WORKING   PAPERS

An Evolutionary Analysis of Insurance Markets under Adverse Selection

Ana B. Ania, Thomas Tröger, Achim Wambach

November, 1998

Working Paper No: 9808

DEPARTMENT OF ECONOMICS

UNIVERSITY   OF   VIENNA

An Evolutionary Analysis of Insurance Markets under Adverse Selection[*]

Ana B. Ania[†], Thomas Tröger[‡] and Achim Wambach[§]

August 1998

Abstract

Since the seminal work by Rothschild and Stiglitz on competitive insurance markets under adverse selection the problem of non-existence of equilibrium has puzzled many economists. In this paper we approach this problem from an evolutionary point of view. In a dynamic model insurance companies remove loss-making contracts from the market and copy profit-making ones. Occasionally, they also experiment, adding new contracts or removing current ones arbitrarily. We show that the Rothschild-Stiglitz outcome arises in the long run if it constitutes an equilibrium in the static framework, but also if it is not an equilibrium, provided that firms only experiment with contracts in the vicinity of their current portfolio.

Journal of Economic Literature Classification Numbers: C70, C72, D82, G22, L1

Keywords: Insurance market, adverse selection, evolutionary game theory, imitation, mutation, withdrawal

0

1 Introduction

Ever since the seminal work by Rothschild and Stiglitz (1976) on competitive insurance markets under adverse selection the problem of non-existence of equilibrium has been one of the major puzzles in insurance economics. The origin of this problem lies in the fact that only separating contracts can constitute an equilibrium in the sense of Rothschild and Stiglitz, while in some cases a single, pooling contract may be preferred by everyone and will, therefore, upset the separating contracts.

There are two main approaches to this problem in the literature. One way out of it is to allow firms to play mixed strategies (Dasgupta and Maskin, 1986).

Another solution is to propose alternative equilibrium concepts. Wilson (1977) introduced the equilibrium requirement that, if a firm wants to offer an additional contract, this one should stay profitable even after loss-making contracts are withdrawn from the market. The result was that, in the case of non-existence of a separating equilibrium, a pooling equilibrium exists. Miyazaki (1977) and Spence (1978) allowed, in addition, for cross-subsidizing contracts. In that case, only separating equilibria exist, but possibly one where one type of contract makes losses and the other one profits.

However, in the Wilson approach, firms must anticipate the reactions of other firms before offering a new contract. This is acceptable if the other firms react sufficiently fast, e.g. in the same period. If there is a delay of at least one period in the reaction time, why should a new firm not enter and make a profit for this single period? The introduction of cross-subsidizing contracts in equilibrium is also subject to criticism. Why should a firm which offers a profit-making contract and a loss-making contract not withdraw the latter and go for the former only?

Riley (1979) introduced the notion of 'reactive' equilibrium. Here, it is the anticipation of entry that deters firms from offering a pooling contract, thus the separating contracts à la Rothschild-Stiglitz constitute the only equilibrium, regardless of the distribution of types.

There exist a few attempts of introducing some form of sequential structure in the model. In Grossman (1979) customers first send some form of signal before firms make

1

their contract offer, while in Hellwig (1987) firms can, after the informed agents have made their choice, decline to serve some contracts. Although in those cases equilibria exist, it is not clear whether these models apply to standard competitive markets.

There are two main features of all the models mentioned above that can be regarded as unsatisfactory. First, none of them provides 'an explicitly dynamic model, which describes how firms adjust their policies over time' (Wilson (1977), p. 205), although the equilibrium concepts proposed have usually some dynamic interpretation. Second, strong rationality assumptions are made in a context of incomplete information, where firms have limited knowledge (if at all) concerning the number of customers' types, their distribution and accident probabilities. In the present work, the non-existence problem is addressed from the perspective of evolutionary game theory. In the spirit of Vega-Redondo (1997), we propose a dynamic model, where the notion of full rationality is dismissed and replaced by simple rules of behavior based on imitation of success.[1]

Nöldeke and Samuelson (1997) already present an evolutionary model for the related case of signaling. In their work, firms use market experience to update their beliefs on the proportion of individuals with different productivity levels in the population. Then wages contingent on observed education levels are determined by the updated productivity. In the case of an insurance market, where insurance firms offer premium-indemnity contracts, this would imply that, if some individuals buy a contract with a given indemnity, firms observe the profit or loss they make with that contract and, for that indemnity, they charge the break-even premium in the next period. This implies Bertrand competition for each indemnity. The result of Nöldeke and Samuelson applied to the case of screening would read as follows: In the long run, the Rothschild-Stiglitz separating contracts are the unique outcome, if, and only if, they constitute a competitive equilibrium. If a competitive equilibrium fails to exist, then the Wilson pooling contract is observed in the long run, but this is not the unique long-run outcome.

The modeling in the present work differs from that of Nöldeke and Samuelson (1997) mainly in two aspects. First, we do not restrict ourselves to a schedule of contracts where the premium is a function of the indemnity; in particular, not every feasible

---

[1]Imitation of success as a behavior of economic institutions was already remarked by Alchian (1950).

indemnity must occur in a firm's menu. Second, we do not explicitly model updating beliefs, although we do let new information come into the system. What we propose is an alternative learning rule on the side of firms, which is also based on their market experience, and allows us to explain how they learn to offer separating contracts for the different types. In our model, in any period, firms offer a menu of contracts, i.e., a set of premium-indemnity pairs. Depending on their risk type, consumers choose the best contract available. After observing what has been offered and the profits obtained by all firms in the market, each firm revises its behavior by removing its loss-making contracts and copying any contract which was offered in the last period and made a profit. This specifies the main driving force of firms' behavior.

Apart from imitating other firms, insurance companies also experiment with their own contracts. They occasionally remove contracts from their own menu or add further contracts which are close to some of the previously offered, that is a new contract's premium and/or indemnity differs from that/those of an existing one by a small amount.[2] This scenario will, in line with the literature, be called local experimentation.[3] Lowering the price or raising the indemnity a little bit might be considered as an attempt by a firm with little information to increase the market share without incurring the risk of a high loss. Experimentation in general, however, may also be seen to implicitly capture the individual learning through the information acquisition or the formation of beliefs, to which the firm may also be subject. Rothschild and Stiglitz already suggest that a way out of the non-existence problem was the definition of a 'local equilibrium' where 'firms experiment with contracts similar to those already on the market' (Rothschild and Stiglitz, 1976, p.646).

With this dynamics of imitation, withdrawal and local experimentation, which is specified more precisely in the next section, we show that, independently of the proportion of high risks in the population, only the separating contracts with the lowest profits to the firm ("Rothschild-Stiglitz" contracts) are offered in the long run. This

---

[2]In this respect we also differ from Nöldeke and Samuelson (1997), who allow experimenting firms to offer any wage-education pair. However in section 3 we discuss the dynamics under global experimentation.

[3]This type of local experimentation has been first introduced by Alós-Ferrer *et al.* (1998).

result points out an instability problem that pooling and cross-subsidizing contracts have: The former are very easily overcome by small changes in the contract structure, while the latter are not robust because firms give up loss-making contracts. This result agrees with the Riley equilibrium but stands in contrast to other equilibrium concepts discussed above.[4]

If experimentation is not local, but global, then the Rothschild-Stiglitz contracts are still the long-run prediction, as long as they constitute a competitive equilibrium in the Rothschild-Stiglitz model. This result is interesting in itself, because it shows that the competitive equilibrium is learned even if the firms have no information about consumers' types and/or preferences. If, however, the competitive equilibrium does not exist, then, under global experimentation, we do not obtain a clear-cut market outcome. In fact, all absorbing states of the dynamics without mutation can be observed with positive probability in the long run.

Although we discuss a model of an insurance market in this paper, the results we obtain apply to a much wider class of competitive markets with adverse selection, for which the single crossing property holds.

The paper is structured as follows. Section 2 specifies the model. In section 3 we carry out the analysis. In section 4 we conclude.

2 The Model

2.1 Rothschild and Stiglitz' model of an insurance market

Consider a large population of individuals facing the risk of losing L. Each individual has an initial wealth W. There are two types in society: High risks (in proportion λ) who have a risk probability π_h and low risks (in proportion $1 - \lambda$) with a risk probability $\pi_l < \pi_h$. Each individual knows her own risk probability but this is not observable in general.

---

[4]Interestingly, it is also in agreement with a recent result by Kahn and Mookherjee (1995). These authors argue, in a cooperative context, that the only outcome which is coalition proof is the "Optimal Non-subsidized Allocation," which is the Rothschild-Stiglitz outcome for the case of two types.

A finite population of insurance companies $j = 1, ..., n$ $(n > 1)$ offer insurance contracts to the individuals described above. An insurance contract $c = (P, I)$ is characterized by the premium P, which has to be paid by the insured, and by the indemnity I, which the insurer pays in case the loss materializes.

Individuals have to decide whether to sign an insurance contract and which one. We will assume that they are risk averse. In particular, the utility that they obtain from money is given by a continuously differentiable and concave function $u : \mathbb{R} \to \mathbb{R}$, such that $u' > 0$. The utility of an individual of type $i = l, h$ after signing an insurance contract (P, I) is given by the expected utility function

$$U_i(P, I) = (1 - \pi_i)u(W - P) + \pi_i u(W - P - L + I).$$

Individuals optimally decide which insurance policy to purchase, if any at all, given the contracts posted by the insurers and their own particular type. If a contract is offered by several firms, individuals choose each with the same probability.

Insurance companies compete for customers. Each company j posts a finite menu of contracts S_j. By $s = (S_1, \ldots, S_n)$ we denote the collection of contracts offered.

Below we will model firms as boundedly rational. In particular, firms do not necessarily know anything about the number of consumers' types or their utility functions. Therefore it seems most natural that firms may offer any menu of contracts.

Insurers are risk neutral, so their expected profit per contract signed by a customer of expected accident probability π is given by

$$V((P, I), \pi) = (1 - \pi)P + \pi(P - I) = P - \pi I.$$

For our dynamics the only important thing will be whether there are profits or losses with any given contract. Therefore, whenever expected profits are positive (resp. negative), we take for granted that also realized profits are positive (resp. negative).[5]

---

[5]If any contract makes nonzero expected profits (losses), then by applying the weak law of large numbers, the probability that such a contract makes a loss (profit) becomes infinitesimally small, and therefore does not interfere with the learning dynamics. Technically, our results hold if, first, the number of individuals goes to infinity, and then the probability of experimentation approaches zero.

Call $\Pi(c, s)$ the profit that any firm j makes with contract $c \in S_j$, which depends on the whole collection of contracts offered s. $\Pi(c, s)$ already accounts for the individuals' reaction to s. The set of contracts which yield positive (resp. negative) profits is then denoted

$$B^+(s) = \{c \in \cup_{i=1}^n S_i | \Pi(c, s) > 0\}, \ B^-(s) = \{c \in \cup_{i=1}^n S_i | \Pi(c, s) < 0\}.$$

Notice that each of these sets may be empty.

We define $c_h^{RS} = (P_h^{RS}, I_h^{RS}) = (\pi_h L, L)$, and $c_l^{RS} = (P_l^{RS}, I_l^{RS})$, such that $U_h(c_h^{RS}) = U_h(c_l^{RS})$ and $V(c_l^{RS}, \pi_l) = 0$. These contracts are called the *Rothschild-Stiglitz contracts.* Call $\pi_{hl} = \lambda \pi_h + (1 - \lambda)\pi_l$. In the Rothschild-Stiglitz model, (c_l^{RS}, c_h^{RS}) is a competitive equilibrium if and only if

$$\mathcal{P} = \{c \in \mathbb{R}^2 | V(c, \pi_{hl}) > 0, \ U_l(c) > U_l(c_l^{RS})\} = \emptyset.$$

2.2 Dynamic behavior of the firms

Suppose the insurance market described in the previous subsection opens every period $t = 1, 2, \ldots$. We are interested in explaining how firms learn which contracts to offer following a dynamic process of imitation, withdrawal and experimentation. For technical reasons we assume the set of possible contracts to be finite.

Assumption 1 *Each contract (P, I) belongs to a finite bidimensional grid $\Gamma = \Gamma_1 \times \Gamma_2 = \{\delta, 2\delta, ..., \bar{P}\} \times \{\delta, 2\delta, ..., \bar{I}\}$, where $\bar{P} > P_h^{RS}$ and $\bar{I} > I_h^{RS}$.*

Note that each firm offers at most $\bar{P}\bar{I}/\delta^2$ contracts. The state of the system at the beginning of any period is denoted $s \in \Omega$, where $\Omega = \{S | S \subseteq \Gamma\}^n$, the *state space,* is the set of all possible collections of contracts.

Imitation and Withdrawal

Let the system be in state $s(t) = (S_1(t), \ldots, S_n(t))$ at the beginning of period t. After contracts have been purchased, damages have occurred, and profits have been realized, each firm revises its set of offered contracts. On making this decision, insurance

companies take into account the profitability of the contracts that were offered in period t. Each firm adds all profit-making contracts observed in the market at t, and removes all loss-making contracts from its menu of period t.[6] This results in the new menu

$$\left(S_j(t) \cup B^+(s(t))\right) \setminus B^-(s(t)).[7]$$

Experimentation/Mutation

After imitation and withdrawal have taken place, each firm experiments with a certain probability $\epsilon \in (0,1)$ independent across them and across time. In the sequel, and following the standard terminology, we will refer to this experimentation also as mutation. For each firm, experimentation consists of either randomly removing some contracts from its own, previously offered menu (these may well be idle contracts, but can also be profit-making ones), or randomly adding to it one or more contracts that lie in neighborhoods of its previously offered ones. This adding works precisely as follows. For any $r > 0$, denote the r-neighborhood of $c \in \mathbb{R}^2$ by

$$\mathcal{N}_r(c) = \{c' \in \Gamma \mid ||c - c'|| < r\},$$

where $|| \cdot ||$ is any norm on $\mathbb{R}^2$. If $S_j \neq \emptyset$, a mutating firm j may add to its menu any subset of

$$\bigcup_{c \in S_j} \mathcal{N}_r(c).$$

---

[6]Alternatively, one can assume that each firm has a probability $0 < \delta \leq 1$ of learning each period, which is independent across them. This would introduce some inertia in the process. Our results would also hold in that case.

[7]Alternatively, one could assume that firms cannot observe the profit that each contract makes *individually*, but only the profits obtained by the *menus*. Our analysis extends to this case as follows. Denote by $B(s)$ be the set of menus with maximal profits in state s. It is straightforward that our theorem still holds under the alternative rule for imitation and withdrawal: any firm j with $S_j(t) \notin B(s(t))$ replaces its menu by

$$\left(S_j(t) \cup \tilde{S}\right) \setminus \left(B^-(s(t)) \cap S_j(t)\right)$$

where $\tilde{S} \in B(s(t))$; those firms j for which $S_j(t) \in B(s(t))$ replace their menus by

$$S_j(t) \setminus \left(B^-(s(t)) \cap S_j(t)\right)$$

If $S_j = \emptyset$, a mutating firm j may add to its menu any subset of Γ.[8] Notice that for r sufficiently large, when firms experiment, they will offer new contracts which can be anywhere on the contract space.

Markov process

As we have described, for each firm j, imitation, withdrawal, and mutation result in some portfolio $S_j(t+1)$. This takes the system from state $s(t)$, at the beginning of period t, to state $s(t+1) = (S_1(t+1), \ldots, S_n(t+1))$, at the beginning of period $t+1$. Given $\epsilon > 0$, at any t, the probability of transition from any state s to any other state s', $P^\epsilon_{ss'}$, is uniquely determined and independent of t, and, thus, $P^\epsilon = (P^\epsilon_{s,s'})_{s,s' \in \Omega}$ defines the transition matrix of a Markov process. Then, the following standard result holds, and is quoted here without proof (see e.g. Freidlin and Wentzell (1984)).

Lemma 1 *If the process is ergodic* [9]*, then, given $\epsilon > 0$, the process P^ϵ has a unique invariant distribution μ^ϵ, which statistically summarizes the behavior of the system along any sample path with probability one independently of initial conditions. Furthermore, the limit invariant distribution $\mu^* = \lim_{\epsilon \to 0} \mu^\epsilon$ exists.*

Any state s with $\mu^*(s) > 0$ is called a *long-run state*. If the mutation probability is sufficiently small, then the invariant distribution concentrates its weight on the long-run states, whereas the probability of all other states tends to zero. Thus, one can take the set of long-run states as a prediction of the long-run behavior if mutation is rare. These states will be characterized using the techniques introduced by Young (1993), and Kandori, Mailath and Rob (1993), as applied in Nöldeke and Samuelson (1993, 1997). We can now turn to the analysis.

---

[8] Our main result still holds if the previous rule is changed to "...may add to its menu any subset of E", where $E \subseteq \Gamma$ contains at least one pair of separating, profit-making contracts that are neighbors to each other. An example that works for all sufficiently small δ is $E = \mathcal{N}_r(0,0)$.

[9] Notice that P^ϵ is not necessarily a positive matrix due to the lack of inertia in the learning process and the fact that mutation is local. This in turn implies that the process might not be irreducible. However, in the appendix, we will show that it has a unique recurrent communication class, which is aperiodic. This is what we mean here by ergodicity.

3 Analysis

In the following we make a simplifying genericity assumption. It requires that purchased contracts yield either profits or losses, and no individual is indifferent between different contracts.

Assumption 2 *For all* $(P, I) \in \Gamma$, *we have* $P/I \notin \{\pi_l, \pi_h, \pi_{hl}\}$. *Furthermore,* $U_i(c) \neq U_i(c')$ *for* $c, c' \in \Gamma$ *with* $c \neq c'$, *and* $i = l, h$.

Before we state our main result, let us denote by $R(c, c')$ the set of states in which *all* firms offer contracts c, c', and (possibly) any set of idle contracts.

Our main result states that, if contracts lie on a sufficiently fine grid, and mutations occur in small neighborhoods of the existing ones, then there exist exactly two profit-making contracts that will be observed in the long run. These are characterized by maximality properties and approximate the Rothschild-Stiglitz contracts. Moreover, if the latter constitute a competitive equilibrium in the static framework, the statements hold even if mutation is global.

Theorem There exist $r_0 > 0$ and a function $\delta_0(\cdot) > 0$ such that, for $r < r_0$ and $\delta < \delta_0(r)$, $\mu^*(s) > 0$ if and only if $s \in R(k_l, k_h)$, where k_l, and k_h are profit-making contracts, uniquely characterized by

$$U_h(k_h) = \max_{c \in \Gamma, \ V(c, \pi_h) > 0} U_h(c), \tag{1}$$

$$U_l(k_l) = \max_{c \in \Gamma, \ V(c, \pi_l) > 0, U_h(c) < U_h(k_h)} U_l(c). \tag{2}$$

For $\delta \to 0$, we have $k_h \to c_h^{RS}$ and $k_l \to c_l^{RS}$. Moreover, if $\mathcal{P} = \emptyset$, then we can choose $r_0 = \infty$.

A proof of this theorem is presented in the Appendix. Here we outline the intuition of the market mechanisms that underlie the result. This will be done in five steps.

First, notice that, in each period, withdrawal eliminates loss-making contracts and imitation only adds profit-making ones already present in the market. Therefore, the number of offered contracts cannot increase by imitation and withdrawal. Furthermore,

once there are no loss-making contracts offered, imitation leads trivially to situations in which all firms are offering the same set of active contracts. Thus, in the long run, the only states that will be potentially observed will be those ones in which all firms are offering the same set of profit-making contracts, and no loss-making ones. These constitute the *absorbing* states of the dynamics without mutation.

Second, any state s in which k_l and k_h are purchased ($s \in R(k_l, k_h)$) is stable in the sense that after any mutation, imitation and withdrawal lead the system back to another state $s' \in R(k_l, k_h)$. To see this, notice first that in the neighborhood of k_l and k_h, there exists no profit-making contract which attracts both risk types or only the high risks, and every profit-making contract which attracts only the low risks must be cross-subsidizing a loss-making contract taken by the high risks (otherwise, both types would be buying the same contract). Even if a firm experiments with such a pair of cross-subsidizing contracts, imitation and withdrawal eventually lead back to the original state (in which k_h and k_l were purchased), because after the loss-making contract is withdrawn, the other mutant contract attracts both risk types and makes itself losses, vanishing then from the market also. If the grid is fine enough, k_l, k_h approximate the Rothschild-Stiglitz contracts.

In the following steps we argue that starting from *any other* absorbing state, a chain of single mutations along absorbing states will lead to a state in $R(k_l, k_h)$. Intuitively, this completes the proof, because, since mutations are rare, two mutations at once (the event necessary to *destabilize* states in $R(k_l, k_h)$) is a much less probable event than a single mutation (the event necessary to destabilize any other state). Moreover, one can connect any state out of $R(k_l, k_h)$ to one in $R(k_l, k_h)$ with a chain of single mutations, but one would have to introduce two mutations at once at the beginning of one of such chains in order to connect any state in $R(k_l, k_h)$ to any other out of it.

Third, suppose the system is in an absorbing state where two arbitrary profit-making contracts are offered, one purchased by the low risks, the other one by the high risks. It may be that in the neighborhood of the low risks' contract there exists another one which does not attract the high risks, but increases the payoff of the low risks and still makes profits. If a firm experiments with such a contract, it will be imitated, implying

10

that the former low risks' contract becomes idle. Repeated application of this kind of competition, and a similar one on the high risks' contracts, leads to the conclusion that all states of this type can be abandoned with a single mutation and that there exists a chain of single mutations along absorbing states which brings the system to a state where no firm can offer a yet more attractive contract to either type and still make profits, i.e., a state in $R(k_l, k_h)$.

Fourth, suppose the system is in a state in which only one profit-making contract is offered which is bought by the high risks only (the low risks are not buying any insurance). Then, price-cutting and/or indemnity-increasing experimentation can lead to a profit-making contract which is bought by both risk types. Therefore, suppose now that the system is in a state in which such a contract is offered. Then, provided that the grid is fine enough, a mutant contract exists which attracts the low risks only. This contract makes profits. However, the incumbent contract is now bought by the high risks only and it can either make profits or losses. In the former case, we are in a state with two profit-making contracts, and the chain of single mutations continues as described above. In the latter, after this contract is discarded, the mutant contract attracts both types, until a new mutant comes in that again attracts the low risks only. Finally, after some mutations of this type, an additional mutant makes losses as soon as it attracts both types. This last contract is discarded, implying that the insurance market has died out.

Fifth, suppose that the system is in a state in which no contract is offered. For small grid size, the contract space contains a pair of profit-making contracts, each taken by a different risk type, such that these contracts are neighbors to each other. One mutation can make the contract for the high risks arise. This one is maybe also taken by the low risks, but makes profits anyway, so an absorbing state is reached after this contract is imitated. The contract for the low risks arises by one further mutation, implying that we have reached a state with two separating, profit-making contracts which was already discussed.

As this intuitive explanation indicates, the theorem holds even if firms only experiment with one new contract at a time. Furthermore, the result also holds if, starting

from a dead market, firms only experiment near the origin of the contract space (see footnote 8).[10]

If the Rothschild-Stiglitz contracts do not constitute a competitive equilibrium in their model, and mutations are global, the evolutionary process does not select any specific set of contracts. Only the loss-making contracts are eliminated.

Corollary Suppose $\mathcal{P} \neq \emptyset$ and $r = \infty$. Then, for all sufficiently small δ, all absorbing states are long-run states.

The idea of the proof (see Appendix) is that all pooling situations can again be destabilized along the same lines of the proof of the theorem , but now there also exists a mutation, which destabilizes the Rothschild-Stiglitz outcome. This mutation is given by a profit-making, pooling contract in the set $\mathcal{P}$, exactly as in their original paper.

4   Conclusions

We have analyzed an insurance market under asymmetric information in an evolutionary context. Our main findings are the following:
- If the separating contracts of the Rothschild-Stiglitz type do constitute an equilibrium in their model, then they are also the only ones purchased in any long-run state of our evolutionary process. This implies that firms, without knowing anything about the risks, utility functions or number of types of their customers, still learn to offer the competitive equilibrium contracts.
- If experimentation is local, then the Rothschild-Stiglitz, separating contracts are still offered in all long-run states, even if a competitive equilibrium of the classic model does not exist.

Our work can be interpreted as follows. Provided that insurance companies decide

---

[10]It is worth mentioning that, for our dynamics, the speed of convergence is high in the sense that the expected waiting time until the system reaches a long-run state is of order ϵ^{-1} (this follows from Theorem 2 in Ellison (1998)).

mainly on the basis of how other firms have performed in the market with their contracts. If, from time to time, they also experiment with new contracts in a cautious manner, in particular only with contracts that are similar to the existing ones. Then, in such a framework, the Rothschild-Stiglitz type of separating contracts will be observed in the long run, even if this does not constitute a competitive equilibrium of the classic, static framework. To destabilize those contracts, one needs firms that can experiment with any contract on the contact space. However, as already Rothschild and Stiglitz (1976, p.646) have noted, 'one would expect that competition would lead to small perturbations'.

References

Alchian, A. (1950): "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy* 58, 211-221

Alós-Ferrer, C., A. B. Ania and F. Vega-Redondo (1998): "An Evolutionary Model of Market Structure," forthcoming in *Theory of Markets and its Functioning,* (G. van der Laan, J.–J. Herings, and D. Talman Eds.). Amsterdam: North Holland.

Dasgupta, P. and E. Maskin (1986): "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications," *Review of Economic Studies* 53(1), 27-41.

Ellison, G. (1998): "Basins of Attraction, Long Run Stochastic Stability, and the Speed of Step-by-Step Evolution," mimeo, MIT.

Freidlin, M. I. and A. D. Wentzell (1984): "Random Perturbations of Dynamical Systems," Berlin: Springer Verlag.

Grossman, H.I. (1979): "Adverse Selection, Dissembling and Competitive Equilibrium," *Bell Journal of Economics* 10, 330-343.

Hellwig, M. (1987): "Some Recent Developments in the Theory of Competition in Markets with Adverse Selection," *European Economic Review* 31, 319-325.

Kahn, C.M. and D. Mookherjee (1995): "Coalition Proof Equilibrium in an Adverse Selection Insurance Economy," *Journal of Economic Theory* 66, 113-138.

Kandori, M., G.J. Mailath and R. Rob (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* 61, 29-56.

Miyazaki, H. (1977): "The Rat Race and Internal Labor Markets," *Bell Journal of Economics* 8, 394-418.

Nöldeke, G. and L. Samuelson (1993): "An Evolutionary Analysis of Backward and Forward Induction," *Games and Economic Behavior* 5, 425-454.

Nöldeke, G. and L. Samuelson (1997): "A Dynamic Model of Equilibrium Selection in Signalling Markets," *Journal of Economic Theory* 73, 118-156.

Riley, J.G. (1979): "Informational Equilibrium," *Econometrica* 47, 331-359.

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

Spence, M. (1978): "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics* 10, 427–447.

Vega-Redondo, F. (1997): "The Evolution of Walrasian Behavior," *Econometrica* 65, 375-384.

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

Young, P. (1993): "The Evolution of Conventions," *Econometrica* 61, 57-84.

Appendix

Proof of the theorem:

For any two states $s, s' \in \Omega$ the *resistence* $r(s, s')$ is the minimum number of mutations in any finite sequence of transitions that brings the system from s to s'. A nonempty set $Q \subseteq \Omega$ is called *absorbing* if Q is minimal with respect to the property that for all $s \in Q$, $s' \notin Q$ we have $r(s, s') \neq 0$. A state s is called *absorbing* if $\{s\}$ is an absorbing set.

The next Lemma characterizes the absorbing sets. Note that in an absorbing state firms may offer different sets of idle contracts.

Lemma 2 *Each absorbing set consists of a single state. A state is absorbing if and only if (1) no firm offers a loss-making contract, and (2) if any firm offers a profit-making contract, this contract is offered by all firms.*

Proof. "if": obvious. "only if": Suppose the system is in any state. Every period, by imitation and withdrawal, the total number of contracts offered decreases or stays constant until it cannot decrease further by subsequent withdrawal, i.e., (1) holds. At this point, imitation implies that (1) and (2) hold. This also shows that each absorbing set consists of a single state.

For any two absorbing states s, s', we write $s \rightarrow s'$ if $r(s, s') = 1$. We write $s \Rightarrow s'$ if there exists a finite sequence of absorbing states $s_1 = s, \ldots, s_q = s'$ $(q \geq 1)$ such that $s_i \rightarrow s_{i+1}$ for $i = 1, \ldots, q-1$ ("*a chain of single mutations*"). We write $s \Longleftrightarrow s'$ if $s \Rightarrow s'$ and $s' \Rightarrow s$. An equivalence class R of $\Longleftrightarrow$ is called a *locally stable component* if there do not exist $s \in R$, $s' \notin R$ with $s \rightarrow s'$. By Nöldeke-Samuelson (1993), Proposition 1, the existence of a unique locally stable component implies that the support of the limit invariant distribution μ^* coincides with the states appearing in that component.

For any $r > 0$, $\delta > 0$, and contracts $c_l, c_h \in \Gamma$, we define

$$\mathcal{U}_l(c_l, c_h) = \{c' \in \mathcal{N}_r(c_l) | U_l(c') > U_l(c_l), \ U_h(c') < U_h(c_h), \ V(c', \pi_l) > 0\},$$
$$\mathcal{U}_h(c_l, c_h) = \{c' \in \mathcal{N}_r(c_h) | U_h(c') > U_h(c_h), \ U_l(c') < U_l(c_l), \ V(c', \pi_h) > 0\}.$$

Claim I. There exists $r_0 > 0$ such that for all $0 < r < r_0$, the following holds: for all sufficiently small $\delta > 0$, a set of states R is a locally stable component if and only if $R = R(k_l, k_h)$ with profit-making contracts k_l and $k_h \neq k_l$ such that $\mathcal{U}_l(k_l, k_h) = \emptyset$ and $\mathcal{U}_h(k_l, k_h) = \emptyset$.

Claim II. For all $r > 0$, the following holds: for all sufficiently small $\delta > 0$, let $k_l = k_l(\delta)$ and $k_h = k_h(\delta)$ be such that $\mathcal{U}_l(k_l, k_h) = \emptyset$ and $\mathcal{U}_h(k_l, k_h) = \emptyset$; then the maximization program given by (1) and (2) and the convergence result formulated in the theorem hold for k_l and k_h.

Remark 1 *The proof of these two claims suffices to prove our Theorem. Claim I shows that the set of long-run states is the unique stable component $R(k_l, k_h)$, i.e., starting from any state in this set, after one mutation, learning will lead the process back into that set; moreover, starting from any other state, there exists a chain of single mutations that leads the process into some state in $R(k_l, k_h)$. Moreover, all states in $R(k_l, k_h)$ can be connected among themselves. At the same time this also shows that all states can be connected to a given one. This immediately implies the existence of a unique recurrent communication class, which includes the states in $R(k_l, k_h)$. If other communication classes exist, they are transient. Aperiodicity follows from the fact that from any absorbing state there is positive probability that the process stays in it. This guarantees that 1 applies. Claim II then refers to the characterization of k_l and k_h.*

Proof of Claim I. If $\mathcal{P} \neq \emptyset$, then let $r_0 > 0$ be the $||\cdot||$-distance between the sets $\{c \in \mathbb{R}^2 | U_h(c) \leq U_h(c_h^{RS})\}$ and $\mathcal{P}$; if $\mathcal{P} = \emptyset$, define $r_0 = \infty$. Now fix $r < r_0$, and let k_l and $k_h \neq k_l$ such that $\mathcal{U}_l(k_l, k_h) = \emptyset$ and $\mathcal{U}_h(k_l, k_h) = \emptyset$. By Claim II, δ can be chosen

so small that $k_l = k_l(\delta)$ and $k_h = k_h(\delta)$ are uniquely characterized by the maximization program given by (1) and (2) in the theorem. Furthermore, we choose δ so small that the $||\cdot||$-distance between the sets $\{c \in \mathbb{R}^2 | U_h(c) \le U_h(k_h)\}$ and

$$\mathcal{P}(k_l) = \{c \in \mathbb{R}^2 | V(c, \pi_{hl}) > 0, \ U_l(c) > U_l(k_l)\}$$

is larger than r (this is possible because, by Claim II, the distance approximates r_0 for $\delta \to 0$.)

It is sufficient to show that for all sufficiently small $\delta > 0$ the following holds: (1) After any single mutation away from a state in $R(k_l, k_h)$, imitation and withdrawal *must* drive the system back to $R(k_l, k_h)$. (2) Starting from any absorbing state there exists a chain of single mutations that leads into some state in $R(k_l, k_h)$. (3) If $s, s' \in \Omega$ are absorbing states with identical nonempty sets of profit-making contracts (s and s' differ only in the idle contracts), then $s \Rightarrow s'$.

(1): We consider three cases. (a) Suppose a firm mutates such that it discards profit-making contracts. At the end of the period, it realizes that the other firms are still making profits with their contracts k_l and k_h, and then it learns to re-offer these contracts. (b) Suppose a firm i mutates such that it invents a set of new contracts, none of which makes profits. One, or both, of the old profit-making contracts may become idle, but none of them can start making losses. Hence, at the end of the period only firm i changes its menu: it discards all loss-making contracts. As a result, some of i's idle contracts may become loss-making or profit-making. If only loss-making contracts come up, our argument can be repeated, while the size of firm i's menu is decreased up to the point where the system is back to a state in $R(k_l, k_h)$. If however, a profit-making contract comes up, we have to consider case (c) below. (c) Suppose a firm mutates such that it invents a set of new contracts N, and there exists a contract $c \in N$ which makes profits. We know that c is not bought by both consumer types, because $c \notin \mathcal{P}(k_l)$. Moreover, by the characterization of k_h, c is not bought by the h-types, but by the l-types. By the characterization of k_l, we have $U_h(c) > U_h(k_h)$, and there exists a contract $d \in N$, bought by the h-types, with $U_h(d) > U_h(c) > U_h(k_h)$. By the characterization of k_h, d makes losses ("cross-subsidization"). All contracts in $N' = N \setminus \{c, d\}$ are

17

idle. Subsequent learning implies that c is imitated and d is discarded. As soon as d has vanished from the market, either c cross-subsidizes another loss-making contract $d' \in N'$, or c makes losses. After some learning, all potential contracts that can be cross-subsidized by c have vanished, and finally also c is discarded. After that, there might exist another profit-making contract $c' \in N$. Repetition of this argument reveals that after some more learning all contracts that are contained in N are discarded or become idle and a state in $R(k_l, k_h)$ is reached.

(2): Assume that the system is in any absorbing state s. First, a chain of single mutations exists such that all idle contracts in s are discarded. Call s' the state reached then. In s' each firm offers the same set S of contracts, and all contracts in S make profits. Three cases are possible: (a) $S = \{c_l, c_h\}$ ($c_l \neq c_h$), (b) $S = \{c\}$, and (c) $S = \emptyset$.

(a): Let c_l denote the contract bought by the l-types. Now suppose that $\mathcal{U}_l(c_l, c_h) \neq \emptyset$. There exists a mutation which adds a contract $c'_l \in \mathcal{U}_l(c_l, c_h)$ to some firm's menu. The contract c'_l makes profits because it attracts only the l-types. Now all firms imitate c'_l, and c_l becomes idle. Thus, the system has arrived in a new absorbing state. Again, by a chain of single mutations the idle contract c_l is discarded, so we are back to case (a) with $S' = \{c'_l, c_h\}$. The utility of the l-types is higher in S' than in S, and the utility of the h-types is unchanged. In a similar way, if $\mathcal{U}_h(c_l, c_h) \neq \emptyset$, then a chain of single mutations leads back to case (a) with the h-types' utility increased and the l-types' utility unchanged. Since the set of possible utility levels is finite for both types, there exists a chain of single mutations which leads the system to $S^* = \{k_l, k_h\}$.

(b): Here we need three facts. The first one is needed due to technical problems which arise near the origin. There exists a $\hat{I} > 0$ such that:

$$\exists \delta_0 > 0 \ \forall \delta < \delta_0 \ c = (P, I) \in \Gamma \ (I \leq \hat{I}, \ U_l(c) > U_l(0, 0)) : \ \mathcal{U}_h(c, c) \neq \emptyset. \qquad (3)$$

Along the rest of the state space, the following statement holds:

$$\exists \delta_0 > 0 \ \forall \delta < \delta_0, \ c = (P, I) \in \Gamma \ (I > \hat{I}, \ V(c, \pi_{hl}) > 0) : \ \mathcal{U}_l(c, c) \neq \emptyset. \qquad (4)$$

(To see (4), note that $P \geq \pi_{hl}\hat{I}$). Facts (3) and (4) hold under our assumptions on the utility function. In particular, $u' > 0$ and continuity of u' guarantee that $[U'_h(c) - U'_l(c)]$

is uniformly bounded from below. The third fact follows from the characterization of k_h given by (1) in the maximization program in the theorem for δ sufficiently small.

$$\exists \delta_0 > 0 \; \forall \delta < \delta_0, \; c = (P, I) \in \Gamma \; (c \neq k_h, V(c, \pi_h) > 0), \; \exists c' \in \mathcal{N}_r(c) : \; V(c', \pi_h) > 0, \; U_h(c') > U_h(c).$$
$$\tag{5}$$

(bα): Consider first the case that both the l- and the h-type buy $c = (P, I)$. If $I \leq \hat{I}$, then, by fact (3), there exists $c' \in \mathcal{U}_h(c, c)$ that attracts the high risks only and makes a profit; c' is imitated and we are in case (a) again. If $I > \hat{I}$, then suppose a mutation occurs which adds a $c' \in \mathcal{U}_l(c, c)$ (existence follows from fact (4)) to one firm's menu. At first c' only attracts the l-types, makes profits and is imitated by all firms. Now, if $V(c, \pi_h) > 0$, then we are in case (a); if, however, $V(c, \pi_h) < 0$, then c makes losses now, it is discarded, and everybody buys c'. If then $V(c', \pi_{hl}) > 0$, we are in case (bα) again with c' instead of c, thereby having increased the payoff of the l-types. If, on the contrary, $V(c', \pi_{hl}) < 0$, then c' makes losses, it is discarded, and we are in case (c) below. Note that we are either in case (a) or in case (c), after a finite number of steps, because there exist only finitely many utility levels for the l-types.

(bβ): Consider now the case that only the h-types buy $c = (P, I)$ and the l-types buy no insurance. If $c \neq k_h$, by fact (5), it follows that there exists a mutant c' that will be imitated and c becomes idle. The construction can be repeated with c replaced by c', until a state is reached with a unique profit-making contract d, where either $d = k_h$ or d is sold to both types. In the latter case, we are in case (bα). To tackle the former, let δ be so small that $U_l(\bar{P}, \delta) < U_l(0, 0)$, and let t_1 be the state where all firms offer the menu $\{(\bar{P}, \delta), k_h\}$. By part (3) below, state t_1 can be reached by a chain of single mutations. Now consider the following sequence of contracts

$$(c_1, c_2, \ldots) = ((\bar{P}, \delta), (\bar{P}, 2\delta), \ldots, (\bar{P}, I_l), (\bar{P} - \delta, I_l), \ldots, (P_l, I_l)),$$

where $(P_l, I_l) := k_l$. Let t_i be the state where all firms offer the menu $\{c_1, \ldots, c_i, k_h\}$, and let k be minimal with $U_l(c_k) > U_l(0, 0)$. By construction, $V(c_k, \pi_l) > 0$ and $U_h(c_k) < U_h(k_h)$. By (3), we have $t_1 \Rightarrow t_k$. After cancelling the idle contracts $\{c_1, \ldots, c_{k-1}\}$ in t_k, we have reached case (a) with the separating pair $\{c_k, k_h\}$.

(c): Let $c_l \in \Gamma$ and $c_h \in \mathcal{N}_r(c_l)$ be profit-making contracts, bought by the l- and

the h-types respectively (such a pair exists if δ is sufficiently small). Define t as the absorbing state where all firms offer the menu $\{c_l, c_h\}$.

In the current state no firm is offering any contract, but there exists a mutation which adds c_h to one firm's menu. After all firms have imitated c_h, another absorbing state is reached. There exists a further mutation which adds c_l to one firm's menu, and imitation leads to state t. Now we are in case (a).

(3): Define $t \in \Omega$ as the state constructed from s by discarding all idle contracts. We have $s \Rightarrow t$. Now consider any idle contract $c = (P, I)$ in state s'. Let $c^* = (P^*, I^*)$ be a profit-making contract in s. We do the case $P \geq P^*$ ($P < P^*$ is similar). Consider the following sequence of contracts,

$$(P^*, I^*), (P^* + \delta, I^*), \ldots, (P, I^*), (P, I^* \pm \delta), \ldots, (P, I).$$

Starting from state t, there exists a mutation which adds $(P^* + \delta, I^*)$ to firm 1's menu, a subsequent mutation which adds this contract to firm 2's menu, and so on. Now there exist mutations which add $(P^* + 2\delta, I^*)$ to all firm's menus, and then there exist mutations which withdraw $(P^* + \delta, I^*)$ from all menus. This process can be continued according to the previous sequence of contracts until a state u is reached which differs from t only by the fact that c is added to all firms' menus. All contracts added and deleted on the way were idle (c^* or c were preferred), and therefore all states on the way are absorbing, which implies $t \Rightarrow u$. In a similar way, all idle contracts in state s' can be added to the firms' menus.

Sketch of proof of Claim II. The convergence result follows easily from the maximization program given by (1) and (2) in the theorem. It remains to show, for any contracts $k_h, k_l \in \Gamma$, that (1) and (2) in that maximization program are satisfied if and only if $\mathcal{U}_i(k_l, k_h) = \emptyset$ ($i = l, h$). "only if" is straightforward from the definitions. To see "if", let $c_h, c_l \in \Gamma$ be a profit-making, separating pair which does not satisfy the maximization program given by (1) and (2) (with k_i replaced by c_i), but satisfies $\mathcal{U}_h(c_l, c_h) = \emptyset$. We have to show that $\mathcal{U}_l(c_l, c_h) \neq \emptyset$.

There can be only three reasons for $\mathcal{U}_h(c_l, c_h) = \emptyset$: (1) in the vicinity of c_h, every contract which is better for the high risks, makes a loss; hence, $c_h = k_h$, and (1) in the

maximization program holds. Then it must be that (2) in that maximization program does not hold, which in turn implies that $\mathcal{U}_l(c_l, c_h) \neq \emptyset$. (2) any contract in the vicinity of c_h which is preferred by the high risks is also preferred by the low risks to their contract c_l; by the single-crossing property and the fact that the low risks prefer c_l to c_h, this can only hold if c_h lies in the vicinity of the upper boundary of the contract space, $P = \bar{P}$, and the low risks just prefer their contract to c_h; by single crossing it then holds that $\mathcal{U}_l(c_l, c_h) \neq \emptyset$. (3) for some contracts in the vicinity of c_h firms would make a loss, for others the contract is preferred by the low risks; this can only hold at a point of overinsurance, and the low risks just prefer their contract to c_h; similar as above, this implies that $\mathcal{U}_l(c_l, c_h) \neq \emptyset$. This completes the proof. □

Proof of the corollary:

We must show that for any two absorbing states s, s' we have $s \Rightarrow s'$. From the theorem, $s \Rightarrow t$ for any $t \in R(k_l, k_h)$. Starting from t, there exists a mutation which invents a $c \in \mathcal{P}$ (to get this, δ must be chosen so small that $\mathcal{P} \cap \Gamma \neq \emptyset$). After all firms have learned this contract, the system is again in an absorbing state. Now a chain of further mutations discards all idle contracts and we are in case $(b\alpha)$ of the proof above. From there, case (c) can be reached. Starting from the dead market, there exists a mutation which adds the set of profit-making contracts in s' to one firm's menu. This firm is then imitated, and we have reached a state s'' with $s \Rightarrow s''$ such that s'' is identical to s' up to idle contracts. By step (3) in Claim I, $s'' \Rightarrow s'$. □