all) firms rather than the consumer. Our insurance illustration assumes that insurers evaluate risk better than policyholders. We explain why certain risky consumers remain unsinsured and why certain market segments are persistently profitable.

*I am endebted to Pierre-André Chiappori for his precise suggestions and constant encouragement. Discussions with Jacques Crémer, Georges Dionne, Nathalie Fombaron, Guido Friebel, Joe Harrington, Charles Kahn, Jean-Jacques Laffont, Lambros Pechlivanos, Jean-Charles Rochet, Bernard Salanié and Désiré Vencatachellum helped me improve the exposition. All errors are mine.

†Institut d'Economie Industrielle and Commissariat à l'Energie Atomique.

1 Introduction

In many situations, a consumer in need of a service faces uncertainty on his true conditions and wants advice, but, unfortunately, the only "experts" able to assess needs are also interested in the provision of the service. The optimistic view is that, by nature, competition enables the consumer to exploit the conflicting interests of the various providers to get the best information, the most appropriate quality, and the lowest price. This intuition is supported by auction theory: setting an auction, in the simple situation where bidding firms are *identical* in terms of technology and information (so that first- and second- prices will be identical), allows the consumer to overcome his inferior information, to ensure efficiency and claim all the surplus.

This paper examines reasons why competitive markets for services may fail to work in the consumer' best interest. Though we do not claim that the market failure is systematic in the real world (after all some people do very well to get cheap insurance or high quality advice), we wish to suggest that information retention, inefficiency and rents have to be expected in many contexts.

We put a general argument in insurance terms, but similar features would emerge in any other market where the experts (i.e. lawyers, doctors, mechanics, coaches, agents, ...) advise and sell. Our model assumes that the prior informational advantage is allocated to the insurers, i.e. they know precisely the structure of their clientele in terms of risks, tastes and beliefs, so as to be able to establish classes of clients and to practice price discrimination.[1] The insured does not know whether he is a low or a high risk. The game is basically the following. Insurers observe, before contracting, the type of the applicant; they make personal offers simultaneously; the individual is able to interpret offers as informative signals, and remains free to accept one or to reject them all. Consistency of the inference of the consumer with actions is guaranteed with a Bayesian equilibrium concept.

A natural prediction in this context is that competition will result in full revelation of available information (his type) and first-best efficiency (full

[1] The right given to the insurers to discriminate is based on the idea that every individual is a different market. Though data on sex, race, religion, etc., cannot be used, or even recorded, many other pieces of information can.

insurance at a fair rate). This article proves that this is not the case:[2] types may be pooled, the high risk consumers may remain without insurance or obtain partial coverage, and profits are not always zero.[3]

The intuition for inefficient though fully informative equilibria is the following. Let us take the case where low risk individuals are fully and fairly insured whereas high risk individuals are not insured at all. Obviously, insurance companies would like to offer full insurance to high risk individuals since this would be profitable and even Pareto improving. But individuals, who do not know their types, can only try to infer information from the offers they receive, and the way they interpret offers play a key role. Any 'claim' by an insurer that the customer is a high risk poses a credibility problem: an insurer would always like to make a low risk think he is a high risk (this increases his willingness-to-pay for insurance), and the consequence is that such claims are not automatically believed. Now consider the following interpretation rule by the customer. Whenever he is offered anything else than no-insurance by one of the insurers, he believes that he is a low risk with probability one. If he is a low risk, fine, competition works well and he gets the first-best. Assume that if he is a high risk, all insurers 'offer' him no insurance. Why is this an equilibrium? If an insurer deviates and offers him, say, the full insurance contract a high risk deserves, the individual thinks he is a low risk and refuses the offer. Deviations being discouraged, the high risk customer finishes up uninsured. The credibility issue is so strong that these beliefs are compatible with strong refinements (such as Cho-Kreps). The key point now is whether individuals believing they have low risk *are* reluctant to accept contracts suitable for high risk customers. This depends on the 'proximity' between types: if types are very different, then contracts that are profitable on high risk types may never be acceptable for customers believing they have a low risk with probability one.

Pooling equilibria (in which no information is revealed) follow a slightly

[2] See Villeneuve [16] (informed insurance monopoly) for comparisons with Stiglitz' [14] (insurance monopoly with adverse selection).

[3] Fagart [6] explored a similar model with the big difference that refusing all offers is not a choice. The forced participation of the agent she assumes implies that information is systematically revealed and that profits are zero. As our results show, non-participation threats explain important features of the equilibrium.

different logic. Take a situation where N insurers offer the same contract to all types. Obviously, this contract have to be acceptable for the average type (the individual who has not learned his type), and profitable on both types. When is this stable? A priori, an insurer is tempted to deviate when a low risk shows up by offering him a neatly better contract: instead of sharing the customer (i.e. earning a large profit with probability $1/N$), it may prefer to earn a lower profit with probability one. Now assume that the consumer believes he his a low risk with probability one whenever he observes a deviation from the collusive offers. Deviations that are accepted by the customer with probability one are not necessarily very profitable, and, consistently, profitable deviations are not accepted... In any case, the collusion is stable. A pooling equilibrium doesn't always exist, but is, surprisingly enough, possible with a infinite number of insurer. Two conditions have to be met: the low risk have to be relatively low (this bonds the willingness-to-pay for insurance of a customer believing he is a low risk), and high risk individuals are relatively more probable (the individuals has a high *prior* willingness-to-pay for insurance).

Predictions that insurance markets do not provide first-best contracts sound sensible. We wish to show that our assumptions find support in recent econometric findings, as well as from casual observation. Adverse selection is undoubtedly a serious threat to insurance, but the success of classification methods may have rendered its effects hardly visible.[4] Chiappori and Salanié [2] offer an important methodological advance on the tests of adverse selection, notably on the necessary precautions to avoid spurious regressions due to endogenous explanatory variables. In their application, their conclude that symmetric information cannot be excluded when tested against adverse selection. Our assumptions have not been tested, and a different dataset could give a different answer, but the results are somewhat encouraging for our view. Indeed, even if the insurer ignores certain individual details, it is able to relate observable data with riskiness by means of statistical methods

[4] It has been noted by Gouriéroux [8] that in the risk classes defined by the insurers in automobile insurance, variability in the choice is observed in a set of classes covering less than 10% of the clientele. If we believe that insurers would propose menus only if they were necessary, this suggests that the hypothesis of adverse selection could be practically dismissed for more than 90% of the clientele.

applied to large samples, whereas the consumer knows all his particulars but relies on too small a reference group to weigh accurately the risk factors. These two information structures are not necessarily comparable in strict mathematical terms, but the following examples suggest that insurers are favored in very common situations. Companies specialized in the insurance of activities where damages are rarely experienced by the insured can reasonably be expected to better evaluate the possible financial consequences of an accident. For instance, in liability insurance, most clients are not really aware of the compensations they could be forced to pay in consequence of their actions since an up-to-date knowledge of the jurisprudence is required. In general, the individual is likely to have chosen his activity or the place where he lives on the basis of comparative informations on safety; still, there may remain aspects of his risk well perceived by insurers that he cannot possibly guess from aggregate data like public statistics or reputations.

Our work has certain links with the literature on informed principals and on credence goods. Leaving the initiative (or negotiation power) to the informed agent is also the core hypothesis of Maskin and Tirole [11]. The essential difference is that here, the uninformed party confronts several (rather than one) informed parties.[5] The strategy space used in Maskin and Tirole [11] is also worth mentioning: the informed party proposes a "mechanism" (a game form); at the price of a certain abstraction, this strategy space improves the efficiency of the equilibria. For the sake of clarity, we base our propositions and interpretations on "contract offers" in the ordinary sense rather than "mechanisms"; the interested readers can find in the Appendix a counter-example proving that our analysis goes through with competition in Maskin-Tirole mechanisms (e.g. no-insurance for all is an equilibrium). A difference with the literature on credence goods (e.g. Emons [5]), where the hidden parameter only affects the consumer's willingness-to-pay, is that for very natural reasons in our model, the cost to the expert is also affected. Our work can be seen as the common value version of this research area. Still, the techniques involved are neatly different. Intuitively, the dependency on type of the expert's profit should facilitate product differentiation, hence, should

[5] An unessential difference is that, in their application to insurance, the informed party is the policyholder. We have in common that there is one policyholder and several insurers.

favor information revelation. This is not the case.

2 The Model

A consumer with wealth W incurs the risk of losing a monetary amount d with a certain probability. Initially, this probability (the *type*) is known to the individual and the insurers only up to its distribution in the population.[6] For the moment, we assume that the consumer is a high-risk type (loss probability p_H) with probability λ_H and a low-risk type (loss probability p_L where $p_H > p_L$) with probability $\lambda_L = 1 - \lambda_H$.[7] As soon as the consumer enters the market, the insurers observe his type strive to provide him with an exclusive contract. The concave VNM utility function of the agent is denoted by u, and the insurers are supposed to be risk-neutral profit maximizers.

The insurers are indexed by a set $I = \{1, 2, \ldots, N\}, N \geq 2$. The steps of the game are the following:

1. The type of the consumer (H or L) is randomly selected.
2. All insurers receive the same perfectly informative signal (the type), but the consumer remains uninformed at this stage.
3. The insurers simultaneously offer contracts.
4. The consumer observes the offers, updates his priors, and selects exclusively the preferred offer. He can prefer the reservation contract.
5. The loss occurs or not. The accepted contract is implemented and final payoffs are realized.

The type of the consumer is "common value" to the principal and the agent: given a contract, the cost to the insurer and the agent's willingness-to-accept are type-dependent. More generally, the ranking between the various

[6] To form a reasonable prior on the hazards incurred, the consumer relies on the publication of epidemiological studies, in scientific or popular form, and on information-based campaigns inciting to insurance purchases.

[7] Extensions to loss distributions with large supports are not substantially different for our main point. Extensions to more than two types, however, necessitate a particular treatment given in Section 5.

contracts is also type-dependent for insurers and dependent on actual beliefs for the individual.

Let Ω be the contract space (specified later), and let us denote by $C^i \in \Omega$ insurer i's offer. The "interpretation mapping" (or simply the beliefs) of the insured, Λ, maps the set of vectors of offers Ω^N into $\Delta(\{H, L\})$, the set of probability distributions over types. Hence, $C^\bullet \equiv (C^i)_{i \in I}$ being a vector of offers, $\Lambda(C^\bullet) = (\lambda_H(C^\bullet), \lambda_L(C^\bullet))$ represents the consumer's updated priors, where $\lambda_K(C^\bullet)$ is the probability (in the consumer's mind) of being of type K (with $K \in \{H, L\}$ and $\lambda_H(C^\bullet) + \lambda_L(C^\bullet) = 1$). Beliefs Λ are endogenous and known to the insurers.

The reservation contract C^0 (here the right to remain uninsured) plays a fundamental role, as we shall see. In accordance with the exclusivity assumption, the decision (or best response to the offers) of the individual will respect the following rules: given the revised priors , (a) any contract which is strictly dominated is never taken; (b) if there are two or more optimal contracts, the individual chooses randomly and symmetrically among optimal offers differing from no-insurance. For example, each insurer gains participation of the agent with probability $1/N$ whenever they all make optimal offers. Notice that our tie-breaking rule ensures that the strategy of the individual is *uniquely* determined by his beliefs. To fix the terminology, an *offer* will be a contract actually proposed by an insurer, whereas the *allocation* will be the (set of) contract(s) actually adopted by the utility-maximizing consumer in equilibrium. No-insurance is not exactly like other contracts: whenever it is the equilibrium allocation of a certain type, (b) ensures that no other contract is taken in equilibrium.

We focus our analysis on the Perfect Bayesian Equilibria (PBE) in pure strategies.[8]

Definition 1 *A PBE here will be: (a) a set of offers $(C_H^{\bullet *}, C_L^{\bullet *})$ where C_K^{i*} is the offer that player i makes when a K shows up and (b) a belief mapping Λ^* such that:*

1. *Actions are sequentially optimal: for any type K, for any insurer i, C_K^{i*}*

[8]Equilibria in pure strategies abound, therefore we leave equilibria in mixed strategies (sometimes called semi-separating equilibria in signalling games) aside to avoid confusion.

maximizes the profit of principal i in Ω given (a) C_K^{-i*} (the actions of the other insurers) and (b) $\Lambda^*(C_K^{-i*},\cdot)$ (its impact on beliefs given *the others' actions).*

2. *Beliefs are consistent with Bayes rule: if $(C_H^{\bullet *}) = (C_L^{\bullet *})$ then $\Lambda^*(C_H^{\bullet *}) = \Lambda^*(C_L^{\bullet *}) = (\lambda_H, \lambda_L)$; else $\Lambda^*(C_H^{\bullet *}) = (1,0)$ and $\Lambda^*(C_L^{\bullet *}) = (0,1)$; they are not restricted elsewhere.*

Three different strategy spaces are investigated: Ω_0 for the simplicity, Ω_1 for the realism, and Ω_2 to show the robustness of the results. Contracts in Ω_0 and Ω_1 are directly interpretable as insurance contracts: C^i in Ω_0 or Ω_1 specifies consumptions W_A^i in case of loss and W_N^i in case of no loss ($C^i \equiv (W_A^i, W_N^i)$, and in particular $C^0 \equiv (W-d, W)$).[9] We only consider "positive" insurance contracts for which $W_A \geq W-d$ and $W_N \leq W$. The full set of insurance contracts is denoted by Ω_1; in particular over-insurance ($W_A > W_N$) or under-insurance ($W_A < W_N$) contracts are not excluded. For simple applications, Ω_0 will denote the subset of full insurance (i.e. constant consumption, or $W_A^i = W_N^i$) contracts; two contracts in this set only differ by the premia.[10] Finally, Ω_2 is the set of mechanisms (or game-forms) as in Maskin and Tirole [11]. Obviously, $\Omega_0 \subsetneq \Omega_1 \subsetneq \Omega_2$.

For any vector of offers $C^\bullet$ in Ω_0^N or Ω_1^N, the linearity of expected utility with respect to probabilities makes the individual evaluate contracts as if his probability of accident were $\tilde{p}(C^\bullet) \equiv \lambda_H(C^\bullet)\, p_H + \lambda_L(C^\bullet)\, p_L$ where obviously $\tilde{p}(C^\bullet) \in [p_L, p_H]$; by extension, $\tilde{p}(C^\bullet)$ being a sufficient statistic for the evaluation of contracts, it will also be called a belief; the weak preference order over Ω_0 over Ω_0 that it implies will be denoted by $\succeq_{\tilde{p}(C^\bullet)}$. In the following, $\tilde{p}_0 \equiv \lambda_H\, p_H + \lambda_L\, p_L$ will denote the expected type at the beginning of the game.

Section 3 solves the model for full insurance contracts (Ω_0). Section 4 gives the general methodology, solves the model for ordinary insurance contracts (Ω_1), and presents the main practical implications of the model. (The

[9] Defining contracts in terms of premium (paid in advance before the loss is known) and indemnity (paid in case of loss) is equivalent.

[10] Notice that $C^0 \notin \Omega_0$. Still, no insurance remains the reservation contract in the simplified model and completing Ω_0 with C^0 would not susbstantially change the results.

Appendix proves that competition in mechanisms—Ω_2—does not ensure efficiency). Section 5 offers an extension to more than two types of individuals, and discusses the contribution of the two main refinement families. The conclusion focuses on competition policy and on the testability of the model.

3 Price Competition

Constraining offers to be full insurance simplifies the exposition, and above all, sets the best conditions for competition to work efficiently: offers differing only in price, comparisons are not type-dependent. However, we show that competition does not lead necessarily to equilibria with discrimination (information is not necessarily revealed) and equilibrium profits are not systematically zero (the Bertrand logic is not fully effective).

A full insurance contract will be simply denoted by the contribution c that is charged to the consumer, the indemnity being equal to d. Five contracts play a particular role: the minimal contributions charged by rational principals type by type ($c_H^{\min} = p_H d$ and $c_L^{\min} = p_L d$); the maximal contribution of a consumer believing that he has, respectively, a high, a low, or an average risk. Formally:

$$u(W - c_K^{\max}) = (1 - p_K)u(W) + p_K u(W - d), K = H, L \tag{1}$$

and

$$u(W - c_{\tilde{p}_0}^{\max}) = (1 - \tilde{p}_0)u(W) + \tilde{p}_0 u(W - d) \tag{2}$$

For non-degenerate distributions: $c_H^{\min} < c_H^{\max}; c_L^{\min} < c_L^{\max}; c_L^{\min} < c_H^{\min}; c_L^{\max} < c_{\tilde{p}_0}^{\max} < c_H^{\max}$.

3.1 Equilibria with Discrimination

We study first equilibria where different types are treated differently.

Proposition 1 *For any given p_H, there exists a threshold $0 < \hat{p}_L < p_H$ such that:*

1. *If $p_L > \hat{p}_L$, then a separating equilibrium is necessarily fair for all types.*

2. *If $p_L \leq \hat{p}_L$, then a separating equilibrium is fair for L, but H can receive either no-insurance or be charged anything between $[c_H^{\min}, c_H^{\max}]$.*

The first important fact is that the actuarially fair full insurance allocation is always an equilibrium: then, competition works and profits are zero. However, other types of equilibria appear, notably those involving the high risk being either exploited (unfair premium), or not insured. The critical comparison here is that of $c_L^{\max}$ with $c_H^{\min}$: are there, or not, contracts acceptable for all beliefs and profitable on all types? The answer depends critically on the support of the types (i.e. p_H, p_L). Notice that the number of insurers is irrelevant.

An interpretation is that when types are very similar, beliefs should not play a major role: low risk individuals are not reluctant to accept contracts suited for high risk, the Bertrand effect (price is equal to cost) is strong and determines entirely the allocation. When types are very dissimilar, beliefs are fundamental to determine whether an offer is acceptable or should be rejected. The high-risk type either receives an inefficient allocation (no-insurance equilibrium) or is charged too high a price. In the latter case, a firm may capture all the economic rent without being seriously threatened by its competitors.

Whenever it is *possible* that the consumer is not at risk ($p_L = 0$ and $\lambda_L > 0$), then the contract that the high risk gets may be anything from a fair contract to a maximal price contract, plus no-insurance. This should be compared to the case where being a low risk is not possible ($\lambda_L = 0$): then, the only equilibrium allocation H is actuarially fair.

3.2 Uniform Pricing in Equilibrium

In a pooling equilibrium, by definition, each insurer charges the same price c^{i*} for all types, though two different insurers could apply different strategies. Clearly, (a) the *allocation* is the same for all types, since all types receive the same offers, and updated beliefs are the same for all; (b) the consumer receives an allocation that differs from no-insurance (whatever his beliefs, a consumer would accept any offer which is profitable and acceptable for the most demanding belief p_L, thus competition ensures that type-L agents cannot be left with no-insurance, and this should be true for all types since the equilibrium pools types together), (c) all insurers offer the same contract c^* (any insurer that would be excluded—and earn no profit—would only have to offer a slightly profitable contract when a L shows up to be better-off and break the equilibrium). Moreover, (d) c^* must profitable on all types, and acceptable for the average type (the consumer keeps his priors), this means $c^* \in [p_H d, c_{\tilde{p}_0}]$, and (e) each insurer compares the value of playing the collusive move c^* (but earning the profit with probability $1/N$) with the profit associated with attracting a type-K consumer for sure, this imposes $\frac{c^* - p_K d}{N} \geq c_L^{\max} - p_K d, K = H, L$.

Pooling equilibria are based on tacit collusion: no insurer wants to undercut price since this would produce adverse effects on the willingness-to-pay of the consumer. Hence, provided there exist contracts satisfying (d) and (e), a pooling equilibrium exists. An example of appropriate beliefs is the following: for vectors of offer differing from $(c^*, c^*, \ldots, c^*)$, let the individual believe that he is a low risk (less extreme beliefs are also possible). This too optimistic attitude discourages deviations since the individual's willingness-to-pay for insurance is minimal. The insurer sees deviations either as non profitable, or, as not acceptable (given the beliefs).

We analyze now the dependency of existence on the following parameters: p_H, p_L, λ_H (or, equivalently, λ_L), and N.

Proposition 2 *Statements are given* ceteris paribus.

1. *The set of pooling allocations enlarges as the proportion of high-risk consumers increases.*

2. *In a given market where $p_L > 0$, there is a finite upper limit on N consistent with pooling equilibria.*

3. *For any N, if the low-risk loss probability and the proportion of low-risk consumers are small enough, there is at least one pooling equilibrium.*

4. *In a given market where $p_L = 0$, if the proportion of low-risk consumers is low enough, then whatever N, there exists a least one pooling equilibrium.*

Though it may seem natural that such allocations cannot be sustained with a large number of insurers, the intuition is counterbalanced by the fact that breaking the collusion for low-profit low-risk consumers is not necessarily attractive. Indeed, the mere *possibility* in the consumer's mind of being a very low probability type always cause existence of a pooling equilibrium. This result may have considerable practical importance since the presence (in whatever proportion) of very low risk individuals is a reasonable scenario in most insurance markets: under these circumstances, non-discrimination and profitable offers have to be expected.

4 A General Methodology

4.1 Dominant Contracts

Contracts which are *obviously* inefficient for the insurer offering them and for the individual should never be proposed. We define two partial orders over the set of contracts (first- and second-degree dominance) that will serve as building-blocks for the characterization of the equilibria. We define first some useful vocabulary.

> $\mathcal{A}(\tilde{p})$: for any belief $\tilde{p}$, $\mathcal{A}(\tilde{p})$ denotes the set of *acceptable* contracts, i.e. preferable in the weak sense to no-insurance (notice that $\mathcal{A}(\tilde{p})$ is strictly increasing with respect to $\tilde{p}$):
>
> $$\mathcal{A}(\tilde{p}) \equiv \left\{ C \mid C \succeq_{\tilde{p}} C^0; C \in \Omega_1 \right\} \tag{3}$$
>
> $\pi_K(C)$: the (expected) profit made on type K if the contract $C = (W_A, W_N)$ is taken:
>
> $$\pi_K(C) \equiv (1 - p_K)(W - W_N) + p_K (W - d - W_A) \tag{4}$$
>
> $\mathcal{P}_K(\pi)$ and $\overline{\mathcal{P}}_K(\pi)$: the open, resp. closed, set of contracts in Ω_1 yielding a strictly, resp. weakly, higher profit than π on type K. Formally:
>
> $$\mathcal{P}_K(\pi) \equiv \{C \mid \pi_K(C) > \pi; C \in \Omega_1\} \tag{5}$$
>
> $$\overline{\mathcal{P}}_K(\pi) \equiv \{C \mid \pi_K(C) \geq \pi; C \in \Omega_1\} \tag{6}$$

C_K^{AF} : the actuarially fair (= zero profit) full-insurance contract of type K.

$C_K^{\max}$: the contract extracting maximal profit under symmetric information (see equation 1) on type K.

The notions defined above are represented in the plane $(W_A, -W_N)$, no-insurance $(W-d, -W)$ being the origin (see Figure). The abscissa is the net indemnity (indemnity minus prepaid premium), and the ordinate is the premium; Ω_1 corresponds exactly to the first quadrant. Whatever his beliefs, the agent's utility increases as we move downwards or rightward. 'U_K' indicates a type-K indifference curve; '$U_{\tilde{p}_0}$' is an indifference curve of the average type; 'K' indicates a type-K isoprofit straight line; $\mathcal{P}_K(\pi_K(C))$ is the open half-plane which is bounded below by the isoprofit line passing through C.

Insert Figure Here

Definition 2 ($\mathcal{F}_K$) *A contract C in Ω_1 is* first-degree dominated *for type K if and only if there exists a contract C' in Ω_1 which (a) is strictly profitable on type K, (b) for all beliefs in the interval $[p_L, p_H]$, is strictly preferred to C, and which (c) is strictly preferred to no-insurance. Formally:*

$$\pi_K(C') > 0 \text{ and } \forall \tilde{p} \in [p_L, p_H] : C' \succ_{\tilde{p}} C, C' \succ_{\tilde{p}} C^0 \tag{7}$$

$\mathcal{F}_K$ is the set of contracts which are first-degree dominant (i.e. not *first-degree dominated) for type K.*

First-degree dominated contracts cannot be offered in equilibrium by an insurer if an other insurer does not make a strictly positive profit. In contrast, a contract in $\mathcal{F}_K$ is robust to external threats: it is a natural candidate as the offer to type-K individuals in an equilibrium where one insurer or any number of insurers are able to capture these consumers with certainty.

Definition 3 ($\mathcal{S}_K$) *A contract C in Ω_1 is* second-degree dominated *for type K if and only if there exists a contract C' in Ω_1 for which (a) the profit on K is strictly larger than $1/N^{th}$ of the profit of C on K, and which, (b) for* *any beliefs in the interval* $[p_L, p_H]$, *is strictly preferred to* C *and (c) is strictly preferred to no-insurance. Formally:*

$$\pi_K(C') > \pi_K(C)/N \text{ and } \forall \tilde{p} \in [p_L, p_H] : C' \succ_{\tilde{p}} C, C' \succ_{\tilde{p}} C^0 \tag{8}$$

$\mathcal{S}_K$ *is the set of contracts which are second-degree dominant (i.e.* not *second-degree dominated) for type* K.

Second-degree dominated contracts cannot be offered in a totally collusive equilibrium (i.e. where the client chooses each insurer with probability $1/N$): any insurer would prefer to deviate and offer C' when a type-K individual shows up. N-vectors of offers in $\mathcal{S}_K$ (provided there exists a certain beliefs for which they are all equivalent) are the natural candidates as offers to type-K individuals in an equilibrium where the insurers share the clients of that type. In other terms, contracts in $\mathcal{S}_K$ are robust to internal threats, i.e. deviations from insurers participating to the collusion.

Notice that first-degree dominant contracts for a certain type are necessarily second-degree dominant for that type ($\mathcal{F}_K \subset \mathcal{S}_K$). The theorems draw general characterization of the equilibria which are applicable to any contract set, provided $\mathcal{F}_K$ and $\mathcal{S}_K$ are adapted. The propositions restate the results in terms of insurance contracts.

4.2 Separating Equilibria

Let C_K^{i*} be the contract offered by insurer i for type-K agents in a separating equilibrium; π_K^{i*} denotes the corresponding *equilibrium profit*.

Theorem 1 *Let* $C_H^{\bullet*}$ *and* $C_L^{\bullet*}$ *be two different vectors;* $(C_H^{\bullet*}, C_L^{\bullet*})$ *is supported in a separating equilibrium if and only if, for any* K, *one of the following two conditions is satisfied:*

1. *(a) The best contract for belief* p_K *in* $C_K^{\bullet*} \cup C^0$ *is profitable, acceptable, and first-degree dominant for type* K. *(b) Either* C^0 *or two contracts (at least) in* $C_K^{\bullet*}$ *are acceptable for type* K *and first-degree dominant for type* K. *(c) Provided (b) is satisfied, the contracts not taken in equilibrium are only restricted to be strictly less desirable for belief* p_K *than the equilibrium allocation.*

2. *All contracts in* $C_K^{\bullet *}$ *are equivalent* and *acceptable for belief* p_K*, and second degree-dominant for type* K.

Moreover, the set of beliefs sustaining the corresponding allocation is convex.

Cases 1 and 2 cover (but do not partition since $\mathcal{F}_K \subset \mathcal{S}_K$) the set of equilibria: either certain or none of the offers are rejected in equilibrium. In each type of equilibrium, one must make sure that the offers accepted with positive probability are sufficiently stable: if the preferred contracts were not first- or second-degree dominant, then "obvious" deviations would be possible for at least one of the insurers (rejected insurers for 1, any insurer in the collusion in 2), and would be accepted whatever the subsequent beliefs. In condition 1, (a) indicates that the allocation must be sufficiently robust to external threats, (b) says that the best offer is blocked by the presence of "second best" offers (no-insurance or non-trivial offers), (c) says that the other contracts are irrelevant. Condition 2 deals with collusive equilibria: no insurer would like to break the profitable arrangement.

If we focus on allocations only, the theorem proves than any contract in Ω_1 which is profitable and acceptable for type K, and first-degree or second-degree dominant for type K is an equilibrium allocation for type K. As a particular case, C^0 is an equilibrium allocation if and only if it is first-degree dominant for type K.

Whenever insurance contracts in the ordinary sense (i.e. contracts in Ω_1) are considered, we can go further by distinguishing three *degrees of proximity* between types:

> **D**: Types are Distant if and only if no profitable contract on H is acceptable for an individual believing he is a L with probability one.
>
> **WR**: Types are Weakly Related if and only if, at the fair price of the high risk, marginal insurance is acceptable for the low risk, but not complete insurance, i.e. $\frac{p_L}{1-p_L}\frac{u'(W-d)}{u'(W)} \geq \frac{p_H}{1-p_H}$ but $C_H^{\mathrm{AF}} \prec_{p_L} C^0$.

CR: Types are Closely Related when the fair full insurance of the high risk is acceptable to the low risk, i.e. $C_H^{\mathrm{AF}} \succeq_{p_L} C^0$.

Notice that **D** implies $\frac{p_L}{1-p_L}\frac{u'(W-d)}{u'(W)} < \frac{p_H}{1-p_H}$ and $C_H^{\mathrm{AF}} \prec_{p_L} C^0$, and **CR** implies $\frac{p_L}{1-p_L}\frac{u'(W-d)}{u'(W)} \geq \frac{p_H}{1-p_H}$. We have now an ordered partition: given p_H, it is necessary and sufficient to make p_L pass from 0 up to p_H for passing from regime **D** to regime **WR** and then to regime **CR**; conversely, given p_L, it is necessary and sufficient to make p_H pass from 1 down to p_L for passing through the three regimes.

The following proposition explains why L is relatively well treated in separating equilibria. Basically, the proof characterizes $\mathcal{F}_K$ and $\mathcal{S}_K$. The proposition states the most apparent results. Readers interested in more details concerning the structure of the equilibria can refer to the developments in the Appendix.

Proposition 3 *In a separating equilibrium:*

1. *Two (or more) offers to the low risk are fair. Both give full or over-insurance and one (at least), say C^i, is taken in equilibrium. The equilibrium allocation contains C^i plus, possibly, other contracts, provided they are acceptable, strictly profitable, and equivalent to C^i for belief p_L.*

2. *For the high risk, either **D** and any contract which is profitable on H and acceptable for belief p_H is allocated to H in an equilibrium, or **WR**, and certain (not all) contracts profitable on H and acceptable for belief p_H are an equilibrium allocation, or **CR**, and everything is as 1 with H instead of L, except that the fair contracts may also provide under-insurance.*

Full fair insurance for all is *always* an equilibrium, but never the only one. In the price competition version of this model, we saw that prices might not be fair because better offers could be seen with distrust, and discouraged. Here inefficiency is aggravated since the quality of coverage may be suboptimal for both types. The restrictions on the allocation to L (and therefore to H in case **CR**) show that if the fair contract in the allocation is not full insurance, then other equivalent contracts offering a better coverage (rather than over-insurance) for a higher price may be in the allocation.

Nevertheless, when types are closely related, beliefs do not matter very much and the allocation is relatively efficient. Case **D** is particularly interesting: if the low risk is sufficiently low, the consumer may become so optimistic that, if he turns out to be a high risk, he spoils the competition between the insurers and ends up with insufficient coverage. In the extension to more than two types (Section 5), we show that H and L can be seen as the extreme types in a large support. In this view, case **D** becomes generic and the strongest inefficiencies (insufficient coverage and capture of the profit by the insurers) have to be expected in most markets.

4.3 Pooling Equilibria

A pooling equilibrium is a *vector* of contracts $C^{\bullet *}$ which are offered to both types without discrimination. Two insurers may offer different contracts. We denote by π_H^{i*} and π_L^{i*} the *equilibrium* profits, which, according to the tie-breaking rule, are such that: if contract C^{i*} is not taken, then, π_H^{i*} and π_L^{i*} are zero; if all contracts are taken with positive probability, $\pi_H^{i*} = \pi_H(C^{i*})/N$ and $\pi_L^{i*} = \pi_L(C^{i*})/N$, from the assumption that the individual chooses randomly when he is indifferent; if contract C^{i*} is exclusively taken, $\pi_H^{i*} = \pi_H(C^{i*})$ and $\pi_L^{i*} = \pi_L(C^{i*})$, etc.

Theorem 2 *The following conditions are necessary and sufficient for $C^{\bullet *}$ to be a pooling equilibrium:*

1. ***Uniform pricing.*** *The equilibrium allocation is the same for all types.*
2. ***Non-triviality.*** *The consumer receives an allocation that differs from no-insurance.*
3. ***Symmetry.*** *The consumer is indifferent between all offers for prior beliefs.*
4. ***Rationality.*** *For any i : C^{i*} is profitable on H and L and C^{i*} is acceptable for the average type.*

5. ***Incentives.*** *For any i, C^{i*} is second-degree dominant for H and L.*

The sets of beliefs sustaining the corresponding allocation is convex.

See Proposition **??** for similar results. Notice that the symmetry property imposes indifference, which is the natural extension of equality to this multi-dimensional version. Concerning the incentives, the appropriate notion is second-order dominance: indeed, if one of the contracts were not second-degree dominant for some type, say $C^{i*} \notin \mathcal{S}_K$, there would exist a contract C' which is a better move for i than C^{i*} at least if a type-K individual shows up.

The comparative statics of existence follows the same logic as for price competition. To avoid repetition, the reader is invited to refer to Subsection 3.2. The paradox is that there always exist pooling equilibria when the loss probability of the low risk is small: the possibility of being a very low risk supports exaggeratedly optimistic beliefs that discourage serious competition.

Insurers would never accept losses on high-risk individuals, therefore the first-best allocation that would provide full insurance at an average fair price is not an equilibrium. Any pooling equilibrium gives a positive profit on both types to the insurers, and the absence of discrimination can only prove the absence of competitive pressure.

5 Extensions

5.1 More Than Two Types

The game is the same but now the support of types is "large". We prove that, for any type, the set of allocations he can receive increases when the support of the prior distribution increases. Retrospectively, H and L in the two-type study can be considered as a particular pair of the more general distribution.

Let us consider a finite support of type, denoted by $\mathcal{K}$, and let J and K be generic types. Let λ_J and λ_K, the prior probabilities of J and K, be strictly positive. Let us consider the two-type "conditional game" where the support of the types is restricted to J and K, and where the proportions of the types are, respectively, $\frac{\lambda_J}{\lambda_J+\lambda_K}$ and $\frac{\lambda_K}{\lambda_J+\lambda_K}$. We denote the set of equilibrium vectors of strategies of this game as $\mathcal{E}(J, K)$. (See Section 4 for details on $\mathcal{E}(J, K)$). A generic element of this set is denoted by $(C_J^\bullet, C_K^\bullet)$, where, e.g., C_J^i would be insurer i's offer if a J enters the market.

Proposition 4 *Let us partition arbitrarily $\mathcal{K}$ in pairs and singletons, and denote by (J_s, K_s) the s^{th} "pair" (a singleton if $J_s = K_s$), with $s \in S$, S being a minimal set of indices. Let us now select an equilibrium strategy of the conditional game for each "pair" of types $(C_{J_s}^\bullet, C_{K_s}^\bullet) \in \mathcal{E}(J_s, K_s)$. Either $(C_{J_s}^\bullet, C_{K_s}^\bullet)_{s\in S}$, or an arbitrarily close approximation, is an equilibrium strategy of the full-support game.*

There are other types of equilibria, where groupings concern three or more types. In spite of the single-crossing property exhibited by the principals' preferences and by the consumers', pooling need not group types that are close in the support. In a four-type economy, for example, extreme types can form one pooling and intermediate ones an other.

The result points at an essential mechanism behind multi-senders equilibria: they entail a high degree of coordination on the type of the consumer, which is a public signal. In consequence, each insurer has a small influence in general, and often none, on the allocation: it is virtually forced to align its offer to the others', even if certain deviations could be very valuable.

5.2 Refinements

Refinements are conceived for eliminating equilibria which are supported by somehow "unreasonable" beliefs. We shall not contribute to the long list of existing refinements: they are all based on two types arguments, the first focussing on the principals' incentives, the second being directly concerned with the revelation of information.

Incentive Based Refinements In these refinements, one starts from the equilibrium allocation and tries to see what types of principals would "gain" from playing a deviation, i.e. a move that is not played in equilibrium. Beliefs associated should not put mass on the types that do not gain. The different versions of this refinement vary on what "gain" exactly means, namely, what sort of reaction from the part of the uninformed agent should be expected to compute the payoff consistently. The example that we will develop is the Intuitive Criterium of Cho and Kreps [3], notoriously efficient in many signalling games (see Villeneuve [16] for the monopoly version of the informed insurers).

We apply the classical definition[11] in the following way: in all equilibria, we verify that the *contribution* to beliefs of any insurer is "reasonable". Otherwise stated, if $C_H^{\bullet *}$ and $C_L^{\bullet *}$ are equilibrium vectors of offers, beliefs $\tilde{p}$ are reasonable if and only if for any i and for any K, beliefs $\tilde{p}(C_K^{-i*}, \cdot)$ are reasonable as far as insurer i's behavior is concerned.

Proposition 5 *A PBE allocation is robust to the Cho-Kreps criterium if and only if (1) it is separating and (2) it assigns the actuarially fair full insurance contract to the low-risk type.*

A certain degree of efficiency is reached via the refinement. Non-discriminating equilibria are eliminated, as well as over-provision of insurance for the low risk in separating equilibria. However, both collusion against H and no-insurance for H are robust against the refinement.

The failure can be understood as follows. Single-crossing arguments are not easily applied in the game: even if the high risk is not insured whereas the low risk is completely insured at a fair price, deviations that are favorable to the high-risk consumer (contracts better than no-insurance) are typically favorable to the insurer whatever the agent's real type.

Revelation Based Refinements This category is particular to multi-principals environments. The idea is that in a given equilibrium, vectors of

[11] Let us take a PBE. Given a type K, some propositions of the insurer would lead it to earn, whatever the subsequent beliefs of the agent provided he acts optimally conditional on them, a profit which is lower (strictly lower for one possible belief at least) than the profit it earns indeed by playing its equilibrium action. The intuitive criterium imposes the following restriction: a "reasonable" belief of the agent associated to such a dominated proposition must put no mass on type K, provided that the proposition gives a higher profit than equilibrium profit on the remaining type. If the condition is not satisfied, the equilibrium is eliminated.

offers for different types are typically different in more than one dimension. The consequence is that no principal's offer is essential for interpreting correctly the message: any short list of offers (even one offer only) is sufficiently informative in most equilibria. Each insurer is tempted to free-ride the revelation part, and earn participation by offering the best deal. However, notice that the reasoning is valid *ex post* only, i.e. too late in a pure PBE.

Bagwell and Ramey [1] proposed that equilibrium beliefs should be unprejudiced in the following sense.[12] To each vector of offers, they associate two numbers: the first (resp. the second) is the number of offers different from the equilibrium pair assigned to the high-risk (resp. the low-risk) type. If these two numbers differ, the belief associated with the examined pair should put probability one on the type corresponding to the strictly smaller number. Applied to our model, their refinement selects as a unique *separating* equilibrium fair full insurance for each type. However, the criterium is incapable of eliminating pooling equilibria.

Should we refine? Supposing perfection of the Bayesian equilibrium as we have done is a refinement, but one which is conceptually clear and which simplifies the resolution of the model. The intuitive criterium is useful against pooling equilibria whereas the Bagwell-Ramey criterium is useful against inefficient separating equilibria. There are other differences: (1) the intuitive criterium does not impose restrictions on the beliefs associated to offers which would all be off-equilibrium; (2) with the intuitive criterium, the dimension of the contract space is crucial for the functioning of sorting-out arguments (applied to price competition, the criterium is incapable of eliminating pooling equilibria); (3) the dimension of the contract space is irrelevant for a correct selection the equilibria sustained by the unprejudiced beliefs. Applying both arguments sequentially leads to the unique allocation that the market would have reached under symmetric information.

If we think that the forces leading to efficiency are strong *in the short*

[12] The economic problem treated in their model (oligopoly pricing when the quality of the product is not observed by the buyer) is very different from ours (no common value between principals); still, as far as refinements are concerned, the game theoretic issue is close.

run,[13] we should only consider as economically meaningful the Bertrand allocation, keeping in mind that the implicit informational and cognitive assumptions are radical and largely obscure, which is a common (and valid) objection against refinements. The message of this paper is that this common sense shortcut does not work well, even if (and especially if) we add a phase of argumentation for explaining why an inefficient "equilibrium" should not be played.

6 Conclusion

Markets for services give a convenient illustration of the reasons why competition does not ensure perfect revelation of the information to the consumers, why, when market interactions reveal information, the allocation may remain inefficient, and finally why a part of the surplus may remain in the hands of some firms. The irony is that the largest rents to the firms are left the most distrustful consumers. Pooling equilibria and separating equilibria where the costly type is exploited are typical of the relationship between distrust and tacit collusion. As refinements indicate, sophisticated consumers are better treated by the market: educating the public's understanding of insurance, not to say of competition, can be at least as effective as regulating tariffs. The efficacy of an education-based policy is even more probable when risks change faster than regulations.

Regulating insurance markets remains a difficult task. *Ex ante*, all equilibria are worse than the allocation that would emerge under symmetric ignorance of the types (namely, average fair full insurance).[14] We encounter a variety of the Hirschleifer paradox, where the value of the better information of the insurers is socially negative. The tragedy is that discrimination techniques are a competitive advantage: one can, for example, propose better terms to low risk consumers, nothing to high risk consumers. The process

[13] This approach could be seen as an illustration of what Hertzendorf and Overgaard [10] call the efficiency argument.

[14] Pooling equilibria in our model are less favorable to the individual since they entail positive profits on both types, i.e. collusion between insurers rather than cross-subsidizing between types. In the separating equilibria, even the best case where each type is fully insured according to its actuarial value, the individual receives *ex ante* less utility than being fully insured under average actuarial terms.

is potentially never ending: insurers seek informational advantages, which are, sooner or later, undermined by imitation. Information becomes symmetric again, but with a higher degree of discrimination, and, again, insurers seek informational advantages. It seems difficult to imagine policies stated in general terms, like anti discrimination laws, that could break this dynamics, notably in markets where adverse selection can suddenly become an issue.

The easiest tests of our theory are based on the predictions that do not contradict our model whereas they falsify that of Rothschild and Stiglitz, and the symmetric information hypothesis. In the absence of a quantitative argument, we want to suggest that the following statements are reasonably attractive: (1) low risk people may be better covered by insurance than high risk people; (2) certain people are not insured in spite of their above average riskiness; (3) in certain niches, policyholders are persistently profitable.[15]

Notice that our statements are not *ceteris paribus.* We base indeed our analysis on the fact that two people may be observationally different for the insurers and the econometrician, and still not be aware of the extent of their differences in riskiness. Compared to the tests of adverse selection, this imposes a different treatment of control variables. In practice, insurers offer menus in which consumers make their choices, but menus are not the same for all individuals (see Puelz and Snow [12] and Chiappori and Salanié [2]). In theory, the most complete view of insurance markets lies between extremes: no party has a superior information and menus signal information from insurers to individuals whereas individuals reveal information to insurers by their choices. Whether the structure of asymmetric information is close to one of the extremes (adverse selection or better informed insurers) is an empirically challenging question. In any case, econometricians have to control for the exact nature of offers and classification techniques in order not to assess asymmetric information where there is none.

[15] These consequences differ from "proselection" (better risks are better covered) as defined in Hemenway [9], since one suspects that individuals there are better informed. He showed that careless people (those who do not fasten their seat belts) tend to have adopted less insurance coverage. The data have been collected at a car rental company by direct observation.

References

[1] K. Bagwell and G. Ramey, Oligopoly Limit Pricing, *RAND Journal of Economics* **22** (1991), 155–172.

[2] P. A. Chiappori and B. Salanié, Testing for Asymmetric Information in Insurance Markets, forthcoming in *Journal of Political Economy.*

[3] I. K. Cho and D. M. Kreps, Signaling Games and Stable Equilibria, *Quarterly Journal of Economics* **102** (1987), 179–221.

[4] K. J. Crocker and A. Snow, The Efficiency Effects of Categorical Discrimination in the Insurance Industry, *Journal of Political Economy* **94** (1987), 444–467.

[5] W. Emons, Credence Goods and Fraudulent Experts, *RAND Journal of Economics* **28** (1997), 107–119.

[6] M. C. Fagart, Compagnies d'assurance informées et équilibre sur le marché de l'assurance, Université de Nanterre, 1996.

[7] D. Fudenberg and J. Tirole, "Game Theory," MIT Press, Cambridge, 1991.

[8] C. Gouriéroux (1997): "Statistics of Insurance," Geneva Lecture, Paris.

[9] D. Hemenway, Propitious Selection, *Quarterly Journal of Economics* **105** (1990), 1063–1069.

[10] M. Hertzendorf and P. B. Overgaard (1997): "Will the High-Quality Producer Please Stand Up? - A Model of Duopoly Signaling," Paper presented at the Econometric Society European Meeting, Toulouse.

[11] E. Maskin and J. Tirole, The Principal-Agent Relationship with an Informed Principal, II : Common Values, *Econometrica* **60** (1992), 1–42.

[12] R. Puelz and A. Snow, Evidence on Adverse Selection : Equilibrium and Cross-Subsidization in the Insurance Market, *Journal of Political Economy* **102** (1994), 236–257.

[13] M. Rothschild and J. E. Stiglitz, Equilibrium in Competitive Insurance Markets : An Essay on the Economics of Imperfect Information, *Quarterly Journal of Economics* **90** (1976), 629–649.

[14] J. E. Stiglitz, Monopoly, Non-Linear Pricing and Imperfect Information : The Insurance Market, *Review of Economic Studies* **44** (1977), 493–510.

[15] M. Spence, "Market Signaling," Harvard University Press, Cambridge, 1974.

[16] B. Villeneuve (1999): "The Consequences for a Monopolistic Insurance Firm of Evaluating Risk Better than Customers: The Adverse Selection Hypothesis Reversed," forthcoming in *The Geneva Papers on Risk and Insurance Theory.*

A Proofs

A.1 Proposition 1

Let $c_H^\bullet$ and $c_L^\bullet$ be the vectors of offers of, respectively, type-H and type-L consumers.

Proposition 6 *A separating equilibrium always exists, moreover:*

1. *The low risk is charged the minimal price $c_L^{\min} = p_L d$. Moreover, at least two of the contracts in $c_L^\bullet$ charge this price.*

2. *If $c_L^{\max} > c_H^{\min}$, the high risk is charged the minimal price $c_H^{\min} = p_H d$. Moreover, at least two of the contracts in $c_H^\bullet$ charge this price.*

3. *If $c_L^{\max} \leq c_H^{\min}$, $[c_H^{\min}, c_H^{\max}] \cup \{C^0\}$ is the set of equilibrium allocations for the high risk. For no-insurance equilibria, the vector of offers is only restricted to be in $(c_L^{\max}, +\infty)^N$. In the other equilibria, the contract c taken in equilibrium is the lowest price in a subset of $[c_H^{\min}, +\infty)^N$ containing c.*

(Proposition 1 is a corollary of this complete version.)

1. and 2. follow the same logic: there exist contracts that are profitable and acceptable whatever the subsequent beliefs of the consumer. These contracts lie in $(c_L^{\min}, c_L^{\max})$ for L, in $(c_H^{\min}, c_L^{\max})$ for H. Competition is therefore total, which explains that two offers at least should and do charge the minimal price (else, the only insurer offering this contract would have some leeway to raise its price). 3. Here the logic is inverted: there does not exist any contract that is both profitable and acceptable for all the possible subsequent beliefs of the consumer. In other words, if we associate to any *profitable* deviation a belief such that the deviant offer is unacceptable (which is always possible by taking, e.g., p_L) and given that the beliefs associated to non-profitable deviations need not be restricted, then deviations are discouraged and the equilibrium is sustained.

A.2 Proposition 2

1. Increasing λ_H increases the set of acceptable contracts for an uninformed agent: $c_{\tilde{p}_0}^{\max} \to c_H^{\max}$ as $\lambda_H \to 1$. This relaxes Condition (e) in the proposition, strictly when the set of pooling is not empty, weakly when it is. 2. N only affects the RHS in Conditions (d) and (e), Proposition **??**; if $p_L > 0$, then (e) becomes violated for a large N. 3. and 4. For $p_L = 0$, existence of a pooling equilibrium is guaranteed, since (e) is satisfied. The proportion of L must be low enough ($c_{\tilde{p}_0}^{\max}$ close enough to $c_H^{\max}$) for the (d) to be met. By continuity, and given that the constraints on pooling equilibria are slack at $p_L = 0$, small non-negative values of p_L and λ_L guarantee existence.

In the general case (contracts in Ω_1), the proof has to be slightly modified. Remark that each C^{i*} is in $\mathcal{A}(\tilde{p}_0) \cap \overline{\mathcal{P}}_L(0) \cap \overline{\mathcal{P}}_H(0) \cap \mathcal{S}_L \cap \mathcal{S}_H$. 1. Proved with the same reasoning with $\mathcal{A}(\tilde{p}_0) \to \mathcal{A}(p_H)$ instead of $c_{\tilde{p}_0}^{\max} \to c_H^{\max}$. 2. If $p_L > 0$, $\mathcal{S}_L \cap \overline{\mathcal{P}}_H(0) = \emptyset$ for N large (any contract profitable on H is second-order dominated for L : deviating is better than sharing). 3. and 4. If $p_L = 0$, no profitable contract can be accepted by an individual believing he is a L. In consequence, no contract is second-order dominated and $\mathcal{S}_L$ contains all the other sets.

A.3 Theorem 1

The necessary conditions on the equilibrium are discussed in the text. The beliefs sustaining the equilibrium satisfy the following restrictions. For equilibrium offers, $\tilde{p}(C_K^{\bullet *}) = p_K, \forall K$. For any K, any i:

$$\forall C \in \mathcal{P}_K(\pi_K^{i*}) : \text{either } C^0 \succ_{\tilde{p}(C_K^{-i*},C)} C \text{ or } \exists j \neq i \text{ s.t. } C_K^{j*} \succ_{\tilde{p}(C_K^{-i*},C)} C \quad (9)$$

In case 1, condition (b) ensures that there always exist appropriate beliefs whenever $\pi_K^{i*} > 0$.

In case 2, π_K^{i*} takes into account the fact that the contract is taken with probability $1/N$. For any possible deviation C by i, second-degree dominance ensures that there exists at least one belief $p_K^i(C)$ thanks to which C_K^{i*} cannot be dominated by C. In consequence the beliefs that are needed may be: (equilibrium beliefs) $\tilde{p}(C_K^{\bullet *}) = p_K, \forall K$; (off-equilibrium beliefs) for any i, j with $j \neq i$, for any $C \neq C_K^j$, $\tilde{p}(C_K^{-j}, C) = p_K^i(C)$. Elsewhere, no restriction is needed.

Given contracts satisfying the conditions, non-emptiness and convexity of the set of beliefs derives from the existence, for each deviation, of upper bounds (one by alternative conditions in equation 9) on beliefs consistent with stability: the larger of these values determine the domain of beliefs rendering the deviation unattractive.

A.4 Proposition 3

$\mathcal{F}_L$ **and** $\mathcal{S}_L$. We prove first that the following conditions are equivalent for any $C \in \mathcal{A}(p_L)$: (a) C is first-degree dominant for L; (b) C is second-degree dominant for L; (c) C is actuarial, and offers at least full insurance, at most the optimal insurance coverage for a high risk who would have access to the low risk's actuarial terms.

Let's prove that (a)⇒(c). Assume that $C = (W_A, W_N) \in \mathcal{A}(p_L)$ is first-degree dominant and strictly profitable on L. For $\epsilon > 0$, $C' = (W_A, W_N + \epsilon)$ is preferred to C and gains participation of the agent (C' offers better conditions than C); by taking small enough an ϵ, the profit generated by C' is necessarily non-negative; therefore C' first-degree dominates C, a contradiction. Let C be an actuarial acceptable contract for L that offers under-insurance;

clearly, C_L^{AF} is strictly better than C for all beliefs (second-order stochastic dominance); it is therefore always possible to find a contract C' first-degree dominating C in a neighborhood of C_L^{AF}. In the same vein, note that the optimal contract for high risks among those which are actuarial for the low risk is preferred by the agent, whatever his beliefs, to any contract offering further over-insurance.

Let's prove now the reciprocal (c)⇒(a). Let C satisfy (c). There exists a belief between p_L and p_H, such that the intersection between profitable contracts and weakly preferred contracts for this belief is $\{C\}$; C cannot be first-degree dominated. To see that (a)⇔(b), remark that when we restrict attention to contracts acceptable for the low risk, first-degree and second-degree dominance are equivalent.

$\mathcal{F}_H$ **and** $\mathcal{S}_H$. We prove the following assertions: (a) **D** if and only if any contract which is (weakly) profitable on H and acceptable for belief p_H is first-degree dominant for type H; (b) **CR** if and only if any contract which is strictly profitable on H and acceptable for belief p_H is first-degree dominated.

(a) **D** if and only if staying uninsured is preferred by low-risk consumers to any contract in $\mathcal{A}(p_H) \cap \mathcal{P}_H(0)$, which, for any contract of this intersection, precludes the existence of a challenging C'. All those contracts belong to $\mathcal{F}_H$. (b) **CR**, if and only if $C_H^{\mathrm{AF}} \in \mathcal{A}(p_L)$, i.e. if and only if for any $C \in \mathcal{A}(p_H) \cap \mathcal{P}_H(0)$, there exists at least one contract in $\mathcal{A}(p_L) \cap \mathcal{P}_H(0)$ in a neighborhood of C_H^{AF} which first-degree dominates C.

A.5 Theorem 2

The necessary conditions are discussed in the text. We prove at the same time non-emptiness and convexity of the set of beliefs supporting an allocation that satisfies the conditions. The equilibrium is supported exactly by beliefs satisfying:

$$\forall K, \forall i, \forall C \in \mathcal{P}_K(\pi_K^{i*}) : \text{either } C^0 \succ_{\tilde{p}(C_K^{-i*},C)} C \text{ or } \exists j \neq i \text{ s.t. } C_K^{j*} \succ_{\tilde{p}(C_K^{-i*},C)} C \tag{10}$$

Each of the two alternative conditions on beliefs furnishes an upper bound on beliefs. The convexity of the beliefs follows. Non-emptiness for contracts satisfying the necessary conditions is a consequence of second-degree dominance.

A.6 Proposition 4

Base case. Let us consider the following three pairs: (a) $C^{\bullet}_{J_s}$and $C^{\bullet}_{J_{s'}}$, (b) $C^{\bullet}_{K_s}$and $C^{\bullet}_{K_{s'}}$, and (c) $C^{\bullet}_{J_{s'}}$ and $C^{\bullet}_{K_s}$. If, for all s, s', the two vectors in the above pairs differ in at least two dimensions (strategies), then $(C^{\bullet}_{J_s}, C^{\bullet}_{K_s})_{s\in S}$ is exactly supported in an equilibrium. To see this, notice that in each of the conditional games, beliefs are restricted only for equilibrium offers (Bayes rule) and for *unilateral* deviations from equilibrium offers (beliefs must be so as to discourage deviations by *one* insurer). Given that beliefs associated with vectors of offers differing in strictly more than one dimension from one of the two equilibrium vector are not restricted at all, the restricted beliefs of all the conditional games can be transplanted in the complete game without interferences. They sustain the allocation.

Other cases. Assume now that one of the conditions above is violated. We deal without loss of generality with the case where $C^{\bullet}_{J_s}$and $C^{\bullet}_{J_{s'}}$ differ by less than two dimensions. Using the fact that the interior of $\mathcal{E}(J_s, K_s)$ and $\mathcal{E}(J_{s'}, K_{s'})$ is not empty, we show that $C^{\bullet}_{J_s}$ and $C^{\bullet}_{J_{s'}}$ can be replaced by arbitrarily close equilibrium allocations so as to eliminate the ambiguity and go back to the base case. We deal separately with two subcases: (a) $(C^{\bullet}_{J_s}, C^{\bullet}_{K_s})$ or $(C^{\bullet}_{J_{s'}}, C^{\bullet}_{K_{s'}})$ is a pooling equilibrium; (b) $(C^{\bullet}_{J_s}, C^{\bullet}_{K_s})$ and $(C^{\bullet}_{J_{s'}}, C^{\bullet}_{K_{s'}})$ are separating.

Subcase (a). For fixing ideas, and without loss of generality, assume that $(C^{\bullet}_{J_s}, C^{\bullet}_{K_s})$ is a pooling. Proposition 2 implies that there is a continuum of pooling equilibria in $\mathcal{E}(J_s, K_s)$. Notice that a pooling equilibrium involving indifference between offers, we have to carefully modify all offers to preserve indifference. This means that there exist $(\varepsilon_i)_{i\in I}$ (with $\varepsilon_i \in I\!R^2$ and $\|\varepsilon_i\| > 0$ as small as necessary), such that, if, for any i, we add vector ε_i to the vectors of conditional transfers (W_A, W_N) implicit in $C^i_{J_s}$ and $C^i_{K_s}$ to eliminate any ambiguity caused by the resemblance with $C^i_{J_{s'}}$ and $C^i_{K_{s'}}$, the modified offers still support a pooling equilibrium. We are back to the base case.

Subcase (b). We have partitioned types, hence $J_s \neq J_{s'}$, and fair contracts are necessarily different for those two types. But due to competitive pressure, in all separating equilibria, low-risk types (in conditional games) are offered at least two (non-trivial) fair contracts (Proposition 3), therefore J_s and $J_{s'}$ cannot be *both* treated as low-risk types in their conditional games. Hence a type, say J_s for fixing ideas, is the high risk in "his" conditional game and is in the situation of Proposition 3, Point 2. In any case, the set of equilibrium allocations for H being a continuum, the offers he gets can be slightly changed to approximate $C^{\bullet}_{J_s}$. We are back to the base case.

A.7 Proposition 5

(1) Let us take a pooling equilibrium. For any i, we denote the proposed contract by C^{i*}, and by π_H^{i*} and π_L^{i*} the associated profits ($\pi_K^{i*} = \pi_K(C^{i*})/N$). Define C^i as the $(N-1,1)$-barycenter of (C^0, C^{i*}); given the linearity of the profit functions: $\pi_K(C^i) = \pi_K^{i*}/N$ for any K. Due to the concavity of the utility function and to the fact that C^{i*} was acceptable for belief $\tilde{p}_0$, C^i is necessarily strictly acceptable for belief p_H, and strictly better than equilibrium offers. Now we invoke a single-crossing argument: according to the refinement, the neighborhood of C^i is cut in four distinct angular sectors in one of which $\tilde{p}(C^{-i*}, \cdot) = p_H$, because there $\pi_H(\cdot) > \pi_H^{i*}/N$ and $\pi_L(\cdot) < \pi_L^{i*}/N$. Insurer i would better offer a contract in this sector: by taking it sufficiently close to C^i, participation of the consumer is guaranteed, and profits are superior. The equilibrium is broken.

(2) Let us fix i and consider $\tilde{p}(C_L^{-i*}, \cdot)$. The refinement impose that for any C^i generating losses on H and profits on L, in particular around $C_L^{\mathrm{AF}}, \tilde{p}(C_L^{-i*}, C^i)$ must be p_L. If the type-L consumer receives a fair but inefficient allocation, then there exist necessarily a deviation by i in the neighborhood of C_L^{AF} that is acceptable, better than the equilibrium offer, and strictly profitable. This eliminates unfair or inefficient allocation to L. We cannot go further. Consider $\tilde{p}(C_H^{-i*}, \cdot)$. Whatever C^i offered by i that would improve its profit on H, C^i would not decrease its equilibrium profit on L (which is zero). In consequence, the beliefs cannot be restricted by the Cho-Kreps criterium and the equilibrium cannot be eliminated.

B Competition in Mechanisms

Maskin and Tirole [11] have somewhat closed a certain type of studies in signalling games by showing that allowing the principal to offer practically unrestricted mechanisms (or game forms), rather than contracts in the narrow sense, guarantees efficiency gains. The following result is mainly conceived for the readers familiar with Maskin and Tirole. It provides a counter-example proving that efficiency gains are not warranted in the multi-principals context.

Proposition 7 *Assume that $p_L = 0$. No-insurance for H and L is an equilibrium allocation of competition in mechanisms.*

Proof. $\mu_H^\bullet$ is the vector of mechanisms offered to H. For any $\mu \in \Omega_2$, any $K, K(\mu)$ denotes the allocation that μ assigns to K. This represents without loss of generality the degree of discretion that the insurer may keep in the offered mechanism. In particular, $L(\mu_H^i)$ is the allocation that a L would receive from i if i proposes μ_H^i to L, and if this offer is accepted by L. This will matter for incentive compatibility.

For any i, any $\mu \in \Omega_2$, we distinguish two cases (a) $L(\mu)$ makes *strictly* positive profits on L if accepted, and (b) $L(\mu)$ makes negative profits on L, if accepted. We check now that the following beliefs sustain the equilibrium allocation (no-insurance for all): if (a) then $\Lambda(\mu_H^{-i}, \mu) = (0, 1)$; if (b) then $\Lambda(\mu_H^{-i}, \mu) = (1, 0)$.

Case (a). μ breaks the equilibrium only if $L(\mu)$ is strictly preferable to no-insurance for belief p_L : indeed, given his beliefs, the individual acts as if he was offered $L(\mu)$. Compared to no-insurance which is the first-best when $p_L = 0, \mu$ would have to improve strictly the insurer's and the consumer's utility. This is impossible.

Case (b). μ breaks the equilibrium only if (i) $H(\mu)$ is strictly preferable to no-insurance for belief p_H (the value of μ for the individual is the value of $H(\mu)$, (ii) $H(\mu)$ is strictly profitable on H (the insurer has to be induced to deviate) and (iii) more profitable than $L(\mu)$ on H (incentive compatibility). Remark that, without loss of generality, $K(\mu)$ can be seen as a random allocation of insurance contracts, "contract" being taken in the ordinary sense. More precisely, all works as if $K(\mu)$ were specifying the random rules according to which the pairs (W_A, W_N) were allocated. Let us define $\tilde{\mu}$ as the insurance (ordinary) contract corresponding to the expected terms of $H(\mu)$ ($\tilde{\mu} = (E(W_A\,|H(\mu)), E(W_N\,|H(\mu))))$. Clearly, due to the risk neutrality of the insurer and risk aversion of the consumer, if μ breaks the equilibrium, then (i) $\tilde{\mu}$ is also strictly preferable to no-insurance for belief p_H, (ii) $\tilde{\mu}$ is also strictly profitable on H, and (iii) still more profitable than $L(\mu)$ on H. If we look at the contracts satisfying these conditions (one can visualize them in the plane $(W_A, -W_N)$), it is clear that $\tilde{\mu}$ (hence $H(\mu)$) necessary generates strictly positive profits on L. On the other hand, incentive compatibility ensures that $L(\mu)$ is more profitable than $H(\mu)$ on L. This is a contradiction with (b), and μ doesn't break the equilibrium. ■

The interpretation that we retain here is that the mere possibility that a consumer is not subjected to the risk is sufficient for high risk to be exaggeratedly distrustful so as to receive no-insurance.

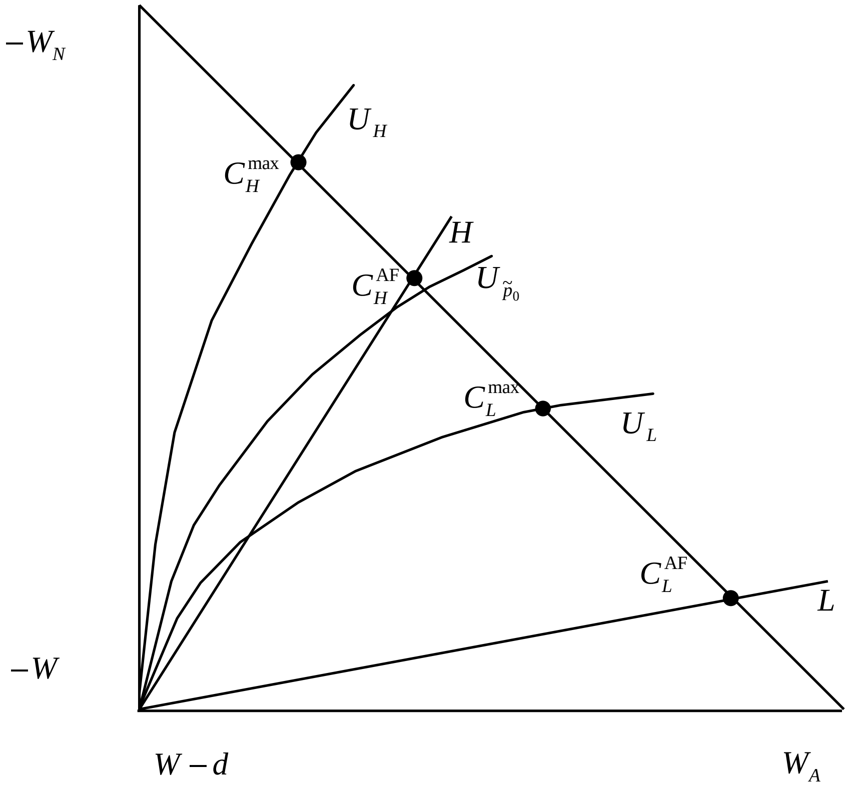

Figure