Characterization, Existence, and Pareto Optimality in Insurance Markets with Asymmetric Information with Endogenous and Asymmetric Disclosures: Revisiting Rothschild-Stiglitz

By JOSEPH E. STIGLITZ, JUNGYOLL YUN, AND ANDREW KOSENKO *

Abstract: We study the Rothschild-Stiglitz model of competitive insurance markets with endogenous information disclosure by both firms and consumers. We show that an equilibrium always exists (even without the single crossing property), and characterize the unique equilibrium allocation. With two types of consumers the outcome is particularly simple, consisting of a pooling allocation which maximizes the well-being of the low risk individual (along the zero profit pooling line) plus a supplemental (undisclosed and nonexclusive) contract that brings the high risk individual to full insurance (at his own odds). We show that this outcome is extremely robust and Pareto efficient. (JEL D43, D82, D86)

Some forty years ago, Rothschild and Stiglitz (1976) characterized equilibrium in a competitive market with exogenous information asymmetries in which market participants had full knowledge of insurance purchases. Self-selection constraints affected individual choices; but unlike the monopoly equilibrium[1], no single firm framed the set of contracts among which individuals chose. There never existed a pooling equilibrium (in which the two types bought the same policy); if there existed an equilibrium, it entailed the high risk getting full insurance, and the low risk individual only getting partial insurance; and under plausible conditions—e.g. if the two types were not too different—a pure strategy equilibrium did not exist. The paper was unsatisfactory not only in its results (equilibrium *seemed* to exist, and often entailed pooling) but on its reliance on a special property, called the single crossing property, whereby the indifference curve of the high risk individual could cross that of the low risk individual only once.[2]

* Stiglitz: University Professor, Columbia University, Uris Hall, Room 212, 3022 Broadway, New York, NY 10027 (email: jes322@columbia.edu, cc: debarati.ghosh@columbia.edu). Yun: Professor, Department of Economics, Ewha University, Seoul, Korea (email: jyyun@ewha.ac.kr). Kosenko: doctoral student, Columbia University, 518 West 111th st, apt 24, New York, NY 10025 (email: ak2912@columbia.edu). We are grateful to Gerry Jaynes for helpful comments on an earlier draft, to Michael Rothschild and Richard Arnott, long time collaborators, to Debarati Ghosh and Lim Nayeon for research and editorial assistance and to the Institute for New Economic Thinking and the Ford Foundation and Fulbright Foundation for financial support.

[1] Stiglitz (1977).

[2] As innocuous as it might seem, it won't be satisfied if the high and low risk individuals differ in their risk aversion;

Since their work, there has been huge literature applying the model to labor, capital, and product markets in a variety of contexts, a large number of empirical applications, and a small literature trying to repair the deficiencies in the underlying framework by formalizing the insurance "game", by changing the information/disclosure assumptions, and by changing the equilibrium concept.

This paper takes an approach that differs fundamentally from this earlier literature by endogenizing the disclosure of information about insurance purchases: each firm and consumer makes a decision about what information to disclose to whom—thus information about contract purchases is not only endogenous but potentially asymmetric. The results were somewhat surprising even to us: (i) asymmetries in information about insurance purchases, especially associated with out of equilibrium moves, do indeed turn out to be important; (ii) there always exists an equilibrium, even when the single crossing property is not satisfied; and (iii) the equilibrium always entails a pooling contract. Indeed, the unique insurance allocation (an insurance allocation describes the sum of benefits and premia for each individual) consists of the pooling allocation which maximizes the well-being of the low risk individual (along the zero profit pooling line) *plus* a supplemental contract that brings the high risk individual to full insurance (at his own odds). While the equilibrium allocation is unique, it can be supported by alternative information strategies.

We begin the analysis by characterizing the set of Pareto efficient (PE) allocations in the presence of a possibly secret contract. We then show that the PE allocation which maximizes the well-being of the low risk individual is the unique equilibrium allocation and can be supported by simple information disclosure strategies.

While the analysis is complex, it is built upon a number of steps, each of which itself is relatively simple. As in RS, insurance firms offer insurance contracts, but now they may or may not decide to reveal information (all or partial) about insurance purchases to other firms. In RS, it was assumed that contracts were exclusive, e.g. implicitly, that if a firm discovered a purchaser had violated the exclusivity restriction, the coverage would be cancelled. Here, we consider a broader range of possible restrictions. Obviously, the enforceability of any conditions imposed is dependent on information available to the insurance firm. Consumers, too, have a slightly more complicated life than in RS: they have to decide which policies to buy, aware of the restrictions

and with multi-crossings, equilibrium, if it exists, can look markedly different.

in place and the information that the insurance firm may have to enforce those restrictions. And they also have to decide on what information to reveal to whom[3].

As in RS, a competitive equilibrium is described by a set of insurance contracts, such that no one can offer an alternative contract or set of contracts and make money. Here, though, a contract is defined not just by the benefit and the premium, but also by the restrictions associated with the contract and the firm's disclosure policy.

The paper is divided into 12 sections. In the first, we set out the standard insurance model. In the second we recall why RS resorted to exclusive contracts. We explain how the existence of a (non-loss making) secret contract offered at the odds of the high risk individual (a) upsets the separating equilibrium; (b) implies that some of the contracts that broke the pooling contract no longer do so; but (c) there always exist some contracts that nevertheless break the relevant pooling allocation. Section 3 then shows that if there is a non-disclosed contract (at the odds of the high risk individual), the Pareto efficient contracts are always of a simple form: pooling plus supplemental insurance purchased only by high risk individuals. Section 4 then defines the competitive equilibrium. Section 5 shows that regardless of the strategies, if there is a competitive equilibrium, the allocation must be the Pareto efficient allocation which maximizes the wellbeing of the low risk individual. Section 6 then describes equilibrium strategies for firms and consumers, shows that the posited strategies support the equilibrium allocation described in the previous section, and are robust against any deviant contract. Section 7 comments on several salient properties of the result and its proof, including that it does not require the single crossing property, but only a much weaker condition. Section 8 and 9 discuss uniqueness of equilibria and show how the equilibrium construct can be extended, for instance to other disclosure strategies and to multiple types of individuals. Sections 10 and 11 relate our results to earlier literature. In particular, section 11 considers the standard adverse selection price equilibrium. We show how our analysis implies that in general a price equilibrium does not exist if there can exist a (non-loss making) insurance contract the purchase of which is not disclosed. Section 12 presents some concluding comments.

1. The model

[3] We assume that consumers can only reveal information to firms, and not to other consumers. Since the game is one of private values, revealing information to other consumers is moot, and therefore we disallow it without loss of generality.

We employ the standard insurance model with adverse selection. An individual is faced with the risk of an accident with some probability, P_i. P_i depends upon the type i of the individual. There are two types of individuals – high risk and low-risk-- who differ from each other only in the probability of accident. The type is privately known to the individual, while the portion θ of H-type is common knowledge. The average probability of accident for an individual is $\bar{P}$, where

$$\bar{P} \equiv \theta P_H + (1-\theta)P_L.$$

An accident involves damages. The cost of repairing the damage in full is d. An insurance firm pays a part of the repair cost, $\alpha \leq d$. The benefit is paid in the event of accident, whereas the insurer is paid insurance premium β when no accident occurs.[4] The price of insurance, q, is defined by $\frac{\beta}{\alpha}$. (In market equilibrium, the amount of insurance that an individual can buy may be limited.) The expected utility for an individual with a contract (α, β) is

$$V_i(\alpha,\beta) = P_i U(w-d+\alpha) + (1-P_i)U(w-\beta) \qquad (1)$$

where the Bernoulli utility function U is quasi-concave and differentiable, with U" < 0 (individuals are risk averse). Sometimes we refer to a contract A ≡ {α, β}, in which case we can refer to the expected utility generated by that contract as $V_i\{A\}$[5]. Under (1), an indifference curve for high-risk individual is steeper than that for low-risk one at any (α, β), satisfying the so-called the single-crossing property. As will be shown later in the paper, however, we can allow for more general preferences, e.g. with a different utility function $U_i(.)$for each type i.[6] In this case, the single crossing property will not be satisfied. The key property of $V_i(\alpha,\beta)$ is that the income consumption curve at the insurance price $\frac{P_i}{1-P_i}$ is the full insurance line,[7] implying that at full insurance, the slope of the indifference curve equals the relative probabilities,

$$\frac{\frac{\partial V_i(\alpha,\beta)}{\partial \beta}}{\frac{\partial V_i(\alpha,\beta)}{\partial \alpha}} = \frac{P_i}{1-P_i}$$

---

[4] This has become the standard formulation since RS. In practice, customers pay β the period before the (potential) accident, receiving back α + β in the event the accident occurs, i.e. a net receipt of α.

[5] Similarly, if the individual purchases policies A and B, we can refer to the expected utility generated as V_i {A + B}.

[6] Indeed, we do not even require preferences to satisfy the conditions required for behavior towards risk to be described by expected utility. We do not even require quasi-concavity.

[7] That is even if the indifference curve is not quasi concave, after being tangent to a given isocline with slope $\frac{P_i}{1-P_i}$, at full insurance, it never touches the isocline again.

so that will full information, equilibrium would entail full insurance for each type at their own odds. We retain this key assumption throughout the paper. There are N firms and the identity of a firm j is represented by j, where j = 1,--, N. The profit π_i of a contract (α, β) that is chosen by i-type (i=H,L) is $\pi_i(\alpha, \beta) = (1 - P_i)\beta - P_i\alpha$. Figure 1 illustrates the zero-profit locus for a firm selling insurance to an i-type or both types of individuals by a line from the origin with the slope being $\frac{P_i}{1-P_i}$ or $\frac{\bar{P}}{1-\bar{P}}$, respectively.

[FIGURE 1 ABOUT HERE]

2. Rothschild-Stiglitz with secret contracts

Central to the analysis of Rothschild and Stiglitz was the assumption that there was sufficient information to enforce exclusivity; the individual could not buy insurance from more than one firm. As RS realized, once we introduce into the RS analysis unobservable contracts in addition to the observable ones, the whole RS framework collapses. Exclusivity cannot be enforced. In this section, we review why they assumed exclusivity; we assume that undisclosed contracts can and will be offered if they at least break-even. In particular, we know that a price contract (where the individual can buy as much of the given insurance at the given price) with a price $\frac{P_H}{1-P_H}$ will at least break even: if it is bought by any low risk individual, it makes a profit.

Breaking a separating equilibrium. When there is secret supplemental insurance, the implicit self-selection constraints change, because whether an individual prefers contract A rather B depends on whether an individual prefers A *plus* the optimally chosen secret contract to B *plus* the optimally chosen secret contract. Thus, in figure 1, the high risk individual prefers the contract which puts him on the highest isocline line with slope $\frac{P_H}{1-P_H}$.

Consider the standard RS equilibrium separating contracts, C and B. C is the full insurance contract for the high risk individual assuming he was not subsidized or taxed and B is the contract on the low risk individual's break-even curve that just separates, i.e. is not purchased by the high risk individual.[8] {B, C} can *never* be an equilibrium if there can be undisclosed contracts, because if there were a secret offer of a supplemental contract at a price reflecting the "odds" of the high risk individual, then the high risk individuals would buy B plus supplemental insurance bringing

[8] In RS, the pair of contracts {B, C} constitutes the equilibrium so long as B is preferred to the contract on the pooling line which is most preferred by the low risk individual. If this is not true, there exists no equilibrium.

him to C'.[9] B and C no longer separate. (Later, we show that there is in fact no alternative set of separating observed contracts.)

Breaking a pooling equilibrium with no disclosure of deviant policy. RS showed that there could be no pooling equilibrium by showing that because of the single crossing property, there always exists contracts preferred by the low risk individual and not by the high risk which lie below the pooling zero profit line and above the low risk zero profit line. But the ability to supplement the breaking contract *may* make the contracts which broke the pooling equilibrium, under the assumption of no hidden contracts, attractive to the high risk individual. Such a contract cannot break the pooling equilibrium.

[FIGURE 2 ABOUT HERE]

Figure 2 provides an illustration. The pooling contract A* is the most preferred policy of the low risk type along the pooling line with slope $\frac{\bar{P}}{1-\bar{P}}$,[10] the only possible pooling equilibrium. Consider the high-risk price line through A*. The high risk individual also purchases the insurance contract A*, thereby obtaining a subsidy from the low risk individual, and supplements it with secret insurance at the high risk odds (represented in Figure 2 by A*C*, where C* is the full insurance point along the line through A* with slope $\frac{P_H}{1-P_H}$.)[11] Consider a policy D_o below the low risk individual's indifference curve through A*, above that for the high risk individual, and which also lies below the zero profit line for high risk individuals through A*. In the RS analysis, *with exclusivity,* D_o would have broken the pooling equilibrium A*. Now, it does not, because the high risk individuals would buy D_o and the (secret) supplemental insurance.[12] And if they do so, then D_o makes a loss, and so D_o could not break the pooling equilibrium.

But the question is, are there *any* policies which could be offered that would break the pooling equilibrium, that would be taken up by the low risk individuals, but not by the high risk individuals *even if they could supplement the contract with a secret contract breaking even.* The answer is yes.

---

[9] This result follows directly from the fact that the implicit price of B is $\frac{P_L}{1-P_L} < \frac{P_H}{1-P_H}$.

[10] Sometimes referred to as the Wilson equilibrium. Obviously, any other posited pooling equilibrium could be broken by A*, since it would be purchased by all the low risk individuals.

[11] Recall that at full insurance, the slope of the indifference curve of the high risk individual is just $\frac{P_H}{1-P_H}$, and full insurance entails $\alpha = d - \beta$.

[12] This is different from the way that the matter was framed by Wilson and Riley, who described the policy A as being withdrawn when a policy such as D_o is offered (which is why their equilibrium concepts are typically described as *reactive*). Here, when D_o is offered, A* is not withdrawn, but nonetheless, because of the secret contract, high risk individuals prefer D_o to A*. See the fuller discussion in the next sections.

There are policies which lie below the zero profit pooling line and above the zero profit line for low risk individuals (that is, would make a profit if purchased only by low risk individuals), below the low risk individual's indifference curve (i.e. are preferred by low risk individuals), and lie *above* the high-risk zero profit line through A* (i.e. even if the high risk individual could have secretly supplemented his purchases with insurance at his actuarial fair odds, he would be worse off than simply purchasing A*). These policies break the pooling contract. In Figure 2, any point (such as D) in the shaded area in the figure, which we denote by z, can thus break the pooling equilibrium. The set z is not empty because the low risk individual's indifference curve is tangent to the pooling line at A*.[13] Formally, for any point such as D, $V_L\{D\} > V_L\{A^*\}$, while $V_H\{D + S_H\} > V_H\{A^*\}$.[14] We collect the results together in

Proposition 1.

(a) *The RS Separating Contracts do not constitute an equilibrium, if firms can offer non-loss making undisclosed contracts.*

(b) *The pooling equilibrium may always be "broken" if there exists undisclosed supplemental insurance and if a deviant firm can choose to keep his offers secret.*

(c) *Some of the contracts that broke the pooling equilibrium in the standard RS equilibrium with exclusivity no longer do so.*

The remaining sections focus on the core issue of an endogenous information structure, with the simultaneous determination of contract offers of firms and with contract purchases and information disclosure by individual customers.

3. Pareto efficiency with undisclosed contracts

In this section, we consider the set of efficient insurance allocations under the premise that there exists a secret (undisclosed) contract being offered at the price $\frac{P_H}{1-P_H}$. We can think of this as a "constrained P.E." allocation—where the constraint is that the government cannot proscribe the

[13] Of course, *if* the offer of the deviant contract were public, sellers of contract A* could make their offer conditional on there not being a contract in z being offered, in which any such contract would lose money. This is in the spirit of Wilson's discussion of "reactive" equilibria, which in turn is *not* in the spirit of competitive equilibria. However, here, firms can chose not to disclose either their offer of insurance or individual's purchase of insurance. (The assumption of non-disclosure of offers is not fully satisfactory in the context of *market* insurance, since if consumers know about a firm selling insurance, presumably so could other insurance firms. But in fact much insurance is non-market insurance (see Arnott and Stiglitz (1991b)), often implicit and not formal, and whether such insurance is available to any individual let alone taken up by him may not be known.

[14] The notation $D + S_H$ refers to the {α, β} associated with the purchase of D plus the optimized value of secret insurance along the price line associated with the high risk individual. Given our assumptions about preferences, we know this brings the high risk individual to full insurance.

secret provision of insurance, unlike the P.E. allocations associated with the RS model, where government could restrain such provision.[15] The difficulties in defining Pareto efficiency in settings of incomplete information are not new[16]; we use the following ex-interim variant of constrained Pareto efficiency[17]:

Definition 1. An allocation $\mathrm{E} = \{(\alpha_i, \beta_i)\}_i$ is constrained Pareto-efficient if the government cannot force disclosure and there does not exist another feasible allocation (i.e. one which at least breaks even), and leaves each type of consumer as well off and at least one type strictly better off.

For simplicity of exposition, in this section we assume (1) is satisfied. We now establish two general properties that a P.E. allocation must satisfy:

Lemma 1. Every Pareto efficient allocation must be a separating allocation (i.e. one where the two types of individuals get different allocations), except possibly for the point along the pooling line providing full insurance.

Any feasible (i.e. making at least zero profit for the firms) pooling allocation must lie on the pooling line. At any point other than full insurance, the utility of the high risk individual will be improved by a pair of allocations (A^* and C^* in Figure 3, for example), that along the pooling line and that bringing the high risk individual to full insurance from there.

[FIGURE 3 ABOUT HERE]

Lemma 2. Every Pareto efficient allocation must entail full insurance for high-risk individuals.

This follows directly from our assumptions on V, quasi-concavity and that at full insurance, the slope equals $\frac{P_H}{1-P_H}$.[18] Define A^* as the point on the pooling line most preferred by the low risk individual, or, more formally, as an allocation $(\bar{\alpha}, \bar{\beta})$ such that

$$\bar{\alpha} = Argmax_\alpha V_L\left(\alpha,\ \frac{\bar{P}}{1-\bar{P}}\alpha\right) \text{ and } \bar{\beta} = \frac{\bar{P}}{1-\bar{P}}\bar{\alpha} \qquad (2)$$

Also, define C^* as a full-insurance point along the line through A^* with slope $\frac{P_H}{1-P_H}$, which can be represented as an allocation (α_H^*, β_H^*) such that

[15] The analysis of P.E. allocations in the RS model is in Stiglitz (2009). The terminology may be confusing. It focuses on the constraints imposed by the government—that it cannot restrict the secret sale of insurance. From the perspective of the market, of course, it is an "unconstrained" equilibrium—they do not face the constraint of disclosing.

[16] See Holmstrom and Myerson (1981)

[17] See also Greenwald and Stiglitz (1986)

[18] It should be clear that these are sufficient conditions. All that is required, as noted above, is that the income consumption curve at the insurance price $\frac{P_H}{1-P_H}$ is the full insurance line. A sufficient condition for this are the restrictions set forth for (1).

$$\alpha_H^* + \beta_H^* = d, \text{ and } \beta_H^* - \bar{\beta} = \frac{P_H}{1-P_H}(\alpha_H^* - \bar{\alpha}) \tag{3}$$

Consider contract pairs {A', C'} in figure 3 where A' lies along the pooling line and C' is the full insurance point along the line through A' with slope $\frac{P_H}{1-P_H}$, or where $A' \equiv (\bar{\alpha}', \bar{\beta}')$ and $C' \equiv (\alpha_H', \beta_H')$ such that

$$\bar{\beta}' = \frac{\bar{P}}{1-\bar{P}}\bar{\alpha}' \tag{4}$$

$$\alpha_H' + \beta_H' = d, \text{ and } \beta_H' - \bar{\beta}' = \frac{P_H}{1-P_H}(\alpha_H' - \bar{\alpha}') \tag{5}$$

All such pairs are feasible outcomes. Then for an allocation {A', C'} such that $\bar{\alpha}' < \bar{\alpha}$, an increase in insurance improves the utility of both the high and low risk individuals, so such allocations cannot be P.E. Consider now a contract pair {A', C'} such that $\bar{\alpha}' > \bar{\alpha}$ as in Figure 3. Given C' and the existence of secret contract, is there an alternative feasible allocation preferred by low risk individuals? Any contract purchases *just* by low risk individuals must lie on or above the line through A' with slope $\frac{P_L}{1-P_L}$, because otherwise it is not feasible; and on or above the line through A' with slope $\frac{P_H}{1-P_H}$, because otherwise it would be chosen by both the high risk and low risk individual. The only contract satisfying these two conditions is A'. On the other hand, any feasible contract purchased by both types must lie along the pooling line. Along the pooling line, any allocation that makes the low risk individual better off (by moving towards A^*) makes the high risk individual worse off. Quasi-concavity of the indifference curves ensures that the low risk individual's indifference curve through A' has a slope that is steeper than $\frac{P_L}{1-P_L}$. Hence, there exists no Pareto improvement over {A',C'}. We have thus fully characterized the set of Pareto efficient allocations.

Proposition 2. The set of P.E. allocations are those generated by an allocation $(\bar{\alpha}', \bar{\beta}')$ (defined by (4)) *along the pooling line, such that* $\bar{\alpha}' \geq \bar{\alpha}$ *and* $\bar{\alpha}' + \bar{\beta}' \leq d$, *for the low risk individual; and by an allocation* (α_H', β_H') (defined by (4) and (5)) *for the high risk individual.*

4. Definition of market equilibrium

In this section, we define the market equilibrium.

4-1. Contract Offers by Firms and Optimal Responses by Consumers

Firms move first, making a set of contract offers.[19] A contract $C_k(=\{\alpha_k, \beta_k, R_k, D_k\})$ offered by a firm k is represented by a benefit α_k, if the accident occurs, a premium β_k, if it does not, a set R_k of restrictions that have to be met for the purchase of (α_k, β_k), and a rule D_k of disclosing information at the firm's disposal, such as about (α_k, β_k) sold to individual i. The restrictions R_k, to be relevant, must be based on observables, i.e. what is revealed to the insurance firm k either by the insured i or by other insurance firms; and we assume that they relate only to the purchases of insurance by the insured; they may entail, for instance, a minimum or maximum amount of insurance obtained from others. The exclusivity provision of RS is an example of a restriction, but there are obviously many potential others.

Two simple disclosure rules would be to disclose the purchase to every other firm, or to disclose the purchase to no firm. The equilibrium disclosure rules to be described below will turn out to be somewhat more complex than these simple rules, but still relatively simple.

Following this, households look at the set of contracts on offer (including the restrictions and disclosure policies) and choose the set of contracts that maximizes their expected utility, given the contract constraints.

Consumers also have an information revelation strategy, e.g. what information (about their purchases) to disclose to whom, taking into consideration disclosure policies and contract offers firms announce. In the central model of this paper, the individual simply reveals the quantity of pooling insurance purchased to those firms from whom he has purchased a pooling contract. In an alternative formulation described briefly in Appendix C, he also tells the price at which he has purchased insurance. Of course, firms anticipate their responses—both their purchases and disclosures.

There is a third period which just entails the "working out" of the consequences of the first two—no new action is taken. The third period takes place in two stages. In the first, firms disclose information according the disclosure rules they announced. In the second, each firm checks to see whether any contract restriction is violated, and if it is, that policy is cancelled. Actually, life is easier than just described, since consumers who always respond optimally to any set of contracts

[19] The firm knows nothing about the individual, other than information about contract purchases. The firm may make inferences about the individual based on the information it has about his purchases.

offered by firms know that if they violate contract provisions, policies will be cancelled[20]; and in this model, there is no strategic value of buying policies which will be cancelled.[21]

4.2. Information Disclosure

As we noted, both consumers and firms disclose information on the contracts they have purchased and sold. We assume that both can withhold information from others.[22] The firm or the consumer can disclose just the amount of insurance (α) or the price or β. Also, as a means of partial revelation of information, a firm might engage in what we call contract manipulation (CM) – dividing its sales to an individual into multiple policies. This would allow a consumer to disclose to others one policy, but to hide the full extent of his insurance purchases. As will be shown below, however, no firm sells an individual multiple contracts in equilibrium, so that no CM occurs in equilibrium.

Suppressing i for notational simplicity, we denote by Ω_k^c and Ω_k^f the information revealed to firm k by consumer i and by the other firms, respectively. The information disclosure rule D_k of a contract specifies what information about individual i firm k reveals to firm j. We assume that the information revealed is a subset of the information Ω_k^c that the firm has on individual i obtained from individual i and the information about its own sale (α_k, β_k) to the individual. Similarly, the decision as to whom to disclose is based upon $\{\Omega_k^c, (\alpha_k, \beta_k)\}$. The disclosure rule of firm k can thus be represented by $D_k(\Omega_k^c; (\alpha_k, \beta_k))$[23][24]. Firms can engage in discriminatory revelation, revealing information to some firms not revealed to others, thus creating an asymmetry of information about the insurance coverage of any individual. If there is discriminatory disclosure, the discrimination has to be based on some information Ω_k^c previously disclosed by the insured to the firm.[25]

[20] That is, no policy is cancelled even out-of-equilibrium as well as in equilibrium.

[21] This is not a repeated game. Consumers are engaging in a “rational expectations best response strategy,” which includes identifying which deceptions are caught out, and since such policies are cancelled, not undertaking them

[22] We assume agents cannot lie; a consumer or his insurer cannot “reveal” that he purchased insurance from a firm when no such purchase happened. In short, they tell the truth, nothing but the truth, but not necessarily the whole truth. We do not analyze the game where firms are free to engage in strategic disinformation. We do allow a contract to be shown with redacted information (the truth, but not the whole truth.)

[23] Note that, as contrasted with Jaynes(1978) and Hellwig(1988), the disclosure rule of a firm is not conditional upon contract offers made by other firms.

[24] In a slightly more general specification of the game, firms can disclose information that is revealed to them by other firms. In this case, the third stage of the game has to be extended, to have a series of rounds of disclosure, i.e. as each firm receives information from other firms (based on their announced disclosure rule), it discloses some or all of what has been disclosed to it.

[25] We do not consider random disclosures.

4.3. Equilibrium

Our equilibrium definition is a straightforward generalization of that of RS, where a set of contracts was an equilibrium if there did not exist another contract (or set of contracts) which could be introduced, be purchased by someone, and make a profit (or at least break even.) Here, contracts are defined by the quadruplet {α, β, R, D}. We denote the set of contract offers of firm k by strategy S_k.

Definition (Equilibrium). An equilibrium is a strategy S_k^* *for each firm k, such that, given the set* $\{S_k^*\}_k$ $(\equiv S^*)$ *of strategies adopted by other firms, there does not exist any other strategy that firm j can adopt to increase its profits, once consumers optimally respond to any sets of strategies announced by firms.*[26]

Firms In Rothschild-Stiglitz, each firm offered only one insurance contract. It turned out that some of the results were sensitive to this somewhat artificial restriction. The results established here do not require that the firm offer a single contract, but the proofs are greatly simplified if we restrict the set of contracts it can offer all to have the same price. In appendix D, we establish the results for the more general case. The set of contracts offered can be discrete, or the firm may offer a continuum of contracts, e.g. any amount of insurance up to some upper bound at a price q.

As the restrictions and the disclosure rules that can be specified by a contract may in general be complex, the strategy space for a firm may also be quite complex. We allow a firm to impose any set of restrictions it wants and to set any disclosure rule it wants. Our purpose, however, is to show that there is a simple strategy that supports the equilibrium allocation, and thus we do not need to consider the most general strategy space possible. We assume that the only information that k takes into account in deciding what information about i to reveal to which other firms is information about purchases of contracts by i.[27] We will focus upon a set of disclosure rules that may discriminate in whom to disclose to but that disclose the same information to all the firms for whom there is disclosure.

[26] We formulate the model with a fixed number of firms, so the deviation occurs on the part of one of those firms. But we could as well have allowed free entry. Note too that the optimal responses of consumers includes responses both about contract choices and disclosures.

[27] This restriction has no consequences. The central theorem established later that all equilibrium allocations must be of a particular form holds regardless of the information strategies. We observe later too that that allocation can be supported by multiple information strategies *within this restricted set of strategies.* We have not investigated whether there exist still other information strategies that support the equilibrium allocation within the more general unrestricted set of strategies.

The disclosure rule in the key theorem will disclose only quantities purchased, and only to those for whom the firm has no information from the consumer that there has been an insurance purchase. In the appendix C, we consider an alternative disclosure rule, disclosing price as well as quantity purchased, which supports the same equilibrium allocation.

5. Equilibrium allocations

In this section, we show that the only possible equilibrium allocation is E* ≡{A*,C*}, the P.E. allocation in the presence of undisclosed insurance which maximizes the well-being of the low risk individual. This is true regardless of the strategies of various firms. The analysis is based simply on showing that for any other posited equilibrium allocation, it is possible for an entrant to attract all of the (low risk) consumers and make a profit; hence that allocation could not be an equilibrium allocation.

The result is almost trivial: assume that there were some other allocation, generated by any set of contracts purchased from any array of insurance firms, that was not P.E. Then there exists a contract A" that a deviant firm could offer (entailing equal or more insurance than A*), selling only one policy to each individual, which would at least break even and be purchased by all individuals, with high risk individuals supplementing that contract with secret insurance to bring the high risk individual to full insurance. The putative equilibrium can easily be broken.

Now assume an equilibrium with a P.E. allocation other than E^*. Then a firm could offer a contract A*, and it would be taken up only by the low risk individual, and so would be profitable. Notice that these results hold regardless of the strategies of incumbent firms. We have thus far established the following Theorem.

Theorem 1: There exists a unique allocation E^* that an equilibrium, if it exists, has to implement.

6. Equilibrium

In establishing the existence of an equilibrium, we will first introduce a posited equilibrium strategy S_k^* and then prove that it supports the equilibrium allocation described above and that it is resilient against any deviancy. We assume that there are a set of firms, k = M+1, -- ,N, that sell the secret contracts at price $q_H (=\frac{P_H}{1-P_H})$. Their strategy is simply to sell to anyone any amount of insurance at the price q_H, without disclosing their sales to anyone.

We now describe the firm strategies S_k^* for the remaining firms, which we refer to as the established firms. (a) They each offer insurance at the pooling price $\bar{q}$ $(=\frac{\bar{P}}{1-\bar{P}})$ with (b) the restriction R_k^* that no individual is allowed to purchase in total (so far as they know) more than $\overline{\alpha}$, the amount of insurance that maximizes the welfare of the low risk individual, i.e., $\alpha_k + \sum_{j \neq k} \tilde{\alpha}_j \leq \bar{\alpha}$, where α_k is the amount of pooling insurance to be purchased from firm k while $\tilde{\alpha}_j$ is the amount of pooling insurance revealed by an individual to have been purchased from firm j . If an individual is revealed to the k^{th} firm to have purchased more than this, the k^{th} firm cancels his policy. (c) Their information disclosure rule D_k^* is equally simple: they disclose *everything* they know about the levels of insurance purchases by individual i to every firm which has not been disclosed to them by individual i as selling insurance to him, and disclose nothing to any firm which has been disclosed by individual i to have sold insurance to him.

Several features of the equilibrium strategy S_k^* are worth noting. First, it is conditional only upon the revealed amount $\tilde{\alpha}_j$ of insurance, not upon the revealed price $\tilde{\beta}_j$ of insurance.[28] Second, it does not entail any latent strategy, that is a strategy that is implemented only in an out-of-equilibrium state. Third, the strategy entails *differential* information disclosure based upon consumer-disclosed information. This is critical in sustaining an equilibrium. Without consumer disclosure in the model, it would be impossible for any Nash disclosure strategy to entail differential information disclosure.[29] And without differential information disclosure, it is impossible to sustain the pooling equilibrium. There has to be some information disclosure to prevent high risk individuals "over-purchasing" the pooling contract. But with full information disclosure (of purchases of pooling contracts), exclusivity can be enforced, and hence the pooling equilibrium can always be broken. We will further emphasize below the importance of asymmetric information disclosure both in implementing E^* and sustaining it against any deviancy.

In showing that the equilibrium strategy S_k^* implements E^*, we first prove the following lemma:

Lemma 3 In equilibrium, no firm sells more than one contract to an individual.

Lemma 3 implies that there is no contract manipulation (CM) in equilibrium. Note first that no low-risk individual would prefer to have multiple contracts from his insurer rather than a single contract, as he purchases the most preferred amount of pooling insurance in equilibrium. It is only high-risk individuals who may want to have multiple contracts from their insurers in order to under

[28] The fact that insurance sales are *conditional* on the sales of other firms does not mean that this is a reactive equilibrium. In the reactive equilibrium, e.g. of Wilson, *offers* of insurance are withdrawn when any other firm makes a particular offer.

[29] See also Hellwig (1988).

report their purchases to other potential insurers, to enable them to purchase more pooling insurance. Knowing this, no firm would offer its customer more than one contract without charging a price at least equal to $\frac{P_H}{1-P_H}$. But high risk individuals would not accept it because they are at least as well of purchasing secret insurance at the price $\frac{P_H}{1-P_H}$.[30]

Given Lemma 3, we can show that consumers' best response to S_k^* consists of no individual buying more than $\overline{\alpha}$,which in turn implies that all purchase just $\overline{\alpha}$.

Lemma 4. With the equilibrium strategy S_k^* *, no individual purchases more than* $\overline{\alpha}$ *from the established firms.*

While a formal proof is given in Appendix A, the intuition is clear. Assume he did. He either fully discloses that he did or does not. If he discloses fully, then given S_k^* all the insurance contracts will be cancelled. So he would not disclose. If he does not disclose some contract, say with firm j, then under S_k^*, all the other firms disclose to j their sales, and j cancels its policy. But the individual would have known that, and so would not have purchased that policy. The one subtlety is the following: Consider a situation with three established firms, A, B, and C. The high risk individual buys ½ $\overline{\alpha}$ from each, discloses its purchases from C to A, from B to C and from A to B. Then A reveals its sales to the individual to B, but B already knew about it, and so on for the others. This is where our assumption that the individual firm reveals *all* of the information at his disposal, not just his direct sales, becomes relevant. A knows about C as well as about its own sales, and thus reveals to B information about C. But then B knows about j's purchases from A, B, and C, i.e. he knows that j has purchased 3/2 $\overline{\alpha}$, and the policy is cancelled. In the appendix, we show that this logic is perfectly general.[31]

---

[30] Of course, high risk individuals (or their insurance firms) do not reveal their purchases of the supplemental policies at the high risk price, because if they did so (truthfully), then all those selling pooling contracts would condition their sales on such supplemental policies not being bought (for such purchases reveal that the individual is high risk).

[31] We have investigated alternative specifications of our model, where a firm discloses just its own sale to its customer, not what the consumer reveals to it. One variant entails insurance being purchased *sequentially,* with sales at any point being conditional on previous purchases. In this setting, a consumer would reveal to his insurer k all of his previous purchases, because otherwise the insurer k will disclose its sale to the previous insurer(s) that were undisclosed to it, who will cancel its policy sold to the consumer. (The only reason that the consumer would not reveal previous purchases was because it had purchased more than $\overline{\alpha}$). That is, in this model, a firm does not need to disclose what its customer reveals to it to prevent its customer from over-purchasing insurance at $\overline{q}$. Also, another formulation that requires a firm to disclose just its sale (but both the quantity of insurance and the price at which it is sold) is a model where firms condition their contract offers upon price information (as well as quantity) revealed by consumers (see Appendix C).

We now prove

Theorem 2: The equilibrium strategy S_k^* *implements the equilibrium allocation* E^*. *An equilibrium always exists*.

The formal proof can be found in Appendix B. The key challenge in formulating the equilibrium strategy was suggested by section 2. With full disclosure (exclusive contracts) one can break any pooling equilibrium. The pooling contract A^* in Figure 2 is sold to both high and low risk individuals, and if it is to be part of the equilibrium it can't be broken. We already established that the only contracts which can break A^* are those in the area labelled z in Figure 2. But if the "established" firms sell to any individual buying such a contract (such as D in Figure 2) a supplemental contract bringing him out of the area z (following the arrow in Figure 2), then that contract will also be bought by the high risk individual. But then the putative contract breaking the pooling equilibrium would lose money.

Given the strategies of all the established firms, they have on offer pooling contracts up to $\overline{\alpha}$. High risk individuals will supplement their purchase of the deviant contract by the pooling contract, and in doing so will find the deviant contract attractive. But if the high risk individuals buy the deviant contract, it loses money.

To see this, observe that the deviant contract *D* either assumes exclusivity (or some restriction to ensure that the individual does not buy enough insurance to take him out of the area z) or does not. The deviant firm knows that given S_k^*, if he does not impose contract restrictions, individuals will buy up to $\overline{\alpha}$, moving him out of the area z. Hence, the deviant firm will impose restrictions. But the consumer knows that the deviant firm cannot enforce those restrictions if the deviant firm doesn't know about his purchases; and he knows that, given the information disclosure rule of (the established) firms, if he reveals his purchases of insurance from the deviant firm to those from whom he has purchased insurance, the firms will not reveal that information. This will be the case regardless of any information disclosure rule the deviant firm adopts. Accordingly, the high risk individual purchases the deviant contract and pooling contracts up to $\overline{\alpha}$ and reveals his purchase of the deviant contract to the sellers of the pooling contract, but not vice versa. He thus moves himself out of the area z, and his new package of policies yields a higher level of utility than the original allocation. Hence the deviant contract loses money and the argument is complete. [32]

[32] This will also be true even when a deviant firm is an entrant firm to whom the established firms never disclose their information. This is because then a high-risk consumer would like to choose the entrant contract all the more as

There is one subtlety that has to be addressed: what happens if the deviant firm offers a menu of policies, in particular one purchased by the high risk individuals, the other by low risk individuals. Is it possible that such a pair of policies—with cross subsidization—could break the equilibrium? In Appendix D, we show that, even when a deviant firm offers multiple contracts at different prices, there still exists an equilibrium.

By making a seemingly weak additional assumption, we can show that our equilibrium can generate full honesty in equilibrium: Assumption A (Truth telling): If individuals are indifferent between telling the truth and not telling the truth, they tell the truth. We have already established that no individual purchases more than $\overline{\alpha}$. Given that that is the case, no individual has an incentive to hide his purchases. It follows that under Assumption A, given the equilibrium strategy S_k^* adopted by the established firms, all individuals reveal the truth about purchases of insurance from other firms except to a deviant firm.

7. Generality of the Result

The existence of equilibrium does not require the single crossing property to be satisfied. First of all, it should be obvious that Theorem 1 on the unique equilibrium allocation can hold for more general preferences so long as the income consumption curve for high-risk individuals is the full-insurance line.

As for Theorem 2: Any cream-skimming strategy must entail a contract preferred by the low risk (diagrammatically, below V_L), and be such that, with whatever supplemental insurance that the high risk individual buys from the established firms, put the individual above the line A*C*--the line through A* with slope $\frac{P_H}{1-P_H}$. The former condition implies that the price of the deviant contract must be below $\overline{q}$. Given the strategies S_k^*, if the deviant contract D entails $\alpha \leq \overline{\alpha}$, the high risk individual tops it up to $\overline{\alpha}$, and it is clear that this allocation is preferred to A*, i.e. D does not cream skim, and loses money.[33] If the deviant contract entails more insurance than $\overline{\alpha}$, it is preferred by V_L, the contract by itself must be below A*C*, i.e. would be purchased by high risk individuals, as is evident in Figure 4 where we have not assumed quasi-concavity. .

8. Extensions: *Non-uniqueness of equilibrium*

---

he can purchase additional pooling insurance from established firms even without disclosing to them his purchase from the entrant firm.

[33] More formally, if the deviant contract entails insurance of α' at price q', then self-selection requires $q'\alpha' + \overline{q}(\overline{\alpha} - \alpha') \geq \overline{q}\overline{\alpha}$, which is never satisfied if α' > 0 and q' < $\overline{q}$.

The equilibrium is not unique: there are other strategies that can sustain the equilibrium allocation E^*. For instance, once we extend the strategy space of firms so that contract sales to an individual can be conditioned on the price as well as the amount of insurance purchased, and information disclosure rules specify the revelation of not just the amounts of insurance, but also the price, we can formulate a slightly different strategy supporting the same equilibrium allocation E^*, as is shown in the Appendix C.[34] In some ways the analysis of the equilibrium is simpler,[35] but it entails using *latent* policies, policies which are *only* sold in response to out of equilibrium purchases from other insurance firms but which are not purchased in equilibrium.

9. Extensions to Cases with Many Types

The result on existence of equilibrium can be extended to the case with many types. (See Stiglitz-Yun (2016).) An equilibrium strategy in a case with the three types, for example, can be described in a similar way to the case with two-types. As illustrated in Figure 5, there is a pooling contract with all three types, contract A, the most preferred by the lowest risk type; and a partial pooling contract B with additional insurance pooling together the two riskiest types, where B is the most preferred along the zero profit line for partial pooling; and finally, a contract C, providing full insurance to the highest risk type. In equilibrium consumers purchase A only (the lowest risk type) or A and B or A, B and C (the highest risk type), depending upon their types.

[FIGURE 5 ABOUT HERE]

There are three types of firms, those selling the full pooling contract, those selling the partial pooling contract, and those selling the price contract to the high risk individuals. They adopt the same information disclosure rule as in the case of two types of individuals.[36] Consumers truthfully fully reveal to the other insurers their information about their purchases of the fully pooling contract A (since all purchase the same amount, such information in equilibrium reveals no information about who they are). Consumers reveal information about their purchases of the partial pooling policies B only to firms not (revealed to be) selling the fully pooling policy.[37] By the same

---

[34] This equilibrium, as well as that discussed in Appendix D, also do not require that the single crossing property be satisfied.

[35] As presented in the Appendix C, this equilibrium may allow for a simpler disclosure rule (than that of S_k^*) of a firm, which is to disclose to others just its own sales, not information revealed by its customers.

[36] That is, revealing information only to firms not revealed to be sellers to individuals.

[37] In fact, in the three-type case, an individual has an incentive to disclose his purchase from a fully pooling seller, because otherwise his potential insurer (or a partially pooling seller) discloses to his fully pooling insurer, who then would cancel (in Stage 3) the contract it sold to him.

reasoning as in the two-type case, there is no room for a cream-skimming deviant contract offering D that profitably attracts only low or medium types, as riskier types are also induced to choose D.[38] This argument can also be applied to the case with a continuum of types as well.

10. Previous literature

In the more than four decades since RS appeared, its disquieting results have given rise to a large literature, which we can divide into a few major strands. The first looked for alternative equilibrium concepts, or game forms, under which equilibrium might always exist, or under which a pooling equilibrium might exist. Hellwig (1987) was the first to provide a game-theoretic framework in a dynamic setting to analyze these equilibria (including RS) and contrast one with another. Rothschild and Stiglitz (1997) reviewed the literature as it existed to that point, suggesting that there had not yet been an adequate alternative resolution as to what a *competitive* market equilibrium should look like in the presence of information asymmetries. For instance, in Wilson's reactive equilibrium (1977), the entry of even a *very* small firm induces all firms to "react," by withdrawing their pooling contracts, making the deviant contract unprofitable and enabling the pooling equilibrium to be sustained.[39] [40]

A second strand more related to the analysis here has explored the consequences of different information structures, in particular, the possibility of non-disclosed contracts.[41] Most notable are a series of papers by Attar-Mariotti-Salanie (2011, 2014, 2016). In the first (which is more akin to Akerlof's model of lemons and has a different scope of applications), they succeed in establishing a condition for existence—the presence of an aggregate capacity constraint, along with latent contracts. Their later (2014) model (which employs preferences that are a generalization of the form considered in this paper) emphasizes the importance of firms being able

[38] By the same token, there is no incentive for contract manipulation.
[39] More recently Netzer and Scheuer (2014) have revived the Wilson-Miyazaki reactive equilibrium. Firms may "opt out" of the market after observing the contract offers of other firms. They show that as long as the costs of opting out are nonzero, but not too large, there is a unique outcome - the Miyazaki-Wilson one.
[40] Mixed strategy equilibria (e.g. studied by Dasgupta and Maskin (1986) and Farinha Luz (TE, forthcoming), while interesting as an analytic solution, are unpersuasive as a description of what any market might look like. The notion that one might go to an insurance firm and choose among lotteries, which would assign probability distributions to benefits or premia, seems largely fanciful. Why that is so may necessitate an enquiry into behavioral economics, or into the economics of trust: how does one know that, say, the contract has been drawn from the purported probability distribution of contracts? One typically only sees one's own outcome.
[41] See also the earlier papers of Jaynes (1978) and Ales-Maziero (2012). Ales-Maziero (2012) focused on the case of adverse selection in a non-exclusive environment, characterizing the conditions for an equilibrium to exist and showing that an equilibrium, if it exists, is a separating one where only the highest-risk type purchase full insurance at the actuarially fair price.

to offer a menu of contracts, but they get existence only under very restrictive conditions—conditions which are never satisfied in our canonical model.[42] In their (2016) model, they allow firms to sell only a single contract, but, again, in general, existence fails. More broadly, we consider a situation that is closely related to those they study—all entail looking for equilibrium in a simple adverse selection model-- but ours is still markedly different from theirs; ours is the natural one relevant in insurance markets, while they employ special assumptions which make their analysis inapplicable to this market.[43]

Their work highlights the important consequences of different information structures. The central objective of this paper, by contrast, is endogenizing the information structure - allowing firms and individuals to decide what information to disclose to whom. The closest works to our paper within the adverse selection literature are Jaynes (1978, 2011) and Hellwig (1988), who analyze a model with a certain type of strategic communication *among firms* about customers' contract information. Jaynes (1978) characterizes an equilibrium outcome that involves a pooling allocation plus supplemental provision at the high-risk price, the allocation which our analysis (as well as that of Attar-Mariotti-Salanie (2016)) showed to be the only possible allocation. However, as Hellwig (1988) clarified, in Jaynes' (1978) 2-stage framework, the strategy of firms including the associated strategic communication is not a Nash equilibrium but a reactive equilibrium, with firms responding to the presence of particular deviant contracts, and thus Jaynes' formulation was subject to the same objections raised earlier. While our work differs from that of Jaynes and Hellwig in several ways[44], perhaps most important is that we consider information revelation strategies by consumers as well as firms. This turns out to be critical in the analysis of the existence of a Nash equilibrium, for it importantly allows the creation of asymmetries of information about insurance purchases between "established" firms and deviant firms. Without that, the pooling contract would not be able to be sustained. As we have noted, there is a delicate balance: on the

[42] Their condition rules out situations in which "both the low-risk and the high-risk agents purchase a basic policy at a medium price, with the high-risk agent purchasing on top of this a supplementary policy at a higher price." By contrast, we characterize precisely such an outcome.

[43] In particular, in our model, preferences in {α,β} are strictly concave.

[44] Importantly, Hellwig's analysis is based on a four-stage game, in which firms decide to whom they send customer information (in stage 3) only after observing contract offers announced by firms (stage 1) and purchased by consumers (stage 2). In other words, their communication strategies are allowed to be conditional upon contract offers made by other firms. While Hellwig (1988) shows that the Jaynes allocation (the equilibrium allocation in our paper) can be sustained as a sequential equilibrium in the four-stage game, Jaynes (2011) attempted to characterize the "Jaynes allocation" as a perfect Bayes equilibrium in a two-stage game, in which firms announce their contract offers and communication strategies simultaneously. His formulation is thus markedly different from that presented here.

one hand, one has to prevent *overinsurance* by high risk individuals purchasing pooling contracts (which requires established firms to know certain information), and on the other hand, one has to prevent a deviant firm from having enough information to enforce an exclusive contract that would break the pooling equilibrium. The consumer and firm information strategies which we describe do this, and it should be apparent that, at least in a simple game form, models relying on just firm information strategies cannot do this, because they do not have the information basis on which to engage in this kind of disclosure discrimination.[45]

11. The no-disclosure limited information price equilibria

A final strand of literature to which this paper is related is that which assumes no disclosure of insurance purchases, implying that the only information which a firm has about the purchases of an individual are the sales the firm of the itself, assuming that there is not anonymity in sales. This literature, however, does not *endogenize* the decision not to disclose, but takes that policy as given. The standard assumption in the adverse selection literature (see e.g. Arrow, 1965) is that insurance firms and individuals simply take the price of insurance as given, and consumers buy as much at that price as they want. Competitive equilibrium requires that there be no profits (on average).

More formally, we denote the purchase by a high risk individual at a price q(P) corresponding to an accident probability P as $\alpha_H(P)$, and similarly for the low risk as $\alpha_L(P)$ where $\mathrm{q(P)} = \frac{P}{1-P}$.

The weighted average accident probability when the price is q is then

$$\hat{P}(\mathrm{q(P)}) \equiv P_H\theta\frac{\alpha_H(q(P))}{\bar{\alpha}^{\mathrm{e}}(q(P))} + P_L(1-\theta)\frac{\alpha_L(q(P))}{\bar{\alpha}^{\mathrm{e}}(q(P))}, \qquad (6)$$

where $\bar{\alpha}^e(q(P)) = \theta\alpha_H(q(P)) + (1-\theta)\alpha_L(q(P))$, and

$$\alpha_L(q(P)) = Argmax\ \ V_L(\alpha,\beta) \quad \text{s.t. } \beta = \frac{P}{1-P}\alpha,$$

and

$$\alpha_H(q(P)) = Argmax\ \ V_H(\alpha,\beta) \quad \text{s.t. } \beta = \frac{P}{1-P}\,\alpha$$

Since at any price, the high risk buy more insurance ($\alpha_H(q) > \alpha_L(q)$), the weighted accident

[45] That is, at least in the initial round of disclosures, firm disclosure can only be based on individual purchases. Assume some firms sold policies which did not disclose their sales. High risk individuals would purchase such insurance beyond $\bar{\alpha}$, and the restriction that they not do so would not be enforceable. Thus, the putative allocation could not be sustained, since the non-disclosure pooling contracts would make a loss. On the other hand, if firms sold only disclosure policies, then a deviant firm offering an exclusive contract in the region z would be able to enforce exclusivity, and this would break the pooling allocation. Hence, again, the putative equilibrium could not be sustained. There has to be some basis on which firms can differentiate among whom to disclose; our consumer revelation mechanism provides this.

probability $\hat{P}(\mathrm{q(P)})$ is higher than the population weighted average $\bar{P}$: $\hat{P}(\mathrm{q(P)}) > \bar{P}$.

Now we define a *(competitive) price equilibrium as* P^e satisfying the following conditions: (a) (uninformed) sellers have rational expectations P^e about the weighted average accident probability of the buyers; (b) with those rational expectations, prices are set to generate zero profits; and (c) at those prices consumers buy the quantities that they wish.[46] Thus, a price equilibrium P^e satisfies

$$P^e = \hat{P}(q(P^e)) \qquad \text{with } \widehat{P}'q' > 0 \qquad (7)$$

Low risk individuals diminish their purchases of insurance as prices increase. This is the well-known *adverse selection* effect. But the value of $\widehat{P}'q'$ depends on the elasticities of demand of the two groups as well as their relative proportions, and so in general *there may be more than one price equilibrium*. A sufficient condition for a unique equilibrium, in which only high risk individuals purchase insurance, is $\alpha_L(q(P_H)) = 0$.[47][48]

Nash equilibrium and non-existence of a partial information-no disclosure price equilibrium. In the no-disclosure price equilibrium, the insurance firms simply take the price as given. However, while a firm doesn't know the size of the policies taken up by an individual from other firms, he knows what he has sold.[49] An insurance firm can offer a large policy - he knows to whom he sells, and can refuse to sell a second policy to the same individual.[50] We define a *partial information-no disclosure (Nash) price equilibrium* as an equilibrium where the insurance firm knows *at least* information about the amount of insurance it sells: a partial information-no disclosure price equilibrium is a set of contracts such that (a) each quantity-contract at least breaks even; (b) there exists a price at which each individuals can buy as much insurance at the price offered at he wishes and at which insurance premiums at least cover pay-outs; and (c) there does not exist any policy which (*given* the information structure) can be offered which will be purchased and make a profit.

[FIGURE 6 ABOUT HERE]

[46] The latter conditions are equivalent to the standard conditions of demand equaling supply for this particular model.

[47] $\alpha_L(q(P_H)) = 0$ implies $\frac{P_L}{1-P_L}\frac{U\prime(W-d)}{U\prime(W)} \leq \frac{P_H}{1-P_H}$.

[48] We could define a price equilibrium in a Nash-Bertrand fashion by adding another condition that each firm, taking the prices of others as given, chooses the price which maximizes its profits. In this case, it can be shown that there exists a unique price equilibrium, the lowest price at which equation (7) is satisfied.

[49] This would not be the case if individuals purchased insurance about an event affecting a third party, and firms sold such insurance without knowledge of the purchaser.

[50] In the context of moral hazard, the implication of this simple observation were explored in Arnott-Stiglitz (1991a, 1987).

Any policy proposing to break a price equilibrium must satisfy two conditions: to be purchased, it has to have a lower price than the market price, but to make a profit, it must have a higher price than that corresponding to the actual pool of people buying the policy. Consider a deviant firm that secretly offered a quantity policy, say the policy which maximizes the utility of the low risk individuals at a price corresponding to P′, with $P^e > \mathrm{P}' > \bar{P}$ such as (α', β') in Figure 6. It sells only one unit of the policy to each individual, and restricts the purchases of all to the fixed quantity policy. Then low-risk individuals will buy the policy, and it will make an (expected) profit. It thus breaks the price-equilibrium. The one case where this argument doesn't work is that where at the pooling price, low risk individuals do not buy any insurance. We have thus established

Theorem 3. There is no partial information-no disclosure price equilibrium where both types of individuals buy insurance.

Put differently, there is no "price equilibrium" when firms can offer an undisclosed quantity contract and ration the sale, say to one policy to a customer[51]. What is remarkable about Theorem 4 is how little information is required to break the price equilibrium: the firm just uses its own contract information to implement the quantity constraint.

It is natural to ask, if there is not a price equilibrium, is there some analogous equilibrium, with say *just* fixed quantity contracts? Consider a case where the two groups are quite similar. Each insurance firm sells insurance in fixed units, say $(\bar{\alpha}, \bar{\beta})$, say the policy which is most preferred by the low risk individual along the break-even pooling line. The high risk individual would not want to buy two units of that insurance. But he would supplement his purchase with the undisclosed insurance at his own price, in an amount that brings him to full insurance. The analysis of this paper has shown that this kind of pooling contract cannot be an equilibrium: there is always a deviant policy that could be offered that would be taken up only by the low risk individuals, given the posited information structure. *In other words, given this partial information structure, there is no equilibrium, ever, where both groups buy insurance*. By contrast, with the more complex

[51] We can also show that *there is a Nash partial information price equilibrium where only the high risk individuals buy insurance if and only if* $\alpha_L(q(\bar{P})) = 0$. This condition is stricter than that in which there exists a price equilibrium with a single type: $\alpha_L(q(P_H)) = 0$. Thus, even a corner price equilibrium may not be a Nash partial information price equilibrium. In a somewhat different set-up, Jaynes (1978) presents a set of similar results. The condition posited here for the existence of a partial disclosure price equilibrium, $\alpha_L((\bar{P})) = 0$, is stricter than that specified by Jaynes (1978), which would be equivalent to $\alpha_L((P_H)) = 0$. Jaynes (1978) shows that a price equilibrium q^* at which each agent purchases his Walrasian demand, which is a no-information equilibrium in our model, cannot be sustained in the presence of a fixed-quantity contract when more than one type of agent purchases insurance at q^*.

endogenous information structure described in the paper, there is *always* an equilibrium.

12. Concluding Remarks

In insurance markets with asymmetric information, firms will use what information is available to make inferences about purchasers of insurance, including information about the amount of insurance purchased. High risk individuals know this, and have an incentive to do what they can to ensure that insurance firms can't tell that they are high risk, and to try to keep any relevant information (such as the amount of insurance purchased) secret, and there may be market incentives for firms to comply.

The earlier work of Akerlof and Rothschild-Stiglitz had, of course, shown the importance of the information structure: information about insurance purchased conveyed important information about the individual's type, and therefore, whether that information was available was central in determining the nature of the equilibrium. The differences between Akerlof and RS reflected differences in assumptions about the information structure, e.g. RS assumed sufficient information to enforce exclusivity. Allowing undisclosed contracts and incorporating realistic assumptions about things that insurance firms know, in particular, that they know the identities of their customers and the quantities purchased, destroys both the Rothschild-Stiglitz and the Akerlof equilibria.

Expanding the equilibrium construct to include endogenous information disclosure rules is complex, but in fact helps resolve some longstanding conundrums in information economics, in particular the general non-existence of pooling equilibria and the possible non-existence even of a screening equilibrium.

When we endogenize information revelation, the unique equilibrium allocation is a *partially* disclosed pooling contract - the pooling contract most preferred by the low risk individual[52] - *plus* undisclosed supplemental insurance for the high risk individuals and no supplemental insurance for the low risk individuals. The equilibrium endogenously creates asymmetries in information about insurance purchases; we show that at least within our framework, such asymmetries are essential to supporting the equilibrium.

In some ways, the equilibrium that arises with endogenous information looks much more like observed equilibria: Equilibrium always exists, and always entails some pooling. Moreover, the

[52] That is, the pooling allocation *at the population weighted accident probabilities* most preferred by low-risk individuals. (This pooling contract is that upon which Wilson (1977) focused.)

analysis and its results do not rely on the highly restrictive single crossing property which has been central in the literature spawned by RS.

The insurance model has proven a useful tool for analyzing more generally markets with asymmetric information, and the papers analyzing imperfect and asymmetric information in that context have spawned a huge literature, with the concepts being applied to a rich variety of institutional structures[53]. The natural information assumptions concerning potentially hidden actions and hidden characteristics differ across markets. This paper has raised questions about both the Akerlof and RS analyses, and by implication, the results in the large literature based on them.

We hope that this paper will, like the earlier RS and Akerlof analyses, spawn further research in the context of other markets in the analysis of market equilibrium with asymmetric information where contracts and the information structure/revelation are endogenously and simultaneously determined.

APPENDICES

Appendix A: Proof of Lemma 4.

Given lemma 3, the consumer purchasing more than $\overline{\alpha}$ must not reveal his full purchases to any firm. Assume there are N purchases and that the firm j to which he is most dishonest has been given information about N-1 purchases, and in particular, he does not reveal purchases from k. Then j reveals to k information about all purchases but that of k; but then k knows about all purchases, and that the individual's total purchases exceed $\overline{\alpha}$. Assume now that the firm j to which he is most dishonest has been given information about N-2 purchases, i.e. the consumer does not reveal purchases from k and k'. Either k knows about k' or not. If k knows about k', then when j reveals all of its information to k, then k knows about all purchases. If k does not know about k', then when k and j reveal all of their information to k', k' knows about all purchases. The argument can be extended to any level of non-disclosure.

Appendix B: Proof of Theorem 2

It is obvious that by Lemma 4, the strategy S_k^* generates the equilibrium allocation E^*. We will now show the strategy S_k^* sustains E^* against any deviant contract. Note first that a deviant firm

[53] It is important to recognize that, for the most part, these models of insurance were not intended to provide a good institutional analysis of the insurance market; rather, the insurance market provided the paradigm for studying behavior in, for example, labor, product, and capital markets because it seemed so simple to strip away institutional details, and study markets unencumbered by them. It was for this reason that these paradigmatic models proved so fruitful. The analysis of this paper should be taken in the same spirit.

cannot make profits by attracting only high-risk individuals in the presence of non-established firms offering any amount of insurances at q_H. This is because then no individual would pay a price higher than q_H since a deviant firm, even with CM, cannot induce the established firms (with S_k^*) to offer more than $\bar{\alpha}$ at $\bar{q}$ under any circumstance. If the deviant attracts both high and low risk, his contract would have to lie on or below the pooling line, and the best that he could be expected to do is zero profits. A deviant firm can thus make positive profits only by attracting only low-risk types.

Lemma 5. A necessary condition for a deviant contract to attract only low risk individuals is that the contract be in the non-empty region z in Figure 2, the set of (α, β)'s such that

$$V_L(\alpha, \beta) > V_L^* \quad \text{and} \quad \beta - \bar{q}\bar{\alpha} \geq q_H(\alpha - \bar{\alpha}),$$

where V_L^* is the expected utility of the low risk individuals in the putative equilibrium.

Clearly, when the first inequality is not satisfied, the low risk individuals will not purchase the policy, and when the second condition is not satisfied, the high risk individual will purchase the policy, supplementing it with the secret insurance.

Consider any policy D$(= (\alpha_D, \beta_D))$ in z (satisfying the above two conditions). Given the equilibrium strategies of the established firms, then high risk individuals will buy D, supplementing it with pooling insurance from the established firms, bringing the entire purchases (of revealed insurance) at least to $\bar{\alpha}$.

Given the conditions imposed on preferences (quasi-concavity, slope of indifference curve equaling q_H with full insurance)[54], high risk individuals will wish to buy as much insurance at the pooling odds as they can. With full disclosure, they can buy $\bar{\alpha}$. Since individuals have a choice of disclosure, they can at least get $\bar{\alpha}$ with full disclosure to established firms but with no disclosure to the firm offering D. Denote by D$'$ total insurance (D plus the pooling contract plus the supplemental secret insurance). It is obvious that $V_H(D') > V_H(A^*)$. With the given consumer and firm disclosure strategies, no firm will disclose to the deviant firm their sales to the insurance, so that the deviant firm cannot enforce the restrictions necessary to prevent consumers from buying supplemental pooling insurance. It follows that there exists no policy breaking the pooling contract A^*.[55]

[54] As discussed in section 7, our results hold even with preferences that are not quasi-concave.

[55] Two minor subtleties: While we showed that in equilibrium, there is no contract manipulation, we have to show that no deviant firm will engage in CM. But it is obvious (by our earlier analysis) that CM is attractive only to high

Appendix C: An alternative equilibrium

In this appendix, we show that the equilibrium allocation can be supported by alternative contracts, entailing different restrictions and disclosures. We now assume that restrictions are based not just upon the amount of insurance purchased but also upon the price (equivalently, on both α and β), and when consumers and firms disclose information, they disclose not just the amount of insurance, but the price at which they purchased insurance. Assume the established firms[56] have a strategy S_k^o which entails the same disclosure rule about to whom to disclose as that of S_k^* [57] while offering

$\bar{\alpha}$ at a price $\bar{q}$ if the individual has no other insurance

0 if the individual has purchased other insurance at a price higher than or equal to $\bar{q}$

$\alpha_k \leq \hat{\alpha}(\mathrm{D}))$ at the price $\bar{q}$ if the individual has purchased elsewhere a contract D that offers insurance α_D at a price $\mathrm{q} < \bar{q}$, where $\hat{\alpha}(\mathrm{D})$ is the maximum amount of insurance that a low risk individual would want to purchase to supplement the contract D at the pooling odds. Because the low risk individual is better off than at A^*, $\hat{\alpha}(\mathrm{D}) + \alpha_D > \bar{\alpha}$, while $\hat{\alpha}(\mathrm{D}) \leq \bar{\alpha}$ with the inequality holding for $\alpha_D > 0$.[58]) In words, the established firms with S_k^o sell the full contract A^* (and only that contract) to an individual with no other insurance (so far as it knows); sells nothing to anyone who has purchased any other insurance at less (or at equally) attractive terms than the pooling equilibrium (it can infer that such a person is a high risk individual); and sells a variable amount of insurance, bringing total insurance purchased up to, at a maximum an amount $\hat{\alpha}(\mathrm{D})$ at the pooling price if the individual has purchased a contract D at a lower price than $\bar{q}$.

The equilibrium looks precisely as before, except now everyone purchases the policy A^* from a single insurance firm. Out of equilibrium behavior entails the use of latent contracts, the policies

risk individuals. This implies that no deviant strategy with CM can attract low-risk types only by charging q $\in \left(\frac{P_L}{1-P_L}, \frac{\bar{P}}{1-\bar{P}}\right)$, since no deviant strategy without CM can. Secondly, our earlier analysis established that the deviant firm would not be able to enforce the exclusivity it needed *regardless of the information disclosure strategy of the deviant firm*.

[56] We also assume, as before, that the other firms (j=M+1, --, N) offer any amount of insurance at a price q_H without disclosure. There is a single deviant firm.

[57] The disclosure rule of S_k^o can be simpler than that of S_k^* ; firms need disclose to others their own sales only, not information revealed by their customers, because every consumer purchases $\bar{\alpha}$ at a price $\bar{q}$ in equilibrium.

[58] $\hat{\alpha}(\mathrm{D}) = Argmax_\alpha\{P_L U(W - d + \alpha + \alpha_D) + (1 - P_L)U(W - \bar{q}\alpha - q\alpha_D)\}$ (where $\frac{P_L}{1-P_L} < \mathrm{q} < \bar{q}$).

Differentiating the f.o.c. for $\hat{\alpha}(\mathrm{D})$ with respect to α_D and α, we have the result that, so long as U" < 0, $\hat{\alpha}(\mathrm{D}) + \alpha_D > \bar{\alpha}$ and $\hat{\alpha}(\mathrm{D}) < \bar{\alpha}$ for $\alpha_D > 0$. Note that this result does not require preferences to satisfy the single-crossing property; the result can hold even when different types of individuals have different utility functions, in which case the single crossing property may not be satisfied.

the sale of which are only triggered when individuals have purchased a deviant contract, D. It should be clear that no low risk individual will buy any policy sold at a price above $\bar{q}$. Accordingly, any policy sold at a price between q_H and $\bar{q}$ loses money. Also, since the amount of additional pooling insurance offered on top of any insurance revealed to be purchased elsewhere is not greater than $\bar{\alpha}$, no high-risk individuals would be willing to pay an average price higher than q_H (getting some part of the package at a price below $\bar{q}$.) to trigger the sale of α_k [59] Thus we can focus on deviant policies sold at a price below $\bar{q}$.

High risk individuals will supplement D, topping up total purchases to $\hat{\alpha}(\mathrm{D})$ of insurance. But that means that expected utility of the high risk individual, supplementing D with the pooling contract (up to $\hat{\alpha}(\mathrm{D})$), and supplementing that with secret insurance (at its own odds) is higher than at the original allocation, i.e. the high risk individual as well as the low risk individual buys D, and that means that D loses money, since D is sold at a price below $\bar{q}$ ($i.e.$ is below the pooling line.) It is thus clear that this simple strategy can support the equilibrium.[60] [61]

Appendix D: Deviants offering Multiple Contracts at Different Prices

In this appendix, we show that our results hold even when firms are allowed to sell multiple contracts at different prices. The central issue is whether this allows a deviant firm to break our putative equilibrium by taking advantage of cross-subsidization. A deviant firm does so to induce self-selection among the applicants - with the self-selection process reducing the costs of the high risk individuals buying insurance from the deviant. We first explain why the set of strategies considered earlier now doesn't “work”. We then describe intuitively the challenges involved in finding an equilibrium strategy. Next we provide the formal analysis, establishing the main theorem of this appendix.

Let {A*, C*} represent the equilibrium allocation described earlier. Now consider the deviant pair of policies {A*B, G} (as depicted in Figure 7), where A*B entails an offer of α_S at $\bar{q}$ without disclosure and G offers α_D at a price q lower than $\bar{q}$ with disclosure and with G being offered

[59] Thus, if the individual chose not to reveal any purchase from the deviant firm, he could have purchased at $\bar{q}$ an amount $\bar{\alpha}$. Earlier, we referred to the kind of deception where an individual purchases two policies (perhaps bundled, as here) and discloses only one as contract manipulation.

[60] As before, it is important that the deviant firm not be able to enforce exclusivity, and the information strategy ensures that this is the case.

[61] In the main text, we showed that the pooling contract cannot be broken except possibly by a contract in the area z, and a straightforward adaptation of the arguments there apply here. The analysis here implies that even contracts in z cannot break the putative equilibrium.

conditional on no additional insurance being purchased. There always exists a continuum of pairs of policies (A*B, G) such that G is chosen by all the low-risk individuals while A*B is chosen by all the high-risk who simultaneously buy A*, that is, the high risk individuals supplement A*B with the pooling insurance A*, i.e. they buy $\bar{\alpha}$ of insurance from the established firms and α_S from the deviant firm. Because the price q is greater than $\frac{p_L}{1-p_L}$, the deviant firm makes a profit on G even though it makes a loss on the contract purchased by the high risk individuals. By carefully choosing {A*B, G} or $\{\alpha_S, (\alpha_D, q)\}$, the deviant firm can make overall positive profits. For instance, this will be so if α_S is small.[62] While there are large total losses associated with the purchase of insurance by high risk individuals, most of those losses are borne by the established firms, who now sell their pooling contract only to the high risk individual. With an appropriately chosen {G}, the deviant firm gets all the low risk individuals for all of their insurance, and the high risk people only for the supplemental amount α_S.

To prevent this type of a deviation, we need to make contract G more attractive to high-risk types by providing more additional insurance at a price $\bar{q}$ than S_k^* does, should a deviant firm try such a strategy, while limiting the total provision by all the firms to $\bar{\alpha}$ in equilibrium. To do this, we need to have a *latent* contract which offers an individual sufficient amount of extra insurance at $\bar{q}$ in the presence of a deviant contract G, so that there can be no profitable self-selection.

More formally, consider a strategy S_k^o which has the same rule about to whom to disclose as S_k^*, but offers the same set of contracts with the same restrictions as S_k^* only when (to its knowledge) the price of insurance purchased elsewhere is not lower than $\bar{q}$ while offering (in the aggregate, among all the established firms) a large policy, say [63] $\hat{\alpha} \leq \bar{\alpha}$, in addition to the policy purchased at $q < \bar{q}$, at a price $\bar{q}$ to those who purchased insurance elsewhere at a price lower than $\bar{q}$. Thus, S_k^o contains a latent contract that is sold only out-of-equilibrium. We can then see that S_k^o supports the allocation E^* in equilibrium as it shares with S_k^* the same set of contracts in equilibrium. But, with the appropriate choice of $\hat{\alpha}, S_k^o$ ensures that the two-contract deviant firm loses money. $\hat{\alpha}$ should be not be greater than $\bar{\alpha}$, because otherwise high-risk

[62] More specifically, a deviant firm can set q ≈ $\bar{q}$, $\alpha_D \approx \bar{\alpha}$ and thus a small α_S, yielding the desired self-selection on the part of high-risk individuals between A*B and G, where G is a contract just below V_L near A*.
[63] It should be clear from the analysis that all that is required is that the latent policies offer a sufficient amount of insurance as to make G attractive to the high risk individual. As in Appendix C, the firms not offering the pooling contract (j=M+1, --, N) offer any amount of insurance at q_H without disclosure.

individuals would be willing to pay an average price higher than q_H, so that through contract manipulation they could purchase $\hat{\alpha}$.[64] We set $\hat{\alpha} = \bar{\alpha}$. Now we will focus upon a cream-skimming strategy G, which offers α_D at price q below $\bar{q}$. A high-risk individual i choosing G would not reveal to the deviant firm d his purchases of pooling insurance from other firms, but has an incentive to reveal to the established firms his purchase of low price insurance, for that triggers the offer of supplemental insurance. But given the strategy S_k^o, that means that the established firms don't disclose their sales to the deviant, which ensures that the exclusivity provision associated with G cannot be enforced. Knowing this, to induce self-selection, a deviant firm offers a "large" contract A*B' – entailing insurance of α_S' without disclosure. Given a choice between G and A*B', all high risk individuals choose A*B' and all low risk individuals choose G. We can then show that any pair of contracts (G, A*B') that induces self-selection makes losses.

To see this, note that if a high risk individual purchases A*B' without disclosure, his total insurance purchased at $\bar{q}$ is $\alpha_S' + \bar{\alpha}$. The high risk individual then supplements this with secret insurance at price q_H bringing him to full insurance. By contrast, with policy G (disclosed) the individual gets $(\alpha_D + \bar{\alpha})$ at a total premium of $(q\alpha_D + \bar{q}\bar{\alpha})$. The high risk individual then supplements this with insurance at price q_H bringing the individual to full insurance. It is easy to show that self- selection requires

$$q_H\{(\alpha_S' + \bar{\alpha}) - (\alpha_D + \bar{\alpha})\} \geq \bar{q}(\alpha_S' + \bar{\alpha}) - (q\alpha_D + \bar{q}\bar{\alpha}) \qquad (8)$$

Condition (8) can be rewritten as

$$\alpha_D \leq (q_H - q)^{-1}(q_H - \bar{q})\alpha_S' \qquad (8')$$

The corresponding profit $\pi\{G, A^*B'\}$ for the deviant firm is

$$\pi\{G, A^*B'\} = -\theta\alpha_S'(q_H - \bar{q}) + (1-\theta)\alpha_D(q - q_L)$$

$$\leq \alpha_S'(q_H - \bar{q})(q_H - q)^{-1}[-\theta(q_H - q) + (1-\theta)(q - q_L)]$$

$$= \alpha_S'(q_H - \bar{q})(q_H - q)^{-1}[q - \bar{q}] < 0,$$

i.e., the total profit for the deviant firm is negative. Alternatively, if the deviant firm fails to

[64] That is, an insurance firm could profitably offer a bundle of two policies, one with a price below $\bar{q}$ but with an average price above q_H, and high risk individuals would purchase such a bundle.

“separate,” so the high risk individuals chooses G, the deviant firm loses money.[65] We have thus established

Theorem 4. If deviant firms are allowed to offer multiple insurance contracts, there always exists an equilibrium strategy that sustains the unique equilibrium allocation E^* .

The Nash equilibrium entails the use of latent contracts, while it does not require preferences to satisfy the single-crossing property.

References

Akerlof, George A. (1970), “The market for ‘lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84(3), pp. 488-500.

Ales, Laurence and Pricila Maziero (2012), “Adverse Selection and Non-exclusive Contracts,” Working Paper, Carnegie Mellon University.

Arnott, Richard J. and Joseph E. Stiglitz (1987, 2013), “Equilibrium in competitive insurance markets with moral hazard.” in *The Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*, Oxford: Oxford University Press, 2013, pp. 660- 689. Edited version of Princeton University Discussion Paper 4, 1987

Arnott, Richard J. and Joseph E. Stiglitz (1991a, 2013), “Price Equilibrium, Efficiency, and Decentralizability in Insurance Markets,” in *The Selected Works of Joseph E. Stiglitz, Volume II: Information and Economic Analysis: Applications to Capital, Labor, and Product Markets*, Oxford: Oxford University Press, 2013, pp. 632-659. Edited version of NBER Working Paper 3642, 1991.

Arnott, Richard J. and Joseph E. Stiglitz (1991b), “Moral Hazard and Nonmarket Institutions: Dysfunctional Crowding Out or Peer Monitoring?” *American Economic Review*, 81(1), pp. 179-90.

Arrow, Kenneth J. (1965), *Aspects of the Theory of Risk-Bearing*, Yrjo Jahnsson Lectures, Helsinki, Finland: Yrjo Jahnssonin Saatio.

Attar, Andrea, Thomas Mariotti, and François Salanié (2011), “Nonexclusive competition in the market for lemons,” *Econometrica*, 79(6), 1869-1918.

[65] Our earlier analysis showed that if high and low risk individuals buy the same policy, a deviant cannot break the equilibrium. We can also verify that a deviant firm cannot profitably attract high-risk individuals through a CM, for at least one of the two contracts has to be at a price below $\bar{q}$, triggering the latent contract.

Attar, Andrea, Thomas Mariotti, and François Salanié (2014), “Nonexclusive Competition under Adverse Selection,” *Theoretical Economics*, 9(1), pp. 1-40.

Attar, Andrea, Thomas Mariotti, and François Salanié (2016), “Multiple Contracting in Insurance Markets,” Working Paper, TSE-532.

Dasgupta, Partha and Eric Maskin (1986), “The Existence of Equilibrium in Discontinuous Games: I,” *Review of Economic Studies*, 53(1), pp. 1-26.

Greenwald, Bruce and J. E. Stiglitz (1986), “Externalities in Economies with Imperfect Information and Incomplete Markets,” *Quarterly Journal of Economics*, 101(2), pp. 229-264.

Hellwig, Martin F. (1987), “Some Recent Developments in the Theory of Competition in Markets with Adverse Selection,” *European Economic Review,* 31(1-2), pp. 309-325.

Hellwig, Martin F. (1988), “A note on the specification of interfirm communication in insurance markets with adverse selection,” *Journal of Economic Theory,* 46(1), pp. 154-163.

Jaynes, Gerald D. (1978), “Equilibria in monopolistically competitive insurance markets,” *Journal of Economic Theory*, 19(2), pp. 394-422.

Jaynes, Gerald D. (2011), “Equilibrium and Strategic Communication in the Adverse Selection Insurance Model,” Working Paper, Yale University.

Riley, John (1979), “Informational Equilibrium,” *Econometrica,* 47(2), pp. 331-359.

Rothschild, Michael and Joseph E. Stiglitz (1976), “Equilibrium in competitive insurance markets: An essay on the economics of imperfect information,” *Quarterly Journal of Economics*, 90, pp. 629-649.

Rothschild, Michael and Joseph E. Stiglitz (1997), “Competition and Insurance Twenty Years Later,” *Geneva Papers on Risk and Insurance Theory*, 22(2), pp. 73-79.

Stiglitz, Joseph E. (1977), “Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market,” *Review of Economic Studies*, 44(3), pp. 407-430.

Stiglitz, Joseph E (2009), “Introduction to Part IIB,” in *Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis*, Oxford: Oxford University Press, 2009, pp. 129-140.

Stiglitz, Joseph and J. Yun (2016), “Equilibrium in a Competitive Insurance Market with Non-exclusivity Under Adverse Selection,” mimeo.

Wilson, C. (1977), “A model of insurance markets with incomplete information," *Journal of Economic Theory*, 16(2), pp. 167-207.

Figures

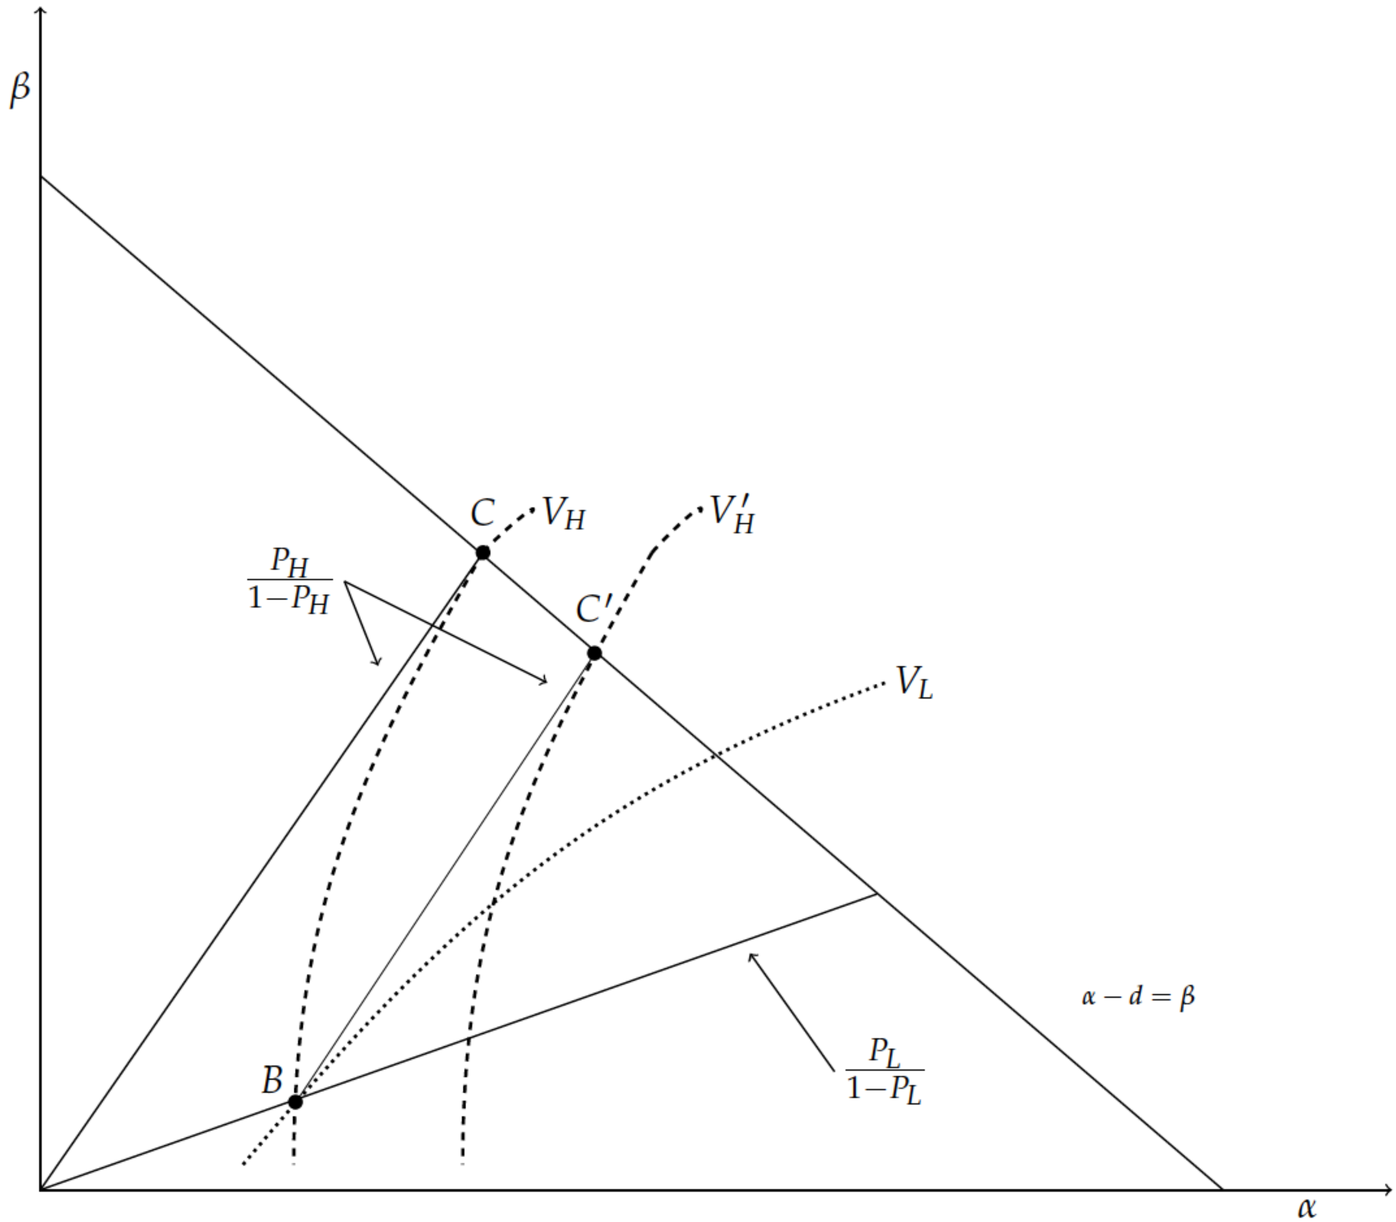

Figure 1: Breaking the RS separating equilibrium in the presence of undisclosed contracts at high-risk odds. Note: in this and subsequent figures ratios of the form $\frac{P_i}{1-P_i}$ denote the slope of the respective lines.

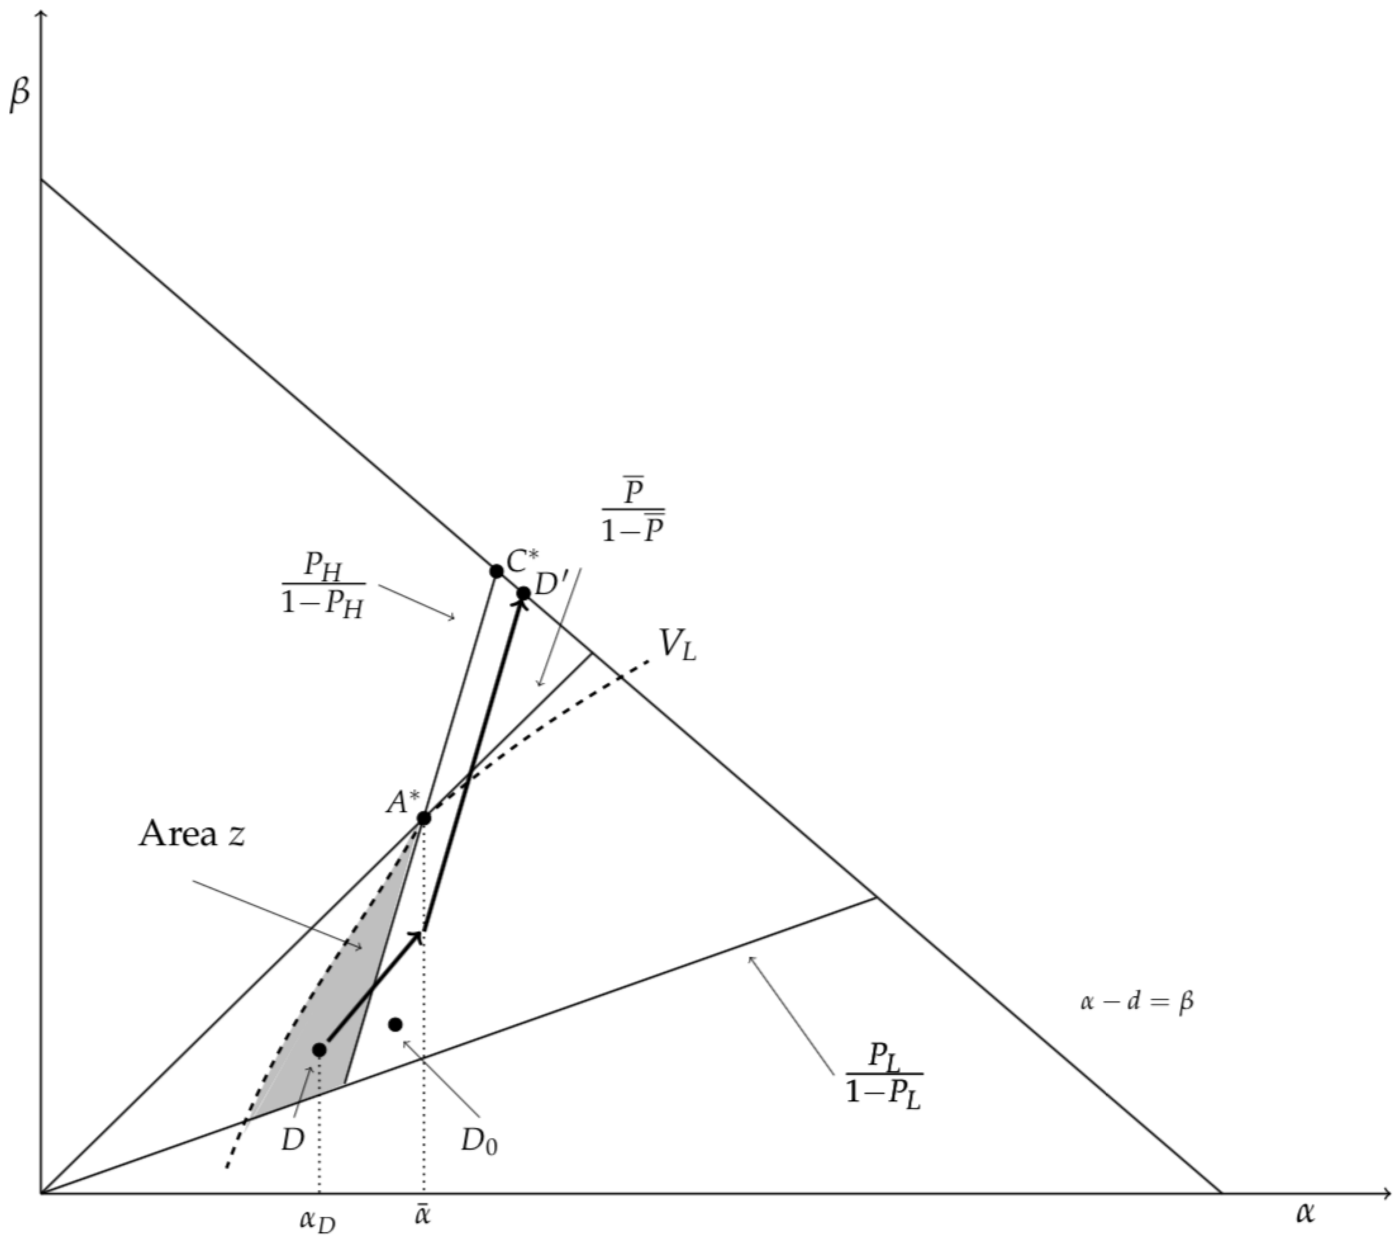

Figure 2: Sustaining an Equilibrium in the presence of a cream-skimming deviant contract D in z.

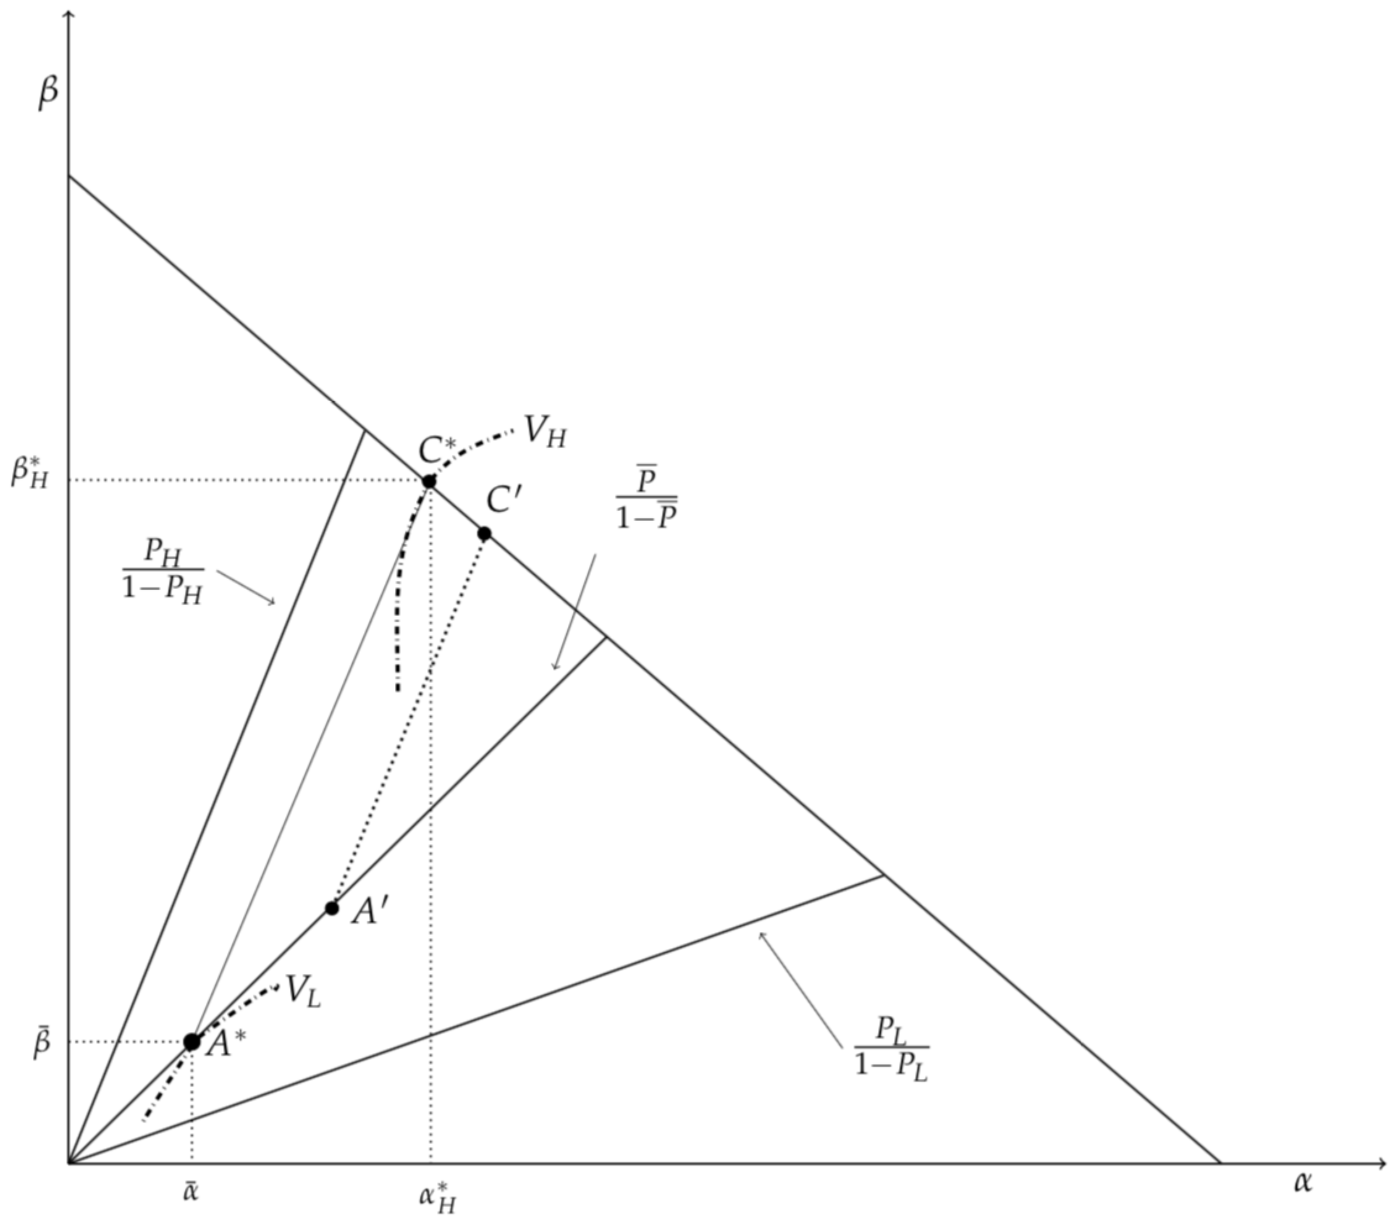

Figure 3: Pareto-efficient allocations $((A^*, C^*), (A', C'))$ and the equilibrium allocation (A^*, C^*).

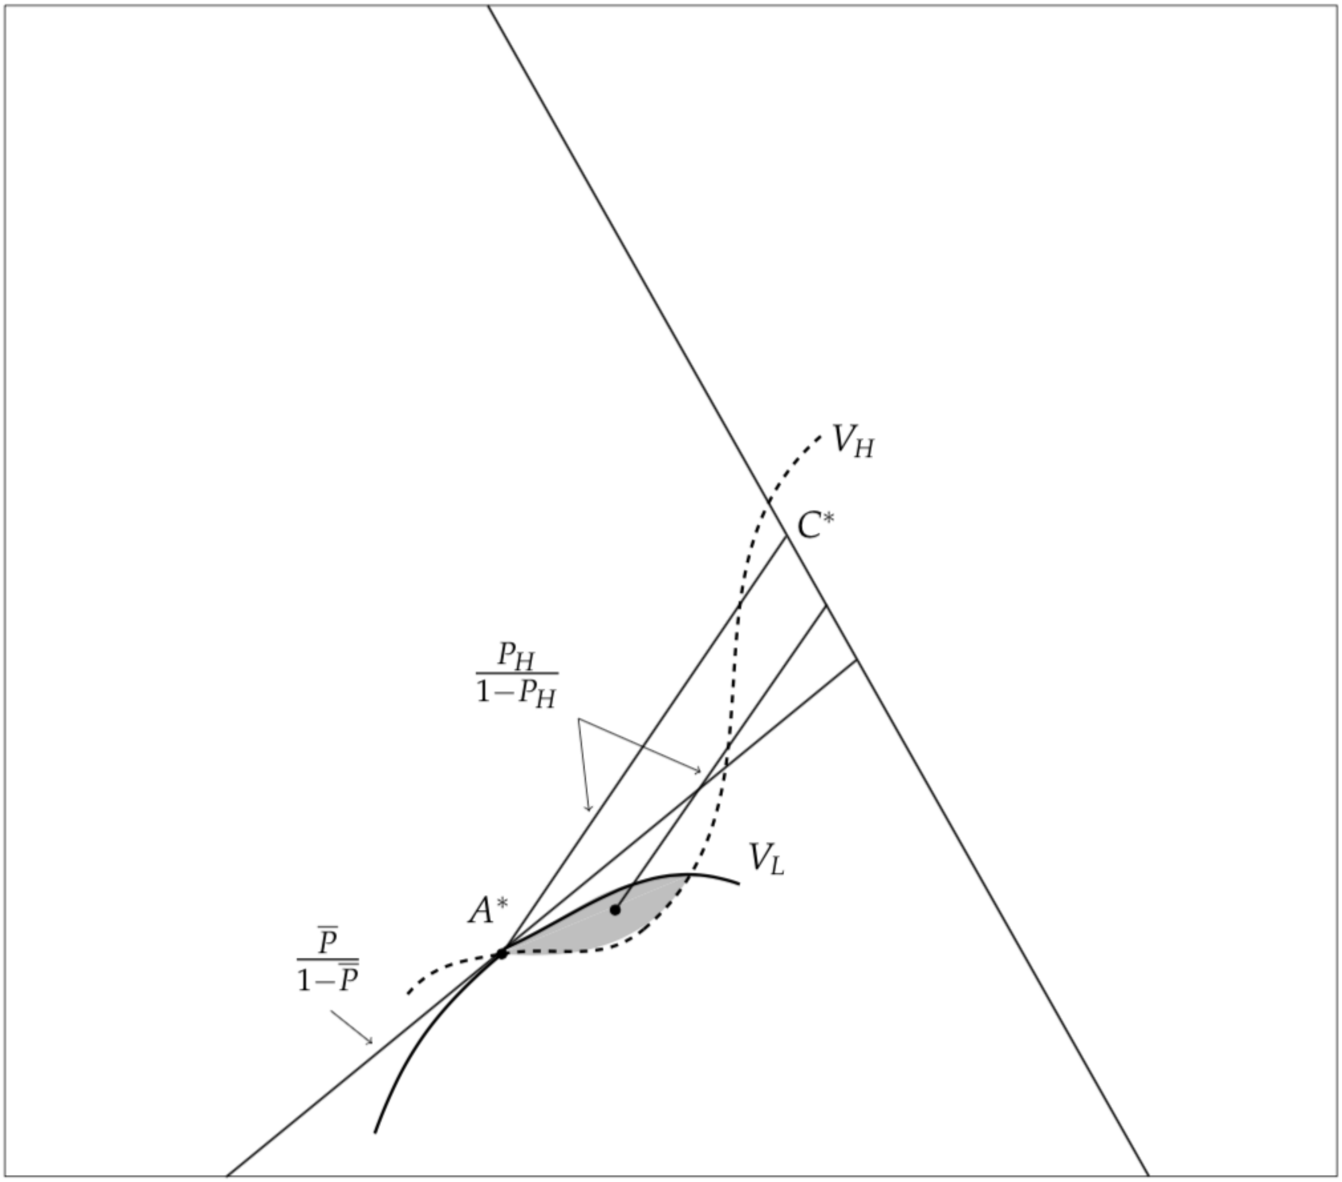

Figure 4: Equilibrium without single-crossing.

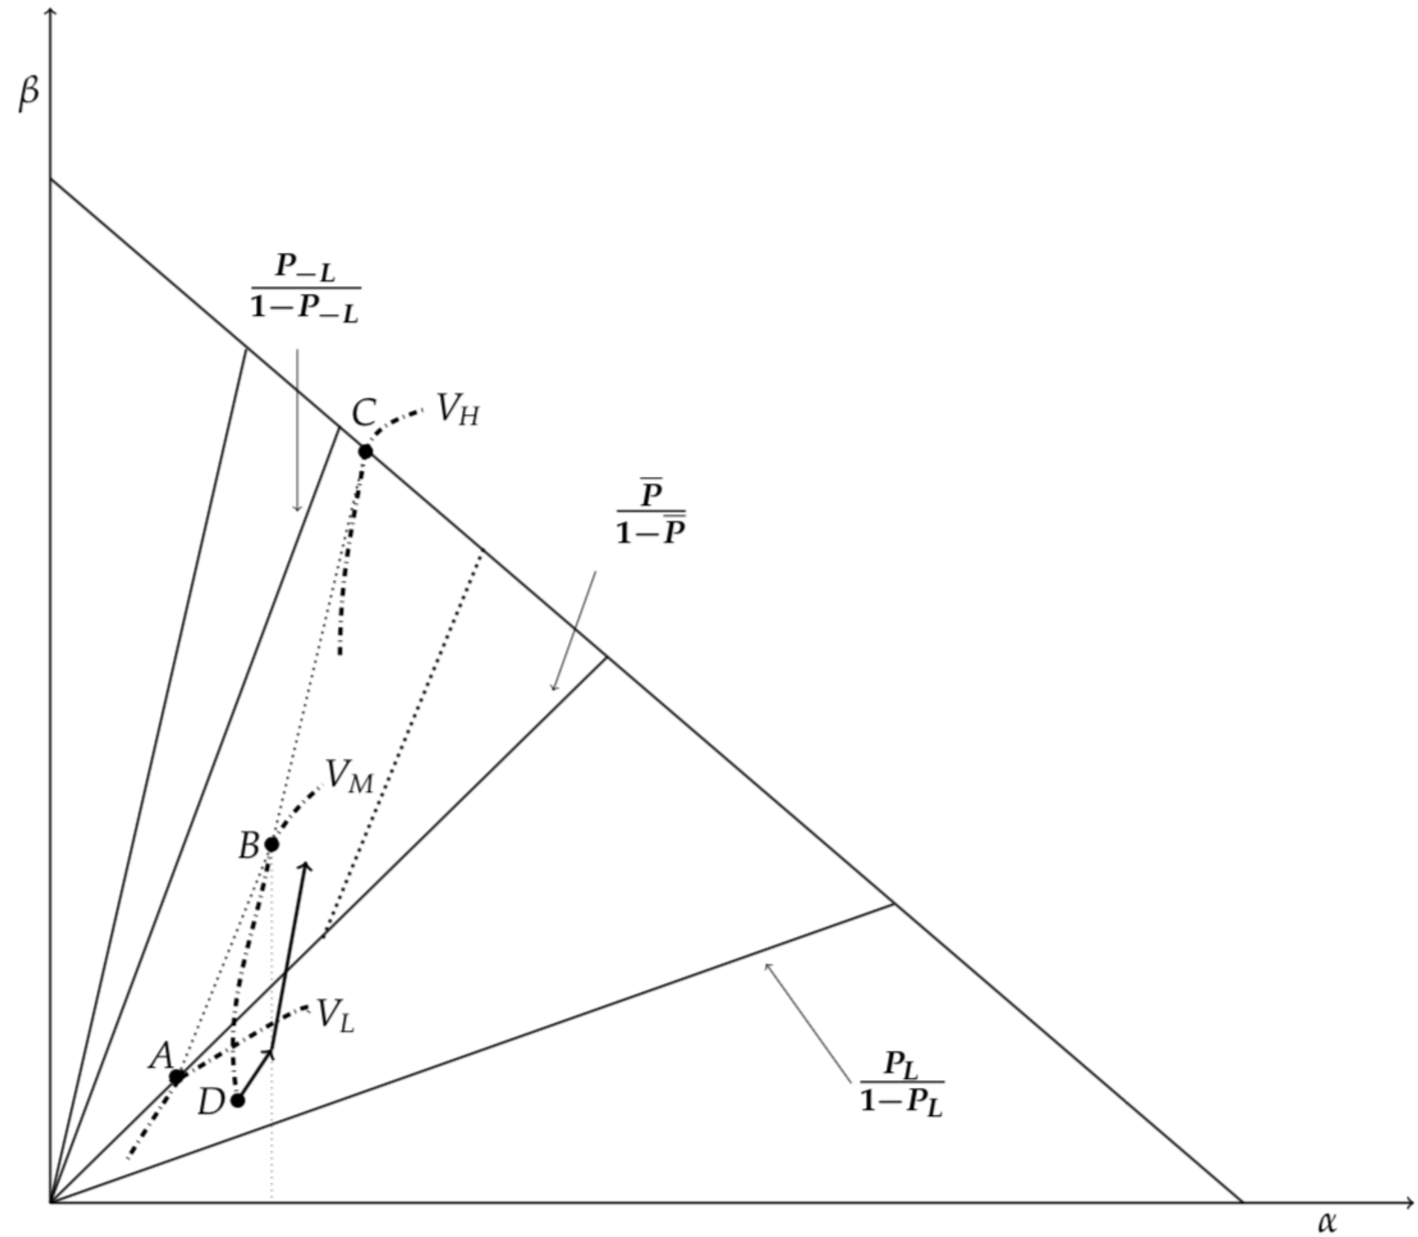

Figure 5: Equilibrium (A, B, C) with three types, which cannot be broken by D as individuals of higher-risk type supplement it by additional pooling insurance (along the arrow) without being disclosed to the deviant firm. P_{-L} denotes the average probability of accident for the two highest risk types, while V_i indicates an indifference curve for i-risk type ($i = H, M, L$).

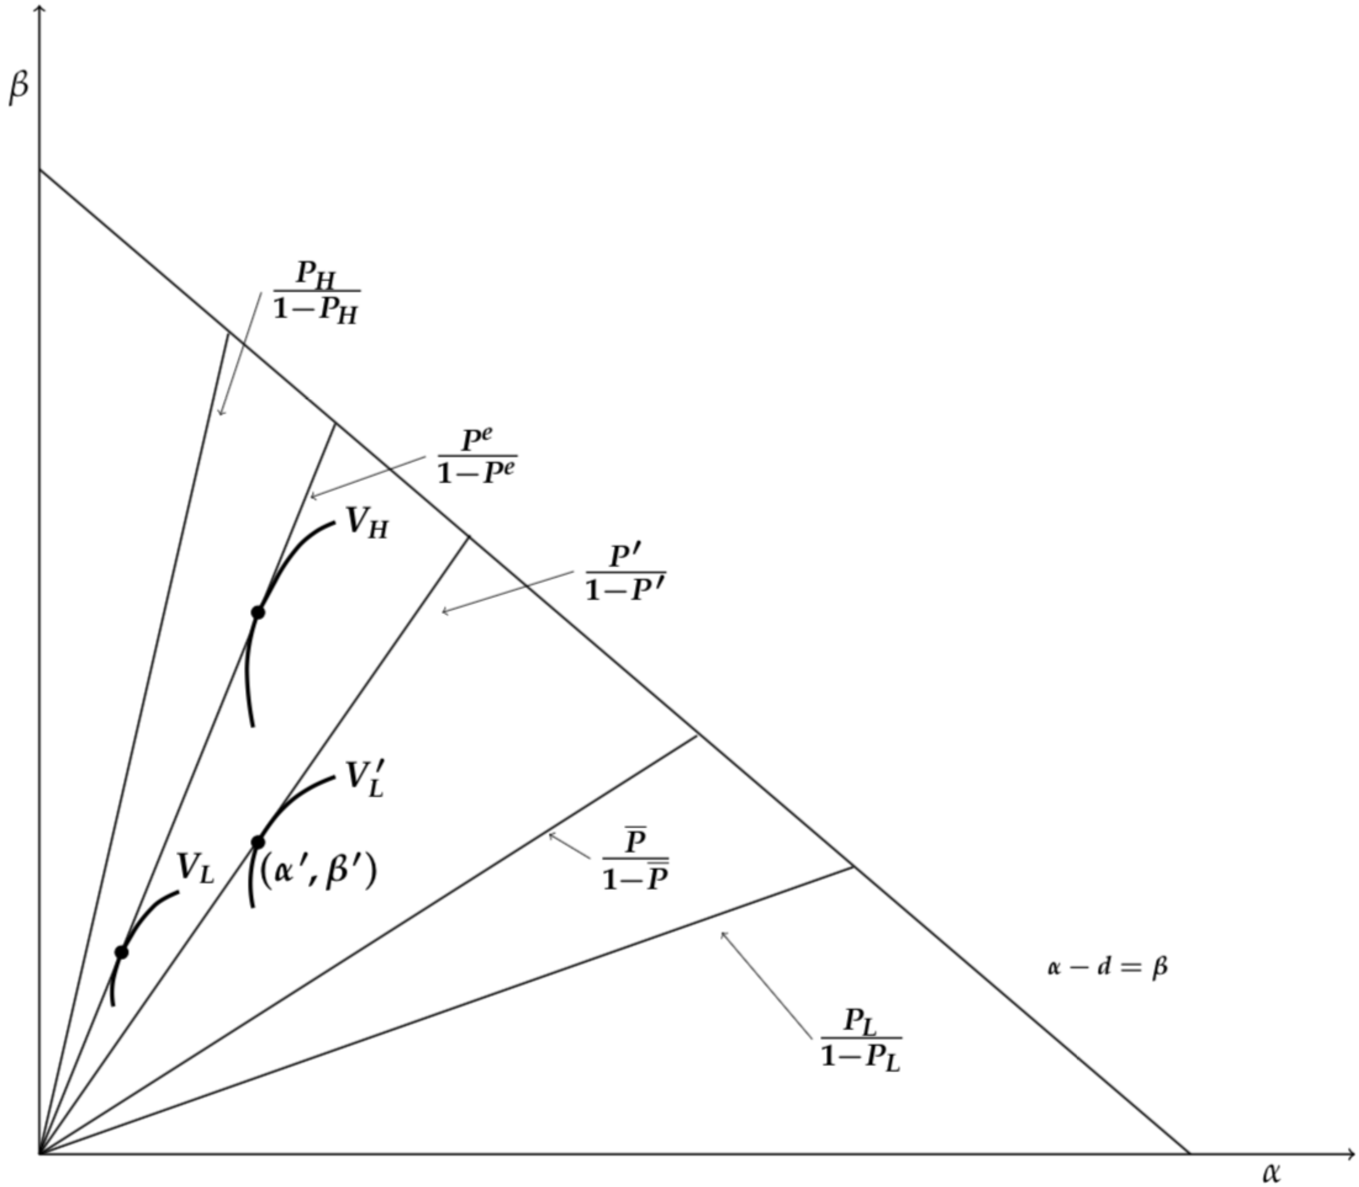

Figure 6: Breaking No-Disclosure-Information Price Equilibrium P^e by a fixed-quantity contract (α', β'), where $P^e > P' > \overline{P}$.

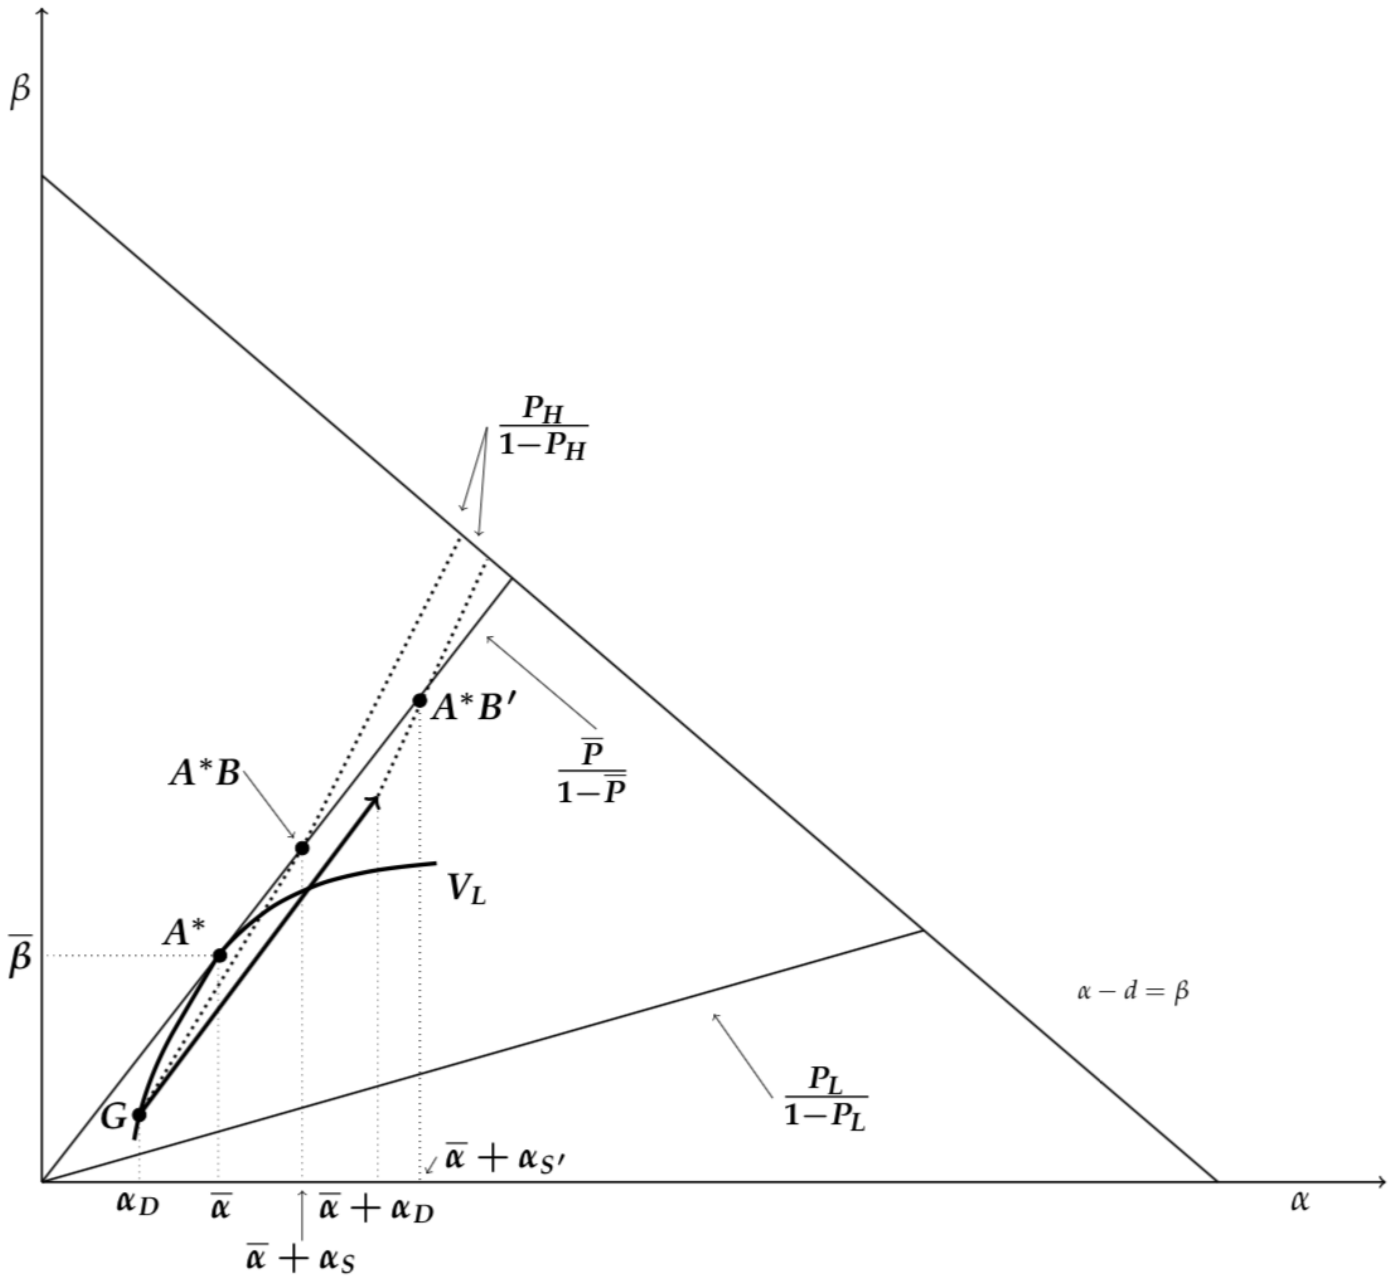

Figure 7: Nash Equilibrium can be sustained against multiple deviant contracts (A^*B, G) or (A^*B', G) offered at different prices as high-risk individuals also choose G (over A^*B) or as (A^*B', G) yields losses for the deviant firm (while inducing self-selection).