Bilateral Information Disclosure in Adverse Selection Markets with Nonexclusive Competition*

Andrew Kosenko,[†] Joseph Stiglitz,[‡] and Jungyoll Yun [§]

September 23, 2022

Abstract

We study insurance markets with nonexclusive contracts, introducing bilateral endogenous information disclosure about insurance sales and purchases by firms and consumers. We show that a competitive equilibrium exists under remarkably mild conditions and characterize the unique equilibrium allocation. With two types of consumers the allocation consists of a pooling contract that maximizes the well-being of the low-risk type (along the zero-profit pooling line) plus a supplemental (undisclosed and nonexclusive) contract that brings the high-risk type to full insurance (at his own odds). We show that this outcome is extremely robust and constrained Pareto efficient. Consumer disclosure and asymmetric equilibrium information flows are critical in supporting the equilibrium.

JEL Classification: D43, D82, D86.

Keywords: Adverse selection, nonexclusivity, nonexclusive competition, Rothschild-Stiglitz, asymmetric information, verifiable disclosure.

---

*We are grateful to several anonymous referees, an anonymous associate editor, and Daniela Puzzello, whose comments significantly improved the paper, to the participants in the seminar at Sciences Po at which an earlier version of this paper was presented, to Gerry Jaynes for helpful comments on an earlier draft, to Michael Rothschild and Richard Arnott, long time collaborators, to Byoung Heon Jun, Debarati Ghosh, Andrea Gurwitt, Parijat Lal, and Lim Nayeon for research and editing, and to the Institute for New Economic Thinking and the Ford Foundation and Fulbright Foundation for financial support. The companion paper "Characterization, Existence, and Pareto Optimality in Insurance Markets with Asymmetric Information with Endogenous and Asymmetric Disclosures; Revisiting Rothschild-Stiglitz" (Stiglitz et al. (2018), NBER Working Paper 24711) contains more results.

[†]Assistant Professor, School of Management, Marist College. 51 Fulton St, Poughkeepsie, NY, 12601. Email: kosenko.andrew@gmail.com.

[‡]University Professor, Columbia University. Kravis Hall, Room 543, 665 West 130th St, New York, NY 10027. Email: jes322@columbia.edu. Corresponding author.

[§]Professor, Department of Economics, Ewha University. 11-1 Daehyun-dong, Seodaemun-gu, Seoul 120-750, Seoul 120, Republic of Korea. Email: jyyun@ewha.ac.kr.

1 Introduction

In 1976, Rothschild and Stiglitz 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 (Stiglitz (1977)), no single firm framed the set of contracts from which individuals chose. There never existed a pooling equilibrium (in which the two types bought the same policy). If an equilibrium did exist, it entailed the high-risk individual getting full insurance, and the low-risk individual getting only partial insurance; and under plausible conditions - for example, if the two types were not too different - a competitive equilibrium did not exist. Finally, the sufficient condition underpinning the existence of a competitive equilibrium when the accident probabilities are different enough, the single-crossing condition, was very restrictive. The results were disquieting, as in reality equilibrium seemed to exist and often entailed pooling.

A vast literature has applied the Rothschild and Stiglitz (1976) model (henceforth RS), to labor, capital, and product markets in a variety of contexts with many empirical applications. A smaller literature focused on remedying the deficiencies in the underlying framework by formalizing the insurance "game", by changing the information assumptions, and by changing the equilibrium concept.

This paper introduces bilateral endogenous information disclosure about insurance purchases. Each firm and each consumer makes decisions about which information to disclose to whom. Thus, information about contract purchases is both endogenous, and potentially asymmetric because a firm may disclose information about a consumer to some firms, but not to others, depending on what the consumer discloses to it. We are motivated by the following observations: The outcomes with full information disclosure (exclusivity is enforceable, so the RS model applies, and a pooling equilibrium is impossible) and with no information disclosure (in which case pooling again cannot be an equilibrium) are known. Without consumer disclosure, any disclosure that firms make has to be symmetric, since they have no basis for differentiation. As we show later symmetric disclosure cannot underpin an equilibrium. The question of what happens if disclosure is bilateral, endogenous, and thus potentially asymmetric, is thus natural.

That asymmetries in information about insurance purchases turn out to be important is perhaps not surprising. What is surprising is that: (i) equilibrium exists under mild assumptions (notably, the single-crossing property need not hold); and (ii) equilibrium always entails a pooling component. The unique insurance allocation (an allocation describes the sum of benefits and premia over all insurance companies for each type of individual) consists of the pooling contract which maximizes the well-being of the low-risk individual subject to the zero-profit constraint, plus, for the high-risk individual, a supplemental contract that brings him to full insurance at his own odds. While the equilibrium allocation is unique, it can be supported by alternative information disclosure strategies.

The allocation we focus on has been the subject of much study in various guises, beginning with Jaynes (1978) and most recently, by Attar, Mariotti, and Salanié ((2020a), (2020b)). Under certain conditions, this allocation has been shown to emerge as the equilibrium in non-competitive models without endogenous information.

One part of our contribution is to show that this allocation can, in fact, be the outcome of a competitive equilibrium with endogenous bilateral disclosure, and that many of the standard but restrictive assumptions are not needed for this result. Previous work has studied some of the efficiency properties of this allocation. We extend that work by showing the context in which the allocation is constrained Pareto efficient (CPE), and then link that property to the set of possible equilibrium outcomes. We explicitly decentralize a CPE allocation as a competitive equilibrium in an insurance setting with adverse selection, nonexclusivity, and endogenous information disclosure.

Crucially, we argue that decentralization via a competitive market in a setting with any fixed information structure is impossible. (Thus, Jaynes (1978), and much of the other literature generating the allocation we focus on is explicitly set in a context of imperfect competition.[1]) Any competitive-like equilibrium requires endogenous and asymmetric information sharing. In RS, Akerlof, and much of the existing related literature, information is endogenous but only affected by actions (e.g., directly observable purchases of insurance). In the analysis here, more in the spirit of much of the mechanism and information design literature that flourished after RS,[2] individuals can disclose (truthful) information, which affects behavior, including others' disclosures. In the equilibrium we construct, firms do not always find out the type of agent they are trading with (as is the case in the RS separating equilibrium), but there are nontrivial information exchanges (between firms, and from consumers to firms) associated with insurance purchases.

We begin by characterizing the set of CPE allocations in the presence of a secret contract. We then show that the CPE allocation which maximizes the well-being of the low-risk individual (supplemented, for the high-risk individuals, by insurance at the high-risk odds, bringing those individuals to full insurance) is the unique equilibrium allocation, and can be supported by simple yet illuminating information disclosure strategies. As in RS, firms offer insurance contracts, but now they have an option to reveal (possibly partial) information about insurance purchases to other firms. In RS, it was assumed that contracts were exclusive, e.g., implicitly firms had full knowledge of individuals' pur-

[1]Evident particularly in the title of Jaynes' (1978) paper: "Equilibria in Monopolistically Competitive Insurance Markets."

[2]See, for instance, Kamenica and Gentzkow (2011), Bergemann and Morris (2019) and Taneva (2019). For the most part, that literature provided limited analyses of market situations, such as the insurance market that is the subject of this paper. Thus, our paper may also be viewed as a contribution to the information design literature.

chases, and 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, and under a broader set of assumptions concerning the information available to insurance firms. Obviously, the enforceability of any conditions imposed is dependent on information available to the insurance firm. Consumers, too, face a more complicated choice: they have to decide which policies to buy, aware of the restrictions in place and the information that the firm may have to enforce those restrictions. They also have to decide what information to reveal to which firms. A competitive equilibrium in this model is a set of insurance contracts, such that no one can offer an alternative contract or set of contracts and make positive profits. A contract is defined by the benefit, the premium, any associated restrictions, and the firm's disclosure policy. And in assessing the consequences of offering an alternative contract, each firm takes into account the consumers' response to the set of contracts on offer, both with respect to insurance purchases and disclosures.

The intuition behind our result is this: in RS a pooling equilibrium can always be broken by a deviant contract that will be purchased only by low-risk individuals and, as a result, is profitable. But that deviant contract will be purchased only by low-risk individuals because the deviant firm can enforce exclusivity. If high-risk individuals can supplement the deviant contract (one breaking the putative pooling equilibrium) with secret insurance at their own odds, that policy may be purchased by high-risk individuals, and thus make a loss. Hence, the deviant contract will not be offered and the pooling contract can be sustained. The trick is to find an information disclosure strategy that ensures that a deviant firm can't enforce exclusivity, but also ensures that the firms selling insurance at the pooling odds (which we refer to as "established" firms) don't "oversell": High-risk individuals would like to buy more insurance at the pooling odds than low-risk individuals. But if they did, the pooling contract would lose money. Accordingly, there has to be sufficient information disclosure to prevent the high-risk individuals from doing so. Thus, supporting the equilibrium allocation requires an *intermediate* amount of disclosure: One needs some information sharing (enough to prevent overselling), but not too much (not enough to enforce exclusivity). Furthermore, disclosure has to be asymmetric in that established firms must have sufficient information, but deviant firms (which, of course, deviate secretly) must not. But firms by themselves have no basis for such asymmetric disclosures. Without further information, they only know whether they themselves have sold insurance to an individual.[3]

This is where consumer disclosure - an essential feature distinguishing our paper from

[3]This is essentially the point that Hellwig (1988) makes in criticism of Jaynes (1978) argument that with endogenous information, there always exists an equilibrium. In contrast, he emphasizes that "...there does not exist a sequential equilibrium for the RS-type specification of the game." He shows that Jaynes' equilibrium requires that each firm's communication strategy be conditioned on the set of contracts that are offered by other firms, making the equilibrium a reactive equilibrium, like that of Wilson (1977), not a competitive equilibrium as in RS.

other work in this area - becomes critical. Firms base the asymmetries in disclosure on consumer-revealed information. The equilibrium firm information disclosure strategy that we analyze induces truth-telling by consumers to established firms, and this in turn enables asymmetries in firm disclosures of information about insurance purchases. Thus, endogenizing consumer disclosure is not just a natural modeling postulate - it is necessary for the outcome we characterize. To the best of our knowledge, these disclosure strategies (expressed simply by saying that "once a consumer makes insurance purchases and freely reveals some subset of his insurers to each firm, each firm reveals information to those firms that the consumer did *not* reveal as his insurers") have not been identified previously.

Our proof strategy is the following. We first arbitrarily divide the set of firms into two groups, one "established," which engages in contracts with disclosure rules, and the other "secret," offering secret contracts. We restrict both the set of policies that any established firm can offer to all having the same price, allowing individuals to purchase any amount of insurance at that price while possibly imposing a constraint on aggregate purchases; and the set of disclosures (in a way we clarify later), establishing that *given these constraints*, there is an equilibrium. We also assume that an insurance firm can sell to any individual only one contract (policy). We then show in section 6.1 that, given these contract offers and disclosures, allowing for *any* contract offer or disclosure, no firm or consumer would deviate from their chosen behavior, i.e., the restrictions imposed in the beginning of the analysis are not in fact binding. Furthermore, no firm that is secret (does not make any disclosures) would want to become an established firm, and vice versa. In appendix B we extend the analysis to the case where deviant firms can offer a menu of contracts.

One can obtain this result by formalizing this setting as a game with appropriately defined strategy spaces and focusing on the outcome in a perfect Bayesian equilibrium; however, doing so would introduce unnecessary complexity,[4] thus detracting attention from the basic insight of our analysis. For this reason, we pursue the route of the original RS paper, positing only the elements that are absolutely necessary to make the point in the simplest possible setting that retains all of the features we are interested in, using an equilibrium concept that is in the spirit of competitive equilibrium.

We view this work as a contribution to the pure theory of competitive equilibrium with asymmetric information; but we believe that some of its main insights are of significant relevance for real-world markets with adverse selection and nonexclusivity. These insights include the robustness of a pooling-plus-separating allocation, and the simple, nontrivial, and illuminating information disclosure strategies which can support such an equilibrium. Towards the end of the paper we illustrate its relevance by considering its

[4]We work through a model along these lines in Stiglitz and Yun (2016).

implications for government provision of insurance under the "public option." We hope that our results provide an impetus for further policy and empirical applications, with insights into why certain markets take the form they do, and how one might improve the design of markets with asymmetric information.

The paper is organized in nine sections. Section two lays out the basic features of our model of insurance with nonexclusive contracts, while section three provides a discussion of nonexclusive ("secret") contracts in workhorse RS and adverse selection models.[5] We characterize the set of CPE contracts in the presence of secret insurance in section four. Section five provides a formal discussion of firms' contracts, information disclosure strategies, the equilibrium concept, and consumer information disclosure and purchase decisions, and shows that there is a unique allocation that an equilibrium - if it exists - has to implement. In section six we first prove that the equilibrium exists under some restrictions (lemma 6.2),[6] and then lift those restrictions to establish the main result of the paper - theorem 6.3 - that the equilibrium established in lemma 6.2 remains an equilibrium without those restrictions. Section seven considers a generalization to the case of many types, while section eight considers the impact of a public insurance option. Section nine relates our results to previous literature. For brevity, we forego a concluding section, and discuss an extension where firms offer menus of contracts (multiple policies with cross-subsidization) in appendix B. Appendix A contains a proof omitted in the text.

2 Model

We employ the standard insurance model with adverse selection. There is a continuum of individuals, each facing the risk of an accident. The two types of individuals - high-risk ($t = H$) and low-risk ($t = L$) - differ only in the probability of accident, P_t, with $P_H > P_L$. The type is privately known to the individual, while the proportion θ of high-risk types is common knowledge. The average probability of accident for an individual is $\overline{P}$, where

$$\overline{P} = \theta P_H + (1 - \theta) P_L \tag{1}$$

An accident involves damages, the cost of repairing which in full is d. An insurance firm pays a part of the repair cost, $\alpha \leq d$ (we disallow negative insurance). The benefit is paid in the event of accident, whereas the insurer is paid an insurance premium β when no

[5] Set against the backdrop of equilibrium nonexistence in those models (which we establish), it seems all the more striking that an equilibrium would exist in a more complex model with endogenous and nontrivial information flows under very mild conditions.

[6] The only restriction on the behavior of the "secret" firms is that they not disclose their sales, a restriction that we remove in section 6.2. While we thus impose restrictions on both the established and secret firms, we impose no restrictions either on the policies or disclosures of deviant firms, those attempting to break the equilibrium, except that a deviant firm can only offer a single policy, a restriction that we remove in appendix B.

accident occurs. Define $q \triangleq \frac{\beta}{\alpha}$ as the price of insurance (payment in the "good state" per dollar received in the case of an accident). We assume the utility of the individual of a given risk type is a function of his consumption if he has an accident ($w - d + \alpha$), where w is the starting wealth level, and if he does not, ($w - \beta$):

$$V_t(\alpha, \beta) \triangleq U_t(w - d + \alpha, w - \beta) \tag{2}$$

We assume that consumer preferences are weakly convex,[7] that V is bounded, continu-

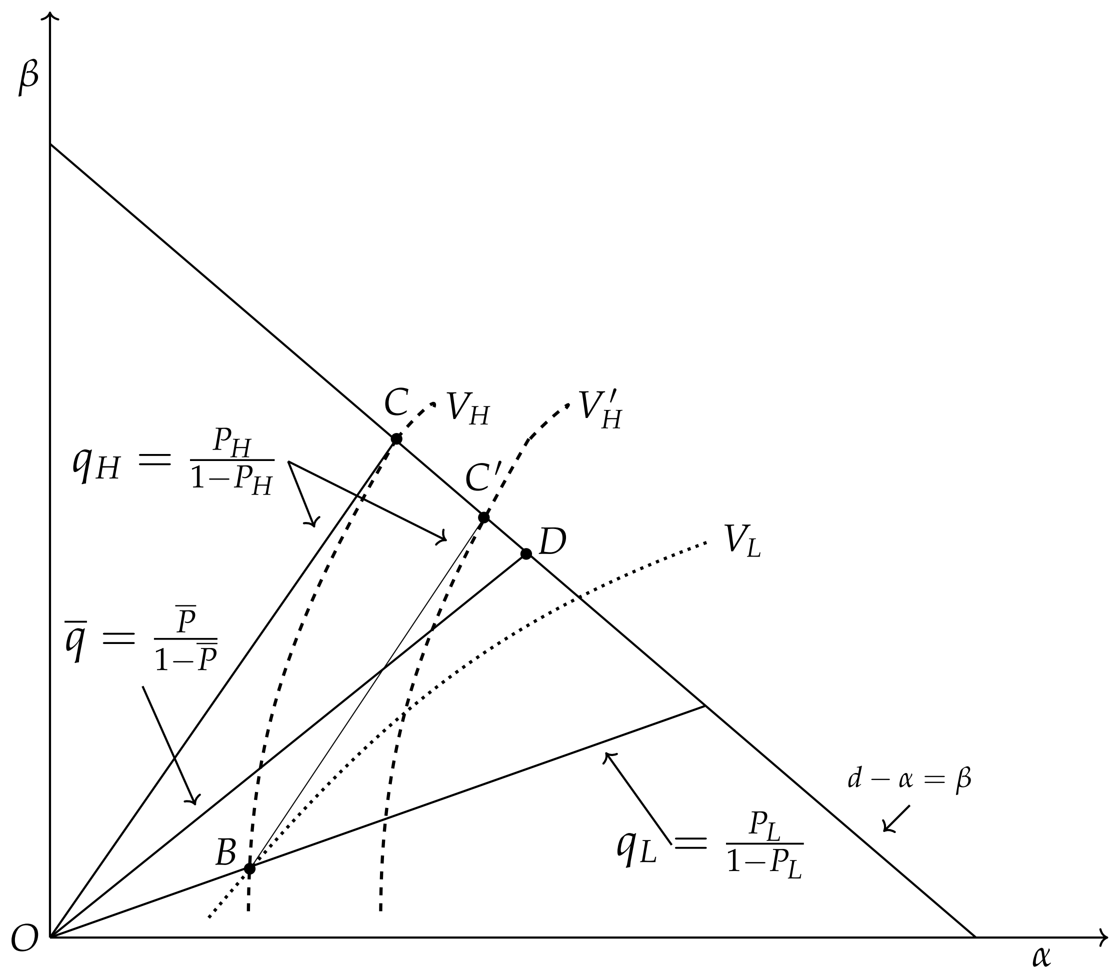

Figure 1: Breaking the RS separating equilibrium (B, C) in the presence of undisclosed contracts at high-risk odds. V_H is an indifference curve of type H, V_L is an indifference curve of type L, the line $d - \alpha = \beta$ is the full insurance line. High-risk individuals will purchase contract B, supplementing it with secret insurance, bringing the individual to full insurance and a higher level of utility than in the separating equilibrium.

ously differentiable, increasing in α, and decreasing in β.

We refer to a policy $A = (\alpha, \beta)$, and to the expected utility generated by that policy

[7] For much of our analysis, we do not even require that, but the exposition is simplified by making this assumption. The assumption plays a role in proposition 2.

as $V_t(A)$. A policy A, with insurance level α and price q can also be described by the vector $(\alpha, \alpha q)$. The key properties of $V_t(\alpha, \beta)$, which we assume are satisfied throughout the paper, are (*i*) at full insurance, the slope of the indifference curve equals the relative probabilities,

$$-\frac{\frac{\partial V_t(\alpha,\beta)}{\partial \alpha}}{\frac{\partial V_t(\alpha,\beta)}{\partial \beta}} = \frac{P_t}{1-P_t} \triangleq q_t \tag{3}$$

and (*ii*) quasi-concavity, so that the indifference curves take on the usual shape. If individuals can purchase as much insurance as they want at a price q_t then, of course, the (absolute value of the) slope of the indifference curve equals q_t at the level of insurance chosen.

These assumptions in turn imply that the income consumption curve at the insurance price $\frac{P_t}{1-P_t}$ is the full insurance line,[8] implying that with full information, equilibrium would entail full insurance for each type at their own odds, i.e., $q_t = \frac{P_t}{1-P_t}$.

A special case of eq. 2 is standard expected utility, with $U_t'' < 0$:

$$V_t(\alpha, \beta) = P_t U_t(w - d + \alpha) + (1 - P_t) U_t(w - \beta) \tag{4}$$

While it is a useful special case, we do not rely on the expected utility formulation for the results. If in eq. 4 the utility functions are the same for high and low-risk individuals, the single-crossing property that was essential in the RS analysis is satisfied. But if, for instance, individuals with a higher probability of an accident are also more risk averse, then the single-crossing property will never be satisfied.[9] None of the results below depend on the single-crossing property being satisfied, so we do not invoke it. Moreover, we do not assume that the utility function is separable across the states.

The profit π_t of a policy (α, β) that is chosen by type t is $\pi_t(\alpha, \beta) = (1 - P_t)\beta - P_t\alpha$. $\pi_t(\alpha, \beta) = 0$ is defined as the t-type's zero-profit locus (the line along which firms selling to type t make zero profit). Figure 1 illustrates the zero-profit locus for a firm selling insurance to a t-type (OB and OC, respectively), or both types (OD) of individuals in proportions θ and $1 - \theta$, by a line from the origin with the slope being $q_t (= \frac{P_t}{1-P_t})$ or $\overline{q} \left(\triangleq \frac{\overline{P}}{1-\overline{P}} \right)$, respectively. The latter is referred to as the *zero-profit pooling line*.

There are $N \geq 3$ firms and the identity of a firm is represented by j, with $j \in \mathcal{M} (\triangleq 1, \ldots, N)$. We also assume initially that these N firms are exogenously sorted into one of two groups: a set $\mathcal{M}^{\mathcal{E}}$ of "established" firms with $|\mathcal{M}^{\mathcal{E}}| = n \geq 2$, and a set $\mathcal{M}^{\mathcal{S}}$ of "secret" firms with $|\mathcal{M}^{\mathcal{S}}| = N - n \geq 1$; the assumption of division into two groups can

[8] This, together with convexity, is the critical property used in proving the main result of the paper. With considerably greater complexity, the critical results can be established without assuming quasi-concavity.

[9] There are other reasons to be concerned about the single-crossing property: in models with moral hazard and adverse selection, where P_t is endogenous, there is also the presumption that the single-crossing property will not be satisfied (see Stiglitz and Yun (2016)).

be dropped, as we have already noted.[10] Of course, $\mathcal{M}^{\mathcal{E}} \cup \mathcal{M}^{\mathcal{S}} = \mathcal{M}$, and thus $|\mathcal{M}^{\mathcal{E}}| + |\mathcal{M}^{\mathcal{S}}| = N$. Whatever the values of n and N, we assume that the firms act competitively, as described below. The firms differ in their information disclosure: The secret firms do not disclose their insurance sales to anyone. While the secret firms do not disclose the insurance policies they sell to anyone, the (possibility of the) existence of such firms is known. Among the secret policies possibly on offer, one plays a critical role in the subsequent analysis: an offer of unlimited insurance at the price q_H. Such a policy must at least break even since the worst that could happen to any insurance firm is that only high-risk individuals buy the policy. The disclosure policy of the established firms is an essential object of study of this paper.

An individual may purchase multiple policies from any subset of firms. A set of benefits and premiums of the insurance policies purchased in the aggregate by each type of individual, denoted by $E = \{(\alpha_t, \beta_t)_{t=L,H}\}$ is called an *allocation*, with $\alpha_t = \sum_j \alpha_t^j$ and $\beta_t = \sum_j \beta_t^j$, where j is the index of the consumer identities.

3 RS and Price-Quantity Equilibria without Exclusivity

Our paper focuses on situations where the information structure is endogenous, in contrast to the earlier competitive equilibrium literature centering around markets with fixed information structures, where firms that disclose their sales always do and those that don't disclose never do. In that earlier literature, insufficient attention was paid either to secret insurance or to the information that a firm might glean from its own sales, even in markets without exclusivity. Before analyzing in the next section the implications for Pareto efficient allocations and competitive equilibrium with secret contracts and endogenous disclosure, we demonstrate in this section that broadening out the information problem in these ways while retaining the assumption of fixed information structures has a devastating consequence: Competitive equilibrium never exists. We begin by looking at the RS equilibrium before turning to a standard price equilibrium in the presence of adverse selection.

3.1 Nonexistence of the RS equilibrium in the Presence of Secret Contracts

Central to RS was the assumption that there was sufficient information to enforce exclusivity; an individual could not buy insurance from more than one firm. As Rothschild and Stiglitz realized, once we introduce into the analysis unobservable contracts in addition to observable ones, the whole RS framework collapses because exclusivity cannot be

[10] If there is only one firm it can trivially impose exclusivity, being a monopolist, so we assume there are at least two established firms. As we discuss below, the same firm can, in fact, be both an established and a secret firm, via-a-vis different customers.

enforced. We begin by recalling the basic definition of equilibrium in RS: a set of contracts such that, given the putative set of equilibrium contracts, there was no contract that could be offered and taken up that would make positive profits.[11] We now ask, what happens if we add to the RS framework the *possibility* of a set of firms offering secret insurance? We assume that each firm can observe the sales of purchases from other established firms to any individual and therefore could enforce exclusivity among the set of established firms - but there may possibly be secret firms. It is easy[12] to show the following:

Lemma 3.1. *Any RS equilibrium, if it exists, entails secret insurance at price q_H.*

We make one more informational assumption: The insurance firms know the identity of the individual to whom they are selling, so that they can ascertain the total quantity of insurance they have sold. This is a natural assumption for individuals purchasing insurance against risks that directly affect only them, such as an accident or their death, but not necessarily individuals purchasing insurance against a market event, like the decrease in the price of a commodity.

Now consider a possible equilibrium where secret insurance is on offer at q_H, which by lemma 3.1 we know must be the case. With this secret insurance on offer, the separating contracts from RS are also not equilibrium contracts, as illustrated in figure 1. The RS separating contracts are (B, C), where C provides full insurance for the high-risk individual at his own odds; and B is the contract at the low-risk individual's odds that just satisfies the self-selection constraints, i.e., it will not be purchased by the high-risk individual. Clearly, if the high-risk individual *can* supplement B with secret insurance at the high-risk odds, he will purchase B. But if high and low-risk individuals both purchase B, it makes a loss. A similar analysis enables us to easily show that there exists no separating equilibrium.

To rule out a pooling equilibrium, suppose that there is one with coverage α_{Pool} at a price greater than $\overline{q}$[13]; then there exists a contract with a price between $\overline{q}$ and that price

[11]In our later discussion with endogenous disclosure, "contracts" are much richer, including disclosure rules. What we call in section 5 a "policy" i.e., a benefit and a premium, RS referred to, given the restricted setting of their model, as a "contract". This terminology should cause no confusion.

[12]Assume that there were a set of contracts with secret insurance at a price greater than q_H, and insurance at price q_H. Clearly, no one would buy the high-price contracts. Now assume that there existed a RS equilibrium in which the secret contract at price q_H is not on offer but there were a set of contracts with secret insurance at a price greater than q_H some of which were purchased by some individuals. Such contracts would make strictly positive profits. Then a firm offering secret insurance at a slightly lower price would still be profitable and attract away all customers from the firms selling secret insurance at a higher price. Assume, on the other hand, that there were a putative RS separating equilibrium without secret insurance. Then a firm that offered secret insurance at a price slightly greater than q_H would be purchased by the high-risk individual and make strictly positive profits, breaking the putative equilibrium. The same would hold for any proposed pooling equilibrium. Thus, if it is possible for there to be secret insurance, any putative equilibrium must have, as part of it, a secret contract at price q_H, and it is the existence of this contract that itself prevents there being either a separating or a pooling equilibrium, as we show below.

[13]It is easy to show that the only possible pooling equilibrium entails an insurance price of $\overline{q}$.

which will make money, because at worst it will attract everyone, and if it attracts everyone it makes a profit. (Recall our assumption that the insurance firm can identify who is buying the insurance, so it will not sell more than one policy to any individual.)

Now assume that the putative equilibrium entails a contract along the pooling line, but that it is not the contract that maximizes the low-risk individual's utility along that line, labeled E_P^* in figure 2. Then, by the reasoning above, a firm that offered a contract just above the pooling line that was arbitrarily close to the one that maximized the low-risk individual's utility would be bought by all low-risk individuals, and make a profit - a large profit if the high-risk did not buy it, possibly a small one if they did.

But E_P^* itself cannot be an equilibrium. From the assumptions on preferences, we know that the high-risk individual's indifference curve through E_P^* has a slope at that point greater than q_H, which is greater than $\overline{q}$, which is the slope of the low-risk individual at that point. Thus, there exists a contract below $V_L(E_P^*)$ (and therefore below the pooling line) but near E_P^*, i.e., above the zero-profit line for the low-risk individual, that gives the low-risk individual a higher level of utility than E_P^*, but which will not be purchased by the high-risk individual, and which accordingly makes a positive profit. This is illustrated in figure 2 and establishes that there cannot exist a pooling equilibrium.

It can similarly be shown that a hybrid equilibrium with partial pooling cannot exist either. Since there cannot exist a pooling equilibrium, a separating equilibrium, or a hybrid equilibrium, it follows that there never exists a RS competitive equilibrium: tautologically, this exhausts the forms that the equilibrium can take.

3.2 Price Equilibrium

A price equilibrium,[14] where insurance firms offer insurance at a fixed (competitive, i.e., zero-profit) price also does not exist, under our earlier assumption that firms know the identity of the individual to whom *they* are selling.

The firm knows purchases from itself but not from other established firms and obviously not from secret firms. The standard price equilibrium in the context of insurance entails a price of insurance reflecting the actuarial risk of the mix of individuals buying the insurance. The presence of adverse selection means that as the price of insurance increases, the actuarial risk increases. A price equilibrium is a price of insurance generating zero profits. The competitive equilibrium price is $q_c > \overline{q}$, reflecting the fact that at any price high-risk individuals buy more insurance than low-risk individuals.[15] Under our informational assumption that the firm can monitor its own sales to an individual, any firm could offer a fixed quantity contract, based only on its own information. Any price

[14]The price equilibrium in insurance markets is the context in which the adverse selection problem was first studied. See e.g., Arrow (1965).

[15]This is the *positive correlation* property that was also identified by Chiappori and Salanié (1997, 2000) in empirical work on this topic.

equilibrium can be broken by a firm offering the quantity contract along a line through the origin just below the line through the origin with slope corresponding to q_c that is most preferred by the low-risk individual (at that price), since every individual will buy it, and that contract will make a profit, *even in the presence of purchases by individuals whose characteristics it cannot observe*. But that contract, supplemented by secret insurance at q_H (purchased only by high-risk individuals) can't be an equilibrium, by the reasoning just given for the RS model. It can be broken by still another contract that will skim off the low-risk individuals, *even in the presence of secret insurance*. Thus, the standard adverse selection price equilibrium does not exist if firms can offer a fixed quantity contract - knowing only the amount of insurance they sell to any individual. To emphasize that we have gone beyond the standard price equilibrium, we refer to this as a price-quantity equilibrium, which is a price equilibrium where any firm has the option of issuing a fixed quantity contract.[16] Summarizing these observations, we have

Proposition 1. *Suppose that an insurance firm knows the identity of the individual to whom it is selling. Then an RS or a price-quantity equilibrium never exists with secret insurance contracts.*

The RS and price equilibria are the two forms of competitive equilibrium in markets with a fixed information structure that have been extensively discussed. An equilibrium does not exist in either case. Thus, if an equilibrium is to exist in the presence of secret contracts, the model must be changed. There are at least two possible directions. One is departing from a competitive framework, e.g. assuming a reactive equilibrium (where incumbent firms respond to offers of an entrant), which may make sense when there are a relatively small number of firms.[17] Another, attempting to find an equilibrium in the spirit of competition, is the approach taken here: we do this by endogenizing the information structure.

4 Constrained Pareto Efficiency with Secret Contracts

We begin by characterizing the set of CPE allocations under the premise that there exists a secret (that is, available for purchase to the consumers, but entailing no disclosure whatsoever) policy, where the constraint is that the government cannot proscribe the secret provision of insurance. This analysis is interesting in its own right, clarifying how the existence of secret contracts constrains the set of feasible allocations. But it will also be useful

[16]While we couch our analysis in terms of insurance, the literature following RS made clear that an analogous analysis applied in many other markets, including credit markets. Enforcing exclusivity in that market has proven to be particularly problematic, with many high-risk individuals over-borrowing. One proposed solution has been a public registry, incentivized by a legal provision that only loans registered would be recognized by courts. But our analysis has made clear the problems with this "solution."

[17]This is the approach taken by Wilson (1977), Riley (1979), Jaynes (1978), and a number of other papers discussed in section nine.

in the next section when we analyze potential equilibrium allocations in the presence of secret policies, showing, importantly though not surprisingly, that any allocation that is not CPE is not entry-proof, and therefore cannot be part of an equilibrium. Formally, we use the following ex-interim variant[18] of constrained Pareto efficiency:

Definition 1. *An allocation $E = \{(\alpha_t, \beta_t)_{t=L,H}\}$ is constrained Pareto-efficient (CPE) if,*

i) in the aggregate, it at least breaks even, and

ii) satisfies the self-selection constraint, i.e., type i weakly prefers the proposed allocation for type i to the proposed allocation for type j; and if each individual prefers the allocation to just buying secret insurance or supplementing the allocation with secret insurance, and

iii) given that the government cannot force disclosure, there does not exist another allocation that (in the aggregate) at least breaks even, and leaves each type of consumer as well off and at least one type strictly better off.

We can think of the government proposing an allocation (aggregate benefits and premia for each of the two types). We will now show that for a proposed allocation to be CPE, it has to satisfy a number of properties.

First, by the reasoning invoked in the previous section, if the government cannot proscribe secret insurance, there will always be on offer an unlimited supply of insurance at price q_H. This greatly constrains the kinds of allocations that the government can propose. Given our assumptions about preferences, it implies that the proposed CPE allocation for the high-risk individual must entail full insurance. (Otherwise, the high-risk individual would want to supplement the proposed allocation with secret insurance.)

The zero-profit constraint means that any losses made from selling to one group must at least be made up for by profits from the other group. That is the case, of course, if the proposed allocation entails, for the low-risk individual, a point along the pooling line, and, for the high-risk individual, that same allocation supplemented by insurance at price q_H bringing the high-risk individual to full insurance.

It is also easy to establish that a CPE allocation must just break even; otherwise the welfare of group or the other can be improved. Moreover, the subsidy to the high-risk individuals has to be non-negative or else a putative allocation entailing a tax on the high-risk individual would be broken by the purchase of just secret insurance at price q_H.

One more necessary condition for a CPE allocation: an allocation with the low-risk individual along the pooling line getting an insurance benefit of α and the high-risk indi-

[18] See also Prescott and Townsend (1984), Hammond (1987), Bisin and Gottardi (2006), and Attar, Mariotti, and Salanié (2019) for important discussions of Pareto efficiency in related contexts; earlier versions of the present paper were (to our knowledge) the first to explicitly consider the "constrained Pareto efficiency" concept introduced here, allowing for secret contracts as well as disclosed contracts.

vidual supplementing that with insurance at q_H to get full insurance must entail $\alpha \geq \overline{\alpha}$, where $\overline{\alpha}$ is the level of insurance along the pooling line that maximizes the low-risk individual's utility. If $\alpha < \overline{\alpha}$, increasing insurance coverage along the pooling line would make the low-risk individual better off, and when the high-risk individual supplemented that with insurance at price q_H, he too would be better off.

We denote the allocation that maximizes the utility of the low-risk individual along the pooling line, and supplements that with insurance at price q_H bringing the high-risk individual to full insurance by $E^* = ((\alpha_L^*, \beta_L^*), (\alpha_H^*, \beta_H^*)) \triangleq ((E_P^*, \overline{q}E_P^*), (E_P^* + S^*, \overline{q}E_P^* + q_H S^*)$, where $S(E')$ is the supplemental policy at q_H ($S(E') \triangleq (\alpha_{S(E')}, q_H \alpha_{S(E')})$) that brings the high-risk individual to full insurance from an allocation E' for the low-risk individual (if E' involves coverage of α', $\alpha_{S(E')} \triangleq \frac{1}{1+q_H}[d + \alpha'(q_H - \overline{q})] > 0$ as long as there is less than full insurance, i.e., $\alpha' + \alpha'\overline{q} < d$), and where $S^* \triangleq S(E_P^*)$. The allocation for the low-risk individual E_P^* is, of course, precisely that discussed in the previous section.

Finally, it is easy to establish that another necessary condition for a CPE allocation is that the self-selection constraint be binding. Assume not. A CPE allocation to the high-risk individuals involves a particular subsidy from each low-risk individual of s^*. We maximize V_L subject to $\beta \geq \alpha q_L + s^*$ and the self-selection constraint. Given our assumptions about preferences, without the self-selection constraint, utility would be maximized at full insurance; but because this is not possible,[19] utility must be maximized where the self-selection constraint is binding.

We will now show that any allocations involving a pooling contract for the low-risk individual with $\alpha \geq \overline{\alpha}$ supplemented by insurance at price q_H bringing the high-risk individual to full insurance is a CPE allocation, and that these are the only CPE allocations.

That such an allocation is CPE is easy to establish. Denote the allocation to the high-risk individual by A_H and that for the low-risk by A_L. We know that a CPE allocation for the high-risk individual must lie along the full insurance line. We also know that a CPE *just* breaking even and giving the high-risk individual the specified utility must lie along a line with slope q_H through A_H. But A_L cannot lie below the pooling line, for any allocation with A_L below the pooling line, given A_H would make a loss. The low-risk individual's utility is maximized at A_L lying along the pooling line - subject to the zero-profit, self-selection, and presence of undisclosed secret insurance at q_H constraints. Thus, the proposed allocation is CPE.

This analysis also shows that there cannot exist any other CPE allocation. We already know that a CPE must have the high-risk individual along the full insurance line, that the self-selection constraint with secret insurance has to be binding, and that the CPE

[19] The high-risk individual would prefer the full insurance allocation for the low-risk individual rather than that intended for the high-risk individual, except trivially in the limiting case with full insurance along the pooling line, where the two are the same. (That point is CPE, that which maximizes the utility of the high-risk individual.)

allocations together must just break even. It follows from this that the CPE allocation to the low-risk individual must lie along the pooling line, and the CPE allocation to the high-risk must be that supplemented with insurance at price q_H.

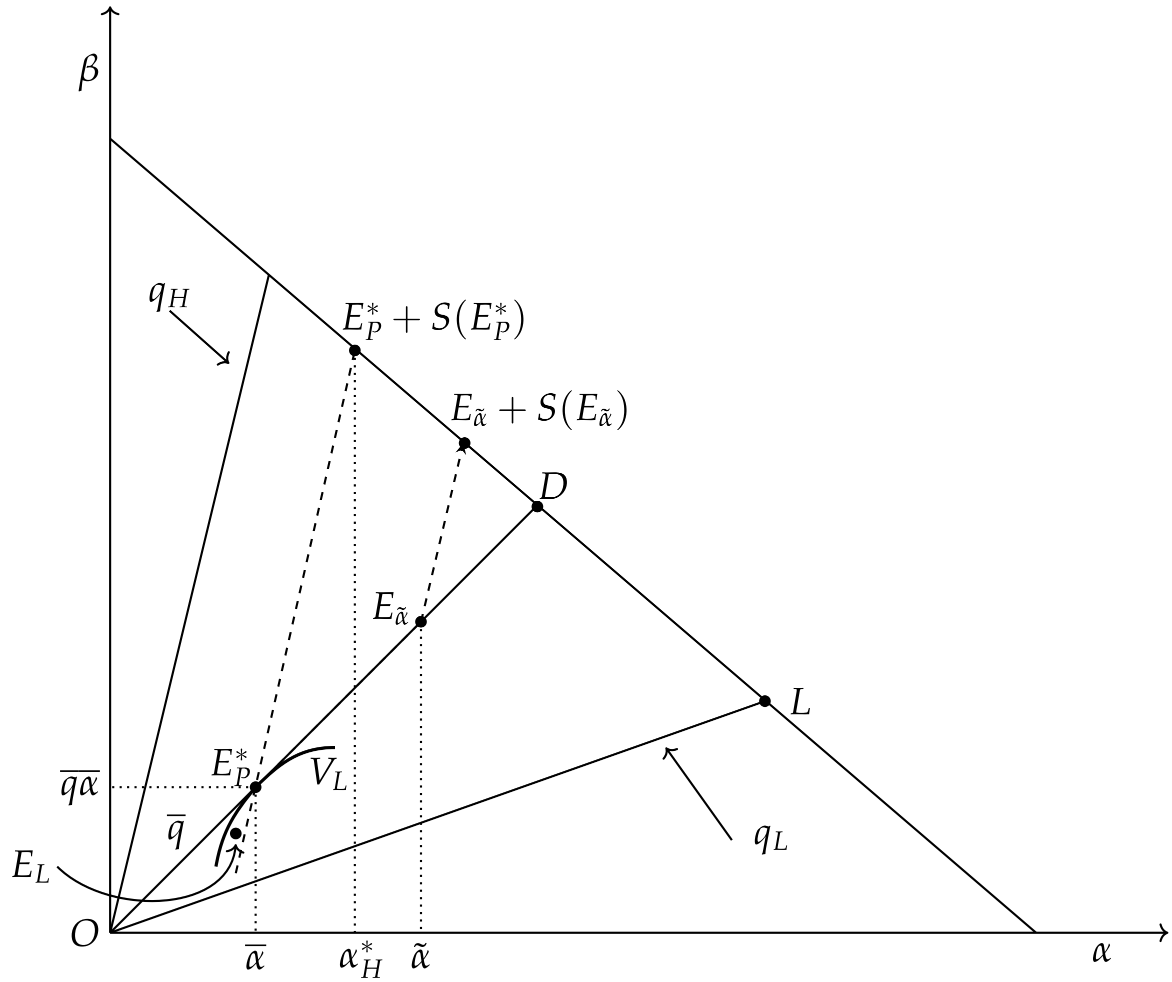

Figure 2: Constrained Pareto Efficient allocations. E_P^* is the allocation that maximizes V_L at price $\overline{q}$; $S(\cdot)$ are supplemental allocations obtainable with secret insurance. $E_{\tilde{\alpha}}$ is a typical CPE allocation for the low-risk type.
The figure also shows why in the presence of secret contracts, the pooling policy corresponding to E_P^* cannot be an equilibrium. Consider E_L, a contract that lies near E_P^* but below line OD, above the line through E_P^* with slope q_H (and so won't be purchased by the high-risk individual), above OL (and therefore makes a profit when purchased only by low-risk types) and below $V_L(E_P^*)$ (and therefore preferred to E_P^* by the low-risk types). Given our assumptions about preferences, there always exist such contracts.

We have thus established:

Proposition 2. *Suppose preferences are convex, and secret contracts cannot be precluded. There is a continuum of CPE allocations, parametrized by $\tilde{\alpha}$, characterized by:*

a) For the low-risk individual, $(\tilde{\alpha}, \overline{q}\tilde{\alpha}) \triangleq E_{\tilde{\alpha}}$, with $\tilde{\alpha} \geq \overline{\alpha}$, where

$$\overline{\alpha} \triangleq \max_{\tilde{\alpha}} \left[\arg\max_{\tilde{\alpha}} V_L(\tilde{\alpha}, \tilde{\alpha}\overline{q}) \right] \tag{5}$$

is the most preferred level of insurance at the pooling price of the low-risk type

b) *For the high-risk individual, $E_{\tilde{\alpha}}$ is supplemented by insurance $S(E_{\tilde{\alpha}}) = \frac{1}{1+q_H}[d + \tilde{\alpha}(q_H - \overline{q})]$ at the price q_H bringing the high-risk individual to full insurance.*

Finally, it is straightforward to show

Proposition 3. *With convex preferences, the CPE allocation that maximizes the utility of the low-risk individual is the unique allocation E^*.*[20]

This allocation will play a critical role in what follows.

5 Contracts and Equilibrium

Individuals may purchase one or more policies from one or more firms. An individual or his insurer may disclose to other firms all or some information about the set of policies purchased or sold, respectively. Information revealed must be truthful, but individuals or firms may choose not to reveal some or all information. What is critical about the information disclosure in the model is that individuals cannot reveal the fact that they have not purchased a particular policy.[21]

As we have noted, an insurance contract consists of two components: (i) a policy, defined by a benefit, α, a price, q, and a set of restrictions that have to be satisfied (as far as the insurer knows) if the policy is to go into effect; and (ii) an information disclosure rule. The set of conceivable contracts is quite rich; all that is required is that firms can only disclose a subset of what they know, and can impose restrictions that can only be implemented based on the knowledge of the insurer. But we show that there exists[22] an equilibrium with a simple set of contracts and disclosure rules.

[20] For the key results below on market equilibrium, we do not actually need the convexity assumption. All the results below hold if we define E_P^* as the (unique) contract along the pooling line that maximizes V_L associated with the highest level of insurance (and therefore the highest level of subsidy to the high-risk individuals). Any other allocation (generating the same level of utility for the low-risk individuals) cannot be CPE, because by switching to E_P^*, the low-risk individuals are unaffected, but the high-risk individuals would be better off with the higher subsidy.

[21] In terms of the literature on strategic communication, this is a setting of verifiable disclosure, or hard information (Milgrom (1981), Grossman (1981)). If individuals could disclose (be asked to show in a verifiable way) that they had *not* purchased any other policy, exclusivity would be enforceable, and we would be back in the RS world. In the case of our model, verification that the individual has the claimed insurance could be demonstrated by showing the insurance policy.

[22] We emphasize that the result that an equilibrium exists is not trivial: Recall our earlier result that if disclosure is not complete but symmetric, there never exists an equilibrium. At the same time, as we elaborate below, we do not provide a complete characterization of all possible equilibria, though we are able to establish a strong result about what the unique equilibrium allocation must look like.

5.1 Setting

We employ a two-stage framework consistent with the conventional setting of a screening model.

- First stage: Each firm offers a set of insurance contracts. Any contract, as we have noted, has two parts. The first is a policy, specifying the quantity of insurance (the benefit) with an associated price (that is, the premium) and a set of restrictions on what insurance individuals can buy from other insurers. The second part is an information disclosure rule, specifying what information the firm will disclose to whom. The implementation of any restrictions depends on the information available to the insurer, which depends on the disclosures of firms and individuals.
- Second stage: Consumers purchase policies and disclose information about them (possibly selectively) to their potential insurers and others, after which each firm executes its contract for its consumer as announced in the first stage by disclosing information as specified by its disclosure rule. Consumers whose insurance purchases are found to be inconsistent with the policy restrictions have their policies cancelled.

Firms disclose their information simultaneously in the second stage, implying that the disclosure rule of a firm may be made conditional only upon consumer-revealed information, in particular, in the equilibrium that we establish, on information about the firm(s) from which individuals have bought insurance.[23] As a policy offer is subject to cancellation once a firm receives information from other firms and from consumers, the enforcement of the restrictions imposed within the policy offers can rely on information disclosed by consumers and firms. After the second stage, there is no further disclosure of information between firms, or further purchases of insurance by consumers.

Consumer disclosure is absolutely essential to our analysis, because it enables the asymmetric disclosure of information that is critical to the existence of equilibrium, as we have noted. Another crucial aspect of this setting is that a contract offered by a firm does not depend on contracts offered by other firms; it is non-reactive, keeping this paper in the spirit of competitive analysis, as in RS. As we have noted, there are other important strands of research in the theory of adverse selection focusing on imperfect competition and reactive equilibria.

5.2 A Simple Illustration of the Equilibrium Contract

Before conducting a formal analysis of an equilibrium we will describe how the set of equilibrium contracts we propose works in a simple context to highlight the core logic

[23]We explore the implications of alternative formulations, including sequential disclosure, in Stiglitz and Yun (2016) and Stiglitz et al. (2019). The present results appear to be quite robust.

of the main argument on the existence of an equilibrium. The equilibrium we propose involves two kinds of firms: a given number n of "established" firms selling insurance at the pooling odds, and $N - n$ "secret" firms, offering an unlimited amount of insurance at price q_H without disclosure. An established firm sells a consumer a single insurance contract at the price $\overline{q}$ with the following restriction on additional insurance purchases, and with the following disclosure rule:

- Restriction: the total amount of revealed purchases is not greater than $\overline{\alpha}$ (the amount most preferred by the low-risk consumers at price $\overline{q}$).
- Disclosure Rule: disclose its sale to all the other firms except those revealed by the consumer to be his insurers.

We denote the equilibrium contract offers for the secret and established firms by C_s^* and C_e^* respectively. For these contracts to sustain the equilibrium allocation, they should be able to do two things: 1) prevent over-purchases by high-risk individuals and 2) deter a cream-skimming deviant contract from breaking an equilibrium, i.e., undermining the pooling contract by offering a contract that would just be purchased by the low-risk individuals. The central result of this paper is to show that the allocation E^* is the only possible equilibrium allocation and can be sustained by the above set of contracts.

In section 6.1 we show that the equilibrium in this simple example with firms offering policies within a restricted set and with disclosure rules within a restricted set is in fact general; that is, if insurance firms are allowed to offer insurance policies and to engage in disclosure rules that were not so restricted, there is still an equilibrium of the form described for this restricted set.[24] Similarly, we show that the assumption dividing firms into established and secret firms is without loss of generality. Any firm could deviate in any way from its offers and disclosures, but in the equilibrium we depict, no firm would want to. So too, the assumption that there are a fixed number of established firms and a fixed number of secret firms can be dropped.

The basic logic of the equilibrium is simple. Assume, for ease of exposition, that individuals honestly reveal to all the firms from which they have purchased insurance all of their purchases from other firms if they do not lose anything by the revelations.[25] In particular, low-risk individuals only buy $\overline{\alpha}$, and have no reason not to disclose it. As we will prove more formally later in lemma 6.1, the above disclosure rule adopted by the established firms leads to at least one firm knowing all the purchases by an individual. This implies that in the equilibrium proposed above, high-risk individuals would not be able to overpurchase insurance from the established firms by withholding some information

[24] The same allocation may be sustained by different disclosure strategies, but all equilibria are allocation-equivalent, modulo information disclosure strategies.

[25] In the later analysis, we both show that this is the case and that our disclosure rule supports the equilibrium.

from them. Thus, this limited version of honesty directly prevents anyone (that is, the high-risk individuals) from overpurchasing the pooling contract.

More subtle is how the asymmetric disclosure rule prevents a deviant contract from breaking the pooling contract. Whenever a deviant firm, say A, offers a quantity of insurance equal to or less than $\bar{\alpha}$, charging a price lower than $\bar{q}$, the policy offered by A is always purchased, regardless of the restriction imposed by A, by both types of consumers, yielding losses for the firm A. This is because any consumer could always purchase the same amount $\bar{\alpha}$ in total from the deviant firm A and another established firm B, hence at an average price lower than $\bar{q}$. (The high-risk individual will always then want to supplement it with secret insurance.) The deviant firm know this, so it has to impose a restriction that the consumer not purchase any supplemental policy from an established firm (a firm selling a pooling contract). But the consumer will always disclose his purchase from A to B, knowing (under B's disclosure policy) that if he doesn't, B will disclose its sales to the consumer to all firms, including the deviant. That means that B does not disclose to A its sale (to that consumer) so that any restriction imposed by firm A can't be implemented. Thus, the asymmetric disclosure rule of the established firms can deter any cream-skimming deviant contract from upsetting an equilibrium while, together with the induced information revelation of consumers, preventing overpurchases of the pooling contract by high-risk individuals.

Asymmetric disclosure by firms based on the consumer disclosure is crucial to our analysis, and this asymmetric disclosure is only possible, as we have noted, because of consumer disclosure. Without consumer disclosure, there would be no basis for the asymmetry of the firm disclosure in a non-reactive framework where a disclosure rule of a firm does not depend upon the offers (policies or disclosure rules) of another firm. On the other hand, if the firm disclosure is symmetric and complete, we would obtain RS results, where an allocation such as we have described cannot be sustained in equilibrium.[26]

5.3 Contracts

Now we formalize these intuitions.[27] To simplify the notation and exposition, we begin by assuming all firms offer insurance with a single price,[28] while possibly imposing a constraint on aggregate purchases. Then in section 6.1 we show that, given the equilibrium contracts described, no firm would want to offer other contracts with any other set

[26] We have already established that if disclosure is not complete but symmetric, there exists no equilibrium.

[27] The formal notation developed in this section is employed only to a limited extent in the subsequent sections. Readers not interested in these formalities may, in a first reading, proceed directly to section 5.4.

[28] See appendix B and Stiglitz et al. (2017) for a generalization to the case with multiple prices and cross-subsidization. The equilibrium outcome is unchanged. Similarly, as we explain below, firms can offer (any set of) fixed quantity contracts, as in RS.

of prices, restrictions or disclosure rules.[29]

1. Policies: (α, β) is given by $(\alpha, \alpha q) \in \mathbb{R}^2_+$. A policy purchased by individual i from firm j is represented by $x^i_j \in \mathbb{R}^2_+$:

$$x^i_j \triangleq (\alpha_j, \alpha_j q_j) \tag{6}$$

while the set of policies purchased from all of the established insurers is denoted by $\hat{X}^i \triangleq \{x_j\}_{j \in K}$ where $K \subseteq \mathcal{M}^{\mathcal{E}}$ is the set of the established insurers from which an individual purchases insurance. Because in the remainder of the paper we focus just on individual i, without confusion, we drop the superscript i. The amount[30] α_j of insurance offered by a firm j may be required to satisfy a restriction, which can be in general be represented by a set of insurance amounts allowed, denoted $\psi_j(X^T_j)$ where $X^T_j \subseteq \hat{X}$, as defined by eq. 11 below, is the total information (about the individual's purchases) available to firm j. The individual (knowing what information the firm will have available to enforce whatever restrictions it imposes) will only choose to purchase an amount of insurance consistent with those restraints (for otherwise, the insurance will be cancelled):

$$\alpha_j \in \psi_j(X^T_j) \tag{7}$$

A policy offer by a firm j may thus be represented by $(q_j, \psi_j(X^T_j))$.

2. Disclosure Rules: an information disclosure rule by a firm j, denoted DIS_j, specifies a set $RE_j(\subseteq M)$ of firms receiving information from j about a particular individual, and information $INF_{jk}(\subseteq X_j)$ to be disclosed to a firm $k(\in RE_j)$, where X_j (defined by eq. 8 below) combines the information the firm has directly about j with the information disclosed by a consumer to firm j about his purchases (including the purchase from j). The information disclosed is obviously a subset of X_j.

   The information disclosed by an individual to his insurer j about purchases from others is denoted by $m_j(\subseteq \hat{X})$, indicating that an individual cannot disclose a policy that he does not purchase[31] although he may withhold from his insurer information about some policies purchased. Thus, the information set of firm j about an individual, X_j, before receiving information from other firms is

[29] Similarly, while we allow disclosure of any information available to firms, the equilibrium entails only disclosure of information revealed to it by its consumers plus what it knows from its own sales. Moreover, the only information disclosed is the quantity of insurance purchases. In more general models with sequential revelation of information (that is, firms can reveal information that they have from other firms to still other firms), disclosure rules can be more complex.

[30] In effect, the firms are offering contracts with linear prices; we show in section 6.1 and in appendix B that this involves no loss of generality.

[31] That is, the individual cannot lie about purchases he has not made in our model.

$$X_j = x_j \cup m_j \tag{8}$$

We suppose that whenever a policy x_j is disclosed, the identity j of the insurer is also disclosed. Thus, the set $I(\subset M)$ of firms (including firm j) disclosed as providing insurance by the consumer is given by:

$$I(X_j) \triangleq \{k \in \mathcal{M} | x(k) \in X_j\} \tag{9}$$

Now a disclosure rule DIS_j of firm j may be represented as follows:

$$DIS_j(X_j) = (RE_j(X_j), INF_{jk}(X_j)) \tag{10}$$

specifying what firms will be disclosed to, and, given that there is some disclosure to firm k, what information is disclosed. Given the disclosure rules $\{DIS_j\}_{j\in\mathcal{M}}$ of all the firms, the aggregate of them will determine the information disclosed to firm j by all the other firms, denoted by X_j^{-j}. Thus, all the information X_j^T available to a firm is that disclosed to firm j by the consumer, by other firms, and what it knows directly from its own sales:

$$X_j^T \triangleq X_j \cup X_j^{-j} \tag{11}$$

A contract C_j offered by a firm j is thus represented by a policy, characterized in turn by a price, a possible constraint on quantities purchased, and a disclosure rule:

$$C_j = \{q_j; \psi_j(X_j^T); DIS_j(X_j)\} \tag{12}$$

The contracts discussed in section 5.2, which we later denote by $\{C_j^*\}_{j\in\mathcal{M}}$, take simple forms. First, the policy offers are of the form:

a) The established firms' price is $\overline{q}$, the offer set is given by

$$\psi_j(X_j^T) = \{[0, \overline{\alpha}] | T(X^T) \leq \overline{\alpha}\} \tag{13}$$

with $T(X^T) = \sum_{x(k)\in X^T} \alpha_k$, the total amount of insurance known to have been purchased. A policy is cancelled if $T(X^T) > \overline{\alpha}$.

b) For the secret firms, $q_j^* = q_H$ with $\psi_j^*(X^T) = \mathbb{R}_+$ (i.e., offering unlimited insurance with no restrictions).

Similarly, the disclosure rule can be simply described: for the established firms,

(a) Disclose to all of the firms that have *not* been disclosed by the consumers as insurers, i.e.,

$$RE_j^*(X_j) = M \setminus I(X_j) \tag{14}$$

(b) All the quantity information that a firm j has about a consumer that it has obtained from the consumer plus its own sales:

$$INF_{jk}^*(X_j) \triangleq X_j \tag{15}$$

The secret firms disclose nothing:

$$RE_j^*(X_j) = INF_{jk}^*(X_j) = \emptyset \tag{16}$$

The contracts defined by eqs. 13-16 with price $\overline{q}$ we refer to as $\{C_j^*\}_{j\in\mathcal{M}}$. As we note in section 6, there are equilibria with the same allocation but with information disclosure rules that include price information.[32] In the next section we also discuss equilibria (again with the same allocation) with more parsimonious equilibrium information disclosure rules.

5.4 Consumer Response

We now analyze consumers' responses to the set of offers. An individual chooses the mix of available contracts and a disclosure policy to maximize his utility, aware of the restrictions and the disclosure rules that may affect the implementation of those restrictions. Formally, given a set $\{C_j\}_{j\in\mathcal{M}}$ of contracts offered by firms, the consumer optimally chooses a set K of established (and K' of secret) insurers from which to purchase insurance, the set $\hat{X}(=\{x_j\}_{j\in K\cup K'})$ of policies to be purchased from them, and disclosure strategies $\{m_j\}_{j\in\mathcal{M}}$ specifying which information about his purchases to disclose to whom. If indifferent across multiple contracts, the consumers randomly choose one.[33] Further, we assume that consumers tell the truth (disclose information) unless it is in their interest not to do so, which we refer to as the assumption of predilection for truth. It is important to emphasize that we do not assume that consumers are always truthful - we only assume that if they are indifferent between truth telling and anything else, they tell the truth. In other words, this is a tie-breaking rule, not an assumption requiring truth-telling.[34]

[32] In the proof of our main result, we only use considerably more coarse information: only the quantities sold: $INF_{jk}^*(X_j) \triangleq \begin{cases} \hat{\alpha}_l \text{ if } l \in I(X_j), \forall k \in \mathcal{M} \\ \emptyset \text{ otherwise} \end{cases}$ suffices for the proof of theorem 6.3. In appendix B we do rely on revelation of both prices and quantities, i.e., eq. 15.

[33] This tie-breaking specification is without loss of generality – the same equilibrium exists under other specifications.

[34] This tie-breaking specification will be used in the proof of lemma 6.1. If consumers pursue another action when indifferent, our equilibrium construction will have to change – in other words, this is a consequential

Consumer disclosure takes the form of messages sent from consumers to firms. Denote the message of consumer i to firm j by m^i_j, and from firm j to firm k about consumer i by $m^i_{j,k}$.[35] For each i, $m^i_{j,k} \subseteq \{\{x_j\}_{j\in\mathcal{M}}\}$. A message can be a statement about how much the individual has purchased from firm k, or the price he has paid for that insurance. In principle, as we have just noted, the message could convey less granular information but aggregate data, e.g., the total amount of insurance purchased, or the total amount purchased from a subset of firms. We require that disclosures have to be truthful, which can be verified by showing an insurance contract. Under these assumptions, the only information we use in our proofs is the set of firms from which i has purchased insurance (which doesn't have to be complete, but which we show is complete) and the aggregate insurance purchased from that set.[36] Thus,

$$m^i_j = \{R^i_j, \{\hat{\alpha}^i_j\}_{j\in R^i_j}\} \tag{17}$$

where R^i_j is the set of identities of firms for which $\alpha^i_j > 0$, and $\hat{\alpha}^i_j(R^i_j) \triangleq \sum_k \alpha^i_k$, for $k \in R^i_j$.

We can formalize the optimization problem for the consumers: each consumer i chooses a set $K \subseteq \mathcal{M}^{\mathcal{E}}$ (and $K' \subseteq \mathcal{M}^{\mathcal{S}}$) of established (and secret) firms, a set of policies $x(k)(= (\alpha_k, \beta_k))$ to purchase from them, and consumer disclosure rules $\{m^i_k\}_{k\in K\cup K'}$ to solve

$$\max_{\{\alpha_k,\beta_k\}_{k\in K\cup K'}} \sum_k V_t\left(\alpha_k, \beta_k\right) \tag{18}$$

$$s.t.\ \alpha_k \in \psi_k(X^T_k), \forall k \in K\cup K' \tag{19}$$

We say that a consumer's choice $\{\{x_j\}_{j\in M}, \{m_j\}_{j\in M}\}$ and disclosure rule is optimal if given $\{x_j\}_{j\in\mathcal{M}}, \{m_j\}_{j\in M}$, and $\{DIS_j(X_j)\}_{j\in\mathcal{M}}$ no policy is ever cancelled, and $\{x_j\}_{j\in\mathcal{M}}$ solves the above problem.

5.5 A Competitive Equilibrium Must Be the CPE Allocation That Maximizes the Utility of the Low-Risk Individual

An equilibrium is defined as follows[37]:

Definition 2. *An equilibrium is a set $\{C_j\}_{j\in\mathcal{M}}$ of contracts offered by firms such that no contract*

tie-breaking assumption.

[35] For clarity and ease of exposition, in this section, focusing on the individual's response, we retain the use of the superscript i to denote individual i.

[36] Because there are no additional costs (in our simplified model) from sending richer messages, nothing changes if the consumer sends all of the information about each of his purchases. Some of this information is, in fact, employed in some alternative equilibria, employing different disclosures; and is used in particular in the equilibrium analyzed in appendix B.

[37] See appendix B and Stiglitz et al. (2017) for a generalization to the case with cross-subsidization. The equilibrium outcome is unchanged. Our definition is the natural generalization of the RS definition of equilibrium to the case with endogenous information.

results in a negative expected profit, and given the contracts offered by other firms $\{C_{-j}\}_{j\in\mathcal{M}}$*, there does not exist any other contract that firm j, for any j, can offer which makes positive profits given consumers' optimal responses to firms' announced contracts.*

We now consider a more general case than the one we have focused on in this section: We consider cases in which firms can offer more than one policy, and can restrict the amount of insurance they offer, i.e., they do not (necessarily) just offer insurance at a given price up to a certain amount. We say that a set of contracts $\{C_j\}_{j\in\mathcal{M}}$ *sustains an allocation E* if when each type chooses the subset of policies that maximizes its utility, supplementing such policies possibly by secret insurance at price q_H, the allocation E is generated, and there is no other contract C' that any firm can offer that would make strictly positive profits. We ask, are there sets of policies with their associated allocations that cannot be part of any equilibrium, whatever the disclosure policies? The answer is that we can rule out any set of policies giving rise to any allocations other than the CPE allocation that maximizes the welfare of the low-risk individual, the allocation we have denoted by E^*.

If each established firm offered only one policy, and they all offered the same one, then each firm would have to break even on that policy. If all firms offered the same policy, it would have to be along the pooling line - with the high-risk individuals supplementing the policy to bring themselves to full insurance (and possibly the low-risk individuals doing so as well). But the only possible such equilibrium allocation is E^*,[38] for if the putative equilibrium set of contracts generated for the low-risk individuals a utility just epsilon worse than $V_L(E_P^*)$, a deviant firm could offer E_P^*, and would clearly attract all the low-risk individuals and would make a profit: it would either attract just the low-risk individuals, in which case it would make a large profit, or both types, and then make just a small profit. Note that this would be true for any disclosure rule, so long as there is some firm offering secret insurance at q_H, which we have argued will always be the case because the effect of a disclosure rule is to provide the deviant firm with more information to assess applications. This either enables the deviant firm to discriminate - so the firm only accepts low-risk individuals, or it doesn't, in which case the firm may get both types. But in either case, the deviant contract is profitable.

If each firm offered only one policy, but different firms offered different policies, the only policy (or sets of policies) that get purchased in equilibrium are those that maximize the utility of the low-risk individual. It should be obvious that the resulting allocation for low-risk individuals cannot yield a utility greater than that of E^*, for if it did, it would also be purchased by high-risk individuals,[39] supplemented by secret insurance; but then

[38] Recall that $E^* = (E_P^*, E_P^* + S^*)$, i.e., it is the allocation that generates for the low-risk individual the highest utility along the pooling line.

[39] The only situation where this is not the case is one where the set of contracts purchased by the high-risk individual gives them an even greater utility, but that would imply that there is an allocation which dominates E^*.

both high and low-risk individuals are better off than in the CPE - which is impossible. If the low-risk individual's utility is lower that at E^*, a firm could offer a policy just slightly worse than E_P^*, make a profit, and attract all the low-risk individuals, and possibly all individuals. By the same reasoning as above, the deviant firm would make a profit regardless of the disclosure rules.

Now consider the case where each firm offered a pair of policies, and without loss of generality we can assume that they all offer the same pair, with the profits of one compensating for the losses on the other. The same reasoning as above shows that, unless the pair of contracts corresponds to the allocation E^*, the putative equilibrium will be broken by a firm offering a single policy "near" E_P^*, with high-risk individuals buying supplemental insurance at q_H.

It is thus apparent that any set of policies that does not correspond to or support E^* cannot be an equilibrium, whatever the disclosure rules for established firms.[40] Thus, we have the following proposition which states that a necessary condition for an equilibrium allocation is that it be CPE:

Proposition 4. *The only allocation that is consistent with a competitive equilibrium is the (unique) CPE allocation that maximizes the utility of the low-risk individual. This holds for any information disclosure strategy of firms and consumers.*

Attar, Mariotti, Salanié (2020b) and Jaynes (1978, 2011) have argued that (in somewhat different contexts) the unique equilibrium allocation with secret contracts must be the allocation that we have labelled E^*. The reason that we should not be surprised is simple: E^* is the CPE, and any allocation that is not E^* is not entry proof in a Nash or putative competitive equilibrium, for the reasons that our analysis makes clear. The only condition under which another allocation might not be upset by E^* is if incumbent firms react to the offer of E^* in ways that make E^* unprofitable, as might conceivably be the case in some reactive equilibria. Of course, if the reactions are not *too* strong,[41] then any allocation that is not the CPE that maximizes the well-being of the low-risk individual would be overturned. Here, however, we focus on competitive equilibria, where it seems inappropriate to assume that existing firms would respond to the entry of a small firm offering an alternative policy or set of policies.

[40] This result holds so long as there is secret insurance at price q_H. But our earlier analysis showed that as long as such insurance cannot be proscribed, e.g., by the government, it will be on offer; any putative equilibrium in which it was not on offer would be broken by one in which such insurance is on offer.

[41] With very strong reactions (where existing firms may pay a price for deterring entry by strong punishments) other equilibria might be sustained. This paper focuses on competitive (non-reactive) equilibria, so we do not pursue the question of whether such equilibria are consistent with standard game-theoretic solution concepts.

6 Proof of Equilibrium Existence

In this section we show that the contracts described in section 5 support the allocation E^* as an equilibrium. There may, of course, exist other equilibrium contracts that differ in the information disclosure elements, or in the number and identity of established and secret firms, but nonetheless result in the same allocation. Our objective is simply to demonstrate the existence of an equilibrium with endogenous information that is implemented using simple, interesting, and illuminating contracts.

In showing that the equilibrium set of contracts $\{C_j^*\}_{j\in\mathcal{M}}$ implements E^*, we first prove the following lemma:

Lemma 6.1. *Given the set of contracts $\{C_j^*\}_{j\in\mathcal{M}}$, no individual purchases more than $\overline{\alpha}$ from the established firms and all individuals reveal their purchases to all established firms from which they buy insurance.*

The intuition behind the lemma (the proof of which is in appendix A) can be illustrated by the following example: assume that the high-risk type tried to over-purchase, for example, by purchasing $\frac{1}{2}\overline{\alpha}$ from 3 different firms, firms A, B, and C; but disclosed only one of his other purchases to each, say only his purchase from B to A, say, but not that from C to A, and symmetrically for the other firms. Then firm A discloses its sales (of $\frac{1}{2}\overline{\alpha}$) and firm B's sale (of $\frac{1}{2}\overline{\alpha}$) to C. But then C knows that the individual's total purchases are $\frac{3}{2}\overline{\alpha}$ and his insurance is cancelled.

Lemma 6.1 generalizes this simple example to *any* pattern of purchases and *any* pattern of disclosures by which an individual might try to hide the fact that his aggregate purchases exceed $\overline{\alpha}$.

The lemma implies that all individuals purchase just $\overline{\alpha}$. Assume an individual purchased more than $\overline{\alpha}$ in the aggregate from the established firms. Given $\{C_j^*\}_{j\in\mathcal{M}}$ he cannot disclose that he has purchased more than $\overline{\alpha}$ (to any of his insurers) because were he to do so the policy would be cancelled. So there must not be full disclosure. If the consumer does not disclose one of his insurers, say purchases from firm j, then all the other insurers disclose to the firm j what they know about the consumer's purchases (i.e., their sales to the consumer, and what the consumer reveals to them), and then the firm knows that the individuals has purchased more than $\overline{\alpha}$, so j cancels its policy. But the individual would have known that, and so would not have purchased a policy from j.

There is one important corollary of lemma 6.1: all individuals reveal their purchases from established firms to all established firms, since they have no reason not to (using the assumption of predilection for truth).

We now prove the central lemma of the paper:

Lemma 6.2. *Under the restricted set of policy offers defined by eqs. 6 and 7 (a single price offer with possible restrictions on aggregate purchases, with a single contract sold to any individual),* *contract offers defined by eqs. 13-16, defined earlier as $\{C_j^*\}_{j\in\mathcal{M}}$, with consumer purchases and disclosures being the best response to the set of contract offers (defined by the solution to eqs. 18 and 19), the equilibrium entails the allocation E^*. This allocation is supported by a policy offer by established firms at price $\overline{q}$, with a restriction on aggregate purchases at $\overline{q}$ of $\overline{\alpha}$, and a disclosure by all established firms of all the information about i's insurance purchases by each firm to all the other firms that that the firm knows not to have been sellers of insurance to i; and the truthful revelation by all individuals of all of their insurance purchases from established firms.*

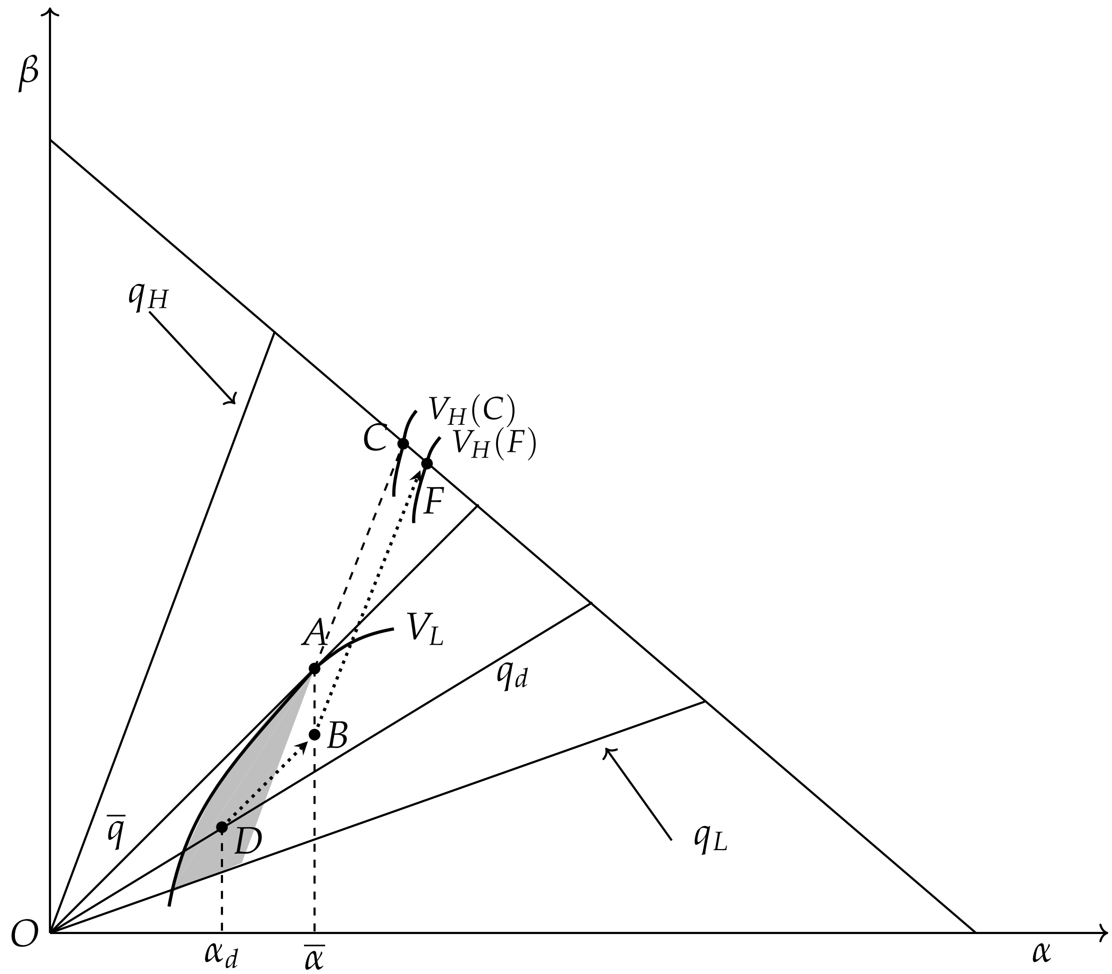

Figure 3: Sustaining an equilibrium against a deviant contract (offering D). High-risk individuals prefer A, supplemented by secret insurance, bringing them to C, to deviant offer D, seeming to break the (partially) pooling equilibrium A. But high-risk individuals would always supplement D with pooling insurance (DB), up to $\overline{\alpha}$ (disclosed only to non-deviant established firms) and secret insurance (BF) bringing the individual to full insurance. Because $V_H(F) > V_H(C)$, the high-risk individual will always purchase D, and the deviant contract, lying below the pooling line, would accordingly make losses.

Proof. It is obvious that by lemma 6.1, the set of contracts $\{C_j^*\}_{j\in\mathcal{M}}$ generates the equilibrium allocation E^*. Because of lemma 6.1 and its corollary, every established firm has the information required to effectively implement the allocation. There is no overinsurance by high-risk individuals. They just purchase $\overline{\alpha}$ from the established firms and supplement it with undisclosed insurance at price q_H, bringing them to full insurance.

We now show the set of contracts $\{C_j^*\}_{j\in\mathcal{M}}$ sustains E^* against any deviant contract, offered by a new entrant or by one of the existing established or secret firms. Note first that a deviant firm, indexed by d, cannot make profits by attracting only high-risk individuals in the presence of firms offering secretly any amount of insurance at q_H. If the deviant firm d attracts both high and low-risk individuals, its policy would have to charge a price q_d equal to or lower than $\overline{q}$, yielding zero profit at best. A deviant firm d can thus make positive profits only by attracting just low-risk types, i.e., only by a cream-skimming contract C_d. We will now show that in the presence of undisclosed insurance at price q_H there is no such cream-skimming contract, that is, the contract C_d always attracts high-risk individuals. The only possible such contracts must be below $V_L(E_P^*)$ and above the line through E_P^* with slope q_H (i.e., in the shaded area in figure 3, or a point like E_L in figure 2). It follows from the fact that $\overline{q} < q_H$ that this area is not empty. We now show that any contract such as D in figure 3 will in fact, under our information assumptions, be purchased by the high-risk individuals.

To attract low-risk individuals, we must have $q_d < \overline{q}$. It is obvious from figure 3 that the high-risk individual, if he could, would purchase the contract OD plus additional pooling insurance DB up to $\overline{\alpha}$ plus supplemental undisclosed insurance (BF in figure 3) at q_H, rendering the deviant contract unprofitable. The deviant firm knows this, and hence must put a restriction on the amount of supplemental pooling insurance that the individual can purchase. But given disclosure rules in $\{C_j^*\}_{j\in\mathcal{M}}$ (eqs. 14-16), no such restriction can be enforced. The high-risk individual obviously will not disclose directly that he has made the supplemental pooling purchases. If the high-risk individual discloses his purchase of the deviant contract to the established firms and limits his total purchases (combining what he has purchased from the deviant firm and amounts purchased from other established firms) to $\overline{\alpha}$, no established firm will cancel insurance that it has sold, and, by its disclosure rule, no established firm will reveal to the deviant firm its sales to the individual. Thus, firm d cannot enforce any restriction on total purchases from itself, plus purchases from the established firms, being less than or equal to $\overline{\alpha}$. Accordingly, high-risk individuals will purchase the deviant contract, and it loses money because the deviant firm could not enforce any restriction on such purchases. Similarly, the deviant firm cannot make its disclosure rule effectively depend on such supplemental insurance, and so can't use its own disclosure rule to deter purchases. That is, a deviant contract would not be able to upset the proposed equilibrium regardless of *any* restriction or any disclosure rule it may take.[42] □

[42]Note that all individuals disclose their purchases of the deviant policy to the established firms (because

6.1 Generalizations

We began the analysis by assuming all firms offer insurance with a single price while possibly imposing a constraint on aggregate purchases. But the proof showed that, given the equilibrium contracts described, no firm would want to offer any other contract(s), so that the equilibrium established in lemma 6.2 holds when firms are not so constrained.[43] This includes an established firm offering only fixed quantity contracts or offering both price and quantity contracts. It could offer a contract at a price higher than $\overline{q}$ but lower than q_H, but such a contract would only be purchased by the high-risk individuals, and so lose money. It could offer a contract at a price lower than $\overline{q}$, and such a contract will attract both all high and low-risk individuals, and thus lose money. (We've already discussed the impossibility of cream skimming in the proof of lemma 6.2.)

By the same token, we began our analysis assuming each firm sold only a single insurance contract to an individual.[44] This ensures that if the individual reveals that he has purchased insurance from firm k he has to reveal all the insurance he has purchased. If the firm sold the individual multiple policies, he could reveal that he purchased some insurance - thus precluding disclosure to that firm - without revealing the full amount of insurance. But it should be obvious that in equilibrium, no established firm sells more than one contract to an individual. It would be only the high-risk individuals that would be interested in purchasing multiple contracts from a firm, because by doing so they might underreport their purchases from one insurer to another established firm (disclosing one policy but not another) to be able to purchase more than $\overline{\alpha}$ at $\overline{q}$. Knowing that the only high-risk individuals would wish to buy multiple policies, an established firm would not sell multiple contracts to an individual without charging a price equal to or higher than q_H, which, however, would not be purchased by any individual.

We can similarly drop the assumption of a fixed number $|\mathcal{M}^{\mathcal{E}}|$ of established firms

otherwise, the established firms disclose their sales to the deviant firm, which would then cancel the deviant policy). But this means that (normally) all individuals purchase the full amount allowed by the established firms, $\overline{\alpha}$, and the established firms then sell exactly the same amount of insurance to the low and high-risk individuals, still breaking even. This will be true if the income consumption curves are negatively sloped (i.e., as individuals get better off, at a fixed price of insurance, they wish to have more consumption in both states). This condition is, for instance, always satisfied in the standard model of separable utility with the utility of consumption in the accident/no accident states being the same. If this condition were not satisfied, then it is possible that the low-risk individual purchase less than the high-risk individual, and the established firms might make a loss. (This problem could be obviated by insurance firms cancelling all policies in which revealed aggregate purchases strictly exceed the observed minimum level of aggregate purchases - which would be that of the low-risk individuals. This would force the high-risk individuals not to purchase more than the low-risk individuals.)

[43]The proof also makes it clear that no firm would want to change its disclosure rule (or to change simultaneously its contract and its disclosure rule). It should be obvious that the results do not depend on convexity of preferences. All that is required is that the income consumption curve for insurance for the high-risk individual at q_H is the full insurance line, i.e., the high-risk individual's indifference curve is tangent to the price line along the full insurance line, and any other point of tangency generates a lower level of utility.

[44]We again require that the insurance firm can tell the identity of the purchaser of the insurance. Obviously, otherwise it could not tell whether it had sold more than one policy to an individual.

and $|\mathcal{M}^{\mathcal{S}}|$ secret firms. There is no incentive either to enter or to change from being an established firm to a secret firm. $|\mathcal{M}^{\mathcal{E}}|$ and $|\mathcal{M}^{\mathcal{S}}|$ are thus perfectly arbitrary (so long as they are large enough that the assumption that any individual firm's action has no effect on the behavior of others is plausible): for any $|\mathcal{M}^{\mathcal{E}}| \geq 2$ and $|\mathcal{M}^{\mathcal{S}}| \geq 1$ there is an equilibrium with the same allocations and the same contracts (policies and disclosure rules).

Corresponding to increased complexity in policy offers (were it desirable for firms to increase that complexity), there can also be increased complexity in disclosure rules. Neither individuals nor firms have any incentives to engage in such more complex rules. We established before (lemma 6.1) under the restricted set of policy offers, that every individual has an incentive to disclose all the firms from which it has purchased insurance and the amount of insurance he has bought. This was true even though it could have revealed a subset of that information, e.g. the sum of the insurance bought from a subset of firms, without revealing the quantities purchased from each. More generally, for an individual to reveal that he had purchased more than $\overline{\alpha}$ would (in equilibrium) reveal the individual to be high-risk, and so no individual would reveal such information. And by the same token, for an individual not to reveal an insurance firm from whom the individual has purchased insurance would, given the equilibrium disclosure policies of firms, lead to the individuals' insurance being cancelled (if he has purchased more than $\overline{\alpha}$), *after the firm revelations that are consequent to that*. These results hold regardless of the complexity of the set of policies on offer and purchased.

We also began our analysis assuming that that each firm sells only one policy. It is easy to extend this, allowing each established firm to sell multiple policies (though by the analysis above, it would sell only one policy to any individual).[45] Denote by $\{C_j^{E*}\}_{j\in\mathcal{M}}$ the set of contracts with the same policy offers and disclosure rules as $\{C_j^{*}\}_{j\in\mathcal{M}}$, but where established firms can offer multiple policies, firms can choose to be established or secret, and can choose arbitrary disclosure policies. We maintain the restriction that a deviant firm can only offer a single contract.

We have thus established the central theorem of this paper:

[45] Without cross subsidization, the analysis is unchanged. As we have already noted, any firm can offer any array of policies; in effect, it can offer every individual a menu of options. Of course, no high-risk individual would ever buy a second policy (at the high-risk odds) from any firm from which it bought a pooling contract, for as we have already noted that the firm would then know that the individual was a high-risk individual, and (in a slight extension of our analysis) it would accordingly cancel the insurance. (That is, an implied condition for the purchase of a pooling contract is that the individual not buy a policy at any price higher than the pooling price - that the insurance firm knows about.) Within the menu of policies that each firm can offer are price-quantity contracts (e.g. of the kind that RS analyzed). Our disclosure strategies both for firms and consumers remain unchanged. Consumers still disclose the amount of insurance that they have purchased and from whom they have purchased it; and firms disclose information about j's purchases to those firms that have not (been disclosed to have) sold insurance to j. The assumptions of our simplified model are chosen to highlight one of the key issues we focus on, the problem of high-risk individuals buying more than $\overline{\alpha}$ of insurance in the context where the individual can purchase small amounts from many providers, and there is not automatic disclosure of information about purchases.

Theorem 6.3. *With convex preferences, there is an equilibrium set of insurance contracts* $\{C_j^{E*}\}_{j\in\mathcal{M}}$*, entailing (a) the equilibrium allocation* E^* *and (b) a disclosure policy by each established firm of disclosing all the information at its disposal to all firms not disclosed by the individual to be a seller to him of insurance. In this equilibrium, all individuals truthfully reveal all the firms from whom they have purchased insurance and the amount of insurance purchased.* E^* *is supported by a contract* C_P^**, the pooling contract maximizing the utility of the low-risk individual at price* $\overline{q}$*, (or a set of contracts from the established firms all at the price* $\overline{q}$ *aggregating to* C_P^**) which sustains* E_P^**, supplemented by secret insurance in the amount* $S(E_P^*)$ *at price* q_H *bringing the high-risk individuals to full insurance.*

Convexity of preferences is a sufficient condition for establishing the existence of equilibrium. Equilibrium exists under the weaker condition that the income consumption curve for insurance for type t individual at price q_t be the full insurance line.[46]

There is one extension that is not so simple: What happens if the deviant firm offers a menu of policies, in particular one purchased by 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 the appendix, we show that, even when a deviant firm offers multiple contracts at different prices, there still exists an equilibrium, and it entails the equilibrium allocation E^* that we have identified, the CPE allocation that maximizes the welfare of the low-risk individual.

6.2 Uniqueness of Equilibrium

Theorem 6.3 shows that the (unique) equilibrium allocation can be implemented by a very simple set of equilibrium disclosure rules. The equilibrium allocation may also be implemented by other disclosure rules, in particular entailing different disclosures off the equilibrium path. But all such disclosure rules must preserve the careful balancing that is central in this paper: asymmetries of disclosure, to prevent, on the one hand, the purchase by high-risk individuals of too much insurance, and on the other, deviant firms from cream-skimming that would break the pooling equilibrium. As an example, if individuals simply disclosed the aggregate of their purchases of insurance from other firms, and the names of firms from which they have purchased and firms then disclosed the aggregate purchases of each individual to all firms who are not disclosed to be sellers - a far more limited set of disclosures than that embodied in lemma 6.2 - then the equilibrium

[46]We employ this assumption for the high-risk individual to ensure that at the pooling contract, the high-risk individual wishes to buy more insurance - to get full insurance. In that sense, it is like a weak, global single-crossing property, much less restrictive that the standard single-crossing property or even the assumption of convexity. The only property on the preferences of the low-risk individual that we employ is that his income consumption curve at price $\overline{q}$ entails strictly less than full insurance, i.e., the CPE allocation E^* for the low-risk individual entails less than full insurance.

allocation would also be implemented.[47]

To see the delicate balance in disclosures required to implement the allocation E^*, assume that the established firms disclose their consumers' purchases to all other firms (not just those that have not been disclosed to have sold insurance to the given individual), whereas secret firms and deviants do not disclose any information. Could the deviant offer a contract C_d to cream-skim the low-risk type? The deviant does not know whether an individual has purchased supplemental secret insurance, but because he knows all the information about purchases from established firms, he can enforce restrictions on purchases of pooling insurance. Such a deviant contract C_d restricting purchases from established firms can cream-skim if (*i*) the low-risk type prefers C_d to the pooling contract C_P^* in E^*; and (*ii*) the high-risk type prefers the pooling contract C_P^* supplemented with the secret contract, and (*iii*) the deviant makes a positive profit, which it will if cream-skimming works, and the price of the insurance is greater than the q_L. Because of the assumption of convexity, there always exists such a contract, which is why the disclosure policy just described cannot implement the equilibrium allocation.[48]

7 Extension to Cases with Many Types

The result on existence of equilibrium can be extended to the case with many types. An equilibrium strategy in a case with three types, (denoting by L the lowest risk type, by M the middle-risk type, and by H the highest risk type) for example, can be described in a similar way to the case with two types. As illustrated in figure 4, there is a pooling allocation with all three types, allocation A, the most preferred by the lowest risk type; and a partial pooling allocation B, supplementing A with additional insurance along the line AB entailing partial pooling, bringing together the two riskiest types, where B is the most preferred allocation by the middle-risk individual along the zero-profit line between the high-risk and middle-risk individuals through A; and finally, a secret allocation C, supplementing B with insurance along the line BC at the high-risk individual's odds, leading to full insurance for the highest risk type. A, B, and C represent the total individual allocations for the three types. In equilibrium, the lowest risk type consumers purchase A only, the M types purchase A and the supplemental policy along the line AB, and the H types purchase A and both supplemental policies along the lines AB and AC.

There are three types of firms, those (denoted by F_A) selling the full pooling contract C_A, those (denoted by F_B) selling the partial pooling contract C_B, and those (denoted by

[47] Similarly, the equilibrium allocation could be implemented by each established firm offering a fixed quantity policy $\overline{\alpha}$ at $\overline{q}$, with the same disclosure rule as in lemma 6.1.

[48] The argument is exactly the same as that given in section 3, showing that there are contracts below the low-risk individual's indifference curve, above the line with slope q_H through E_P^* (and so won't be chosen by high-risk individuals), but near E_P^*, so that it is above the zero-profit line for the low-risk individual, and thus makes strictly positive profits.

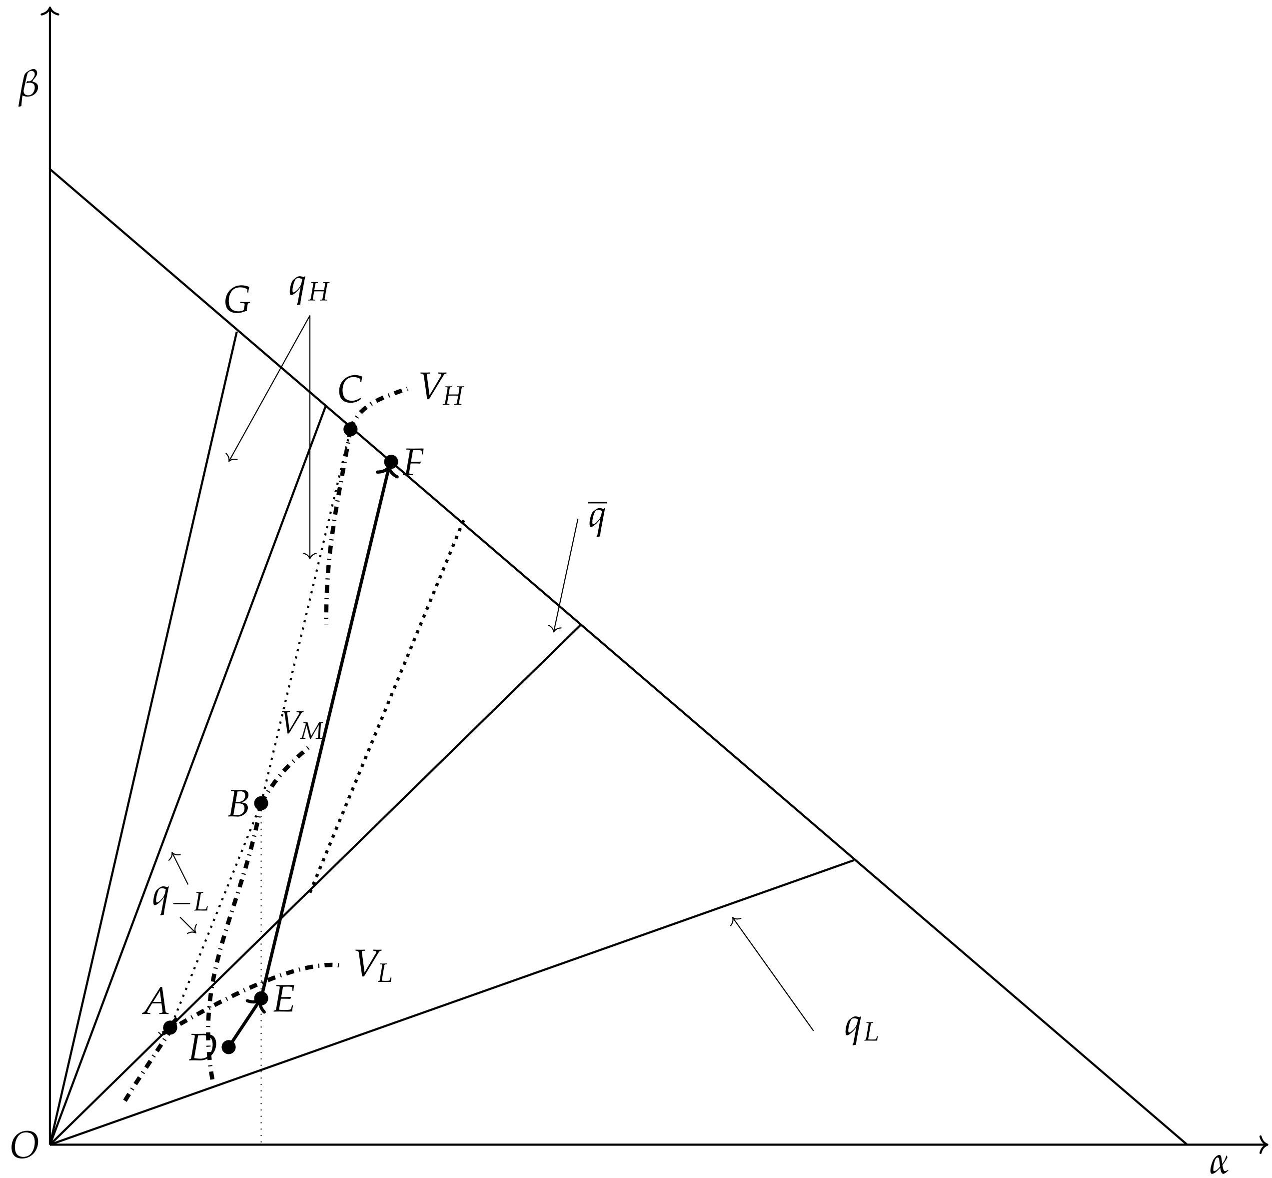

Figure 4: Equilibrium allocation (A, B, C) with three types, which cannot be broken by D as individuals of higher-risk type supplement it by additional pooling insurance (along the arrow EF) without being disclosed to the deviant firm. The slopes of the various lines are indicated by the q's. P_{-L} denotes the average probability of accident for the two highest risk types (and $q_{-L} = \frac{P_{-L}}{1-p_{-L}}$).

F_C) selling the secret insurance contract (C_C at price q_H) to the H types. Thus, *f*irms F_A are selling *c*ontracts C_A that lead to allocation A. They adopt an information disclosure rule analogous to that in the case of two types of individuals.[49] Consumers truthfully fully reveal to the other insurers their information about their purchases of the fully pooling contract C_A (since all purchase the same amount, such information in equilibrium reveals

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

no information about who they are).[50] Consumers reveal information about their purchases of the partial pooling contract C_B only to firms not selling them the fully pooling contract (for if they disclosed that, the firm would know that they were of one of the two riskier types). At the same time, they do reveal the firm from which they bought C_A (the say, F_B firms know that they are selling insurance to middle or high-risk individuals, so there is no news in the fact that the individual has also bought C_A. But that means that F_B does not reveal that he has sold C_B to L types. So F_A can't cancel its policy. Moreover, even if F_A reveals to F_B that the individual has bought the full pooling insurance, it reveals no information: F_B knows that all individuals do so.

By the same reasoning as in the two-type case, there is no possibility of M or H types buying more of the fully pooling contract or H types buying more of the partially pooling contract and no room for a cream-skimming deviant contract.

Figure 4 illustrates with a deviant contract offering D that attempts to profitably attract only low and medium types, but riskier types are also induced to choose D as they can purchase additional insurance along the lines DE and EF; and by the same argument used earlier, the partially pooling contract B can thus be sustained. This reasoning allows us to extend the analysis not only to any finite number of types, but to a continuum of types.

8 The Public Option

Our equilibrium analysis highlights the importance of asymmetric disclosure of information about insurance purchases. While there is some complexity in the formal description of the market equilibrium it is actually simple and intuitive. But that raises the question, is there some other way to robustly implement a CPE equilibrium? Sustaining the CPE equilibrium has distinct welfare benefits, with those who are fortunate in having a low probability of the occurrence of the insured-against event subsidizing those with a high probability. Could government regulation, such as a public register, help?

Our analysis has provided some caution to such approaches. Recall that the "trick" in sustaining the pooling contract is creating asymmetries in information about insurance purchases, as part of the market response to the natural asymmetries of information about risk types. If the government required disclosure of information, made that information public, and could enforce such a requirement for all firms, we would be back in the RS world.

More likely, while the government can force disclosure of "established" firms, it cannot

[50] Accordingly, no individual is worse off revealing his purchases of the full pooling contract A than not fully revealing his purchases. In fact, in the three-type case, an individual buying insurance from other than a fully pooling seller has an incentive to disclose his other purchases to the seller of the partial pooling policy, because otherwise that insurer discloses to his fully pooling insurer his sales, and then the seller of that policy would cancel the contract it sold to him.

force disclosure of informal insurance and insurance purchased from small firms. With disclosure of established firms accompanied by secret insurance, we are in the world of section 3, where no equilibrium exists, either in the quantity-constrained world or in the price-cum-quantity-constrained world (the price-quantity equilibrium).

An alternative is a well-designed public option, accompanied by appropriate disclosure policies. The government offers insurance at the market odds, requiring individuals to disclose whether they have purchased insurance at that price from anyone else, and restricting (as in our model) total purchases to be less than or equal to $\overline{\alpha}$. But the government discloses its information only to those firms that have not been disclosed to be sellers of insurance to the individual. Then our reasoning shows that the public option cannot be undercut. It is sustainable. Without such disclosure, there will be over insurance - the problem with the standard price equilibrium; with full disclosure, there is cream-skimming *à la* RS, and the only possible equilibrium is the separating equilibrium - and there may be no equilibrium. With full disclosure from established firms but still secret insurance from others, there is never an equilibrium - cream-skimming breaks the pooling equilibrium and the ability to purchase supplemental policies breaks the separating equilibrium. Our disclosure rule under the public option prevents both overinsurance from established firms and cream-skimming by deviant firms.

One of the intents of the public option is to encourage competition in lowering transactions costs - though there is little evidence that the private sector can come anywhere near the costs for the government. But we can allow for this. The government can still prevent cream-skimming. If the entrant is truly more efficient, it will still constrain its own sales to $\overline{\alpha}$ and displace government sales. If it is not more efficient, it will lose money because if it attempted to charge a price equal to or greater than its costs (greater than $\overline{q}$ but less than or equal to q_H) it would only attract high-risk individuals and so couldn't make a profit.

9 Previous Literature

In the more than four decades since RS appeared, its disquieting results have given rise to several large literatures.

Varying Equilibrium Concepts and Game Forms

The first strand looked for alternative equilibrium concepts or game forms under which equilibrium might always exist or under which a pooling equilibrium might exist. Rothschild and Stiglitz (1997) and Mimra and Wambach (2014) reviewed the literature as it existed to those points, with Rothschild and Stiglitz (1997) suggesting that proposed seeming resolutions of their non-existence result contravened plausible specifications of what a competitive market equilibrium should look like in the presence of information asymmetries. For instance, in the "reactive" equilibrium of Riley (1979) contracts are added in response to out-of-equilibrium offers, while in "anticipatory" equilibrium (Wilson (1977)), the entry of even a very small firm induces all firms to withdraw their pooling contracts making the deviant contract unprofitable. Such deterrence enables the pooling equilibrium to be sustained. Miyazaki (1977) (in the case of two types) and Spence (1978) (in the case of n types) extend this reactive equilibrium concept to allow for menus of contracts; the Miyazaki-Wilson-Spence (MWS) outcome entails separating, jointly zero-profit contracts with cross-subsidization.

Formal game theory literature since then has supported the MWS equilibrium under various conditions, with Mimra and Wambach (2017) endogenizing capital level choice before playing the RS game, Mimra and Wambach (2019) relying on latent contracts, and Netzer and Scheuer (2014) showing that the MWS outcome is a "robust" equilibrium when there are small costs associated with withdrawing from the market.[51] However, none of these extensions have fully resolved the inherently non-competitive nature of the proposed equilibria.

Consequences of Different Information Structures

One important strand of research is focused not on different equilibrium constructs but on the consequences of different information structures, allowing for nondisclosed contracts but *not endogenizing disclosure*. Most notable are the series of papers by Attar, Mariotti, Salanié (2011, 2014, 2016), employing a variety of assumptions about consumer preferences, firm behavior, and market structure. While a complete explication of the differences among these papers and between these papers and ours would take us beyond this paper, we note some key salient differences.

The 2014 model, employing strictly convex preferences, provides necessary conditions for the existence of an equilibrium in a much more general setting than discussed here. Applied to the insurance market, the equilibrium (when it exists) turns out to be the allocation where no one but the highest-risk individuals purchase insurance. The difference between their results and ours, where we have focused on *endogeneity* of information disclosure, are marked and obvious. In their (2016) model, they allow individuals to buy insurance from multiple insurers without disclosure.[52] Within their equilibrium construct (distinct from ours), they prove a parallel result, that any competitive equilibrium must entail the allocation E^*. Our analysis in section 3 has explained why this should not be a surprise: E^* is the CPE allocation that maximizes the welfare of the low-risk individuals.

[51] Another strand has focused on a mixed strategy equilibrium (Dasgupta and Maskin (1986)); Farinha Luz (2017) provides a full characterization of equilibria in this setting.

[52] See also Ales and Maziero (2014).

They are able to establish the existence of equilibrium, using latent contracts, but only under a very restrictive set of preferences, more restrictive even than the single-crossing property.

Attar, Mariotti, and Salanié (2020b) construct an equilibrium employing a sophisticated auction scheme under the assumption of a fixed information structure. Though the equilibrium is Nash, the structure of the model has an important reactive element: Firms can observe and punish deviators. The resulting allocation is again E^*, but that allocation is sustainable for very different reasons. Ours is a *competitive* theory, where the market does not respond to the entry of a new firm or a deviation in the contract offers of a single deviant firm. Theirs is in the long tradition of trying to build in the possibility of large responses to such occurrences into the response function (so, for example, the entry of a new firm does not alter the response function), and it takes those large responses which enable the equilibrium to be sustained, e.g. entry to be deterred.

Furthermore, the key issue posed in our analysis of nonexclusivity is that firms can offer contracts that are not observed by other firms, and that in the presence of such contracts overinsurance by the high-risk types is difficult to deter. Thus, in our model, markets are nonexclusive not only in the sense that buyers can trade with multiple sellers (which is the sense of "nonexclusivity" employed by Attar, Mariotti, and Salanié (2020b)) but also in a much stronger sense - sellers cannot *observe* purchases made by consumers from others.[53] Attar, Mariotti, and Salanié (2020b) prevent such overinsurance by a combination of elements – utilizing a sequential auction where offers cannot be withdrawn, using a rationing rule for sales, leveraging the fact that deviations by sellers are observable (and therefore can be punished) - which together obviate the issue of unobservable offers. Thus, the most important difference between our analysis and theirs is how equilibrium is sustained. In ours, endogenous and asymmetric information disclosure is crucial in establishing the competitive equilibrium; in theirs, it is the combination of elements just described, which result in their essentially non-competitive equilibrium. Given the markedly different settings of the analyses, it is not surprising that we require different assumptions to ensure the existence of equilibrium: Attar, Mariotti, and Salanié (2020b) employ strong assumptions on preferences (not just strict convexity and differentiability, but also a condition stronger than single-crossing). By contrast, as we have noted, we don't even need the single-crossing property, and while our analysis has employed convexity, in footnotes we have indicated how the results hold under weaker conditions.

We have focused on competitive equilibrium and how the CPE allocation we have identified can be decentralized with endogenous information disclosure; theirs entails a fixed information structure, where the CPE is decentralized through a particularly struc-

[53]This usage is also consistent with that of Arnott and Stiglitz (1990) who examined the welfare economics of nonexclusivity in much simpler settings.

tured auction.

There are therefore, several senses in which our contributions differ: 1) our notion of nonexclusivity is stronger, embracing not just the possibility of buying from multiple sellers but of each seller not knowing about the sales of others, 2) in our setting firms cannot observe the offers of other firms (and therefore, cannot directly detect cream-skimming deviations) whereas in their setting they can, 3) they decentralize the allocation using an auction scheme, whereas we use information disclosure in a way that seems perhaps more illuminating and more in the spirit of most competitive markets, like insurance, in which there are asymmetries of information - they are not, in fact, conducted as auctions, 4) the convexity of the price schedule is a consequence of the auction rules assumed, whereas we obtain it without requiring the schedule to have any properties in advance, and 5) we have a simultaneous, competitive setting with endogenous information revelation, whereas they have a sequential setting of almost perfect information, with a specific, fixed, information structure that generates far more complete information than that which arises in our model.

Thus, it turns out that both of our models decentralize the allocation E^*, but in markedly different ways under markedly different assumptions. The key differences, however, are the role of endogenous information disclosure in our model, and the reactive nature of their equilibrium. As they write (italics added): the allocation E^* "...emerges as the essentially unique outcome of competition *when each seller can quickly react to his competitors' offers.*" Which formulation is more natural depends, of course, on applications.

Endogenous Information

The closest works to our paper within the adverse selection literature are those papers that attempt, in one way or another, to endogenize the information available to sellers of insurance, including Pauly (1974) and especially Jaynes (1978, 2011) and Hellwig (1988), who analyze a model with a kind of strategic communication *among firms* about customers' contract information. Jaynes (1978) analyzes the same allocation E^* that we do. However, as Hellwig (1988) clarified, in Jaynes (1978)'s two-stage framework, the strategy of firms, including the associated strategic communication, is a reactive equilibrium, with firms responding to the presence of particular deviant contracts, and thus Jaynes (1978)'s formulation was subject to the same objections to reactive equilibria raised earlier.[54] Hellwig (1988) formulated a four-stage game, in which E^* emerges as the sequential equilibrium, but as he emphasizes, it has the unattractive property that firm behavior (in the final two stages) is conditioned on knowing the offers of all firms, including the deviant firm. Thus, in contrast to our model, a firm cannot offer a contract in secret. Moreover, as Hell-

[54] As we noted earlier, even in the title of Jaynes' (1978) paper makes clear that he was not trying to formulate a competitive equilibrium.

wig (1988) observes, "...it is not the endogenous treatment of interfirm communication that solves the existence problem of Rothschild, Stiglitz, and Wilson. Instead the existence problem is solved by the sequential specification of firm behavior which allows each firm to react to the other firms' contract offers."[55]

While our work differs from that of Jaynes (1978, 2011) and Hellwig (1988) in several ways, perhaps most important is that we consider information revelation by consumers as well as firms.[56] This allows the creation of asymmetries of information about insurance purchases between established firms and deviant firms, which, in turn, enables the pooling contract to be sustained *in the context of a competitive non-reactive equilibrium*. As we have noted, there is a delicate balance: On one hand, one must 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 we describe achieve this. In contrast, at least in a simple setting, models relying on just firm information strategies cannot do this, because they do not have any basis on which to engage in this necessary kind of selective disclosure. Moreover, it is natural to allow consumer revelation of information. Such revelation is a standard feature in markets with asymmetric information, and especially when dishonest disclosure can be punished, it occurs naturally in such markets, because less risky individuals attempt to distinguish themselves from the more risky. Furthermore, the distinction that is central to our paper, and much of the literature, between a requirement to "tell the truth and nothing but the truth," and the requirement to "tell the whole truth" is also a natural one. Individuals can be punished for lying (e.g. when the accident occurs, the individual cannot collect on the benefits if it is ascertained that he lied); but there may be implicit insurance, e.g. from one's family, which is not disclosed and the existence and non-disclosure of which is not punishable.

Welfare Economics of Adverse Selection

While our main focus has been on the analysis of equilibrium in markets with adverse selection, section 4 addressed issues of the welfare economics of adverse selection, introducing the concept of constrained Pareto optimality. Our work can thus be viewed as a

[55]More generally, changing the sequence of behavior, e.g. which side of the market moves first, can have a significant effect on the market equilibrium. While in the context of insurance markets, it is natural to have firms move first (making offers), in the context of other adverse selection models, e.g. the labor market where individuals have to choose a level of education and firms have to decide the wages to pay to those with different levels of education, there is more ambiguity. See Stiglitz and Weiss (2009).

[56]By the same token, while the equilibrium allocation we identify shares some features with that of the limit-order book studied by Glosten (1994), the context is different; most importantly, there is not the endogenous determination of information sharing between consumers and firms that is the central feature of our analysis. It is important not to confuse the similarity of the allocations in these different settings with the marked differences in the settings and the specification of the full equilibrium.

continuation of that of Bisin and Gottardi (2006) who prove versions of the first and second welfare theorems in a setting with adverse selection, but with exclusivity.[57] Our result in effect shows that a limited analogue of the second welfare theorem with nonexclusivity (in the way we use the term) also holds: The (constrained) Pareto efficient equilibrium that maximizes the utility of the low-risk individual can be implemented by a market mechanism; and we show explicitly how it can be done. Moreover, in parallel to the first welfare theorem, we have established that if there exists a market equilibrium it must be CPE.

Attar, Mariotti, Salanié (2020a) take a normative point of view, similar to our analysis in section 4. The constraint on which they focus is that the planner is unable to prevent consumers from *trading* with a third firm. This is related to, but not the same as, the constraint on which we focus. Our definition of constrained Pareto optimality here and in our earlier papers (2016, 2018) centers on the inability of government to force full disclosure, though it may (as in our equilibrium construct) induce some disclosure.[58] Thus, while the allocation they identify is the same the logic is different.[59] Still, it is not surprising that the results are parallel.

Concluding Remark

Rothschild and Stiglitz (1976) and Akerlof (1970) showed that market equilibrium with information asymmetries look markedly different from those without such asymmetries. The difference in the characterization of the equilibrium in the two papers highlighted the importance of information about purchases (quantities). Models that assume full disclosure to everyone, as RS and its decedents do, and models that assume no disclosure to anyone, as in the standard adverse selection price equilibrium models, do not create asymmetries of information about the amounts of insurance purchased. It is hardly a surprise that such asymmetries matter. What is perhaps a surprise is how much they matter: with appropriately structured asymmetries, which arise *endogenously* as part of a natural definition of equilibrium in competitive markets with endogenous disclosure,

[57] Earlier, Greenwald and Stiglitz (1986 and Arnott, Greenwald, and Stiglitz (1994) had provided a more general analysis of the first welfare theorem, showing that, in general, if and only if there is a single commodity does the first welfare theorem hold. Arnott and Stiglitz (1990, 1991) analyzed non-exclusivity in competitive insurance markets with a fixed information structure.

[58] One could, for instance, force disclosure but not restrict trade. It is, of course, difficult to restrict trade without information about what trades occur. Thus, there is some presumption that the constraint restricting trade is a more binding constraint. Moreover, as the discussion of section 8 makes clear, forcing disclosure will not result in the asymmetries of information about insurance purchases required to sustain in a competitive equilibrium the allocation E^*.

[59] Attar, Mariotti, and Salanié (2020a) in discussing an earlier version of our paper, observe: "...the logic of our approach is entirely different. First, these authors allow firms to react to the information disclosed by their competitors by possibly enforcing exclusivity clauses, which is at odds with the very notion of side trading that we emphasize. Second, [. . .] we are interested in the normative implications of side trading and not in characterizing the equilibrium of a given extensive-form game."

equilibrium always exists, even without the single-crossing property being satisfied, and entails partial pooling. Finally, it is extremely surprising (at least to us) that among all possible information revelation structures in our setting, there exists a very simple structure that, without restricting what firms and consumers share with whom, endogenously implements the equilibrium outcome.

Appendix A: Proof of lemma 6.1

Proof. Given the equilibrium contract, a consumer purchasing more than $\overline{\alpha}$ must not reveal his full purchases to any firm from whom he has purchased insurance. We first prove the following result: given the equilibrium disclosure rules, in spite of this non-disclosure by consumers, there is at least one firm that knows all the firms from whom the individual has purchased insurance. Assume a consumer purchases more than $\overline{\alpha}$ from K firms, and suppose the consumer makes any set of disclosures. Pick up first the firm that is the most informed (by the consumer) , say firm $j_1(< K)$, who knows about the consumer's purchases from firms $1, \ldots, j_1$ (including his own sales) and does not know about his purchases from firms $j_1 + 1, \ldots, K$, a group of firms undisclosed to j_1. (When there is a tie in which firm is the most informed, choose any of those; $j_1 = 1$ if a consumer does not disclose anything to any firm). Focus then upon the firms $(j_1 + 1, \ldots, K)$ undisclosed to j_1, and consider a firm who is the most informed of the purchases from those firms, say j_2, who knows about the purchases from $j_1 + 1, \ldots, j_2$. Similarly, we consider the most informed of the firms undisclosed to j_2 and j_1, say j_3. We can continue until we get j_k, where $k = K$. Then, clearly, the purchase from firm j_k is undisclosed to firms $j_1, j_2, \ldots, j_{k-1}$. Now consider the disclosures by firms. As a firm discloses to any other firm that is undisclosed by the consumer as his insurer, all the firms $j_1, j_2, \ldots, j_{k-1}$ (at least) will disclose to the firm j_k their own sales and information received from the consumer, implying that the firm j_k knows all of the K purchases.

The result of lemma 6.1 is now immediate: since that firm knows all of the individual's purchases, it knows that the individual has purchased more than $\overline{\alpha}$, and so cancels the policy. But the individual would not make those purchases, knowing that they would be cancelled. □

Appendix B: Multiple Contracts and Cross-subsidization

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. A deviant firm does so to induce self-selection among the applicants - with the self-selection process designed to reduce the costs of the high-risk individuals buying insurance from the deviant.

The definition of equilibrium is modified to read:

Definition 3. *An equilibrium is a set $\{C_j\}_{j \in \mathcal{M}}$ of contracts offered by firms such that no contract results in a negative expected profit, and given the contracts offered by other firms $\{C_{-j}\}_{j \in \mathcal{M}}$, there does not exist any other set of contracts that any firm j can offer that would make positive profits given consumers' optimal responses to firms' announced contracts.*

This definition highlights that the deviant firm may offer a pair of contracts, losing money on one but more than making up for it on the other.[60]

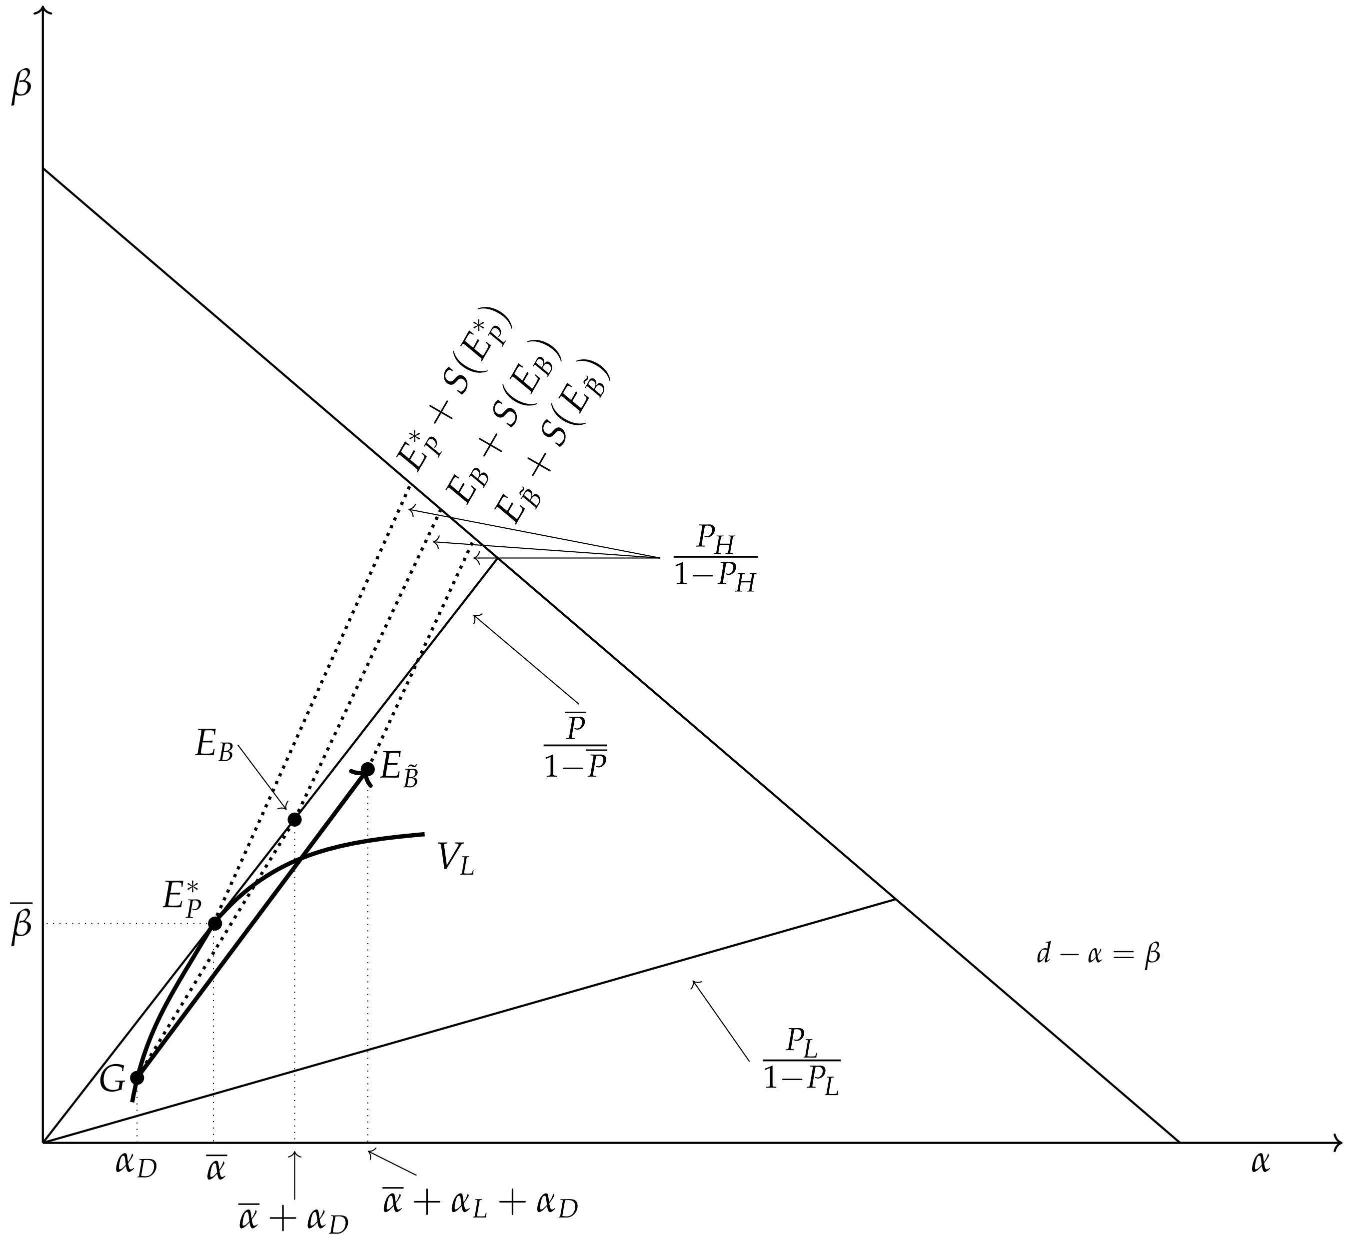

Figure 5: Equilibrium can be sustained against a deviant firm offering multiple contracts (E_B, G) or $(E_{\tilde{B}}, G)$ offered at different prices as high-risk individuals also choose G (over E_B), as $(E_{\tilde{B}}, G)$ yields losses for the deviant firm (while inducing self-selection). In breaking the original equilibrium the high-risk individual purchases $\overline{\alpha} + \alpha_D$. In keeping the deviant firm from breaking the equilibrium the individual purchases $\overline{\alpha} + \alpha_D + \alpha_L$.

[60]By the same token, we can extend the definition still further allowing in equilibrium cross subsidization. It is easy to show, however, that in our formulation, such cross subsidization cannot be sustained in equilibrium: a firm could always attract just the low-risk individuals by offering a policy that they would prefer slightly and that would not be preferred by high-risk individuals.

We first discuss why the analysis presented in the text no longer works. $(E_P^*, E_P^* + S(E_P^*))$ in figure 5 represent graphically the equilibrium allocation described earlier. Now consider the deviant pair of policies $(E_P^* E_B, G)$, where $E_P^* E_B$ is a supplemental policy offered at price $\overline{q}$ without disclosure (which, when combined with E_P^* leads to the allocation E_B) and G is offered at a price lower than $\frac{\overline{P}}{1-\overline{P}}$ but greater than $\frac{P_L}{1-P_L}$ without disclosure, and with G being offered conditional on no additional insurance being purchased. There always exists a set (E_B, G) such that G is preferred by all the low-risk individuals while E_B is preferred by all the high-risk, with the high-risk individuals supplementing E_P^* with $E_P^* E_B$, yielding total insurance in figure 5 of $\overline{\alpha} + \alpha_D$. Because the price of insurance 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 the policy G and the pooling supplemental policy which just separates (i.e., will be purchased by the high-risk individual), the deviant firm can make overall positive profits. In particular, this will be so if the supplemental policy is small. While there are large total losses associated with the purchase of pooling 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. The deviant firm gets all the low-risk individuals for all of their insurance, and the high-risk people only for the supplemental amount E_B.

To prevent this type of deviation, we need to make the choice of G more attractive to high-risk types, for instance, by providing more additional insurance at $\frac{\overline{P}}{1-\overline{P}}$ than the original equilibrium does, while limiting the total provision by all the firms to $\overline{\alpha}$ in equilibrium. One way of doing this is to have a latent contract,[61] which offers an individual sufficient amount of extra insurance at $\frac{\overline{P}}{1-\overline{P}}$ in the presence of a deviant contract G that the high-risk individual purchases G. The established firms announce that they will sell to anyone who purchases insurance at a price lower than $\frac{\overline{P}}{1-\overline{P}}$ additional insurance in fixed quantity α_L, which they will not disclose, and which the high-risk individuals then supplement with secret insurance.[62] An individual 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 that means that the established firms don't disclose

[61] The equilibrium allocation may be supported in other ways, but investigating that (both policy offers and disclosure rules) would take us beyond the scope of this paper.

[62] Of course, the latent contracts, being taken up only by high-risk individuals, would make a loss. We emphasize however, that the equilibrium concept upon which we focus (the competitive equilibrium defined above) does not require sub-game perfection, and does not require the firm to make non-negative profits along any out-of-equilibrium path. Moreover, it is easy to construct commitment devices, which entail payments from the established firms should it fail to offer the promised latent policy, which ensure that established firms will indeed offer it, should any deviant firm try to cream skim. Note that the disclosure strategies now also have to be somewhat more complex than before: consumers need to disclose not only the quantity of insurance that they have purchased from each firm, but also the price; and firms then disclose this information fully to those firms who have not been disclosed to be sellers of the insurance to the given individual.

their sales to the deviant, which ensures that the exclusivity provision associated with G cannot be enforced.

References

[1] Akerlof, George A., 1970. The market for 'lemons': quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), pp. 488-500. https://doi.org/10.2307/1879431

[2] Ales, Laurence and Pricila Maziero, 2014. Adverse selection and non-exclusive contracts. Working Paper, Carnegie Mellon University.

[3] 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.

[4] Arnott, Richard J. and Joseph E. Stiglitz, 1991, 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.

[5] Arnott, Richard J. and J. E. Stiglitz, 1990. The welfare economics of moral hazard. In: H. Louberge (ed.). Risk, Information and Insurance: Essays in the Memory of Karl H. Borch. Norwell: Kluwer Academic Publishers, 1990, pp. 91-122.

[6] Arnott, Richard J. and Joseph E. Stiglitz, 1991, 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.

[7] Arnott, R, B. Greenwald, and J. E. Stiglitz, 1994. Information and Economic Efficiency. *Information Economics and Policy*, 6(1), March 1994, pp. 77-88. https://doi.org/10.1016/0167-6245(94)90037-X

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

[9] Attar, Andrea, Thomas Mariotti, and François Salanié, 2011. Nonexclusive competition in the market for lemons. *Econometrica*, 79(6), 1869-1918. https://doi.org/10.3982/ECTA8665

[10] Attar, Andrea, Thomas Mariotti, and François Salanié, 2014. Nonexclusive Competition under Adverse Selection. *Theoretical Economics*, 9(1), pp. 1-40. https://doi.org/10.3982/TE1126

[11] Attar, Andrea, Thomas Mariotti, and François Salanié, 2016. Multiple Contracting in Insurance Markets. Working Paper, TSE-532.

[12] Andrea Attar, Thomas Mariotti, and François Salanié, 2019. The Social Costs of Side Trading. TSE Working Paper, n. 19-1017, June 2019, revised October 2019.

[13] Attar, Andrea, Thomas Mariotti, François Salanié, 2020a. The Social Costs of Side Trading. *The Economic Journal*, 130(630), pp.1608–1622. https://doi.org/10.1093/ej/ueaa041

[14] Attar, Andrea, Thomas Mariotti, François Salanié, 2020b. Entry-Proofness and Discriminatory Pricing under Adverse Selection. Forthcoming, *American Economic Review*. https://doi.org/10.1257/aer.20190189

[15] Bisin, A., and Gottardi, P., 2006. Efficient Competitive Equilibria with Adverse Selection. *Journal of Political Economy*, 114(3), 485-516. https://doi.org/10.1086/503754

[16] Chiappori, Pierre-Andre and Bernard Salanié, 1997. Empirical contract theory: the case of insurance data. *European Economic Review*, page 943-950. https://doi.org/10.1016/S0014-2921(97)00052-4

[17] Chiappori, P., and Salanié, B., 2000. Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy*, 108(1), 56–78. https://doi.org/10.1086/262111

[18] Dasgupta, Partha and Eric Maskin, 1986. The Existence of Equilibrium in Discontinuous Games: I. *Review of Economic Studies*, 53(1), pp. 1-26. https://doi.org/10.2307/2297588

[19] Farinha Luz, Vitor, 2017. Characterization and uniqueness of equilibrium in competitive insurance. *Theoretical Economics* 12, pp. 1349–1391. https://doi.org/10.3982/TE2166

[20] Glosten, L.R., 1994. Is the Electronic Open Limit Order Book Inevitable? *Journal of Finance*, 49(4), 1127–1161. https://doi.org/10.1111/j.1540-6261.1994.tb02450.x

[21] Greenwald, Bruce, 1979. *Adverse Selection in the Labor Market*, Routledge: New York.

[22] Greenwald, Bruce, 1986. Adverse Selection in the Labour Market. *The Review of Economic Studies*, Volume 53, Issue 3, p. 325. https://doi.org/10.2307/2297632

[23] Greenwald, Bruce, J. E. Stiglitz, and A. Weiss, 1984. Informational Imperfections in the Capital Markets and Macroeconomic Fluctuations. *American Economic Review*, 74(2), May 1984, pp. 194-199. https://doi.org/10.7916/D8KD27WR

[24] 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. https://doi.org/10.2307/1891114

[25] Hammond, Peter J., 1987. Markets as Constraints: Multilateral Incentive Compatibility in Continuum Economies. *The Review of Economic Studies*, 54(3), pp. 399–412. https://doi.org/10.2307/2297566

[26] 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. https://doi.org/10.1016/0014-2921(87)90046-8

[27] 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. https://doi.org/10.1016/0022-0531(88)90155-X

[28] Holmstrom, B., and Myerson, R., 1983. Efficient and Durable Decision Rules with Incomplete Information. *Econometrica*, 51(6), 1799-1819. https://doi.org/10.2307/1912117

[29] Jaynes, Gerald D., 1978. Equilibria in monopolistically competitive insurance markets. *Journal of Economic Theory*, 19(2), pp. 394-422. https://doi.org/10.1016/0022-0531(78)90102-3

[30] Jaynes, Gerald D., 2011. Equilibrium and Strategic Communication in the Adverse Selection Insurance Model. Working Paper, Yale University.

[31] Jaynes, Gerald D., 2018. Endogenous Beliefs and Institutional Structure in Competitive Equilibrium with Adverse Selection. Working Paper, Yale University.

[32] Grossman, S. J., 1981. The Informational Role of Warranties and Private Disclosure about Product Quality. *The Journal of Law and Economics*, 24(3):461–483. https://doi.org/10.1086/466995

[33] Milgrom, P. R., 1981. Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12(2):380–391. https://doi.org/10.2307/3003562

[34] Mimra, W., and Wambach, A., 2014. New Developments in the Theory of Adverse Selection in Competitive Insurance. *The Geneva Risk and Insurance Review*, 39(2), 136-152. https://doi.org/10.1057/grir.2014.11

[35] Mimra, W., Wambach, A., 2017. Endogenous Insolvency in the Rothschild–Stiglitz Model. *Journal Risk and Insurance*, 86: 165-181. https://doi.org/10.1111/jori.12206

[36] Mimra, W., Wambach, A., 2019. Contract withdrawals and equilibrium in competitive markets with adverse selection. *Economic Theory* 67, 875–907. https://doi.org/10.1007/s00199-018-1101-4

[37] Netzer, N. and Scheuer, F., 2014. A Game Theoretic Foundation of Competitive Equilibria with Adverse Selection. *International Economic Review*, 55: 399-422. https://doi.org/10.1111/iere.12054

[38] Miyazaki H., 1977. The rat race and internal labor markets. *Bell Journal of Economics*. 8:394–418. https://doi.org/10.2307/3003294

[39] Pauly, M., 1974. Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection. *Quarterly Journal of Economics*, 88(1), 44-62. https://doi.org/10.2307/1881793

[40] Prescott, E., and Townsend, R., 1984. Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard. *Econometrica*, 52(1), 21-45. https://doi.org/10.2307/1911459

[41] Riley, J., 1979. Informational Equilibrium. *Econometrica*, 47(2), 331-359. https://doi.org/10.2307/1914187

[42] 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. https://doi.org/10.2307/1885326

[43] 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. https://doi.org/10.1023/A:1008607915478

[44] Spence, M., 1978. Product Differentiation and Performance in Insurance Markets. *J. Public Econ*. 10(3), 427–447. https://doi.org/10.1016/0047-2727(78)90055-5

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

[46] 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.

[47] Stiglitz, Joseph E. and Andrew Weiss, 2009. Sorting Out the Differences Between Signaling and Screening Models. In Selected Works of Joseph E. Stiglitz, Volume I: Information and Economic Analysis, Oxford: Oxford University Press, 2009, pp. 223-231. Originally NBER Technical Working Paper 93, November 1990.

[48] Stiglitz, Joseph and J. Yun, 2016. Equilibrium in a Competitive Insurance Market with Non-exclusivity Under Adverse Selection. Mimeo.

[49] Stiglitz, Joseph, Jungyoll Yun, and Andrew Kosenko, 2017. Equilibrium in a Competitive Insurance Market under Adverse Selection with Endogenous Information. NBER Working Paper No. 23556, June.

[50] Stiglitz, Joseph, Jungyoll Yun, and Andrew Kosenko, 2018. Characterization, Existence, and Pareto Optimality in Insurance Markets with Asymmetric Information with Endogenous and Asymmetric Disclosures: Revisiting Rothschild-Stiglitz. NBER Working Paper No. 24711, June.

[51] Stiglitz, Joseph, Jungyoll Yun, and Andrew Kosenko, 2019. Characterization, Existence, and Pareto Optimality in Insurance Markets with Asymmetric Information with Endogenous and Asymmetric Disclosures: Basic Analytics of Revisiting Rothschild-Stiglitz. NBER Working Paper No. 26251, September.

[52] Wilson, C., 1977. A Model of Insurance Markets with Incomplete Information. *Journal of Economic Theory*, 16(2), pp. 167-207. https://doi.org/10.1016/0022-0531(77)90004-7