Universität Trier

Adverse selection and heterogeneity of demand responsiveness

Normann Lorenz

Research Papers in Economics
No. 2/14

Adverse selection and heterogeneity of demand responsiveness

Normann Lorenz*

January 27, 2014

Abstract

This paper analyzes the distortions of (health) insurers' benefit levels due to adverse selection if individuals' responsiveness to differences in contracts is heterogeneous. Within a discrete choice model with two risk types and imperfect competition the following results are shown: In the pooling equilibrium, a positive correlation of low risk and high responsiveness (e.g., younger individuals being both healthier and faster to switch insurers than older individuals) increases the distortion of the uniform benefit level if the share of low risks is small; if the share of low risks is large, the reverse holds, but only if the average level of responsiveness is high. In the separating equilibrium, a positive correlation increases the distortion of the contract for the low risks, unless the number of insurers offering the contract for the high risks is very small or a large share of the high risks chooses the contract designated for the low risks. These results imply that the welfare effects of a policy intervention of making individuals more responsive crucially depend on which risk types' responsiveness is increased more. The results also have implications for the estimation of the level of risk aversion and of the welfare effects of adverse selection.

JEL-**classification:** I 13.

Keywords: Adverse selection, discrete choice.

*Universität Trier, Universitätsring 15, 54286 Trier, Germany; e-mail: lorenzn@uni-trier.de; phone: 0049-651-2012624. I thank Friedrich Breyer, Mathias Kifmann and Esther Schuch for helpful comments and suggestions.

1 Introduction

Adverse selection in (health) insurance markets has recently attracted renewed attention, as it has been observed that heterogeneity in other dimensions than risk may be of similar importance as differences in risk and either mitigate or exacerbate the distortions caused by adverse selection (Cohen and Siegelman 2010). Important examples for such additional dimensions are risk aversion (Cutler et al. 2008), cognitive ability (Fang et al. 2008), switching costs (Handel 2013) or income (Johar and Savage 2012).

Yet another dimension that exhibits a considerable degree of heterogeneity is individuals' general responsiveness to differences in the contracts offered by insurers. Some individuals are rather attentive when buying insurance; these individuals respond even to small differences in contracts like additional benefits or price discounts. For others, a particular contract has to yield considerably higher utility than all the other contracts before it is chosen with high probability. For health insurance, one example of an observable variable that is correlated with responsiveness is age: younger individuals usually belong to the first group, and older individuals to the second.[1] Since age is also a determinant of expected expenditures, this would imply a positive correlation of low risk type and high responsiveness. Ericson and Starc (2012b) have shown that such a correlation influences the effects of the modified community rating regulation of the Massachusetts Health Insurance Exchange, and Bijlsma et al. (2011) have shown that it alters the optimal design of a risk adjustment scheme.

So far it has neither been demonstrated empirically nor derived in a theoretical model how a correlation of risk type and responsiveness affects the distortions caused by adverse selection in the absence of such regulatory means.[2] The general notion, however, seems to be that a positive correlation of low risk and high responsiveness exacerbates the distortion: if addressing the preferences of the low risks attracted many of these individuals (because they are very responsive), this would increase the incentive to distort the contract. However, this argument could also be reversed, i.e. argued that a negative correlation exacerbates the distortion: if addressing the preferences of the low risks induced many of the high risks to choose another insurer (if they were the ones who are very responsive), this would also increase the incentive to distort the contract.

Within a discrete choice model with two unobservable risk types we show that either of the two cases can occur: In the pooling equilibrium, a positive correlation increases the distortion if the share of low risks is small; if the share of low risks is large, the reverse holds, but only if the average level of responsiveness is high. In the separating equilibrium, a positive correlation increases the distortion of the contract for the low risks, unless the number of insurers offering the contract for the high risks is very small or a large share of the high risks chooses the contract designated for the low risks.

These results have two main implications: First, from a policy perspective, they indicate that the welfare effects of increasing individuals' responsiveness by, e.g., providing easy

[1] See Ericson and Starc (2012b), who show that the higher premiums older individuals have to pay cannot entirely be explained by higher costs, but are also due to the lower responsiveness of older individuals which allows insurers to charge higher premiums.

[2] See also Einav et al. (2010, p. 333), who explicitly raise the question whether a correlation of the amount of consumer search or consumer interest in plan switching and risk might affect competition.

access to information about insurers' offers (as with the government run websites for the Health Insurance Exchanges in the U.S.), depend crucially on which of the two risk types' responsiveness is increased more.

The second implication concerns the distortions in the separating equilibrium. Some studies have used the set of contracts offered in such equilibria to estimate an underlying preference parameter of risk aversion; they have also been used to determine an estimate of the welfare losses caused by adverse selection.[3] This study shows that using these contracts may entail something akin to a 'measurement error', and that the size of this measurement error depends on the level of correlation. Ignoring this measurement error results in a downward biased estimate of the level of risk aversion and of the welfare losses caused by adverse selection.

We derive our results within a discrete choice model, namely the (conditional or mixed) logit. The logit model has been extensively used for empirical analyses of health insurance choice;[4] here it is used to capture different degrees of responsiveness in a theoretical model of adverse selection.[5]

There are three important aspects of using this discrete choice model for a theoretical analysis: First, it endogenizes whether a pooling or separating equilibrium occurs: if the average level of responsiveness is low, the former, if it is high, the latter equilibrium emerges.

Secondly, it allows to capture the fact that some individuals 'make mistakes' when choosing their (health) insurance contract.[6] If some individuals are less responsive to differences in the contracts offered, they have a higher probability of making such a mistake. This higher probability of choosing the 'wrong' contract can easily be depicted graphically by introducing the concept of an 'indifference curves area'. Such an 'indifference curves area' provides an intuitive understanding of the economic forces driving the additional distortions caused by a correlation of risk type and responsiveness.

Thirdly, this discrete choice model relaxes the assumption of a strong demand asymmetry that is implicit in the studies which so far have analyzed health insurance choice under imperfect competition: Most of these studies consider a setting where – for the case of two risk types – each insurer offers two contracts so that the incentive compatibility constraint is satisfied; imperfect competition is then captured by a Hotelling-model.[7] These models imply the following strong asymmetry of demand responses: Consider a group of individuals holding a contract from a particular insurer. A new contract, yielding slightly higher utility than the contract they currently hold, would attract all these individuals, if offered by the same insurer, but only a small share of them, if offered by a different insurer. This would be a reasonable assumption if individuals were perfectly informed about all the offers of their insurer, but not about those of the other insurers; it would also be reasonable if switching to

[3]See Einav et al. (2010) for an overview of such studies.

[4]See, e.g., Royalty and Solomon (1999), Harris et al. (2002), Keane (2004) and Ericson and Starc (2012a).

[5]The model is similar to the one employed by Lorenz (2013), who – in a setting without heterogeneity in responsiveness – has examined the impact of imperfect competition on the effectiveness of a risk adjustment scheme.

[6]See Handel and Kolstad (2013) and Sinaiko and Hirth (2011) for empirical evidence.

[7]See, e.g., Biglaiser and Ma (2003), Jack (2006), Olivella and Vera-Hernandez (2007) and Bijlsma et al. (2011).

another insurer incurred much higher transaction costs (like filling out an application form) than switching to another contract of one's current insurer. However, with more and more individuals using the internet, it seems more appropriate to consider getting informed about (and switching to) an alternative contract from either one's current or from a different insurer as equally difficult (or easy).[8] We therefore relax the assumption of the strong demand asymmetry, and we do that in the easiest way possible in that we assume that each insurer offers only one contract.

We will present the model in the way that the risk the individuals face is to develop an illness which requires to be treated. All individuals are obliged to buy health insurance, where each insurer offers a benefit level and charges a premium. We thus analyze benefit level-premium-bundles and present the results graphically in a benefit-premium-diagram.[9] The model and all results can, however, easily be transferred into the premium-deductible-space or income when sick-income when healthy-space (Rothschild and Stiglitz 1976); the results are therefore not confined to a health insurance setting.

Also, we will present the model under the assumption that both the risk type and the level of responsiveness are unobservable (at the individual level). A different setting to which the model applies, is that risk type and responsiveness are observable, but that insurers are not allowed to charge type-specific premiums because of community rating.

The remainder of this paper is organized as follows: In Section 2 we present the model, introduce the concept of the 'indifference curves area' and discuss the three types of equilibria that can occur (a pooling equilibrium and two types of separating equilibria). The impact of a correlation of risk type and responsiveness on the pooling equilibrium is analyzed in Section 3, and on the separating equilibrium in Section 4. In Section 5 the policy implications of the results and some of the assumptions are discussed. Section 6 concludes.

2 The Discrete Choice Model

2.1 Basic model

Each individual may suffer from an illness that occurs with probability p. If it occurs, utility changes by $v(m)$, where m is the medical services (measured in monetary terms) provided by an insurer, who in return charges a premium R. $v(m)$ is increasing at a decreasing rate, i.e. $v'(m) > 0$ and $v''(m) < 0$; the efficient level of m is implicitly defined by $v'(m^{FB}) = 1$. There are two unobservable risk types $r = L, H$, with $p^L < p^H$; the share of L-types is λ. Individuals' preferences are therefore given by

$$u = p^r v(m) - R. \tag{1}$$

As already stated in the introduction, a different (and may be more traditional) representation of preferences would be

$$u = p^r v(y^s) + (1 - p^r) v(y^h), \tag{2}$$

[8] E.g., for Massachusetts, Ericson and Starc (2013) report that most individuals enrolled through the Health Insurance Exchange's website.

[9] We therefore adopt the representation of preferences of Zweifel et al. (2009), chapter 7.

with $y^s = y - D - R$ and $y^h = y - R$, where y^s represents income when sick, y^h income when healthy, y initial income, and D the deductible.[10] All results could just as well be derived in y^s-y^h-space using (2), but in this paper we will refer to the representation of preferences as given by (1) and depict all results in m-R-space (as in Zweifel et al. (2009), chapter 7).

There are n insurers j, each offering a contract $\{m^j, R^j\}$. Individual i's utility when choosing an insurer j not only depends on this benefit-premium-bundle, but also on an insurer specific utility component ε_{ij}, which captures all the influences on the choice of an insurer that are independent of m and R, like, e.g., perceived friendliness of personnel, location or which insurer was recommended by family and friends; it may, however, also be unfounded and thus represent 'decision mistakes'. The utility of an individual i that is of risk type r and chooses insurer j therefore is

$$u_i(m^j, R^j) = p^r v(m^j) - R^j + \varepsilon_{ij}. \tag{3}$$

ε_{ij} is distributed according to the extreme value distribution with $Var(\varepsilon_{ij}) = \sigma_s^2 \frac{\pi^2}{6}$.[11] There are two unobservable σ-types $s = C, I$, with $\sigma_C < \sigma_I$; the share of the σ-type C is η. For individuals with a small variance of ε_{ij}, i.e. for σ-type C, all the additional utility components ε_{ij} are very similar and thus only have a small influence on which insurer is chosen. Individuals of σ-type C are therefore very responsive to differences in the contracts offered, i.e. 'careful' or 'conscientious' to choose the benefit-premium-bundle which provides the highest utility. On the other hand, individuals of σ-type I are rather 'insensible' to or 'ignorant' about (small) differences in the contracts; as for these individuals ε_{ij} assumes large (positive and negative) values, the additional utility components have a much larger impact on the decision of which insurer to choose than for the individuals of σ-type C.[12]

Table 1: Shares μ_{rs} of the four types of individuals; a positive correlation of low risk and high responsiveness is captured by $\delta > 0$.

	p^L	p^H	
σ_C	$\mu_{LC} = \lambda\eta + \delta$	$\mu_{HC} = (1-\lambda)\eta - \delta$	η
σ_I	$\mu_{LI} = \lambda(1-\eta) - \delta$	$\mu_{HI} = (1-\lambda)(1-\eta) + \delta$	$1-\eta$
	λ	$1-\lambda$	

The shares of the four types of individuals, μ_{rs}, are given in Table 1, where $\delta > 0$ captures the case of a positive correlation of low risk type and high responsiveness.[13] Increasing δ increases this positive correlation without altering the shares of the two risk types, λ and $(1-\lambda)$, and the shares of the two σ-types, η and $(1-\eta)$. Because these shares have an

[10] In this case the deductible would equal the (fixed) cost of treatment m^{FB} minus the indemnity, m, i.e. $D = m^{FB} - m$.

[11] Note that it is common to state the variance of ε_{ij} as a multiple of $\frac{\pi^2}{6}$ for the extreme value distribution, see Train (2009, p. 24).

[12] A low level of σ therefore corresponds to a high level of responsiveness, and a high level of σ to a low level of responsiveness.

[13] In the following, we will often use the term responsiveness instead of 'general responsiveness to differences in the contracts offered', and the term positive correlation instead of 'positive correlation of low risk type and high responsiveness'.

influence on the equilibrium, it is important to hold them constant when changing the level of correlation.

Denote by V_r^j the utility of contract j for an individual of risk type r without the additional utility component ε_{ij}:

$$V_r^j = p^r v(m^j) - R^j.$$

An individual i (being of risk type r) will choose an insurer k, if this insurer provides the highest level of overall utility, i.e. if

$$V_r^k + \varepsilon_{ik} > V_r^l + \varepsilon_{il} \quad \forall\, l \neq k. \tag{4}$$

With ε_{ij} distributed extreme value, it follows that the probability of individual i choosing insurer k is[14]

$$Prob(i \text{ chooses } k | i \text{ is of risk type } r \text{ and } \sigma\text{-type } s) = \frac{e^{\frac{V_r^k}{\sigma_s}}}{\sum_j e^{\frac{V_r^j}{\sigma_s}}}. \tag{5}$$

Denote this probability by P_{rs}^k; it is also insurer k's market share among the group of individuals of risk type r and σ-type s.

As it will turn out easier to provide an intuitive explanation of the results, we formulate insurer k's objective in terms of $\{m^k, V_L^k\}$ instead of $\{m^k, R^k\}$. Graphically, in m-R-space, insurer k chooses an indifference curve $\mathcal{I}^{V_L^k}$ for the L-types associated with the utility level V_L^k, and a benefit level m^k along this indifference curve. For $\{m^k, V_L^k\}$, the utility level of the H-type is

$$V_H^k = V_L^k + (p^H - p^L)v(m^k). \tag{6}$$

Normalizing the mass of individuals to one and assuming profit maximization, the objective of insurer k is

$$\max_{m^k, V_L^k} \pi^k = \sum_r \sum_s \mu_{rs} P_{rs}^k \pi_r^k, \tag{7}$$

where π_r^k denotes insurer k's profit per individual of risk type r, which is given by

$$\pi_r^k = p^L v(m^k) - V_L^k - p^r m^k. \tag{8}$$

2.2 The equilibrium with one risk type and one σ-type

We briefly discuss the equilibrium with only one risk type and one σ-type to introduce the concept of the 'indifference curves area' and to show the impact of different levels of σ.[15] With only one risk type and one σ-type, in this section we can skip the indices r and s and write insurer k's objective as

$$\max_{m^k, V^k} \pi^k = P^k \pi_i^k = \frac{e^{\frac{V^k}{\sigma}}}{\sum_j e^{\frac{V^j}{\sigma}}} \left(pv(m^k) - V^k - pm^k \right), \tag{9}$$

[14] See Train (2009, p. 40).

[15] This section is similar to Lorenz (2013).

where π_i^k is insurer k's profit per individual. Using the property of P^k that its derivative with respect to V^k can be expressed in terms of P^k itself in a simple way,

$$\frac{\partial P^k}{\partial V^k} = \frac{P^k(1-P^k)}{\sigma_s}, \tag{10}$$

the FOCs for insurer k's objective are

$$\frac{\partial \pi^k}{\partial m^k} = P^k\left[pv'(m^k) - p\right] = 0 \tag{11}$$

$$\frac{\partial \pi^k}{\partial V^k} = \frac{P^k(1-P^k)}{\sigma}\pi_i^k - P^k = 0. \tag{12}$$

Condition (11) yields $v'(m^k) = 1$, so m^k is chosen efficiently. Condition (12) requires $\pi_i^k = \frac{\sigma}{1-P^k}$. As it can be shown that the only equilibrium is a symmetric one, all insurers choose the same level of utility $V^j = V^* \ \forall j$. Since, in this case, $P^k = \frac{1}{n}$, in equilibrium profit per individual is

$$\pi_i^k = \frac{n}{n-1}\sigma, \tag{13}$$

and total profit per insurer is

$$\pi^k = \frac{\sigma}{n-1}. \tag{14}$$

As is to be expected, both, profit per individual and total profit per insurer increase in σ and decrease in n. If σ is small, offering a higher utility level yields a large increase in the share of individuals, because individuals are responsive even to small differences in contracts. This raises the incentive to offer a higher utility level, thereby reducing profits in equilibrium.

If n is large, each insurer's market share is small. Offering a higher utility level then attracts individuals from a large 'external' market share $1 - P^k$. This again raises the incentive to offer higher utility levels, lowering profits. We refer to this as the 'more competition due to a larger external market share'-effect. This effect plays an important role in the separating equilibrium.

2.3 Graphical representation of the equilibrium with one risk type and one σ-type

As P^k denotes the share of all individuals choosing insurer k and depends on V^k, it can be considered a distribution function $P^k(V^k)$. In equilibrium, when all the other insurers offer the same level of utility V^*, this distribution function is given by

$$P^k = P^k(V^k|\sigma, V^*) = \frac{e^{\frac{V^k}{\sigma}}}{e^{\frac{V^k}{\sigma}} + (n-1)e^{\frac{V^*}{\sigma}}}. \tag{15}$$

The shape of this distribution function and of the corresponding density $P^k(1-P^k)\frac{1}{\sigma}$ is shown in Figure 1 for two different values of σ; (ignore the curves labeled 'average' and the letters A and B at this point).

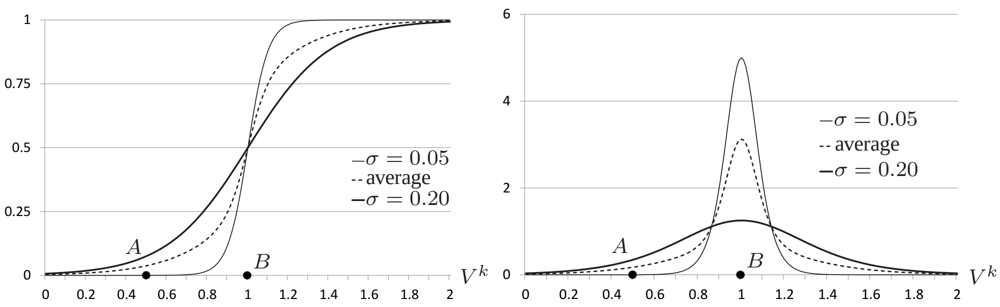

(a) Distribution function (b) Density function

Figure 1: Distribution function $P^k(V^k|\sigma, V^*)$ and density function $P^k(1-P^k)\frac{1}{\sigma}$ with $n = 2$ and $V^* = 1$ for $\sigma_C = 0.05$ and $\sigma_H = 0.20$ (solid curves); the dashed curve represents the average of the two functions for $\eta = 0.5$.

This distribution function can be depicted in the m-R-diagram that shows the equilibrium where all insurers offer $\{m^*, V^*\}$ by drawing a shaded area around the $\mathcal{I}^{V^*}$-indifference curve representing the corresponding density $P^k(1-P^k)\frac{1}{\sigma}$, see Figure 2; the different levels of darkness of this shaded area are a measure of the level of this density.[16]

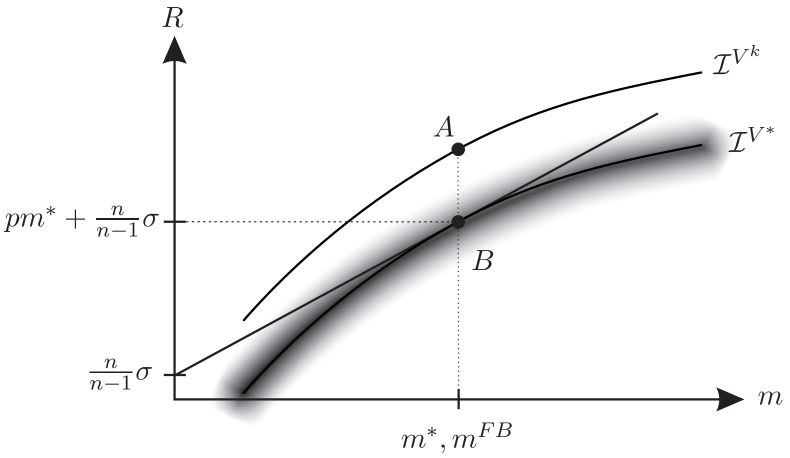

Figure 2: Equilibrium contract B if there is one risk type and one σ-type. The shaded area around the indifference curve $\mathcal{I}^{V^*}$ represents the density $P^k(1-P^k)\frac{1}{\sigma}$ of the distribution function P^k.

Above the shaded area, P^k and the corresponding density are zero. The density is also zero below the shaded area, where $P^k = 1$. In Figure 2, a contract like A (with V^k considerably below V^*) would therefore not attract any of the individuals; in Figure 1, this contract could

[16]As a technical detail, note that for $n = 2$, the maximum of this density is at $V^k = V^*$, but for $n > 2$, it is at $V^k > V^*$. Therefore the 'center' of the shaded area is at the $\mathcal{I}^{\tilde{V}}$-indifference curve for $n = 2$, and somewhat below it for $n > 2$. To simplify the exposition in the graphs, we will always draw the center of the shaded area at V^*.

be, e.g., at $V^k = 0.5$ (for $\sigma = 0.05$).[17]

For the following reason, this shaded area could be referred to as an 'indifference curves area': Consider the case that $n = 2$, so that there is only one other insurer l that offers V^*. Insurer k, to be chosen by individual i, has to offer a benefit-premium-bundle which yields utility

$$V^k > V^* + (\varepsilon_{il} - \varepsilon_{ik}).$$

For some individuals, $\varepsilon_{il} - \varepsilon_{ik} > 0$; the indifference curve insurer k must offer to make such an individual indifferent between the two insurers is somewhere below $\mathcal{I}^{V^*}$. On the other hand, for those with $\varepsilon_{il} - \varepsilon_{ik} < 0$, it suffices to offer an indifference curve above $\mathcal{I}^{V^*}$. From the perspective of insurer k, i.e. taking into account all the additional utility components ε_{ij}, the shaded area therefore also represents the whole set of the indifference curves of all individuals, or, an 'indifference curves area'.[18]

2.4 The impact of the level of σ on the equilibrium with one risk type and one σ-type

The equilibrium is affected by an increase of σ in two ways. First, the iso-profit line associated with the equilibrium contract is shifted upward as less responsive individuals allow insurers to charge a higher premium, see condition (13).

Secondly, the shaded area around the $\mathcal{I}^{V^*}$-indifference curve changes as follows: It is straightforward to show that the distribution function P^k as stated in (15) increases for $V^k < V^*$ and decreases for $V^k > V^*$; the corresponding density decreases around V^* and increases in the tails; see Figure 1, where the distribution and density function are drawn for $\sigma = 0.05$ and $\sigma = 0.20$. If σ increases, the distribution is spread out (over a wider range); in Figure 2, this can be depicted by a wider (and lighter) shaded area around the $\mathcal{I}^{V^*}$-indifference curve.

In the full model, each risk type consists of two unobservable σ-types, so from the perspective of an insurer, the average distribution and density functions are relevant. In Figure 1, such an average is shown for an equal share of both σ-types, i.e., for $\eta = 0.5$. If η increases, the distribution gets closer to the one of σ-type C, which, for the density, implies an increase at the mode and a decrease in the tails.

The changes in the shape of the average distribution and density functions are the reason why the equilibrium depends on the level of correlation. If the correlation increases, the share of the σ-type C increases for the L-types and decreases for the H-types. Therefore, the (average) density of the L-types increases at the mode and decreases in the tails, while for the H-types, the reverse holds: the density decreases at the mode and increases in the tails. We will now show how the equilibrium is affected by these changes, but before give an overview of the types of equilibria that can occur.

[17] Of course, strictly speaking, $P^k > 0 \ \forall \ V^k$, see (15), but above the shaded area, both P^k and the density $P^k(1 - P^k)\frac{1}{\sigma}$ are extremely small and almost equal to zero.

[18] If there is more than one other insurer, the argument is the same if ε_{il} is replaced by $max_{l \neq k} \varepsilon_{il}$.

2.5 Types of equilibria with two risk types

If there are two unobservable risk types (and one σ-type), three types of equilibria can be distinguished:[19] For low levels of σ, (i.e., a high degree of responsiveness), a separating equilibrium very similar to the Rothschild-Stiglitz-equilibrium emerges, where H-types receive the efficient level of m, while L-types receive a benefit level $m < m^{FB}$, so that the H-types are (about) indifferent between the two contracts.[20] For intermediate levels of σ, the separating equilibrium is of a different type as both benefit levels are distorted, i.e., $m^L < m^H < m^{FB}$. Finally, if the level of σ is high enough, a pooling equilibrium emerges, where all insurers offer the same benefit-premium-bundle.

We will analyze these three cases in turn: We begin with the pooling equilibrium in Section 3. We consider the separating equilibrium where only the benefit level of the L-types is at an inefficient level in Section 4.1; the separating equilibrium where both benefit levels are distorted is analyzed in Section 4.2.

3 The pooling equilibrium

In the full model, there are two risk types and two σ-types; the FOCs for this full model are stated as conditions (24) and (25) in Appendix A.1. As in this section we consider the pooling equilibrium where all insurers offer the same contract, in the following we will skip the index k. Also, to simplify the notation, we use $\widetilde{\sigma}$ to represent the harmonic mean of σ_C and σ_I:

$$\widetilde{\sigma} = \frac{1}{\frac{\eta}{\sigma_C} + \frac{1-\eta}{\sigma_I}}. \tag{16}$$

We first consider the case without correlation, i.e. $\delta = 0$.

3.1 The equilibrium without correlation

In equilibrium, individuals are distributed equally among the insurers, so that for all market shares we have $P_{rs}^k = \frac{1}{n}$. With $\delta = 0$, the FOC with respect to V_L, condition (24), then simplifies to

$$\lambda\pi_L + (1-\lambda)\pi_H = \frac{n}{n-1}\widetilde{\sigma}; \tag{17}$$

as before, (average) profit increases in $\widetilde{\sigma}$ and decreases in n. The FOC with respect to m, condition (25), can be written as

$$\lambda\left[p^L v'(m) - p^L\right] + (1-\lambda)\left[p^L v'(m) - p^H\right] + \left[(1-\lambda)\frac{n-1}{n\widetilde{\sigma}}(p^H - p^L)v'(m)\right]\pi_H = 0. \tag{18}$$

[19] See Lorenz (2013).

[20] See Zweifel et al. (2009), chapter 7.

Solving (17) for π_H by using $\pi_L = \pi_H + (p^H - p^L)m$, and substituting in (18) yields

$$\left[1 - \frac{n-1}{n\widetilde{\sigma}} \frac{(p^H - p^L)^2}{\overline{p}} \lambda(1-\lambda)m^*\right] v'(m^*) = 1. \tag{19}$$

Because the bracket is smaller than one, $v'(m^*) > 1$, i.e. $m^* < m^{FB}$. As is to be expected, the equilibrium level of m decreases in the difference of the probabilities, $p^H - p^L$. It also decreases in n and increases in $\widetilde{\sigma}$. The higher the average level of responsiveness (i.e. the lower $\widetilde{\sigma}$), the more distorted the benefit level is.

The equilibrium contract is shown in Figure 3: Each insurer offers utility V_L (depicted by the indifference curve $\mathcal{I}^{V_L}$) and medical services m^*, which determine the utility level for the H-types, V_H. Average profit per individual equals $\frac{n\widetilde{\sigma}}{n-1}$, represented by the distance of contract B to the pooling zero profit line $\bar{p}$.

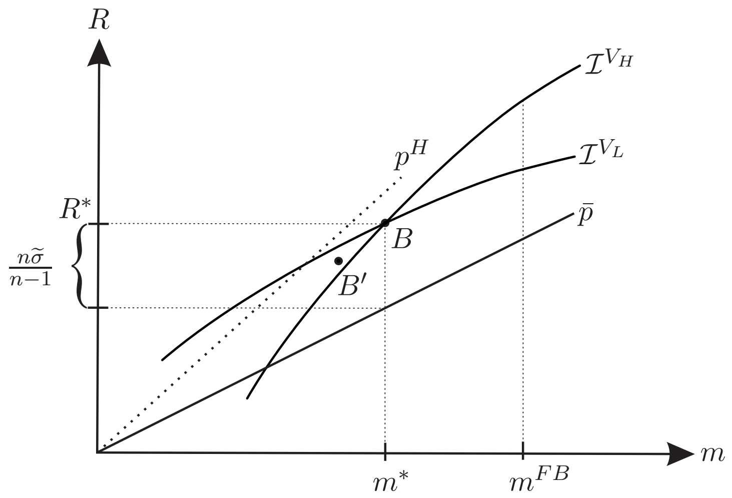

Figure 3: Equilibrium with $\pi_H < 0$. For larger levels of $\widetilde{\sigma}$, contract B is above the p^H-line; then $\pi_H > 0$.

Similar to Figure 2, one could draw the shaded areas around the two indifference curves $\mathcal{I}^{V_H}$ and $\mathcal{I}^{V_L}$, representing the density of the two risk types. It would then be straightforward to derive the demand responses of the two risk types if one of the insurers deviated from contract B. Moving, e.g., along the $\mathcal{I}^{V_H}$-indifference curve to the right would keep the share of H-types constant and decrease the share of L-types as indicated by the darkness of the shaded area around $\mathcal{I}^{V_L}$.

The shaded areas also make clear why it is not possible to destroy the pooling equilibrium by offering a contract like B', as would be the case under perfect competition, where B' would be chosen by all the L-types and none of the H-types. With a low degree of responsiveness, the shaded areas around both indifference curves are wide and overlap at B', so that an insurer offering B' instead of B would only attract a few more L-types and a few less H-types. Because profit per L-type, π_L, is considerably lower at B' than at B (as B' is well below the iso-profit-line for the L-types through B which has slope p^L), offering B' does not yield a higher total profit than B.

The economic forces determining the degree of the distortion of the equilibrium level of m are easiest to interpret using condition (18), where each of the three summands captures one of the forces: With $v'(m) > 1$, the first bracket, representing the L-types, is positive: Increasing m along the $\mathcal{I}^{V_L}$-indifference curve increases profit per L-type, because the indifference curve has a larger slope than the iso-π_L-line for $m < m^{FB}$.

The sign of the second bracket, representing the H-types, is indeterminate: increasing m (along $\mathcal{I}^{V_L}$) may either increase or decrease profit per H-type, depending on whether the $\mathcal{I}^{V_L}$-indifference curve at m^* has a larger slope than the iso-π_H-line.

The last term captures the effect that increasing m (along $\mathcal{I}^{V_L}$) increases the share of H-types choosing this insurer,[21] where this increase is given by the density of the H-types at contract B; in condition (18), this density is captured by the last bracket $[\cdot]$. Weighting these additional H-types by π_H then yields the effect on total profit.[22] For the equilibrium level of m, these three effects have to cancel out.

3.2 The dependence of the equilibrium on the level of correlation

Replacing μ_{rs} by the respective values as given in Table 1, the FOC with respect to V_L, solved for π_H yields

$$\pi_H = \frac{n}{n-1}\widetilde{\sigma} - \lambda(p^H - p^L)m - \delta\widetilde{\sigma}\left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right)(p^H - p^L)m, \tag{20}$$

and shows that the direct effect of an increase of the correlation (keeping m constant) is a decrease of π_H, and thereby also of π_L: If δ increases, the share of the σ_C-types among the L-types increases, so that the L-types become more responsive (on average). As it is the L-types, and not the H-types, insurers compete for, this creates an incentive to provide higher utility V_L, which reduces profits.

The FOC with respect to m can be simplified to

$$\begin{aligned}\lambda\left[p^L v'(m) - p^L\right] + (1-\lambda)\left[p^L v'(m) - p^H\right]& \\ + (1-\lambda)\frac{n-1}{n\widetilde{\sigma}}(p^H - p^L)v'(m)\left[1 - \frac{\delta}{1-\lambda}\widetilde{\sigma}\left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right)\right]\pi_H = 0,&\end{aligned} \tag{21}$$

where the second line, except for π_H, represents the density of the H-types. An increase of δ decreases this density: Because the H-types, for a positive correlation, consist of a larger share of σ_I-types which are less responsive, increasing m does not attract as many H-types as before.

The overall effect of an increase in δ therefore depends on how the product of the density of the H-types and the profit per H-type changes. For this overall effect we can state the following:

[21] Increasing m along the $\mathcal{I}^{V_L}$-indifference curve, of course keeps the number of L-types choosing this insurer constant.

[22] If $\pi_H > 0$, insurers like to have more H-types, so the third effect reduces the distortion; if $\pi_H < 0$, insurers try to avoid being chosen by the H-types, which increases the distortion.

Proposition 1. *In the pooling equilibrium,*

$$\frac{\partial m^*}{\partial \delta} \gtreqqless 0 \quad \textit{for} \quad \lambda \gtreqqless \frac{1}{2} + \frac{n\widetilde{\sigma}}{2(n-1)(p^H - p^L)m} - \delta\widetilde{\sigma}\left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right). \tag{22}$$

Proof. See Appendix A.2. □

Condition (22) defines a threshold level of λ below which m decreases in δ. Consider first the case that $\delta = 0$ (so that the last term in (22) can be ignored), i.e., that a positive (or negative) correlation is introduced beginning from no correlation. The threshold level then is larger than $\frac{1}{2}$; it is below 1 if $\widetilde{\sigma}$ is small, and above 1, if $\widetilde{\sigma}$ is large.

If the share of L-types is below the threshold, m decreases in δ. This is the case confirming the general notion: If low risks are more responsive, this increases the distortion. However, if λ is above the threshold, the reverse holds and m increases in δ. This case is at odds with the general notion, but it may in fact be the more relevant of the two cases for health insurance markets: If H-types represent the chronically ill, and L-types are those who have not yet developed this illness, then $1 - \lambda$ represents the prevalence rate: For most illnesses, the prevalence rate is small, so λ will be close to one.[23]

The following intuitive explanation can be given for why the relationship of the correlation and the distortion reverses if the share of L-types is large, but only, if $\widetilde{\sigma}$ is small: We showed that the distortion of m depends on the level of correlation only via the product of π_H and the density of the H-types, where the direct effect of an increase in δ is a decrease in π_H (see (20)), and a decrease of the density (see (21)). For the effect on the product of the two terms, two cases have to be distinguished:

If $\pi_H > 0$, the product of the two terms – which are both positive and decreasing – decreases. If $\pi_H < 0$, the product of the two terms is negative, and what is important then is how the two factors change in relative terms. If π_H decreased relatively little, but the density decreased relatively much, the product would increase (get closer to zero). This is exactly the case if the share of L-types is large and $\widetilde{\sigma}$ is small:

First, if λ is large, the difference in risk type specific profits, $\pi_L - \pi_H = (p^H - p^L)m$, is large, because m^* increases in λ for $\lambda > (1 + \sqrt{p^L/p^H})^{-1}$, see (19).[24] Secondly, if $\widetilde{\sigma}$ is small, both π_L and π_H are small, see condition (17). Therefore, the combination of a large level of λ and a low level of $\widetilde{\sigma}$ leads to the smallest profit per H-type (i.e., the largest loss); then the decrease of π_H due to an increase of δ is small in relative terms. In addition, if λ is large, the relative change of the density caused by an increase in δ is large: see the term $\frac{\delta}{1-\lambda}$ in the last brackets of (21).

So far we considered the case that $\delta = 0$. If $\delta > 0$, the threshold level of λ is smaller since $\left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right) > 0$. If δ increases (above zero), the share of σ-type C among the L-types increases: Because the L-types become more responsive, V_L^A increases, which decreases profits, so for a given level of λ, π_H is more negative. Therefore, with $\delta > 0$, λ can be somewhat smaller for π_H to be still negative enough so that the change of π_H due to an increase in δ is small in relative terms.

[23]E.g., for diabetes, $1 - \lambda$ is about 6% for most OECD countries; see OECD (2011).

[24]Note that λ also enters $\bar{p}$.

To sum up: If the share of L-types is large and the average responsiveness is high (i.e. $\widetilde{\sigma}$ is small), H-types entail a particularly large loss. If δ increases, the relative change of this loss is small, but the relative change of the responsiveness of the H-types is large. Increasing m then attracts a substantially lower share of H-types which incur only slightly higher losses, so in equilibrium, the distortion is reduced. This case seems to be the one most relevant for health insurance markets: the share of H-types is small, but each H-type entails a large loss.

3.3 Example

We will illustrate all results, both for the pooling and the separating equilibrium, with an example, for which we assume $n = 10$, $\Delta\sigma = \sigma_I - \sigma_C = 0.05$, $p^L = 0.2$, $p^H = 1$ and $v(m) = \ln(m)$, so that one of the risk types is chronically ill and $m^{FB} = 1$; varying $\widetilde{\sigma}$ then yields the three different types of equilibria.

For $\widetilde{\sigma} = 0.20$, a pooling equilibrium emerges. Assuming an equal share of both σ-types, i.e. $\eta = 0.5$, then requires $\sigma_C = 0.178$ and $\sigma_I = 0.228$.

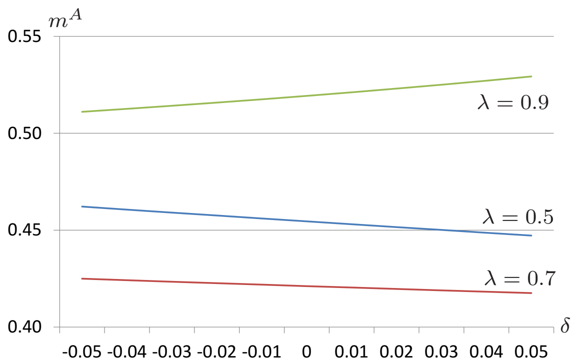

Figure 4: Example with $n = 10$, $\Delta\sigma = \sigma_I - \sigma_C = 0.05$, $p^L = 0.2$, $p^H = 1$ and $v(m) = \ln(m)$. m^A for $\widetilde{\sigma} = 0.20$ and $-0.05 \leq \delta \leq 0.05$.

Figure 4 shows the equilibrium levels of m as a function of δ for three different values of λ. For $\lambda = 0.5$ and $\lambda = 0.7$, an increase in the correlation reduces m, but for a high level of L-types ($\lambda = 0.9$), it increases m.[25] Table 2 shows the corresponding profits for L- and H-types: In all three cases, profit per L-type decreases in δ because insurers – competing more heavily for the L-types if they become more responsive – provide higher utility V_L by charging a lower premium. Whether H-types incur a profit or a loss depends on λ: For $\lambda = 0.5$, $\pi_H > 0$, so m decreases in δ. For $\lambda = 0.7$, $\pi_H < 0$, but is close to zero, so that the decrease of π_H due to an increase in δ is large in relative terms. Finally, for $\lambda = 0.9$, π_H is negative enough, so that its decrease is small in relative terms and more than compensated by the decrease in the density of the H-types.

[25]For $\lambda = 0.9$, δ is confined to the interval $[-0.05, 0.05]$, so that all $\mu_{rs} \geq 0$; this is why we present the results only for $-0.05 \leq \delta \leq 0.05$.

Table 2: Example with $n = 10$, $v(m) = \ln(m)$, $p^L = 0.2$, $p^H = 1$, $\eta = 0.5$, $\sigma_C = 0.1781$, $\sigma_I = 0.2281$, for different values of λ and δ.

λ	δ	m^*	π_L	π_H
0.5	0.00	0.455	0.4040	0.0404
0.5	0.05	0.447	0.3967	0.0389
0.7	0.00	0.421	0.3233	-0.0136
0.7	0.05	0.418	0.3183	-0.0157
0.9	0.00	0.519	0.2638	-0.1517
0.9	0.05	0.529	0.2594	-0.1641

4 The separating equilibrium

In the separating equilibrium, two types of insurers can be distinguished. We denote insurers offering a contract for the L-types as insurers of type A, and insurers offering a contract for the H-types as insurers of type B. The number of insurers is n^A and n^B respectively, with $n^A + n^B = n$. We denote insurer type by A and B, and not by L and H, because the contracts offered by insurers of type A, although designated for the L-types, may be chosen by both risk types; the same applies to the contracts offered by insurers of type B.

Because insurers of type B offer a contract designated for the H-types, we express their objective in terms of $\{V_H^B, m^B\}$ instead of $\{V_L^B, m^B\}$. The FOCs for the two types of insurers are explicitly stated as conditions (24)-(27) in Appendix A.1. In addition, as insurers can decide whether to be of type A or type B, the following profit equality condition,

$$\pi^A = \sum_r \sum_s \mu_{rs} P_{rs}^A \pi_r^A = \sum_r \sum_s \mu_{rs} P_{rs}^B \pi_r^B = \pi^B, \tag{23}$$

has to be satisfied. This condition implicitly defines n^A and n^B. However, since n^A and n^B have to be integer, it is only an approximation. As it is not important for deriving the results, we refrain from elaborating on a formula that determines whether n^A as given by (23) has to be rounded up or off.[26]

Because there is an intuitive graphical derivation for why the impact of a correlation of risk type and responsiveness in general is ambiguous for the separating equilibrium, we focus on this graphical derivation and only refer to the FOCs to confirm the results. We begin with the case of a low level of $\widetilde{\sigma}$, so that only the benefit level of the L-types is distorted (Section 4.1); the case of an intermediate level of $\widetilde{\sigma}$ where both benefit levels are distorted is analyzed in Section 4.2.

[26]Note that the requirement of n^A and n^B to be integer can, for some parameter settings, cause the non-existence of an equilibrium: For some values of n^A and n^B, it may be profitable for an insurer of type B to enter the market for the L-types and become an insurer of type A; but after the new 'equilibrium' has been attained, where π_i^A is decreased and π_i^B increased (because the increase of n^A increases competition among insurers of type A, while the decrease of n^B decreases competition among insurers of type B), the same insurer may then find it profitable to become of type B again. This problem of the existence of an equilibrium is discussed in greater detail in Lorenz (2013).

4.1 The separating equilibrium for a high level of responsiveness

4.1.1 The equilibrium without correlation

We first give a brief intuitive explanation for why, in Figure 5, contracts B and A_3 constitute the equilibrium.[27]

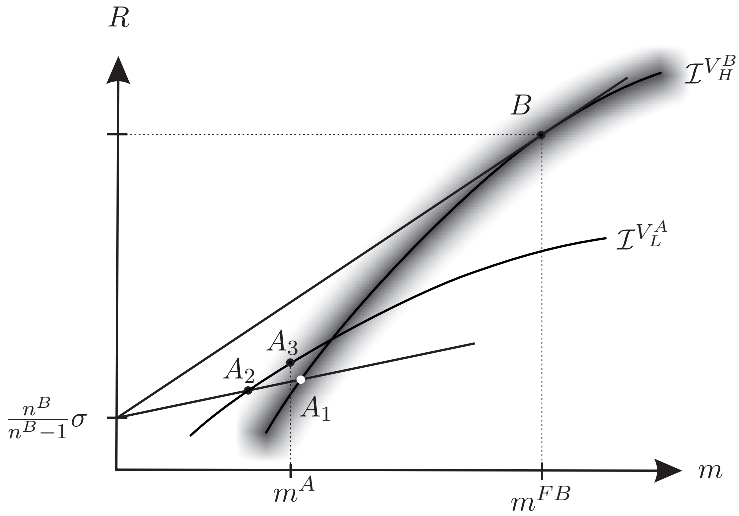

Figure 5: Separating equilibrium: contracts B and A_3 are offered.

With perfect competition, and applying the equilibrium concept of Rothschild and Stiglitz (1976), the equilibrium consists of contract B, chosen by the H-types, and contract A_1, chosen by the L-types (with both iso-profit lines shifted downward so that they pass through the origin).[28] However, as the shaded area around the $\mathcal{I}^{V_H^B}$-indifference curve shows, with less than perfect responsiveness, a considerable share of the H-types chooses contract A_1. Because these H-types, when choosing A_1, entail a large loss, insurers of type A have to shift their contract to the left.

Assume, that it is shifted to A_2, where (almost) none of the H-types choose this contract. But then an insurer of type A could move its contract along the $\mathcal{I}^{V_L^A}$-indifference curve to the right: This would leave the number of L-types choosing this insurer unaffected, but increase profits per L-type, π_L^A.[29] It would also increase the number of H-types choosing this insurer, but since the average density of the H-types is (almost) zero at contract A_2, at the boundary of the shaded area this effect is of second order. The third effect when moving along $\mathcal{I}^{V_L^A}$ is ambiguous: Depending on whether the slope of the $\mathcal{I}^{V_L^A}$-indifference curve is smaller or larger than the slope of the iso-profit lines for the H-types, this will increase or decrease profit per H-type, π_H^A.

[27] This section is again similar to Lorenz (2013).

[28] See Zweifel et al. (2009), chapter 7. In Figure 5, we assume $\lambda = 0.5$, so that $n^A = n^B$. For $n^A \neq n^B$, the iso-profit lines start at different points on the ordinate.

[29] This is because the slope of the $\mathcal{I}^{V_L^A}$-indifference curve is larger than the slope of the iso-π_L^A-lines for all contracts with $m^A < m^{FB}$.

Insurers of type A will move their contract to the right until these three effects – the increase of π_L^A, the increase of P_H^A, and the change of π_H^A – cancel out, which will be at a contract as indicated by A_3.

Therefore, in equilibrium, a small share of the H-types chooses the contracted designated for the L-types. This contrasts with the contract offered by insurers of type B: As contract B is far away from the shaded area that can be drawn around the $\mathcal{I}^{V_L^A}$-indifference curve, none of the L-types choose contract B. As there is no interference of the L-types, contract B is at the efficient level, as in the case of perfect competition.

4.1.2 The dependence of the equilibrium on the level of correlation

We will now analyze how this equilibrium is affected by a correlation of risk type and σ-type. We will first show that an increase in the correlation leads to a 'shift' of the indifference curves $\mathcal{I}^{V_L^A}$ and $\mathcal{I}^{V_H^B}$, and then how these shifts together with the change of the shaded area around $\mathcal{I}^{V_H^B}$ affect the distortion of m^A.

We begin with the insurers of type B, which are chosen by the HC- and HI-types. Replacing μ_{rs} in the FOCs (24)-(27) by the respective values as given in Table 1, it can first be shown that for the FOC with respect to m^B the effects of δ cancel out, so that $v'(m^B) = 1$ still holds.[30] As contract B is chosen only by the H-types, there is no reason to offer an inefficient benefit level m^B, irrespective of the level of correlation.

We next consider the direct effect of a (positive) correlation on V_H^B:

As is apparent from the shapes of the two densities which constitute the shaded area around $\mathcal{I}^{V_H^B}$, the H-types choosing insurer A consist primarily of σ-type I, so $P_{HI}^A > P_{HC}^A$; see Figure 1, where contract A_3 could be, say, at 0.5. Those for which the additional utility component plays a larger role when choosing an insurer are the ones who have a higher probability of choosing the 'wrong' benefit package. Insurers of type B therefore lose a larger share of the HI-types than of the HC-types to insurers of type A, so $P_{HC}^B > P_{HI}^B$.

The 'number' of individuals of the two σ-types choosing an insurer of type B are $\mu_{HC} P_{HC}^B$ and $\mu_{HI} P_{HI}^B$. If δ increases, μ_{HC} decreases and μ_{HI} increases. The direct effects of these changes, holding P_{HF}^B and P_{HS}^B constant, can be found in Table 3, which is to be read as follows:

For insurers of type B, the number of HC-types decreases and the number of HI-types increases; the total number of H-types decreases, and the share of the σ_I-types among the H-types increases (see column two of Table 3).[31]

These changes have the following direct effects on V_H^B: First, the decrease in the total number of H-types choosing an insurer of type B creates an incentive to increase V_H^B: This is the "more competition due to a larger external market share"-effect: The larger the share of individuals who have chosen another insurer (in this case, an insurer of type A),

[30] See condition (31) in Appendix A.3.

[31] The total number of H-types decreases because $\Delta H = P_{HC}^B \Delta\mu_{HC} + P_{HI}^B \Delta\mu_{HI} = -P_{HC}^B \Delta\delta + P_{HI}^B \Delta\delta < 0$ for $P_{HC}^B > P_{HI}^B$.

Table 3: Direct effect of positive correlation holding P^k_{rs} constant. The first entry is to be read as $\frac{\partial}{\partial\delta}(\mu_{HC}\overline{P}^B_{HC}) < 0$; the next entry in the same row as $\frac{\partial}{\partial\delta}(\mu_{LC}\overline{P}^A_{LC}) > 0$, and so on.

	insurer B		insurer A				
	H-types		L-types		H-types		
C	$-$		$+$		$-$		
I	$+$		$-$		$+$		
$C+I$	$-$	$V^B_H \nearrow$	0		$+$	$V^A_L \searrow$	$m^A \searrow$
$\frac{I}{C+I}$	$+$	$V^B_H \searrow$	$-$	$V^A_L \nearrow$	$+/0$	$V^A_L(\searrow)$	$m^A(\searrow)$

the larger the demand response when V^B_H is increased. Secondly, the increase of the share of σ_I-types among the H-types decreases the density of the H-types at contract B, which creates an incentive to decrease V^B_H: As the H-types become less responsive, this reduces the incentive to provide a higher utility.

In general, the aggregate of these two effects is indeterminate. However, the first effect can only be important if the number of insurers of type B is small. If it was large, each insurer of type B would have a small market share, so the external market share would be close to one. In this case, any increase in the external market share would be small, and therefore also the "more competition due to a larger external market share"-effect, which would then be dominated by the second effect.

Lemma 1. *If n^B is large enough, the direct effect of an increase in δ is a decrease of V^B_H.*

Proof. See Appendix A.4. □

Consider now the insurers of type A, which are chosen both by the L-types and the H-types. The direct effects of the changes in μ_{rs} due to the increase of δ, holding P^A_{rs} constant, can be found in the third and fourth column of Table 3. As can be seen from column three, the number of LC-types increases and the number of LI-types decreases; with $P^A_{LC} = P^A_{LI}$, the total number of L-types is unaffected. The only effect on V^A_L caused by the L-types is due to the lower share of σ-type I, which increases the density of the L-types at contract A: As the L-types become more responsive, the incentive to offer a higher utility increases, so V^A_L is increased.

Insurers of type A are, however, also chosen by the H-types. The last column of Table 3 shows that the number of HC-types decreases and the number of HI-types increases. The total number of H-types increases, as does the share of the σ-type I (unless $P^A_{HC} = 0$, in which case it stays constant); this raises the average density of the H-types at contract A, which creates an incentive to reduce V^A_L. However, this second effect is small and dominated by the first effect if P^A_{HI} (and therefore also P^A_{HC}) is small; then an increase of μ_{HI} increases the total number of H-types choosing an insurer of type A only to a small degree.

Lemma 2. *If P^A_{HI} is small enough, the direct effect of an increase in δ is a decrease of V^A_L.*

Proof. See Appendix A.5. □

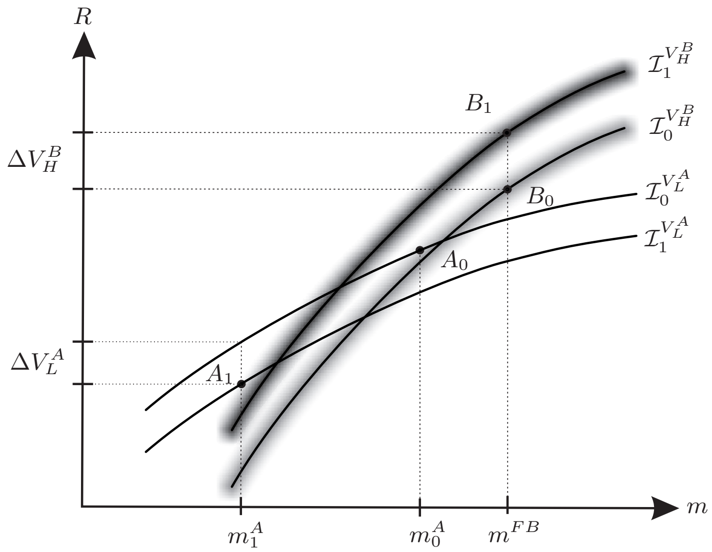

Figure 6: Equilibrium before (B_0, A_0) and after (B_1, A_1) an increase of the correlation of risk type and responsiveness.

If n^B is large enough and P_{HI}^A is small enough so that the main effects dominate, V_H^B is decreased (because the H-types become less responsive) and V_L^A is increased (because the L-types become more responsive). These changes in V_H^B and V_L^A can be depicted by an upward shift of the $\mathcal{I}^{V_H^B}$-indifference curve and a downward shift of the $\mathcal{I}^{V_L^A}$-indifference curve; see Figure 6, where rather narrow shaded areas have been drawn so that these shaded areas do not overlap. Clearly, if m^A did not change, a higher share of the H-types would choose the insurers of type A. This creates an incentive to reduce m^A.

In addition, there is a second effect on m^A that is independent of the changes of V_H^B and V_L^A: As δ increases, the average density of the H-types increases in the tails. This increase of the density of the H-types at contract A creates a second incentive to reduce m^A. In Figure 6, this is reflected by the darker boundaries of the upper of the two shaded areas, and by a distance to the indifference curves of the H-types that is larger for contract A_1 than for A_0.

Proposition 2. *If the number of insurers offering a contract for the H-types is large enough and the share of H-types choosing a contract designated for the L-types is small enough, then m^A decreases if the correlation of low risk type and high responsiveness increases:*

$\frac{\partial m^A}{\partial \delta} < 0$ *for P_{Hs}^A small enough and n^B large enough.*

Proof. See Appendix A.6. ☐

We comment on the implications of this result in the discussion section.

If the conditions of Proposition 2 are not satisfied, m^A may decrease in δ, as we now show with the second of the following two examples.

4.1.3 Example

We now illustrate the results of Proposition 2 using the example introduced in Section 3.3. Here we assume a low level of $\widetilde{\sigma}$, so that the separating equilibrium emerges. With $\widetilde{\sigma} = 0.05$, we have $\sigma_C = 0.0354$ and $\sigma_I = 0.0854$. The equilibrium values of m^A are shown in Figure 7; the equilibrium values of some of the other variables can be found in Table 4, where we show the results always for the lowest level of δ for which each combination of n^A and n^B occurs.[32] For $n^A = n^B = 5$, we also present the results for a few more values of δ so that the effect on V_L^A, V_H^B and the market shares can be seen. In addition, the Rothschild-Stiglitz-equilibrium is given in the first row.

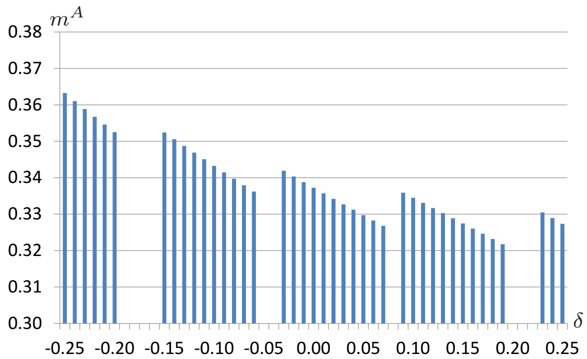

Figure 7: m^A for $-0.25 \leq \delta \leq 0.25$.

For all levels of δ, m^B is at the efficient level $m^{FB} = 1$.[33] As can be seen from Figure 7, m^A in general decreases in δ, but there are a few upward jumps, which occur whenever n^A decreases by one.[34] If we compare the different levels of m^A from Figure 7 with the Rothschild-Stiglitz-equilibrium where $m^A = 0.398$, we see that depending on the level of δ, the distance of m^A to its level under perfect competition can differ by a factor of two.[35] The distortion of m^A clearly depends on the level of correlation of risk type and responsiveness.

For δ increasing from -0.03 to 0.07, we see the opposite effects on V_L^A and V_H^B: While insurers of type A increase V_L^A from -0.363 to -0.358, insurers of type B decreases V_B^H from -1.058 to -1.068. Accordingly, profits for insurers of type A decrease from 0.0068 to 0.0055, while profits for insurers of type B increase from 0.0054 to 0.064. If δ increases above 0.07, there exists an incentive for one of the insurers of type A to offer the contract designated for the H-types and become an insurer of type B; therefore, n^B increases.

Comparing the utility levels provided by the two insurers, we see that the difference for the L-types, $V_L^A - V_L^B$, is much larger than the difference for the H-types, $V_H^B - V_H^A$. Accordingly, none of the L-types choose contract B ($P_{Ls}^B = 0$), but some of the H-types choose contract A, where the share of the σ-type I is larger than the share of the σ-type C

[32] E.g., the smallest level of δ so that $n^A = 6$ is $\delta = -0.15$.

[33] We therefore omit m^B from Figure 7 and Table 4.

[34] For those values of δ for which an entry is missing in Figure 7, an equilibrium does not exist, see footnote 26. If we determined n^A and n^B according to formula (23), i.e. as a real instead of an integer number, all equilibria exist and m^A would be strictly decreasing in δ for all levels of δ.

[35] $\Delta m^A(\delta = 0.19)/\Delta m^A(\delta = -0.25) = (0.398 - 0.322)/(0.398 - 0.363) = 2.17$.

Table 4: Example with $n = 10$, $v(m) = \ln(m)$, $p^L = 0.2$, $p^H = 1$, $\lambda = 0.5$, $\eta = 0.5$, $\sigma_C = 0.0354$, $\sigma_I = 0.0854$, for different values of δ. The first row contains the Rothschild-Stiglitz-equilibrium (RS).

δ	n^A	n^B	m^A	V_L^A	V_L^B	V_H^A	V_H^B	P_{HC}^A	P_{HI}^A	P_{Ls}^B	π^A	π^B
RS	-	-	.398	-.264	-1.00	-1.00	-1.00	-	-	-	-	-
-.25	7	3	.363	-.393	-1.05	-1.20	-1.052	.0045	.041	.00	.0080	.0084
-.15	6	4	.352	-.375	-1.05	-1.21	-1.052	.0028	.032	.00	.0072	.0062
-.03	5	5	.342	-.363	-1.06	-1.22	-1.058	.0019	.026	.00	.0068	.0054
.00	5	5	.337	-.361	-1.06	-1.23	-1.060	.0016	.024	.00	.0064	.0057
.03	5	5	.333	-.359	-1.06	-1.24	-1.063	.0014	.023	.00	.0060	.0059
.07	5	5	.327	-.358	-1.07	-1.25	-1.068	.0011	.021	.00	.0055	.0064
.09	4	6	.336	-.356	-1.07	-1.23	-1.068	.0018	.023	.00	.0072	.0053
.23	3	7	.330	-.354	-1.09	-1.24	-1.093	.0022	.024	.00	.0088	.0062

($P_{HI}^A > P_{HC}^A > 0$); e.g., for $\delta = 0.00$, about 12% of the HI-types (2.4% for each of the five insurers), but only about 0.8% of the HC-types choose an insurer of type A.

We now alter two of the parameters of the example to show that the main result may be reversed, i.e., that m^A may increase in δ. In Figure 8, we depict the increase in m^A for an increase of δ from 0.00 to 0.01, for different levels of λ (the different curves) and η (on the abscissa). For most levels of λ, m^A decreases in δ irrespective of the level of η. However, if both λ and η are high (e.g. $\lambda = 0.9$ and $\eta = 0.9$), m^A increases in δ. If λ is large, n^B is small (because there are not many H-types). In addition, if η is large, the share of the σ_I-types is small; then, if δ increases, the increase of the HI-types (who are the ones choosing an insurer of type A) is large in relative terms. Therefore, if both λ and η are large, there is a particularly strong 'more competition due to a larger external market share'-effect, so that V_H^B increases, which leads to a decrease of m^A.

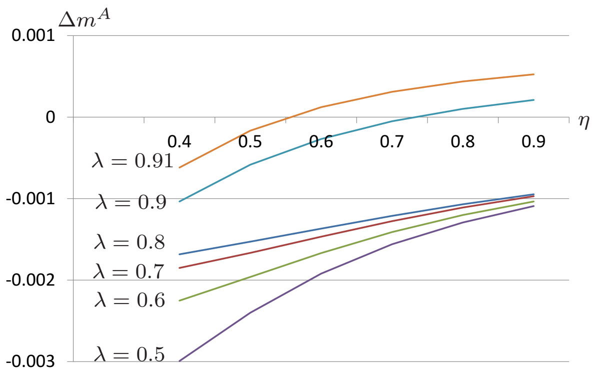

Figure 8: Increase of m^A for an increase of δ from 0.00 to 0.01 for different levels of λ and η, i.e., $\Delta m^A = m^A(\delta = 0.01|\lambda, \eta) - m^A(\delta = 0.00|\lambda, \eta)$.

4.2 The separating equilibrium for an intermediate level responsiveness

4.2.1 The equilibrium without correlation

In the previous section the separating equilibrium for a low level of $\widetilde{\sigma}$, where only the contract for the L-types is distorted, has been analyzed. If $\widetilde{\sigma}$ increases, at some point insurers of type B begin to distort the contract designated for the H-types, and if $\widetilde{\sigma}$ is large (but still below the level for which the pooling equilibrium emerges), the distortion of m^B can be substantial.

The reason for the distortion of m^B is the following. For a low level of $\widetilde{\sigma}$, the shaded area around the indifference curve of the L-types is so narrow that contract B_0 is far away from it. If $\widetilde{\sigma}$ increases, the distribution functions P^k_{rs} are spread out, so that the shaded area around $\mathcal{I}^{V^A_L}$ becomes wider; at some point, it 'reaches' contract B_0, and if $\widetilde{\sigma}$ increases further, contract B will be inside this shaded area, see Figure 9.

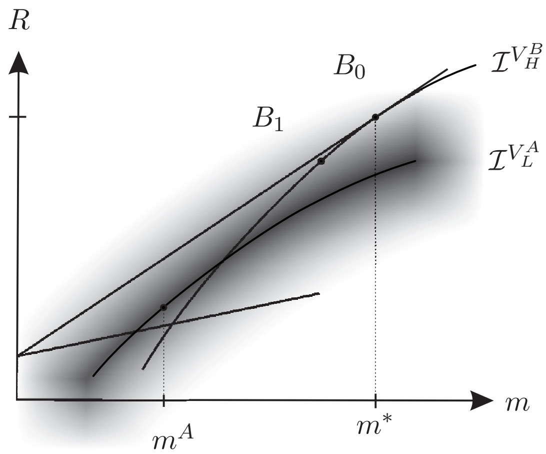

Figure 9: Separating equilibrium with two unobservable risk types; $\widetilde{\sigma}$ large: Contract B distorted from B_0 to B_1.

If contract B_0 is inside the shaded area around $\mathcal{I}^{V^A_L}$, an insurer of type B can increase its profit by moving its contract along the $\mathcal{I}^{V^B_H}$-indifference curve to the left: This leaves the share of H-types choosing this insurer unaffected, but increases the share of L-types.[36] Of course, moving along the $\mathcal{I}^{V^B_H}$-indifference curve also reduces profit per H-type, π^B_H, but initially, as m is close to m^{FB}, this effect is of second order. Contract B_0 will therefore be shifted to the left, until these three effects cancel out.

4.2.2 The dependence of the equilibrium on the level of correlation

It follows immediately how the distortion of m^B is affected by a positive correlation of risk type and responsiveness if the two main effects described in the previous section dominate,

[36]This is because the slope of the $\mathcal{I}^{V^B_H}$-indifference curve is larger than the slope of the iso-P^B_{Ls}-curves, which are identical to the slope of the $\mathcal{I}^{V^A_L}$-indifference curve.

so that V_H^B is decreased and V_L^A is increased: Both the upward shift of $\mathcal{I}^{V_H^B}$ and the downward shift of $\mathcal{I}^{V_L^A}$ move contract B closer to the boundary of the shaded area around $\mathcal{I}^{V_L^A}$: This decreases the density of the L-types at contract B, so the incentive to distort m^B is reduced. With the lower density of the L-types, insurer B now attracts fewer L-types when moving contract B along $\mathcal{I}^{V_H^B}$ to the left, so that the countervailing effect of the decrease in profits for the H-types will stop the distortion when m is closer to m^{FB}.

Like in the previous Section 4.1.2, there is an additional effect that reduces the distortion of m^B, irrespective of the change of V_H^B and V_L^A. As shown above, if δ increases, the shape of the density represented by the shaded area around $\mathcal{I}^{V_L^A}$ increases at the mode and decreases in the tails. Because the density of the L-types decreases at contract B, the incentive to distort m^B is reduced again. We can therefore state the following result:

Proposition 3. *In the separating equilibrium for intermediate levels of $\widetilde{\sigma}$ where m^B is below the efficient level, m^B increases in the correlation if the main effects dominate so that V_H^B decreases and V_L^A increases:*

If $m^B < m^{FB}$, then $\frac{\partial m^B}{\partial \delta} > 0$ for $\frac{\partial V_L^A}{\partial \delta} > 0$ and $\frac{\partial V_H^B}{\partial \delta} < 0$.

4.2.3 Example

We illustrate this result using the same example as before. Here we set $\widetilde{\sigma} = 0.17$, which requires $\sigma_C = 0.1486$ and $\sigma_I = 0.1986$. We refrain from presenting the results as detailed as in Table 4, but only show the levels of m^A and m^B in Figure 10.

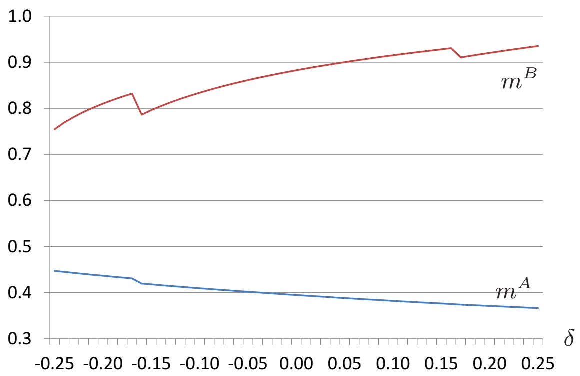

Figure 10: m^A and m^B for $\sigma_C = 0.1486$, $\sigma_I = 0.1986$ and $-0.25 \leq \delta \leq 0.25$.

Again, the downward jumps of m^B occur for all levels of δ, for which an increase in n^B emerges.[37] As can be seen, there is a small and steady decrease of m^A, but a considerable increase of m^B from 0.755 to 0.935 as δ increases from -0.25 to 0.25. We comment on the implications of this result in the following section.

[37] In this example, the equilibrium exists for all levels of δ, so there is no entry missing in Figure 10.

5 Discussion

5.1 Implications of the results

The purpose of this paper is to show that the distortion of the benefit levels caused by adverse selection not only depend on the shares of risk types and the average level of responsiveness (via the type of equilibrium that emerges), but also on the level of correlation of risk type and responsiveness.

We think that the results derived have three main implications. The first one regards the policy intervention of making individuals more responsive by, e.g., providing easy access to information about the contracts offered by insurers. One example of such a policy intervention are the government run websites about the Health Insurance Exchanges in the U.S. We do not question the usefulness of these websites, but want to point to a side effect such information provision may have. For the pooling equilibrium we showed that the effect of such a policy intervention not only depends on how the average responsiveness changes, but also on which of the two risk types is affected more. Interestingly, for the parameter setting which is likely to be the most relevant one for health insurance markets – a small share of H-types which incur a large negative profit – increasing the responsiveness of the H-types (e.g., the chronically ill) has the most detrimental effect on the distortion, and therefore on welfare. This shows that an information campaign about the quality of insurers' benefit packages that is specifically targeted at the chronically ill would have the largest negative side effect on welfare.

The second implication regards the separating equilibrium for a high level of average responsiveness, which is similar but not identical to the equilibrium under perfect competition. Some studies have used the sets of contracts offered in separating equilibria to estimate individuals' risk aversion (the curvature of the indifference curves). Our analysis shows that with less than perfect responsiveness, the 'average' H-type – i.e., average with respect to the additional utility components – is not indifferent between the two contracts. Instead, the contract for the L-type is located to the left of the indifference curve of the H-type (in the benefit-premium-diagram). This distance between the contract and the indifference curve could be considered a 'measurement error'. It would lead to a downward biased estimate of the level of risk aversion, as the indifference curve for the H-type appears to have a lower curvature than it actually has. Our analysis shows that the size of this measurement error is largest for a positive correlation.[38] It may therefore be important to take this measurement error into account when estimating the level of risk aversion.

The third implication regards the separating equilibrium for an intermediate level of responsiveness where both benefit levels are distorted. For such a separating equilibrium, even the benefit level of the H-types depends on the level of correlation of risk type and responsiveness; as the example given in Section 4.2.3 showed, the impact of the correlation on this benefit level can be substantial. The contract for the H-types may therefore be a (severely) biased indicator of the efficient benefit level. If this bias is large, this may have to be taken

[38] In the example of Section 4.1.3 we found a factor of two when comparing the largest and the smallest distance of m^A to the respective level under perfect competition.

into account when estimating the overall welfare effects of adverse selection. In fact, recent empirical studies have found surprisingly low estimates of the welfare consequences of adverse selection caused by inefficient pricing of a *given* set of benefit packages.[39] However, as explicitly stated by Einav, Finkelstein, and Levin (2010), the welfare losses due to an inefficient set of benefit packages may be much larger, and our analysis shows that the level of these additional welfare losses depends on the correlation of risk type and responsiveness.

5.2 Robustness of results

In the discrete choice model we analyzed, we made a number of simplifying assumptions. Several of these assumptions have been discussed in Lorenz (2013), to which we refer. There it has been shown that the model can also be applied if insurers can offer more than one contract or if the premium is set by a regulator and insurers offer multidimensional benefit packages; there it has also been discussed why the IIA-property does not cause a problem for this model, so that the conditional logit is more appropriate than the nested logit model, and that the equilibrium is very similar for other distributional assumptions for ε_{ij} than the extreme value distribution.

Here, we discuss the simplifying assumption that there are only two types of responsiveness to differences in the benefit-premium-bundles offered. We think that the results are very similar if there are more than two σ-types or if there is a continuous distribution of σ-types: In this case, the expression for $\widetilde{\sigma} = \left(\frac{\eta}{\sigma_C} + \frac{1-\eta}{\sigma_I}\right)^{-1}$ would have to be augmented by the additional risk types or replaced by the respective integral terms, but the main results should still hold. For the separating equilibrium, the impact of a correlation depends on how this affects the utility levels (V_L^A and V_H^B) and the shape of the shaded areas around the indifference curves: A positive correlation will increase the responsiveness of the L-types and decrease the responsiveness of the H-types also in this more general case; then, as long as a positive correlation raises the density close to the boundaries of the shaded area around the indifference curve of the H-types and lowers it for the shaded area around the indifference curve of the L-types, the distortion of the contract for the L-types should increase, and (if a distortion existed) decrease for the H-types. For the pooling equilibrium it was shown that the impact of a correlation depends on whether the share of low risks is large and $\widetilde{\sigma}$ is small, so that the H-types incur a large negative profit. As the effect of a (positive) correlation only depends on the product of this negative profit and the density of the H-types, where the relative changes of these two factors mainly depend on the share of L-types and the average level of responsiveness, the result for the pooling equilibrium should also not be affected by the number of σ-types.

6 Conclusion

We have analyzed how the distortions of (health) insurers' benefit levels caused by adverse selection are influenced by a correlation of risk type and responsiveness to differences in the

[39] See Einav, Finkelstein, and Cullen (2010), Bundorf et al. (2012) and Handel (2013).

contracts offered by insurers. Within a discrete choice model which endogenizes whether a separating or a pooling equilibrium emerges, we showed the following main results: For the pooling equilibrium, the effect of such a correlation depends on the share of low risks and the average level of responsiveness. If the share of low risks is small, a positive correlation increases the distortion; if the share of low risks is large, the reverse holds, but only if the average level of responsiveness is high. For the separating equilibrium the effects of a positive correlation are ambiguous. If the main effects dominate, i.e., if the share of H-types choosing the contract designated for the L-types is small enough and the number of insurers offering the contract for the H-types is large enough, a positive correlation will increase the distortion of the benefit level for the L-types, while the distortion of the benefit level of the H-types – if such a distortion exists – will decrease.

Regarding policy implications we discussed that selectively increasing the responsiveness of the H-types (by, e.g., providing information about insurers' offers only to the chronically ill) has the largest welfare-decreasing effect in the pooling equilibrium. For the separating equilibrium we showed that it may be important to take into account that the benefit levels of the two contracts are affected by the level of correlation of risk type and responsiveness when estimating the degree of risk aversion or the welfare effects of adverse selection.

A Appendix

A.1 First order conditions for the full model

The FOCs in the full model with two risk types and two σ-types are:

$$\frac{\partial \pi^k}{\partial V_L^k} = \sum_r \sum_s \mu_{rs} \left[\frac{P_{rs}^k (1 - P_{rs}^k)}{\sigma_s} \pi_r^k - P_{rs}^k \right] = 0 \tag{24}$$

$$\begin{aligned} \frac{\partial \pi^k}{\partial m^k} &= \sum_r \sum_s \mu_{rs} P_{rs}^k \left[p^L v'(m^k) - p^r \right] \\ &+ \sum_s \left[\mu_{Hs} \frac{P_{Hs}^k (1 - P_{Hs}^k)}{\sigma_s} (p^H - p^L) v'(m^k) \pi_H^k \right] = 0. \end{aligned} \tag{25}$$

For the separating equilibrium where we distinguish insurers of type A and type B, in (24) and (25), k has to be replaced by A; in addition, the FOCs for insurers of type B are:

$$\frac{\partial \pi^B}{\partial V_H^B} = \sum_r \sum_s \mu_{rs} \left[\frac{P_{rs}^B (1 - P_{rs}^B)}{\sigma_s} \pi_r^B - P_{rs}^B \right] = 0 \tag{26}$$

$$\begin{aligned} \frac{\partial \pi^B}{\partial m^B} &= \sum_s \mu_{Ls} P_{Ls}^B \left[p^H v'(m^B) - p^L \right] - \sum_s \mu_{Ls} \frac{P_{Ls}^B (1 - P_{Ls}^B)}{\sigma_s} (p^H - p^L) v'(m^B) \pi_L^B \\ &+ \sum_s \mu_{Hs} P_{Hs}^B \left[p^H v'(m^B) - p^H \right]. \end{aligned} \tag{27}$$

A.2 Proof of Proposition 1

The only terms in (21) containing δ are the last two factors, i.e., $\left[1 - \frac{\delta}{1-\lambda} \widetilde{\sigma} \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I} \right) \right] \pi_H$. Substituting (20) for π_H, this expression equals

$$\begin{aligned} &\frac{n\widetilde{\sigma}}{n-1} - \lambda \left(p^H - p^L \right) m + \frac{\delta^2 \widetilde{\sigma}^2}{1-\lambda} \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I} \right)^2 (p^H - p^L) m \\ &\qquad + \frac{\delta \widetilde{\sigma}}{1-\lambda} \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I} \right) \left[(2\lambda - 1)(p^H - p^L) m - \frac{n\widetilde{\sigma}}{n-1} \right]. \end{aligned} \tag{28}$$

Taking the derivative of (28) with respect to δ and dividing by $\frac{\widetilde{\sigma}}{1-\lambda} \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I} \right)$ yields

$$(2\lambda - 1)(p^H - p^L) m - \frac{n\widetilde{\sigma}}{n-1} + 2\delta\widetilde{\sigma} \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I} \right) (p^H - p^L) m. \tag{29}$$

This expression is $\gtreqless$ zero for

$$\lambda \gtreqless \frac{1}{2} + \frac{n\widetilde{\sigma}}{2(n-1)(p^H - p^L) m} - \delta\widetilde{\sigma} \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I} \right). \tag{30}$$

Therefore, (28) decreases in δ if λ is smaller than the right hand side of (30); then $v'(m)$ has to increase, i.e. m has to decrease so that (21) is still satisfied.

A.3 FOC of insurers of type B with respect to m^B if $\widetilde{\sigma}$ is small

With $P^B_{LC} = 0$ and $P^B_{LI} = 0$, condition (27) simplifies to

$$\left[\mu_{HC} P^B_{HC} + \mu_{HI} P^B_{HI}\right] \left[p^H v'(m^B) - p^H\right]. \tag{31}$$

Therefore, $v'(m^B) = 1$, irrespective of μ_{HC} and μ_{HI}, and therefore irrespective of δ.

A.4 Direct effect of an increase in δ on π^B_H

Using the $P^B_{LC} = 0$ and $P^B_{LI} = 0$ and replacing μ_{rs} by the respective values of Table 1, condition (26) solved for π^B_H yields:

$$\pi^B_H = \frac{(1-\lambda)\left[\eta P^B_{HC} + (1-\eta) P^B_{HI}\right] + \delta(P^B_{HI} - P^B_{HC})}{(1-\lambda)\left[\frac{\eta P^B_{HC}(1-P^B_{HC})}{\sigma_C} + \frac{(1-\eta)P^B_{HI}(1-P^B_{HI})}{\sigma_I}\right] + \delta\left[\frac{P^B_{HI}(1-P^B_{HI})}{\sigma_I} - \frac{P^B_{HC}(1-P^B_{HC})}{\sigma_C}\right]}. \tag{32}$$

Taking the derivative with respect to δ we have

$$\frac{\partial \pi^B_H}{\partial \delta} = \frac{(1-\lambda) P^B_{HI} P^B_{HC}\left(\frac{1-P^B_{HC}}{\sigma_C} - \frac{1-P^B_{HI}}{\sigma_I}\right)}{D^2}, \tag{33}$$

where D is the denominator of (32). Therefore, the direct effect of an increase in δ is an increase in π^B_H if n^B is large enough so that both $1 - P^B_{HC}$ and $1 - P^B_{HI}$ are close to one.

A.5 Direct effect of an increase in δ on π^A_L

Taking the FOC of insurer A with respect to V^A_L, solving for π^A_L and taking the derivative with respect to δ holding all P^A_{rs} constant yields that the sign of this derivative is equal to the sign of

$$\begin{aligned}
&\frac{n-1}{n^2}(p^H - p^L) m^A \Bigg[\left(\frac{P^A_{HI}(1-P^A_{HI})}{\sigma_I} - \frac{P^A_{HC}(1-P^A_{HC})}{\sigma_C}\right)\frac{\lambda}{\widetilde{\sigma}} \\
&\qquad - \left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right)(1-\lambda)\left(P^A_{HC}(1-P^A_{HC})\frac{\eta}{\sigma_C} + P^A_{HI}(1-P^A_{HI})\frac{1-\eta}{\sigma_I}\right)\Bigg] \\
&-\frac{n-1}{n^2}\left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right)\left[\frac{\lambda}{n} + (1-\lambda)\left(\eta P^A_{HC} + (1-\eta) P^A_{HI}\right)\right] \\
&+\frac{\lambda}{\widetilde{\sigma}}\frac{n-1}{n^2}(P^A_{HI} - P^A_{HC}) \\
&+(1-\lambda)(P^A_{HI} - P^A_{HC})\left(P^A_{HC}(1-P^A_{HC})\frac{\eta}{\sigma_C} + P^A_{HI}(1-P^A_{HI})\frac{1-\eta}{\sigma_I}\right) \\
&-\left(\frac{P^A_{HI}(1-P^A_{HI})}{\sigma_I} - \frac{P^A_{HC}(1-P^A_{HC})}{\sigma_C}\right)\left[\frac{\lambda}{n} + (1-\lambda)\left(\eta P^A_{HC} + (1-\eta) P^A_{HI}\right)\right].
\end{aligned} \tag{34}$$

The last line of (34) is negative, but the first line is positive and may, depending on the levels of λ and η, dominate the negative second line; likewise, the positive fourth and fifth line may dominate the negative third line. However, if the share of H-types choosing an insurer of type A is small enough, so that P^A_{HC} and P^A_{HI} are close to zero, we have

$$\text{sign}\left(\frac{\partial \pi^A_L}{\partial \delta}\right) = \text{sign}\left(-\frac{n-1}{n^2}\left(\frac{1}{\sigma_C} - \frac{1}{\sigma_I}\right)\frac{\lambda}{n}\right) = -1. \tag{35}$$

A.6 Effect of a (positive) correlation on m^A in the separating equilibrium

Replacing both P^A_{LC} and P^A_{LI} by $\frac{1}{n^A}$, condition (25) can be simplified to

$$\begin{aligned}
&\frac{\lambda}{n^A}\left(p^L v'(m^A) - p^L\right) \qquad (36)\\
&+\left[(1-\lambda)(\eta P^A_{HC} + (1-\eta)P^A_{HI}) + \delta(P^A_{HI} - P^A_{HC})\right]\left(p^L v'(m^A) - p^H\right)\\
&+\Bigg[(1-\lambda)\left(\eta\frac{P^A_{HC}(1-P^A_{HC})}{\sigma_C} + (1-\eta)\frac{P^A_{HI}(1-P^A_{HI})}{\sigma_I}\right)\\
&\qquad +\delta\left(\frac{P^A_{HI}(1-P^A_{HI})}{\sigma_I} - \frac{P^A_{HC}(1-P^A_{HC})}{\sigma_C}\right)\Bigg](p^H - p^L)\pi^A_H v'(m^A) = 0.
\end{aligned}$$

If the main effects dominate, an increase of δ leads to a decrease of V^B_H and an increase of V^A_L; this leads to an increase of both P^A_{HC} and P^A_{HI}. This increases the bracket $[\cdot]$ in the second line of (36), so v' has to increase, because $p^H > p^L$. It also increases the bracket $[\cdot]$ of the third and fourth line, so v' has to increase again, because $\pi^A_H < 0$. In addition, the bracket in the second line also increases due to the increase in δ itself, so v' has to be increased again. Finally, an increase in δ increases the last bracket of the third and fourth line itself, so v' has to increase again. Therefore, if δ increases, m^A decreases.

References

BIGLAISER, G., AND C. MA (2003): "Price and quality competition under adverse selection: market organization and efficiency," *RAND Journal of Economics*, 34, 266–286.

BIJLSMA, M., J. BOONE, AND G. ZWART (2011): "Competition leverage: How the demand side affects optimal risk adjustment," *TILEC Discussion Paper*, 2011-039.

BUNDORF, K., J. D. LEVIN, AND N. MAHONEY (2012): "Pricing and welfare in health plan choice," *American Economic Review*, 102(7), 3214–3248.

COHEN, A., AND P. SIEGELMAN (2010): "Testing for Adverse Selection in Insurance Markets," *Journal of Risk and Insurance*, 77(1), 39–84.

CUTLER, D. M., A. FINKELSTEIN, AND K. MCGARRY (2008): "Preference heterogeneity and insurance markets: Explaining a puzzle of insurance," *American Economic Review: Papers and Proceedings*, 98(2), 157–162.

EINAV, L., A. FINKELSTEIN, AND M. R. CULLEN (2010): "Estimating Welfare in insurance markets using variation in prices," *Quarterly Journal of Economics*, 123(3), 877–921.

EINAV, L., A. FINKELSTEIN, AND J. LEVIN (2010): "Beyond Testing: Empirical Models of Insurance," *Annual Review of Economics*, 2, 311–336.

ERICSON, K. M. M., AND A. STARC (2012a): "Heuristics and Heterogeneity in Health Insurance Exchanges: Evidence from the Massachusetts Connector," *American Economic Review: Papers and Proceedings*, 102(3), 493–497.

——— (2012b): "Pricing Regulation and imperfect competition on the Massachusetts Health Insurance Exchange," *NBER Working Paper*, 18089.

——— (2013): "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange," *NBER Working Paper*, 19527.

FANG, H., M. P. KEANE, AND D. SILVERMAN (2008): "Sources of advantageous selection: Evidence from the Medigap insurance market," *Journal of Political Economy*, 116(2), 303–350.

HANDEL, B. R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 103(7), 2643–2682.

HANDEL, B. R., AND J. K. KOLSTAD (2013): "Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare," *unpublished*.

HARRIS, K. M., J. SCHULTZ, AND R. FELDMAN (2002): "Measuring consumer perceptions of quality differences among competing health benefit plans," *Journal of Health Economics*, 21(1), 1–17.

JACK, W. (2006): "Optimal risk adjustment in a model with adverse selection and spatial competition," *Journal of Health Economics*, 25(5), 908–926.

JOHAR, M., AND E. SAVAGE (2012): "Sources of advantageous selection: Evidence using actual health expenditure risk," *Economics Letters*, 116(3), 579–582.

KEANE, M. P. (2004): "Modeling Health Insurance Choices in "Competitive" Markets," *unpublished*, pp. 1–57.

LORENZ, N. (2013): "Adverse selection and risk adjustment under imperfect competition," *Trier University Working Paper*, 13/05.

OECD (2011): "Diabetes prevalence and incidence," in *Health at a Glance 2011: OECD Indicators*, pp. 42–43. OECD Publishing, http://dx.doi.org/10.1787/health_glance-2011-13-en.

OLIVELLA, P., AND M. VERA-HERNANDEZ (2007): "Competition among differentiated health plans under adverse selection," *Journal of Health Economics*, 26(2), 233–250.

ROTHSCHILD, M., AND J. STIGLITZ (1976): "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information," *Quarterly Journal of Economics*, 91, 629–649.

ROYALTY, A. B., AND N. SOLOMON (1999): "Health Plan Choice: Price elasticities in a managed competition setting," *Journal of Human Resources*, 31(1), 1–41.

SINAIKO, A. D., AND R. A. HIRTH (2011): "Consumers, health insurance and dominated choices," *Journal of Health Economics*, 30(2), 450–457.

TRAIN, K. E. (2009): *Discrete Choice Methods with Simulation*. Cambridge University Press.

ZWEIFEL, P., F. BREYER, AND M. KIFMANN (2009): *Health Economics*. Springer.