Discussion Papers
Department of Economics
University of Copenhagen

No. 13-15

Monopoly Insurance with Endogenous Information

Johan N. M. Lagerlöf and Christoph Schottmüller

Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K., Denmark
Tel.: +45 35 32 30 01 – Fax: +45 35 32 30 00
http://www.econ.ku.dk

ISSN: 1601-2461 (E)

Monopoly Insurance with Endogenous Information

Johan N. M. Lagerlöf
University of Copenhagen and CEPR

Christoph Schottmüller*
University of Copenhagen and TILEC

November 26, 2013

Abstract

We study a monopoly insurance model with endogenous information acquisition. Through a continuous effort choice, consumers can determine the precision of a privately observed signal that is informative about their accident risk. The equilibrium effort is, depending on parameter values, either zero (implying symmetric information) or positive (implying privately informed consumers). Regardless of the nature of the equilibrium, *all* offered contracts, also at the top, involve underinsurance. The reason is that underinsurance at the top discourages information gathering. We identify a *sorting effect* that explains why the insurer wants to discourage information acquisition. Moreover, a public policy that decreases the information gathering costs can hurt both parties. Lower information gathering costs can harm consumers because the insurer adjusts the optimal contract menu in an unfavorable manner.

Keywords: asymmetric information, information acquisition, insurance, screening, adverse selection
JEL codes: D82, I13

*Address to both authors: Department of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark; email: `johan.lagerlof@econ.ku.dk` and `christoph.schottmueller@econ.ku.dk`. We are grateful for helpful comments from seminar audiences at the University of Copenhagen and at EARIE 2013 (in Évora, Portugal).

1. Introduction

It is often said that markets work better with well-informed consumers. For example, it is not uncommon to hear statements from policymakers that link more information for consumers with higher welfare.[1] However, there are situations where the gathering of information may exacerbate informational asymmetries between buyer and seller. If better information of consumers implies more informational asymmetry, it is no longer clear that markets work better with more informed consumers. In this case, a public policy that facilitates information acquisition might worsen market outcomes and potentially be detrimental to welfare.

One area in which a public policy that leads to more well-informed consumers could exacerbate informational asymmetries is the health insurance market. Suppose a consumer, prior to purchasing a health insurance, has the opportunity to acquire information about his health risk. The consumer knows that one grandparent and one of his aunts are affected by a particular genetic disorder, but his parents are not. To understand how these circumstances impact on the consumer's own health risk, he must learn about how the inheritance pattern of the genetic disorder works.[2] If the consumer indeed learns about this, then he will end up with private information about his risk as long as the insurance company does not know the health history of the consumer's family. As another example, a consumer may know, in broad terms, that his smoking and exercise habits affect his risk of developing a cardiovascular disease. But in order to learn about how substantial the increase in risk is if he smokes, say, ten cigarettes a day, the consumer would have to gather more information. Again, if the consumer does gather the required information and if the insurer does not know about the details of the consumer's life style and habits,[3] then the consumer will end up with private

[1]For example, the chair of Consumer Focus, a public consumer organization in the U.K., argues that "Well informed consumers make markets work better. This is good for consumers, good for economic growth and good for Britain." See `http://www.consumerfocus.org.uk/news/greater-clarity-on-consumer-rights-welcomed-by-consumer-focus` (accessed September 20, 2013). Similarly, the European Commission's Directorate-General for Health and Consumers (DG-SANCO) writes that one of their focus points is: "EU-wide collation and analysis of health data [...] to provide citizens and health professionals with the information they need". See `http://ec.europa.eu/dgs/health_consumer/general_info/mission_en.html` (accessed September 20, 2013).

[2]If, in the example, we are talking about Huntington's disease, this particular health history would be favorable for the consumer, as Huntington's disease has an autosomal dominant inheritance pattern. If, however, we are talking about Wilson's disease, then the consumer would be at an increased risk as the inheritance pattern is autosomal recessive.

[3]It may also be that although having some information about those things, the insurer is legally prohibited from making the insurance policy contingent on them.

information about his health risk.

Thus suppose that a public policy that leads to more well-informed consumers indeed exacerbates informational asymmetries. The presence of these asymmetries is likely to distort the insurance policy that the consumer is offered, at least if the insurance company has some market power. In particular, the distortion could manifest itself in a relatively low insurance coverage, which is costly for a risk averse individual. On the other hand, the fact that the consumer has private information can help him earn a rent that outweighs the adverse effects on consumer surplus. All in all, it is not a priori clear if the net effect on consumer welfare of a policy that facilitates information gathering by the consumer is positive or negative.

In this paper, we examine the questions discussed above in a formal analysis. We do this by studying a monopoly insurance model in which consumers have the opportunity to, privately and covertly, gather information about their (health) risks. In particular, after having observed the insurance company's offered menu of insurance policies, a consumer makes a continuous effort choice that determines the informativeness of a signal about his true risk (which is either "low" or "high"). The consumer observes the signal and then either chooses a policy from the menu or decides to remain uninsured. We characterize the equilibrium menu of insurance policies and we perform comparative statics exercises with respect to the consumer's information acquisition cost. We think of these exercises as a way of exploring the welfare effects of a public policy that facilitates consumer learning with respect to (health) risks. Examples of public policies that we have in mind include informational campaigns, the launch of a website, and the funding of phone lines with free expert advice.

The equilibrium of the model belongs to one of three categories: *pooling*—the consumer is induced to choose a zero effort, which means that there is effectively only one type of agent in the model; *exclusion*—the consumer chooses some positive effort and then purchases an insurance only if observing a high-risk signal; and *separation*—the consumer chooses some positive effort and then buys an insurance with high (low, respectively) coverage if observing a high-risk (low-risk) signal. We show, by means of examples, that there are parameter values for which an equilibrium belonging to each one of these categories exists.

The first main result of our analysis is that, regardless of which of the three categories the equilibrium belongs to, *all* contracts in the offered menu involve underinsurance. This means, in particular, that the famous "no distortion at the top" property does not hold in our model: Lowering the coverage of the high-coverage contract makes the offered insurance contracts more similar and, therefore, discourages information gathering. Relatively badly informed consumers are beneficial for the insurer's profit.

This is because of something that we call the *sorting effect*: As the precision of the consumer's signal drops, the probability that a consumer whose true risk is high receives a low-risk signal (and therefore buys a low coverage contract) increases; similarly, a lower signal precision also leads to a higher likelihood that a low-risk consumer receives a high-risk signal (and therefore buys a high coverage contract). This kind of sorting increases the insurer's profit because consumers buying the high-coverage (low-coverage) contract have a lower (higher) risk. Hence, the indemnities that the insurer must pay are lower in expectation.

In our second main result, we show that our model with endogenous information acquisition contains, as limit cases, (i) the model of Stiglitz (1977) with exogenous asymmetric information and (ii) a model of symmetric information. This implies, for example, that the results of the Stiglitz model hold approximately if information acquisition is possible and the costs of gathering information are sufficiently low.

The third main result concerns the welfare effect of a reduction in the cost of information gathering. We show that such a cost reduction can, contrary to naive intuition, hurt both the consumer and the insurer. The result that the consumer is hurt is the most surprising one. The reason is that the distortion in the offered contracts changes as the costs of information gathering change. Technically speaking, lower information gathering costs exacerbate the consumer's threat to acquire more information (which corresponds to a binding constraint in the insurer's maximization problem). This threat is mitigated by distorting the high-coverage contract which can harm consumer surplus.

Our paper is related to a relatively small but growing literature on mechanism design with endogenous types. This literature has analyzed procurement settings—see, for example, Crémer and Khalil (1992) and Crémer et al. (1998)—but also auctions (Persico, 2000; Shi, 2012) and implementation of efficient allocations à la Vickrey-Clark-Groves (Bergemann and Välimäki, 2002). For a survey of this literature, see Bergemann and Välimäki (2006). Closest to our setup is Szalay (2009), who analyzes a procurement setting in which a firm, by exerting effort, can choose the extent to which a privately observed signal is informative about the firm's marginal cost. As in our model, the effort choice is continuous. An important difference between his procurement and our insurance setting is that it is only in the former case that knowledge about the agent's type is required to achieve the first best allocation. In an insurance setting, the first best allocation is always, for all types, full coverage, which implies that exerting effort is wasteful from a first best point of view. Szalay uses a first order approach—that is, he focuses on situations where the agent's effort choice problem has an interior solution. Given that the first best effort is zero in the insurance setting, we must allow for the possibility that also the optimal (second best) effort is zero, which means that we have to take global constraints into account and cannot restrict ourselves to a pure first order approach. The wastefulness of effort also gives an intuitive explanation for why the equilibrium value of information for the agent is positive in Szalay (2009) but zero in our setup.

Doherty and Thistle (1996) study a question related to our paper. They model a perfectly competitive (health) insurance market in which some consumers do not know their risk type but can learn about this by taking a test. One example that Doherty and Thistle suggest is HIV testing. The focus of their paper is the effect of observability and verifiability of test taking and test outcome. This leads to policy questions concerning the regulation of contractible information in health insurance contracts. Our paper studies a different kind of model (the insurer having market power, the consumer's effort choice being continuous) and it focuses on another policy question, namely, the welfare effects of facilitating information acquisition. The model of Doherty and Thistle (1996) has been extended by adding prevention decisions (Bardey and De Donder, 2012) and early treatment possibilities (Peter et al., 2012).[4]

Finally, in a companion paper (Lagerlöf and Schottmüller, 2013), we study an alternative setting in which the information gathering decision is binary (like deciding whether or not to take a test). The results of that paper, and the logic behind them, differ in several ways from the results of the present paper. In the concluding section, we discuss our companion paper in greater detail and relate it to the present one.

The remainder of the paper is organized as follows. The next section describes our model. We then solve the model with the help of backward induction. Section 3 analyzes the contract choice of the agent, Section 4 looks at the agent's optimal effort choice, and Section 5 examines the optimal menu design problem of the insurer. In Section 6, we return to the policy questions raised in the text above and show, inter alia, that a lower information acquisition cost can lead to both lower consumer surplus and lower profits. Section 7 concludes. An appendix consists of proofs and examples.

2. Model

The principal of the model is a risk neutral and profit-maximizing insurance monopolist offering a menu of insurance contracts. The agent is a risk averse consumer who faces an accident risk and maximizes expected utility. Before presenting the details of the

[4]Broadly related to these papers is also Ligon and Thistle (1996), which is a monopoly model of health insurance in which consumers can take preventive effort. While information acquisition is not modeled, the authors compare a scenario where agents know their risk with a scenario where they do not know this.

model, we explain the timing of events.

1. Nature determines the value of the accident probability $\theta \in \{\theta^l, \theta^h\}$ according to the prior distribution $\Pr\left[\theta = \theta^h\right] = \alpha$. Neither the insurer nor the agent can observe the realization of this draw.

2. The insurer chooses a menu of insurance contracts, which is then observed by the agent.

3. The agent chooses an effort level $e \geq 0$ and then observes a signal $\sigma \in \left\{\sigma^l, \sigma^h\right\}$. The insurer cannot observe the effort level or the signal.

4. Given the signal σ and the effort choice e, the agent uses Bayes' rule to form an interim belief about θ and then picks one insurance contract from the menu (or remains uninsured).

We denote the consumer's initial wealth by $w > 0$ and assume that an accident leads to (the equivalent of) a monetary loss, or damage, $D \in (0, w)$. An insurance contract specifies a premium p and an indemnity R that is paid in case the loss occurs.

The accident occurs with the exogenous probability $\theta \in (0, 1)$. This probability is either high ($\theta = \theta^h$) or low ($\theta = \theta^l$, with $\theta^l < \theta^h$), and initially the value is unknown to both the insurer and the agent: Principal and agent share a common prior $\alpha = \Pr\left[\theta = \theta^h\right]$, where $\alpha \in (0, 1)$. The agent, however, privately observes a binary signal $\sigma \in \left\{\sigma^l, \sigma^h\right\}$, which may be informative about θ. The informativeness of the signal is determined by the agent's effort $e \geq 0$. The larger is e, the higher is the correlation between the signal and the true accident probability.

More precisely, let $\alpha_i(e)$ denote the agent's updated belief that the accident probability is high, after having chosen the effort e and observed the signal σ^i (for $i \in \{l, h\}$): $\alpha_i(e) \equiv \Pr\left[\theta = \theta^h \mid \sigma = \sigma^i, e\right]$. We can then write the agent's expected accident probability, given a signal σ^i and an effort level e, as $\beta_i(e) \equiv \alpha_i(e)\,\theta^h + [1 - \alpha_i(e)]\,\theta^l$. We will define our signaling technology in terms of properties of the functions $\beta_l(e)$ and $\beta_h(e)$. First β_i is assumed to be two times continuously differentiable with $\beta_l'(e) < 0$ and $\beta_h'(e) > 0$. This captures the notion that a higher effort leads to a more informative signal. We also assume that

$$\beta_l(0) = \beta_h(0) = \alpha\theta^h + (1 - \alpha)\theta^l \equiv \beta.$$

That is, the signal provides no additional information beyond the prior if the agent does not exert any effort at all. Furthermore, we make the assumption that, for all $e \geq 0$, the unconditional probability of a high signal equals the prior: $\Pr\left[\sigma = \sigma^h\right] = \alpha$. Together with Bayes' rule, this assumption implies that

$$\alpha\beta_h(e) + (1-\alpha)\beta_l(e) = \beta \qquad \text{for all } e \geq 0. \tag{1}$$

By differentiating both sides of the identity in (1) with respect to e, it is straightforward to verify that $\alpha\beta_h'(e) = -(1-\alpha)\beta_l'(e)$, which we will make use of in the subsequent analysis.[5]

Finally we make the following assumption:

Assumption 1. *The signaling technology exhibits decreasing returns from effort; that is, $\beta_h''(e) \leq 0$ for all $e > 0$ (and therefore $\beta_l''(e) \geq 0$ for all $e > 0$).*

This assumption will restrict the number of local maxima in the agent's effort choice problem (see Section 4). In Appendix A, we provide examples of signaling technologies that satisfy all our assumptions.

The agent maximizes expected utility. His Bernoulli utility function u is assumed to be twice continuously differentiable with $u' > 0$ and $u'' < 0$. Exerting positive effort leads to a disutility $c(e)$, which is assumed to enter additively in the agent's payoff.[6] The cost function c is strictly increasing, strictly convex, twice continuously differentiable and satisfies $c(0) = c'(0) = 0$. An agent's expected utility after exerting effort e, receiving signal σ^i and buying the insurance contract (p, R) is therefore

$$U = \beta_i(e)u(w - p - D + R) + (1 - \beta_i(e))u(w - p) - c(e).$$

The solution concept we employ is sequentially rational Nash equilibrium.[7] Put differently, we will solve the game by backwards induction starting with the agent's contract choice.

3. Contract choice

At stage 2 of the game, the insurer chooses a menu of contracts, which the agent then takes as given when making his contract choice. Lemma 1 below tells us that we can without loss of generality rule out menus consisting of more than two contracts: Any

[5] More generally, both the first and second derivatives of $\beta_l(e)$ and $\beta_h(e)$ must have opposite signs: $sgn\left[\beta_l'(e)\right] \neq sgn\left[\beta_h'(e)\right]$ and $sgn\left[\beta_l''(e)\right] \neq sgn\left[\beta_h''(e)\right]$.

[6] In the supplementary material to this paper, we explore a setup where the costs of effort are monetary and enter, therefore, in the argument of the utility function; see the discussion in Section 7.

[7] Subgame perfect equilibrium has no bite since the game itself is the only subgame. Our equilibrium is equivalent to a sequential equilibrium. As consistency of beliefs is unproblematic, we do not want to clutter notation with it.

behavior on the part of the agent that the insurer can induce with such a menu can just as well be induced with the help of only one or two contracts.

Lemma 1. *For any equilibrium with more than two contracts in the menu, there is a corresponding equilibrium in which (i) the insurer offers at most two contracts and (ii) the expected profits of the insurer and the expected utility of the agent are the same as in the original equilibrium.*

Proof. See Appendix B.

The lemma holds true because of the binary nature of the signal. There will be at most two types of agents—those receiving a high signal and those receiving a low signal—and there is no need to offer more contracts than there are types.[8] In the remainder of the paper we will, justified by Lemma 1, focus on menus with at most two contracts. Also we ignore—without loss of generality—dominated contracts; that is, if the insurer offers two contracts and contract 2 has the lower indemnity ($R_2 < R_1$), then contract 2 must have the lower premium ($p_2 < p_1$) as well.

When making his contract choice, the agent thus has up to three options (up to two contracts in the menu and the outside option of no insurance). What will determine the agent's choice between these? The options will differ from each other in terms of how much coverage they involve. In the remainder of this section, we will formalize the intuitive statement that the higher is the agent's expected accident risk $\beta_i(e)$, the greater is his tendency to choose an option with relatively high coverage.

Suppose that there are indeed two contracts in the chosen menu and that these are distinct from each other and indexed by $j \in \{1, 2\}$. In addition, refer to the option of not purchasing any insurance as "contract 0" and accordingly write $p_0 = R_0 = 0$. Let $\underline{u}_j$ denote the agent's utility if having chosen contract j and then having an accident. Similarly, let $\overline{u}_j$ denote the agent's utility if having chosen contract j and then *not* having an accident. Formally, $\underline{u}_j \equiv u\,(w - D - p_j + R_j)$ and $\overline{u}_j \equiv u\,(w - p_j)$ for $j \in \{0, 1, 2\}$.

Given some chosen effort e and some observed signal σ^i, the agent weakly prefers contract 1 to contract 2 if, and only if,

$$\beta_i(e)(\underline{u}_1 - \underline{u}_2) + [1 - \beta_i(e)](\bar{u}_1 - \bar{u}_2) \geq 0. \tag{2}$$

Assume, without loss of generality, that $p_1 \geq p_2$ or, equivalently, that $\bar{u}_1 \leq \bar{u}_2$, which implies $\underline{u}_1 \geq \underline{u}_2$ as contract 1 would be dominated otherwise. By the assumptions that

[8] Of course, there are many more equilibria. For example, the insurer could offer a menu with a third contract that is strictly worse for both types of agent and which is therefore not purchased in equilibrium.

the contracts are distinct and none is dominated by the other, both inequalities must hold strictly: $\bar{u}_1 < \bar{u}_2$ and $\underline{u}_1 > \underline{u}_2$. We will refer to contracts 1 and 2 as the high and the low coverage contract, respectively. Rewriting inequality (2) yields

$$\beta_i(e) \geq \frac{\bar{u}_2 - \bar{u}_1}{(\bar{u}_2 - \underline{u}_2) - (\bar{u}_1 - \underline{u}_1)} = \frac{\bar{u}_2 - \bar{u}_1}{\Delta_2 - \Delta_1}, \tag{3}$$

where we introduced the shorthand notation $\Delta_j \equiv \bar{u}_j - \underline{u}_j$. Note that Δ_j measures the extent to which contract j involves underinsurance (we have indeed that $\Delta_2 > \Delta_1$). The condition in (3) says that the agent prefers the high coverage contract only if the expected accident probability is high enough. The agent also, of course, has the option of remaining uninsured (i.e., to pick contract 0). Following a procedure analogous to the one above, it is easy to see that the agent weakly prefers contract 1 to both the other alternatives if, and only if,

$$\beta_i(e) \geq \max\left\{\frac{\bar{u}_0 - \bar{u}_2}{\Delta_0 - \Delta_2}, \frac{\bar{u}_2 - \bar{u}_1}{\Delta_2 - \Delta_1}\right\} \equiv \tilde{\tilde{\beta}}.$$

Similarly, the agent weakly prefers no insurance (i.e., contract 0) if, and only if,

$$\beta_i(e) \leq \min\left\{\frac{\bar{u}_0 - \bar{u}_2}{\Delta_0 - \Delta_2}, \frac{\bar{u}_0 - \bar{u}_1}{\Delta_0 - \Delta_1}\right\} \equiv \tilde{\beta}.$$

Contract 2 is of course strictly preferred by the agent if, and only if, both the above inequalities are violated. Note that $\tilde{\beta} \leq \tilde{\tilde{\beta}}$, with strict inequality whenever there exists a belief $\beta_i(e)$ at which contract 2 would be strictly preferred to contract 0 and 1.

We can conclude from the above analysis that the larger is the expected accident risk $\beta_i(e)$, the stronger is the agent's preference for a high coverage contract. In particular, if we let $\beta_i(e)$ gradually increase (and assume that its initial value is sufficiently low), the agent with signal i will move, in turn, from contract 0 to contract 2 to contract 1.

4. Effort choice

The agent's expected utility from choosing contract (p_j, R_j), conditional on the effort level e and the signal σ^i, equals (gross of the effort costs $c(e)$)

$$\beta_i(e)\,\underline{u}_j + [1 - \beta_i(e)]\,\bar{u}_j. \tag{4}$$

Let $m \in \{0, 1, 2\}$ denote the contract that maximizes the expression in (4) across all available contracts j (including contract 0, the outside option).[9] Moreover, let $\underline{u}^*(\sigma, e)$ denote the agent's utility if having chosen that optimal contract and then having an accident: $\underline{u}^*(\sigma, e) \equiv u(w - D - p_m + R_m)$. Similarly, let $\overline{u}^*(\sigma, e)$ denote the

[9] If there are more than one such contract, then let m denote any one of them.

agent's utility if having chosen contract m and then *not* having an accident: $\overline{u}^* (\sigma, e) \equiv u(w - p_m)$.

We can now write the agent's expected utility, at the point in time when he is about to choose the effort level e, as follows:

$$U(e) = \sum_{i \in \{l,h\}} \Pr\left[\sigma = \sigma^i\right] \left[\beta_i(e)\, \underline{u}^*\left(\sigma^i, e\right) + (1 - \beta_i(e))\, \overline{u}^*\left(\sigma^i, e\right)\right] - c(e), \tag{5}$$

where we recall that $\Pr\left[\sigma = \sigma^h\right] = \alpha$. Notice that the function $U(e)$ is not necessarily everywhere differentiable with respect to e, as it may have a kink at a point where the agent prefers to move from one contract to another (although $U(e)$ is continuous for all $e \geq 0$). However, for each of the two signals there are at most two such points with a possible kink; this follows from the analysis in Section 3 and the assumption that $\beta_l(e)$ and $\beta_h(e)$ are monotone in e. Everywhere else the expression is indeed differentiable, and we can calculate

$$\begin{aligned} U'(e) &= \sum_{i \in \{l,h\}} \Pr\left[\sigma = \sigma^i\right] \left[\beta_i'(e)\, \underline{u}^*\left(\sigma^i, e\right) - \beta_i'(e)\, \overline{u}^*\left(\sigma^i, e\right)\right] - c'(e) \\ &= \alpha \beta_h'(e) \left(\Delta_l^* - \Delta_h^*\right) - c'(e), \end{aligned} \tag{6}$$

where $\Delta_i^* \equiv \overline{u}^*(\sigma^i, e) - \underline{u}^*(\sigma^i, e)$ measures the extent to which the *optimally chosen* contract (given a signal σ^i and an effort e) involves underinsurance. The second equality follows from $\alpha\beta_h'(e) = -(1-\alpha)\beta_l'(e)$.

Thus an increase in the agent's effort has potentially two effects on his utility. First, it always increases the information gathering cost $c(e)$. This effect is negative and it is captured by the last term in (6). Second, if the agent chooses different contracts after having observed the two possible signals, then a larger e increases the probability of a correct contract choice. This effect on the agent's utility is positive and it is captured by the first term in (6) (note that we always have $\Delta_l^* \geq \Delta_h^*$, with the inequality being strict if the contracts are different).

It follows that if the agent for some interval of effort levels chooses the same contract regardless of the signal, then $U'(e) = -c'(e) < 0$. For example, if the agent, given his prior (that is, with zero effort), prefers one contract strictly over the others, then $U(e)$ will be strictly decreasing in the effort level for e small enough. The function $U(e)$ is also strictly concave in that interval (since $c''(e) > 0$). Moreover, if the agent for some interval of effort levels chooses different contracts for different signals, then his expected utility is possibly increasing in e. The implication of these observations is that, despite Assumption 1 and the fact that the cost function is strictly convex, the expected utility function $U(e)$ is not guaranteed to be quasiconcave.

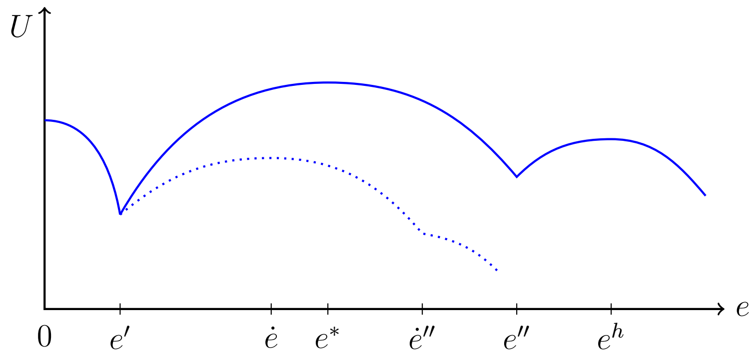

Figure 1: Utility as function of effort given optimal contract choice

We illustrate the lack of quasiconcavity with the help of Figure 1. Suppose, as before, that the insurer has offered a menu with two distinct contracts, 1 and 2. Also suppose that for $e = 0$ the agent strictly prefers contract 2, the one with low coverage, to contracts 0 and 1. We know that then $U(e)$ must be strictly decreasing at $e = 0$, as is shown in the figure. As we gradually increase e, starting at zero, $\beta_h(e)$ becomes larger and $\beta_l(e)$ becomes smaller. Therefore, at some point (say, at $e' > 0$) the agent will be indifferent between contracts 1 and 2 when he receives a high signal, and at some point (say, at $e'' > 0$) the agent will be indifferent between contracts 0 and 2 when receiving a low signal. In Figure 1, we have $e' < e''$ (although the reverse relationship is possible in principle). For effort levels immediately above e', $U(e)$ could be increasing (if $c'(e')$ is not too high). Eventually, however, $U(e)$ will decrease again, because the cost effect of increasing e becomes larger ($c'' > 0$) while the information effect becomes smaller ($\beta_h'' \leq 0$). At e'', the agent with a low signal decides to remain uninsured instead of buying contract 2. Again, $U(e)$ may be increasing for effort levels immediately above e'', but eventually it will decrease again. If the utility function is the solid curve in Figure 1, optimal effort is $e^* > 0$. If the utility function is the dotted line (from e' onwards), the optimal effort is zero.

Hence, there can be several interior local maxima as well as a local maximum at the boundary, $e = 0$. Assumption 1 and the convexity of the cost function ensure that $U(e)$ is strictly concave on $[0, e']$ and on $[e', e'']$ as well as on $[e'', \infty)$. Hence, $U(e)$ can have up to three local maxima if two contracts are offered.

Suppose the insurer wants to implement an outcome where $e > 0$ and where the agent buys contract 1 if he gets a high signal and contract 2 if he receives a low signal. Then the agent's optimal effort is the unique solution to the first order condition

$$\alpha\beta_h'(e)(\Delta_2 - \Delta_1) - c'(e) = 0. \tag{7}$$

With a slight abuse of terminology, we will refer to the solution of (7) as *optimal interior* *effort*. This leads to the following useful result:

Result 1. *The optimal interior effort choice depends only on the coverage levels Δ_i and not directly on the utility levels.*

5. Contract design

We now turn to the profit maximization problem of the monopolist insurer. Depending on what agent behavior the insurer optimally induces, the solution to this problem will belong to one of the following three categories:

1. *Pooling.* The agent does not gather information ($e = 0$) but does purchase an insurance contract (independently of any signal as there is no informative signal with $e = 0$).

2. *Exclusion.* The agent gathers information ($e > 0$) and purchases an insurance only if observing the signal σ^h (hence the low type agent remains uninsured).

3. *Separation.* The agent gathers information ($e > 0$) and purchases insurance contract 1 (i.e., the one with high coverage) if observing the signal σ^h and contract 2 (i.e., the one with low coverage) if observing the signal σ^l.

5.1. Constraints in contract design

In this subsection, we explain the constraints that appear in the profit maximization problem of the insurer. We also show which of these are binding at the optimum. In the text below we will, to begin with, discuss the constraints that are relevant if the insurer's optimal menu involves separation. After having done this, however, we will show that this setup includes the cases of pooling and exclusion as special cases. Hence, we will be able to deal with all possible contract designs within one single formulation of the insurer's problem.

Thus suppose that the insurer wants to induce separation. The insurer's problem is then to maximize his expected profits subject to altogether nine constraints. Four of the constraints concern the agent's behavior at the interim stage where he has observed the signal; these are identical to the constraints in standard formulations of the screening problem, namely, two individual rationality constraints and two incentive compatibility constraints (one of each for each type). As in the standard screening problem, it is straightforward to show that only the low type's individual rationality constraint and the high type's incentive compatibility constraint can be relevant. These constraints read:

$$\beta_l(e)\underline{u}_2 + (1 - \beta_l(e))\bar{u}_2 \geq \beta_l(e)\underline{u}_0 + (1 - \beta_l(e))\bar{u}_0, \tag{IRl}$$

$$\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1 \geq \beta_h(e)\underline{u}_2 + (1-\beta_h(e))\bar{u}_2. \tag{ICh}$$

The remaining five constraints concern the agent's behavior at the ex ante stage, prior to the observation of the signal. First, the agent must prefer to exert the effort level e that the insurer wants to induce rather than exerting zero effort and to remain uninsured. We will refer to this as the agent's *ex ante individual rationality constraint*:

$$\alpha[\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e)$$
$$\geq \beta\underline{u}_0 + (1-\beta)\bar{u}_0. \quad \text{(EAIR)}$$

The incentive compatibility constraints at the interim stage ensure that, after having observed the signal, the agent indeed chooses the contract aimed at his particular type (as the left hand side of (EAIR) above assumes). The next two ex ante constraints require that the agent prefers to exert the effort level e that the insurer wants to induce rather than exerting zero effort and then always choose one of the two contracts in the menu (always contract 2 or always contract 1, respectively). We will call these two conditions the *information gathering constraints*. They read:

$$\alpha[\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e)$$
$$\geq \beta\underline{u}_2 + (1-\beta)\bar{u}_2, \quad \text{(IGl)}$$

$$\alpha[\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e)$$
$$\geq \beta\underline{u}_1 + (1-\beta)\bar{u}_1. \quad \text{(IGh)}$$

The fourth ex ante constraint requires that the agent does not prefer an effort level that is (significantly) *higher* than is desired by the insurer. In particular, the agent may want to choose some effort e^h (cf. Figure 1) and then purchase contract 1 if receiving a high signal and contract 0 if receiving a low signal. For the agent not to have such an incentive, the following condition must hold (we will refer to this as the *effort-high constraint*):[10]

$$\alpha\left[\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1\right] + (1-\alpha)\left[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2\right] - c(e)$$
$$\geq \alpha\left[\beta_h(e^h)\underline{u}_1 + (1-\beta_h(e^h))\bar{u}_1\right] + (1-\alpha)\left[\beta_l(e^h)\underline{u}_0 + (1-\beta_l(e^h))\bar{u}_0\right] - c(e^h), \quad \text{(EH)}$$

[10]Note that it is not rational to exert some effort e^l and buy contract $(\bar{u}_2, \underline{u}_2)$ if the signal is σ^h and no insurance if the signal is σ^l. If $e^l < e$, then $\beta_l(e) < \beta_l(e^l)$ and therefore buying the low coverage contract would be utility maximizing even with a low signal after e^l because of (IRl). If $e^l \geq e$, then $\beta_h(e^l) > \beta_h(e)$ and by (ICh) it is then rational to buy contract 1 instead of contract 2 in case of a high signal.

where the effort level e^h satisfies the following first order condition:

$$c'(e^h) = \alpha\beta_h'(e^h)\left(\bar{u}_0 - \underline{u}_0 - \bar{u}_1 + \underline{u}_1\right). \tag{8}$$

Because of Assumption 1 and $c'' > 0$, this first order condition is necessary and sufficient for e^h to maximize the agent's expected utility, given that he chooses contract 1 after a high signal and contract 0 after a low signal.

The fifth and final ex ante constraint is a *local information gathering constraint*. It requires that the effort level e that the insurer wants to induce is a local optimum to the agent's effort choice problem. In particular, it says that the first order condition in (7) must be satisfied. This condition is necessary and sufficient for e to maximize the agent's expected utility, given that he chooses contract 1 after a high signal and contract 2 after a low signal (again, because of Assumption 1 and $c'' > 0$).

The constraints that we have discussed are able to deal also with exclusion and pooling. To see this, first suppose that the insurer wants to induce pooling, which means making contracts 1 and 2 identical ($\bar{u}_1 = \bar{u}_2$ and $\underline{u}_1 = \underline{u}_2$) and providing the agent with incentives to choose $e = 0$. The same interim constraints as above are required, although several of them are trivially satisfied when the two menu contracts are the same. For the agent not to have an incentive to gather information, (EH) must be satisfied; this constraint ensures that the agent does not want to choose the effort level e^h and then rejects the single contract in the menu when observing the low signal (this is the only way in which the agent would be able and potentially willing to deviate to some $e > 0$). Finally, (EAIR) must hold. The remaining constraints that are required to induce separation—namely, (IGl), (IGh), and the first order condition in (7)—will under pooling be trivially satisfied with equality (recall our assumption that $c'(0) = 0$).

The case with exclusion is similar. Here the insurer makes contracts 0 and 2 identical ($\bar{u}_0 = \bar{u}_2$ and $\underline{u}_0 = \underline{u}_2$) and provides the agent with incentives to choose $e = e^h > 0$. The agent's incentives to choose that effort level will be guaranteed by the first order condition in (7) and by (IGh). The constraints (IGl) and (EAIR) will coincide and therefore both act as the ex ante individual rationality constraint. The constraint (ICh) is effectively the high type's (interim) individual rationality constraint, and (IRl) and (EH) are trivially satisfied.

When formulating the insurer's profit maximization problem we will not have to consider all the constraints discussed above. The following lemma tells us that (IGl) and (EH) are binding in equilibrium while all other constraints that were stated as inequalities can be neglected. (The local information gathering constraint (7) and (8) are obviously also binding.)

Lemma 2. *Constraints (ICh), (IRl) and (EAIR) are implied by (EH) and (IGl). In equilibrium, (EH) and (IGl) bind while (IGh) is slack.*

Proof. See Appendix B.

Lemma 2 is not only technically important for solving the contract design problem but has also a straightforward economic implication. Consider the cases of separation and exclusion (i.e., the ones where the agent chooses a positive effort level). Since (IGl) is binding the agent will be indifferent between information acquisition and no information acquisition when facing the optimal menu. Doherty and Thistle (1996) refer to the utility difference between acquiring information (and choosing the best contract given the acquired information) and not acquiring information (and choosing the best contract given the prior) as the *value of information* to the agent. To indicate that we refer to this utility difference given the optimal menu, we prefer to call this concept the *equilibrium value of information.* We can conclude the following:

Corollary 1. *The equilibrium value of information to the agent is zero.*

We want to point out that, despite the corollary, effort can be positive in equilibrium. The corollary only states that, in equilibrium, the agent will be indifferent between the optimal positive effort and no effort (only with pooling will the optimal effort be zero). In this respect, our model differs from other models of endogenous information acquisition (e.g., Crémer et al. (1998), Szalay (2009), and Doherty and Thistle (1996)), in which agents can have strict preferences over effort in equilibrium. The reason why our result is different is that in our model the agent's effort choice problem is (i) continuous and (ii) non-quasiconcave (in particular, the agent's effort choice problem has several local maxima and one of those is the boundary solution $e = 0$).

5.2. Profit maximization program and distortion results

We now turn to the formulation of the insurer's profit maximization problem. As is often the case in the formal insurance literature, it will be more convenient to work with the inverse of the utility function than with this function directly. We write $h(u)$ for the inverse of the utility function (i.e., $u^{-1} = h$) and note that $h'(u) > 0$ and $h''(u) > 0$. By Lemma 2, the optimal menu of contracts solves the following program:[11]

$$\begin{aligned} \max_{\bar{u}_1, \underline{u}_1, \bar{u}_2, \underline{u}_2, e, e^h} \; & w - \beta D - \alpha \left[\beta_h(e) h(\underline{u}_1) + (1-\beta_h(e)) h(\bar{u}_1)\right] \\ & -(1-\alpha)\left[\beta_l(e) h(\underline{u}_2) + (1-\beta_l(e)) h(\bar{u}_2)\right] \end{aligned} \tag{9}$$

[11]The objective function in (9) is an equivalent way of writing $\alpha\left[p_1 - \beta_h(e) R_1\right] + (1 - \alpha)\left[p_2 - \beta_l(e) R_2\right]$, i.e., the insurer's expected profits.

subject to the constraints

$$\alpha[\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e) \\ = [\alpha\theta^h + (1-\alpha)\theta^l]\underline{u}_2 + [\alpha(1-\theta^h) + (1-\alpha)(1-\theta^l)]\bar{u}_2, \tag{10}$$

$$\alpha[(\beta_h(e) - \beta_h(e^h))\underline{u}_1 + (\beta_h(e^h) - \beta_h(e))\bar{u}_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e) \\ = (1-\alpha)\beta^l(e^h)\underline{u}_0 + (1-\alpha)(1-\beta_l(e^h))\bar{u}_0 - c(e^h), \tag{11}$$

$$c'(e) = \alpha\beta_h'(e)\left(\bar{u}_2 - \underline{u}_2 - \bar{u}_1 + \underline{u}_1\right), \tag{12}$$

$$c'(e^h) = \alpha\beta_h'(e^h)\left(\bar{u}_0 - \underline{u}_0 - \bar{u}_1 + \underline{u}_1\right), \tag{13}$$

$$e^h \geq e \geq 0. \tag{14}$$

The equalities in (10) and (11) are the binding (IGl) and (EH) constraints, respectively.

In Proposition 1 below, we state our first main result, which concerns the allocative efficiency of the equilibrium contracts. The result says that *all* equilibrium contracts must involve underinsurance. This result stands in sharp contrast to the usual "no distortion at the top" property that holds in standard screening models in which the information structure is exogenous.

Proposition 1. *All insurance contracts in the equilibrium menu have less than full coverage (i.e., $\underline{u}_j < \bar{u}_j$ for all j).*

Proof. Here in the main text, we will prove that contract 1 cannot have full coverage when the equilibrium menu involves separation. In Appendix B, we show that the result holds also with pooling and exclusion and that overinsurance is never optimal. Since contract 2 by construction has less coverage than contract 1, also contract 2 must involve underinsurance.

The proof will proceed in three steps. We will show that for any full coverage contract aimed at the high type, there is a partial coverage contract that (i) yields higher profits for the insurer, (ii) is preferred by the high type, and (iii) is feasible. We begin by computing the slope of the isoprofit curve in the $(\underline{u}_1, \bar{u}_1)$ plane. This slope will tell us how much $\bar{u}_1$ must be changed to keep profits constant if $\underline{u}_1$ is increased marginally (while keeping $\underline{u}_2$ and $\bar{u}_2$ fixed). Using the profit expression stated in (9) and invoking the implicit function theorem (acknowledging that also e depends on $\underline{u}_1$ through the optimality condition in (12)), we obtain the following expression for the slope of the isoprofit curve:

$$\begin{aligned}
&\left.\frac{d\bar{u}_1}{d\underline{u}_1}\right|_{\pi=\bar{\pi}} \\
= \;& -\frac{\alpha\beta_h(e)h'(\underline{u}_1) - \alpha\beta_h'(e)\frac{de}{d\underline{u}_1}[D + h(\underline{u}_1) - h(\bar{u}_1)] - (1-\alpha)\beta_l'(e)\frac{de}{d\underline{u}_1}[D + h(\underline{u}_2) - h(\bar{u}_2)]}{-\alpha[1-\beta_h(e)]h'(\bar{u}_1) - \alpha\beta_h'(e)\frac{de}{d\bar{u}_1}[D + h(\underline{u}_1) - h(\bar{u}_1)] - (1-\alpha)\beta_l'(e)\frac{de}{d\bar{u}_1}[D + h(\underline{u}_2) - h(\bar{u}_2)]} \\
= \;& -\frac{\beta_h(e)h'(\underline{u}_1) + \beta_h'(e)\frac{de}{d\underline{u}_1}(R_1 - R_2)}{[1-\beta_h(e)]h'(\bar{u}_1) - \beta_h'(e)\frac{de}{d\underline{u}_1}(R_1 - R_2)} \\
< \;& -\frac{\beta_h(e)h'(\underline{u}_1)}{[1-\beta_h(e)]h'(\bar{u}_1)}. \qquad (15)
\end{aligned}$$

The second equality above makes use of $\alpha\beta_h'(e) = -(1-\alpha)\beta_l'(e)$ and $D + h(\underline{u}_i) - h(\bar{u}_i) = R_i$ as well as $-de/d\bar{u}_1 = de/d\underline{u}_1 > 0$ (see (12) for the latter). The inequality follows from the fact that $\beta_h'(e) > 0$, $de/d\underline{u}_1 > 0$ and $R_1 > R_2$.

Our next step is to calculate the slope of the high type agent's indifference curve in the $(\underline{u}_1, \bar{u}_1)$ plane. The high type's expected utility from contract 1 equals $EU^{high} = \beta_h(e)\underline{u}_1 + (1-\beta_h(e))\,\bar{u}_1$. By implicitly differentiating both sides of this identity with respect to $\underline{u}_1$, while acknowledging that also e depends on $\underline{u}_1$ through the optimality condition in (12), we obtain the following expression for the slope:

$$\begin{aligned}
\left.\frac{d\bar{u}_1}{d\underline{u}_1}\right|_{EU^{high}=const} &= -\frac{\beta_h(e) + \beta_h'(e)(\underline{u}_1 - \bar{u}_1)\frac{de}{d\underline{u}_1}}{1-\beta_h(e) + \beta_h'(e)(\underline{u}_1 - \bar{u}_1)\frac{de}{d\bar{u}_1}} \\
&= -\frac{\beta_h(e) + \beta_h'(e)(\underline{u}_1 - \bar{u}_1)\frac{de}{d\underline{u}_1}}{1-\beta_h(e) - \beta_h'(e)(\underline{u}_1 - \bar{u}_1)\frac{de}{d\underline{u}_1}},
\end{aligned}$$

where the second equality uses $de/d\bar{u}_1 = -de/d\underline{u}_1$. It follows that at the point where there is full coverage (i.e., where $\bar{u}_1 = \underline{u}_1$), the isoprofit curve is steeper than the indifference curve (both are negatively sloped). The implication is that, by marginally moving the high type's contract along the indifference curve and toward the region with underinsurance, the insurer's profits increase. That is, given some arbitrary full insurance contract for the high type, there exist partial insurance contracts that are more profitable for the insurer and preferred by the agent.[12] Finally we consider the question whether these partial insurance contracts are feasible. As they give a higher utility to the high type, offering such a contract instead of a full coverage contract relaxes (IGl). Moreover, as the partial insurance contracts reduce coverage (i.e., lower $\underline{u}_1$ and higher $\bar{u}_1$) also (EH) is relaxed. □

What is the intuition behind the result that there is underinsurance also for the high type? The main driving force behind the result is the following *sorting effect*: It

[12] Using the envelope theorem, one can derive $\left.\frac{d\bar{u}_1}{d\underline{u}_1}\right|_{EU=const} = \frac{\beta_h(e)}{1-\beta_h(e)}$. Hence, the ex ante utility of the agent increases when changing to such a partial insurance contract.

is beneficial for the insurer if the agent acquires less information. The reason is that the agent tends to choose the "wrong" contract if he is badly informed. In particular, an agent whose true accident probability is high will, if the signal is noisy, often receive a low signal and thus choose the low coverage contract; similarly, an agent whose true accident probability is low will, if the signal is noisy, often receive a high signal and thus choose the high coverage contract. Both these kinds of mistakes are beneficial for the insurer's profit: A high-risk agent has a high expected loss. If this agent buys a contract with less coverage, the insurer needs to bear less of that loss. Similarly, a low-risk agent has a low expected loss. It is, therefore, better for the insurer if the low-risk agent purchases the high coverage contract than if the high-risk agent buys this contract.[13]

The sorting effect is the main reason why the insurer optimally underinsures also the high type. By reducing the coverage in the high type's contract (i.e., by increasing Δ_1), the insurer induces a lower effort choice (see the local information gathering constraint (7)). This is beneficial for the insurer's profits, due to the sorting effect. Given that the adjustment of contract 1 is made starting out from a situation with no distortion, the cost of distorting the contract slightly will be of second order magnitude only. The profit enhancing sorting effect, on the other hand, is of first order magnitude.

The distortion result of Proposition 1 does not appear in the previous literature that has endogenized the information structure in screening models.[14] There are two main reasons why we obtain a different result in our model. First, our effort variable is continuous. Most of the previous literature assumes a binary effort decision (see, e.g., Crémer et al. (1998)); in such a setting, the logic that gives rise to distortion at the top in our setting cannot arise: There is no possibility of marginally reducing information acquisition by distorting the top contract. Second, we consider an insurance problem. In a standard procurement problem, there is no sorting effect because the principal's payoff does not directly depend on the agent's type (it only depends indirectly through the contract choice) while the agent's type determines the expected payout in an insurance problem.

We now return to the insurer's profit maximization problem. The constraints (10)–(13) are all linear in the agent's ex post utility levels. It is therefore straightforward to

[13]In terms of the objective function in (9), the sorting effect is simply the effect of an increase in the effort that goes directly through the variable e (keeping all the ex post utility levels fixed). The partial derivative of the objective with respect to e is $\beta_h'(e)\left([h(\bar{u}_1) - h(\underline{u}_1)] - [h(\bar{u}_2) - h(\underline{u}_2)]\right)$, which is negative for $\underline{u}_2 < \underline{u}_1 < \bar{u}_1 < \bar{u}_2$.

[14]The notable exception that we are aware of is a result in an extension of Szalay's (2009) main analysis in which he allows for a moving support—that is, the choice of effort affects the possible range of expected marginal costs. Distortion at the top will then affect the effort choice. This affects the support which in turn affects the agent's expected rent and therefore the principal's payoff.

solve the four constraints for the four utility levels as functions of e and e^h. This leads to the following expressions:

$$
\begin{aligned}
\bar{u}_2(e,e^h) &= \bar{u}_0 - \beta_l(e^h)\frac{c'(e^h)}{\alpha\beta_h'(e^h)} + \beta_l(e)\frac{c'(e)}{\alpha\beta_h'(e)} - \frac{c(e^h)-c(e)}{1-\alpha}, && \text{(16a)}\\
\underline{u}_2(e,e^h) &= \underline{u}_0 + \left(1-\beta_l(e^h)\right)\frac{c'(e^h)}{\alpha\beta_h'(e^h)} - (1-\beta_l(e))\frac{c'(e)}{\alpha{\beta_h}'(e)} - \frac{c(e^h)-c(e)}{1-\alpha}, && \text{(16b)}\\
\bar{u}_1(e,e^h) &= \bar{u}_0 - \beta_l(e^h)\frac{c'(e^h)}{\alpha\beta_h'(e^h)} - [\beta_h(e)-\beta_l(e)]\frac{c'(e)}{\alpha\beta_h'(e)} - \frac{c(e^h)-c(e)}{1-\alpha} + \frac{c(e)}{\alpha}, && \text{(16c)}\\
\underline{u}_1(e,e^h) &= \underline{u}_0 + \left(1-\beta_l(e^h)\right)\frac{c'(e^h)}{\alpha\beta_h'(e^h)} - [\beta_h(e)-\beta_l(e)]\frac{c'(e)}{\alpha\beta_h'(e)} - \frac{c(e^h)-c(e)}{1-\alpha} + \frac{c(e)}{\alpha} && \text{(16d)}
\end{aligned}
$$

Plugging these expressions into the objective function (9) leads to a maximization problem over the two variables e and e^h, where the only constraints are $e^h \geq e \geq 0$. An equilibrium with pooling results if $e = 0$. An equilibrium with exclusion occurs if $e^h = e$. This is summarized in the following proposition.

Proposition 2. *The optimal contract menu is given by (16), where e and e^h solve the following maximization problem:*

$$
\begin{aligned}
\max_{e,e^h}\ w - \beta D - \alpha\left[\beta_h(e)h(\underline{u}_1(e,e^h)) + (1-\beta_h(e))h(\bar{u}_1(e,e^h))\right]&\\
-(1-\alpha)\left[\beta_l(e)h(\underline{u}_2(e,e^h)) + (1-\beta_l(e))h(\bar{u}_2(e,e^h))\right]&\\
s.t.:\ e^h \geq e \geq 0.&
\end{aligned}
$$

The optimal contract menu satisfies the first order conditions:

$$
\begin{aligned}
\frac{\partial\pi(e,e^h)}{\partial e} &= -\alpha\beta_h'(e)\left[h(\bar{u}_2) - h(\underline{u}_2) - h(\bar{u}_1) + h(\underline{u}_1)\right]\\
&\quad -(1-\alpha)\beta_l(e)(1-\beta_l(e))[h'(\bar{u}_2) - h'(\underline{u}_2)]\frac{c''(e)\beta_h'(e) - \beta_h''(e)c'(e)}{\alpha(\beta_h'(e))^2}\\
&\quad +\alpha\left[\beta_h(e)h'(\underline{u}_1) + (1-\beta_h(e))h'(\bar{u}_1)\right][\beta_h(e)-\beta_l(e)]\frac{c''(e)\beta_h'(e) - \beta_h''(e)c'(e)}{\alpha(\beta_h'(e))^2}\\
&\leq 0 \qquad \textit{with ``='' if } e > 0; && \text{(17)}\\
\frac{\partial\pi(e,e^h)}{\partial e^h} &= \left[\alpha\left(-\beta_h(e)h'(\underline{u}_1)(1-\beta_l(e^h)) + (1-\beta_h(e))h'(\bar{u}_1)\beta_l(e^h)\right)\right.\\
&\quad \left.+(1-\alpha)\left(-\beta_l(e)h'(\underline{u}_2)(1-\beta_l(e^h)) + (1-\beta_l(e))h'(\bar{u}_2)\beta_l(e^h)\right)\right]\times\\
&\quad \frac{c''(e^h)\beta_h'(e^h) - \beta_h''(e^h)c'(e^h)}{\alpha(\beta_h'(e^h))^2}\\
&\leq 0 \qquad \textit{with ``='' if } e^h > e; && \text{(18)}\\
e^h &\geq e \geq 0, && \text{(19)}
\end{aligned}
$$

where the $\bar{u}_i$ and $\underline{u}_i$ are given by (16).

To illustrate some of the insurer's incentives when designing the optimal menu, consider the effect on the profits of an increase in the effort level e, i.e., the left hand side of (17). The first part of the derivative, on the first line, is strictly negative as long as the coverage of contract 1 is strictly higher than that of contract 2. This part represents the sorting effect discussed above. The remaining part of the expression, on the second and third lines, represents the effect on the insurer's profit that goes via changes in the ex post utility levels in (16). For example, a higher e requires an increase in $\Delta_2 - \Delta_1$ (see the local information gathering constraint in (7)). And a larger $\Delta_2 - \Delta_1$, in turn, relaxes (IGl), which enables the insurer to lower the utility levels in contract 1 and thus increase profits. This second part of (17) can therefore be positive, making it profitable for the insurer to induce a positive effort level.

Evaluated at $e = 0$, the derivative in (17) is always non-positive:

$$\frac{\partial \pi(0, e^h)}{\partial e} = -(1-\alpha)\beta(1-\beta)[h'(\bar{u}_2) - h'(\underline{u}_2)]\frac{c''(0)}{\beta_h'(0)} \leq 0.$$

Indeed, if $c''(0) > 0$, then the derivative is strictly negative and $e = 0$ must thus be a local maximum. However, as we have already noted, the optimization program above is not quasiconcave and therefore the globally profit maximizing menu of contracts can lead to positive effort.

The lack of quasiconcavity of the maximization problem implies that the first order conditions allow for multiple solutions. This is true also for simple functional forms—for example, for a linear signaling technology and a quadratic cost function. However, as the insurer's problem can be written as the problem of maximizing a continuous function over a compact and convex set, the problem is easy to solve numerically.[15] We show in a numerical example (see Appendix C) that each equilibrium type (pooling, exclusion, separation) exists for some parameter values.

5.3. Limit results

This subsection relates our analysis to the insurance model in Stiglitz (1977). In this classic model, the agent does not exert effort in order to learn about his accident risk. Instead, each agent type is assumed to know his risk from the outset of the game. We call the optimal contract menu in this model the *Stiglitz menu*.

Knowing one's risk is intuitively very similar to being able to learn one's risk at zero costs. To formalize this intuition, we denote the effort cost function employed in this subsection not by $c(e)$ but by $\gamma c(e)$, with $\gamma > 0$. Effort costs are therefore zero if the

[15] Compactness of the domain is not apparent from Proposition 2. However, it is easy to bound e^h from above by some $\bar{e}$: For example, $\bar{e}$ defined by $c(\bar{e}) = u(w - \beta D) - (\beta\underline{u}_0 + (1-\beta)\bar{u}_0)$ is such an upper bound, as the right hand side is an upper bound on the benefit an agent can get from insurance.

parameter γ is zero. The following proposition states, inter alia, that the equilibrium contract in our model converges to the Stiglitz contract as γ becomes arbitrarily small.

Proposition 3. *As $\gamma \to 0$, the optimal contract menu is either separating or exclusionary and, therefore, induces a strictly positive effort level. If $\lim_{e\to\infty} \beta_h(e) = \theta^h$, the optimal contract converges to the Stiglitz menu as $\gamma \to 0$.*[16] *As $\gamma \to \infty$, the equilibrium coverage of all offered contracts converges to full coverage (i.e., $\Delta_1 \to 0$ and $\Delta_2 \to 0$).*

Proof. See Appendix B.

The continuity result in the first part of Proposition 3 implies that the results of the classic Stiglitz model are robust. The proposition can thus serve as a justification for the common practice in the literature of neglecting the information acquisition decision: As long as the costs of information gathering are small, the results of the Stiglitz model and those of the richer model with information acquisition roughly coincide. The second part of Proposition 3 deals with the case where exerting effort is becoming arbitrarily costly. The limit $\gamma = \infty$ is intuitively equivalent to a situation where the agent simply cannot learn about his type. This means that the agents, throughout the game, are homogeneous and have a common perceived accident risk that equals the prior β. In this situation, a monopolist insurer has no incentive to distort the coverage. The reason is that, in a model with homogeneous agents, only their (common) individual rationality constraint matters, which cannot be relaxed by distorting the coverage. Proposition 3 confirms this intuition: As γ becomes arbitrarily large, the optimal contract menu converges in a continuous fashion to the limit outcome with full insurance.

6. Policy intervention: Information provision

In this section, we address the policy questions that we raised in the Introduction. As we explained there, it is common to hear claims from policymakers and commentators that well-informed consumers make markets work better and are therefore good for welfare. Often, the ambition of these policymakers is to *empower* consumers, and key elements of this empowerment are to provide information to and facilitate information acquisition by consumers. For example, the E.U. Consumer Policy Strategy 2007-2013 states in its objectives that "Empowered consumers need real choices, *accurate information*, market transparency and the confidence that comes from effective protection and solid rights" (our italics).[17] We will here, within the framework of the model studied in the previous

[16] If $\lim_{e\to\infty} \beta_h(e) \neq \theta^h$, the optimal contract menu will converge to the Stiglitz menu of a model in which the two types' risks are $(\lim_{e\to\infty} \beta_h(e), \lim_{e\to\infty} \beta_l(e))$ instead of (θ^h, θ^l).

[17] European Commission (2007, p. 13). See also the quotes in footnote 1 in the Introduction.

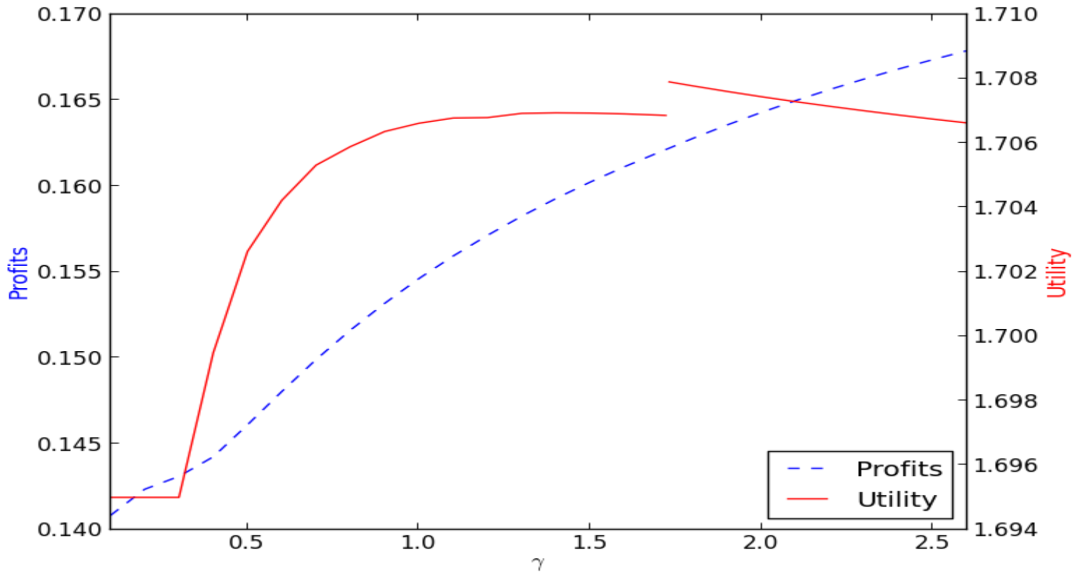

Figure 2: Consumer surplus and profits (numerical example)

sections, investigate the welfare implications of a policy that lowers the consumers' cost of gathering information about their health risks.

As we did in Subsection 5.3 above, we will denote the effort cost function not by $c(e)$ but by $\gamma c(e)$, with $\gamma > 0$. We will think of "a policy that facilitates for consumers to gather information" as a policy that lowers γ. The question we will ask is: Given the equilibrium menu of insurance policies characterized in Section 5, what is the effect on consumer surplus and on the insurer's profits of an (exogenous) reduction in γ?[18]

A first answer to that question is given by a numerical example (see Appendix C for details). Figure 2 plots the values of the insurer's profits and the consumer surplus that the example gives rise to, for a range of different values of the cost parameter γ (from $\gamma = 0.05$ to $\gamma = 1.3$).[19] The example in Figure 2 shows that, in general, the effect of a reduction in γ on consumer surplus is ambiguous: It can be either positive or negative. Figure 2 shows that both the consumer and the insurer can be hurt by a reduction in the cost. The more surprising of these results is probably the one saying that consumer surplus can decrease as a consequence of a reduction in γ. As information acquisition

[18]Note that this means that we abstract from any "investment costs" that the policymaker may have to incur in order to facilitate for consumers to gather information (e.g., costs of funding an information campaign, leaflets, or a website). If we find that the net welfare effect is negative when not counting such investments costs, then it is of course even more so if including those costs.

[19]The example assumes the functional forms $u(x) = \sqrt{x}$ and $c(e) = \gamma e^4$. See also Appendix C.

becomes less difficult, consumers benefit directly from that through the effect on their costs. However, there is also an indirect effect that goes through the conditions of the optimally chosen insurance policy. As it becomes easier for the consumer to learn about his risk, the constraints in the insurer's maximization problem are modified. In particular, the constraint (EH) will be harder to satisfy. To deal with this, the insurer adjusts the menu of insurance policies, which sometimes can have an adverse effect on the consumer surplus. This is especially apparent around $\gamma = 0.86$ in the figure, where the equilibrium type changes from separating to pooling (i.e., at that value of γ, the global maximum of the principal's objective function changes from one local maximum to another). Hence, contracts and consumer surplus change discontinuously in γ at this point.

Although the effects of a change in γ on consumer surplus and profits are ambiguous in general, clear-cut results emerge in special cases.

Proposition 4. *If the optimal contract menu is pooling, profits are increasing in γ. If the optimal contract menu is exclusionary, consumer surplus is constant in γ.*

Proof. See Appendix B.

The second part of Proposition 4 says that in an equilibrium with exclusion (meaning that the consumer exerts some positive effort and then buys an insurance only if receiving a high-risk signal) a change in the information acquisition cost has no impact on the consumer's expected utility. The reason for this is that in this kind of equilibrium the constraints (IGl) and (EAIR) coincide (see our discussion in Section 5.1), which means that here (EAIR) binds and the consumer does not receive any ex ante rents. Since the consumer's outside option is independent of γ, a change in this parameter cannot affect his expected equilibrium utility. In other words, while the reduction in γ lowers the consumer's effort costs, the insurer also makes the conditions in the insurance contract less favorable to the consumer, leaving the net effect equal to zero.

The effect of a change in γ on profits consists, in general, of three parts. First, a lower γ tightens the constraint (EH): Deviating to the high effort level e^h becomes more attractive when the cost is relatively low. Second, a lower γ relaxes the constraint (IGl): While the utility from simply buying contract 2 (with zero effort) is not affected, the costs of exerting effort e are lower and therefore the expected utility from the contracts in the menu is higher. Third, due to the sorting effect that we have discussed above, a reduction in γ lowers profits: A lower value of γ leads to higher effort e and therefore a more precise signal, which is bad for the insurer's profits. In an equilibrium with pooling, it is only the first one of these three effects that is non-zero. As a consequence, a reduction in γ decreases profits.

7. Concluding discussion

In this paper, we have studied a monopoly insurance model with endogenous information acquisition. We assumed that the insurer and the consumer are initially symmetrically informed about the latter's health risk. However, the consumer can, privately and covertly, learn more about his risk by exerting effort. This effort choice is continuous. In spite of our assumptions that the effort cost function is strictly convex and the signaling technology is concave, the consumer's expected utility at the point in time when he makes his effort choice is not quasiconcave. This means that a first order approach is not valid and we must impose global "information gathering constraints" on the insurer's problem of designing the optimal menu of contracts. Two of these global constraints turn out to be binding in equilibrium and they are important for our results. Also the principal's profit maximization problem has several local maxima. Nevertheless, we are able to solve the model and to derive some qualitative results.

Among our results are: (i) Depending on parameter values, the equilbrium of the model may involve zero effort (meaning symmetric information) or a positive effort (which implies a privately informed consumer). (ii) Regardless of the nature of the equilibrium, all contracts, also at the top, involve underinsurance. (iii) In the limit where the information acquisition cost approaches zero, the equilibrium outcome approaches the one in Stiglitz's (1977) classic model. Thus, the results of the latter model are robust to the introduction of endogenous information gathering if information gathering costs are sufficiently small. (iv) An exogenous reduction in the information acquisition cost has, for certain parameter values, a *negative* effect both on profits and consumer surplus. (v) For other parameter values, a reduction in the cost can be beneficial; so welfare results are ambiguous in general.

We also identify a *sorting effect* that implies that the insurer wants the consumer to be badly informed. This effect is an important reason behind our result that there is underinsurance also at the top. The logic is that a consumer with a relatively uninformative signal is more likely to choose the "wrong" contract—the low-coverage contract if the true health risk is high, and the high-coverage contract if the true health risk is low. From the insurer's point of view, it is a *good* thing if the consumer tends to choose the wrong contract! The reason is that this lowers the expected value of the indemnities that the insurer must pay to the consumer: It is preferable if the high coverage contract is purchased by the low-risk consumer since that consumer will need compensation with a smaller likelihood. We are not aware of any discussion of the sorting effect in the previous literature. The effect should appear in insurance settings quite generally, whenever the precision of the insuree's information is a continuum. However, the sorting effect does not appear in, for example, a procurement setting because the agent's type does not affect the principal's payoff directly in such a setting.

In our model, we assumed that the consumer's information acquisition cost entered his payoff additively, rather than being subtracted from his wealth and thus entering the argument of the utility function. One interpretation of this model specification, which would justify our modeling choice, is that the cost is a disutility of effort, as opposed to a monetary expenditure. The assumption that the cost is additive has been made by other authors (see Doherty and Thistle, 1996) and it makes our analysis much more tractable than it would otherwise be. Still, a natural question is whether, or to what extent, our results would be altered if we assumed that the cost of effort were monetary and therefore entered the argument of the consumer's utility function. We explore this question in the supplementary material to this paper (available on our websites). We can show that our main results hold for the case of CARA preferences (which imply that there are no income effects on risk aversion). For other utility specifications, the analysis becomes untractable as the endogenous extent of information acquisition can change the degree of risk aversion and therefore the demand for insurance.

The insurance company in our model is a monopolist. While admittedly this does not reflect the reality in all insurance markets, insurance companies tend to be big due to scale economies in risk diversification. Recent empirical evidence (Dafny, 2010) shows that health insurance companies in the U.S. indeed have market power. In smaller countries with more public involvement in health care than in the U.S., health insurers will typically have even more market power. In Denmark, for example, a public system that is funded by taxes covers basic health care. However, dental care, physiotherapy, chiropractors and care in private hospitals are not covered by the public system. Instead there is a private insurance company (*Sygeforsikringen "danmark" g.s.*) that offers additional health insurance for these services. This company is indeed a monopolist in its market and it offers a menu of insurance policies with different levels of coverage.

We also assumed that the consumer's information acquisition decision is continuous. We believe this is a natural modeling choice for many economic environments—for example, if the accuracy of the information that the consumers acquire depends on the amount of time they spend on searching. However, one can also think of situations where the information acquisition decision is essentially binary—for example, if consumers have the opportunity to take a test that tells them whether they carry some particular gene associated with an increased risk for certain diseases.

In a companion paper (Lagerlöf and Schottmüller, 2013), we study a similar question to the one modeled here but with a binary information acquisition decision. The binary decision naturally fits to the application of genetic and other testing where consumers decide to take a test or not. The continuous effort choice fits more naturally to the idea of spending time in order to search for information on the internet. The qualitative results of the two papers differ in several respects. In a situation where there is no information acquisition, the differences are minor: In both settings, a local reduction in costs hurts the insurer and might benefit the consumer. In an equilibrium with a positive level of information acquisition, there are bigger differences. While in the binary setting both parties always gain from a marginal cost reduction, this is not guaranteed in the continuous setting, as we have seen above. For example, as we showed in Figure 2, it can be that both parties lose from the cost reduction. An important reason behind these differences is that in the continuous setting, but only there, the consumer can deviate in his effort choice either downwards (to zero) or upwards (to e^h). The two constraints that ensure that such deviations do not pay off are both binding in equilibrium, and they force the insurer to distort the contracts. Most importantly, the sorting effect, as defined in this paper, appears only in the continuous setting. The sorting effect also leads to the result that there is always underinsurance, also at the top. This "distortion-at-the-top" result does not hold in the binary setup. Intuitively, in the continuous-effort setting the insurer distorts the contract with the objective of inducing "a bit less" information acquisition. If the effort choice is binary, this is clearly not possible.[20]

Summing up, our welfare analysis in this paper sends a message of caution. In an insurance market, a public policy that facilitates consumer learning can lead to quite intricate market responses. It can even lead to an outcome where all parties are worse off.

[20]Distorting the contracts so much that the consumer gathers no information at all is usually not profitable for the principal. Technically speaking, such huge distortions will lead to first order losses for the principal that were not present in the proof of Proposition 1.

Appendix

A. Signal technology examples

We here provide two examples of signaling technologies that satisfy all the assumptions that we made in the model description. The first example leads to a linear $\beta_h(e)$, whereas the second one leads to a strictly concave $\beta_h(e)$. Let $g(e)$ be an increasing function satisfying $g(0) = 0$ and $g(e) \leq 1$ for all $e \geq 0$. Both our examples assume the following: Given an effort e, the signal reveals the true state ($\sigma^i = \theta^i$) with probability $g(e)$; with the complementary probability, $1 - g(e)$, the signal is drawn from the prior ($\Pr\left[\sigma = \sigma^h \mid \theta = \theta^h\right] = \Pr\left[\sigma = \sigma^h \mid \theta = \theta^l\right] = \alpha$) and is therefore completely uninformative. These assumptions, together with Bayes' rule, yield (we here treat e as a fixed parameter):

$$\begin{aligned}
\alpha_h(e) &\equiv \Pr[\theta = \theta^h|\sigma = \sigma^h] = \frac{\Pr[\theta = \theta^h]\Pr[\sigma = \sigma^h|\theta = \theta^h]}{\Pr[\sigma = \sigma^h]} \\
&= \frac{\alpha\left[\alpha + (1-\alpha\ )\, g\,(e)\ \right]}{\alpha} = \alpha + (1-\alpha\ )\, g\,(e)\,.
\end{aligned}$$

Thus, $\beta_h(e) \equiv \alpha_h(e)\theta^h + [1 - \alpha_h(e)]\,\theta^l = \alpha\theta^h + (1-\alpha)\theta^l + (1-\alpha)(\theta^h - \theta^l)g\,(e)\,.$ To obtain our first example we set $g(e) = \min\{e, 1\}$. This gives us, for $e < 1$,

$$\beta_h(e) = \alpha\theta^h + (1-\alpha)\theta^l + (1-\alpha)(\theta^h - \theta^l)e,$$

$$\beta_h'(e) = (1-\alpha)(\theta^h - \theta^l).$$

Hence, $\beta_h(e)$ is linear in e for $e \in [0, 1]$ and flat (equal to θ^h) for higher values of e. Consequently, β_h is weakly concave. To obtain our second example we set $g(e) = 1 - \frac{1}{e+1}$. This gives us

$$\beta_h(e) = \alpha\theta^h + (1-\alpha)\theta^l + (1-\alpha)(\theta^h - \theta^l)\left(1 - \frac{1}{e+1}\right),$$

$$\beta_h'(e) = \frac{(1-\alpha)(\theta^h - \theta^l)}{(e+1)^2} > 0,$$

$$\beta_h''(e) = -\frac{2(1-\alpha)(\theta^h - \theta^l)}{(e+1)^3} < 0.$$

B. Proofs

Proof of Lemma 1: Suppose that the insurer offers more than two contracts in equilibrium. We will show that there exists another equilibrium in which the insurer offers only two contracts and earns the same profits as in the original equilibrium.

As a first step, delete all contracts from the menu that are bought with zero probability under equilibrium play. Second, denote a pure strategy of the agent by (e, i, j), where e is the chosen effort, i is the contract in the menu that is chosen if the signal is high, and j is the contract chosen if the signal is low.[21] The pure strategy (e, i, j) induces some particular expected profit of the insurer. Let (e^*, i^*, j^*) denote the pure strategy that, among the pure strategies that the agent uses with positive probability in the original equilibrium, yields the highest expected profits. There must exist an equilibrium in which the insurer offers only the contracts i^* and j^*. This is because, by definition, offering only these contracts yields expected profits that are at least as high as the ones in the original equilibrium. Moreover, the agent will not have an incentive to deviate from such a new equilibrium, as he played (e^*, i^*, j^*) with positive probability in the original equilibrium (when his choice set was larger). We can also conclude that the insurer's expected profits cannot be strictly higher in the new equilibrium than in the original one. Because if they were, then the original situation would not have been an equilibrium (offering only i^* and j^* would be have been a profitable deviation for the insurer). Finally, if in the original equilibrium the agent uses more than one pure strategy, then he must be indifferent between them. Therefore the expected utility of the agent is the same in the original and in the new equilibrium. $\square$

Proof of Lemma 2: First, (ICh) is implied by (IGl). Recall that $\beta = \alpha\beta_h(e) + (1 - \alpha)\beta_l(e)$. Consequently, (IGl) can be rewritten as

$$\beta_h(e)\underline{u}_1 + (1 - \beta_h(e))\bar{u}_1 - \beta_h(e)\underline{u}_2 - (1 - \beta_h(e))\bar{u}_2 \geq c(e)/\alpha,$$

which is (weakly) more stringent than (ICh) as $c(e)/\alpha \geq 0$.

Second, (IRl) is slack. Suppose (IRl) was binding. We will show that in this case (EH) would be violated. Given that (IRl) holds with equality, (EH) can be written as $f(e) \geq f(e^h)$ where

$$f(e) = \alpha[(\beta_h(e) - \beta_h(e^h))\underline{u}_1 + (\beta_h(e^h) - \beta_h(e))\bar{u}_1] + (1 - \alpha)[\beta_l(e)\underline{u}_0 + (1 - \beta_l(e))\bar{u}_0] - c(e).$$

Since f is strictly concave in e and since e^h is defined as the maximizer of f, the inequality $f(e) \geq f(e^h)$ cannot hold and therefore (EH) would be violated if (IRl) was binding. Consequently, (IRl) cannot bind.

Third, (EAIR) is implied by (IGl) in combination with (IRl). As $\beta_l(e) \leq \beta$ and $\underline{u}_0 \leq \underline{u}_2$ and $\bar{u}_0 \geq \bar{u}_2$, (IRl) implies[22]

$$\beta\underline{u}_2 + (1 - \beta)\bar{u}_2 \geq \beta\underline{u}_0 + (1 - \beta)\bar{u}_0.$$

[21] Possibly, i and j are the same contract.

[22] Note that in an equilibrium with separation or exclusion the inequality below holds strictly, as $e > 0$ in such an equilibrium. This will play a role below.

Plugging this last inequality into (IGl) yields

$$\alpha\left[\beta_h(e)\underline{u}_1+(1-\beta_h(e))\bar{u}_1\right]+(1-\alpha)\left[\beta_l(e)\underline{u}_2+(1-\beta_l(e))\bar{u}_2\right]-c(e) \\ \geq \beta\underline{u}_0+(1-\beta)\bar{u}_0$$

which is (EAIR).

Fourth, (IGl) is binding in equilibrium. Note that the statement is tautological in an equilibrium with pooling, as in that case (IGl) reduces to an identity. Thus consider an equilibrium with separation or exclusion and suppose that (IGl) is not binding. If so, the insurer can decrease $\underline{u}_1$ by $\varepsilon>0$, which increases the profits.[23] This deviation relaxes (or does not affect) all the other potentially binding constraints. The change will decrease the expected utility of the agent by $\alpha\beta_h(e)\varepsilon$.[24] This is less than $\alpha\beta_h(e^h)\varepsilon$ with $e^h>e$. Consequently, (EH) is relaxed by the change. Similarly, (IGh) is relaxed as the effect on the right hand side of (IGh) is $\beta=\alpha\beta_h(e)+(1-\alpha)\beta_l(e)>\alpha\beta_h(e)$. As (EAIR) cannot bind in a separating/exclusion equilibrium (see above), the change does not violate (EAIR) for $\varepsilon>0$ small enough. Therefore, (IGl) must bind.

Fifth, (IGh) is slack in the optimal contract. Suppose (IGh) was binding. We will show that the principal can deviate to a pooling contract that satisfies all constraints and increases profits. We distinguish two cases.

In the first case, assume $\bar{u}_1>\underline{u}_1$. In this case, we claim that profits can be increased by only offering contract 1, i.e., by dropping contract 2 from the menu. Since (IGh) was binding, the agent can achieve the same utility as before by buying contract 1 without exerting effort. This must be his optimal choice, as reducing his choice set cannot result in higher ex ante utility. In particular, (EAIR) and (EH) are not affected and (IGl) is irrelevant in the pooling situation. Now we must show that profits are increased:

$$\begin{aligned}
\pi^p &= w-\beta D-\beta h(\underline{u}_1)-(1-\beta)h(\bar{u}_1)\\
&= w-\beta D-\alpha\left[\beta_h(e)h(\underline{u}_1)+(1-\beta_h(e))h(\bar{u}_1)\right]-(1-\alpha)\left[\beta_l(e)h(\underline{u}_1)+(1-\beta_l(e))h(\bar{u}_1)\right]\\
&> w-\beta D-\alpha\left[\beta_h(e)h(\underline{u}_1)+(1-\beta_h(e))h(\bar{u}_1)\right]-(1-\alpha)\left[\beta_l(e)h(\underline{u}_2)+(1-\beta_l(e))h(\bar{u}_2)\right]
\end{aligned}$$

where the inequality follows from (ICl) and the strict convexity of h. More specifically, the line between $h(\underline{u}_2)$ and $h(\bar{u}_2)$ is strictly above the line connecting $h(\underline{u}_1)$ and $h(\bar{u}_1)$, because h is strictly convex and $\underline{u}_2<\underline{u}_1<\bar{u}_1<\bar{u}_2$. By (ICl), $\beta_l(e)\underline{u}_2+(1-\beta_l(e))\bar{u}_2\geq \beta_l(e)\underline{u}_1+(1-\beta_l(e))\bar{u}_1$. Therefore, $\beta_l(e)h(\underline{u}_2)+(1-\beta_l(e))h(\bar{u}_2)>\beta_l(e)h(\underline{u}_1)+(1-\beta_l(e))h(\bar{u}_1)$. This concludes the proof for the case $\bar{u}_1>\underline{u}_1$.

[23]The indirect effect of this change is that the optimal e will be decreased. This indirect effect increases profits as well.

[24]The indirect effect through e is negligible for ε small enough, as the agent maximizes utility over e; that is, there is no first order effect.

For the second case, assume $\bar{u}_1 \leq \underline{u}_1$. Then a full coverage pooling contract that gives the same ex ante utility to the agent increases profits and is feasible. Let $u_p = \alpha[\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e)$ be the utility level of the full coverage deviation contract. The agent can achieve the same utility level as before by buying the pooling contract without exerting effort. Hence, (EAIR) is satisfied (and (IGl) is irrelevant in pooling). Because (IGh) was binding initially by assumption, we have $u_p = \beta\underline{u}_1 + (1-\beta)\bar{u}_1$. As $\beta_h(e^h) > \beta$ and $\bar{u}_1 \leq \underline{u}_1$, it follows that $u_p < \beta_h(e^h)\underline{u}_1 + (1-\beta_h(e^h))\bar{u}_1$. Therefore, (EH) is relaxed by the deviation. Finally, we show that profits are higher under the deviation contract:

$$
\begin{aligned}
\pi_p &= w - \beta D - h(u_p) \\
&= w - \beta D - h\left[\alpha(\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1) + (1-\alpha)(\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2) - c(e)\right] \\
&\geq w - \beta D - h\left[\alpha(\beta_h(e)\underline{u}_1 + (1-\beta_h(e))\bar{u}_1) + (1-\alpha)(\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2)\right] \\
&> w - \beta D - \alpha\left[\beta_h(e)h(\underline{u}_1) + (1-\beta_h(e))h(\bar{u}_1)\right] - (1-\alpha)\left[\beta_l(e)h(\underline{u}_2) + (1-\beta_l(e))\bar{u}_2\right],
\end{aligned}
$$

where the first inequality follows from $h' > 0$ and $c(e) \geq 0$ and the second inequality follows from $h'' > 0$.

Last, (EH) must be binding. If not, $\bar{u}_2$ and $\underline{u}_2$ could both be decreased by $\varepsilon > 0$. This would not affect the optimal choice of e and it would increase the principal's profit. The binding constraint (IGl) would also be relaxed by this decrease. □

Proof of Proposition 1 (continued): In an exclusionary equilibrium the same proof as in the main text holds. The only difference is that we must use the subscript 0 instead of 2. As $R_0 = 0$, the R_2 term is not present in the slope of the indifference curve. However, this does not change (15) and therefore the argument goes through also in this case. The relevant constraint in an exclusion equilibrium is the ex ante individual rationality constraint, which is relaxed by giving the high risk type a higher utility.

Before turning to pooling equilibria, we show that overinsurance is not optimal in a separating or exclusion equilibrium. Suppose there is overinsurance: $\bar{u}_1 < \underline{u}_1$. Call the optimal effort level under the (supposedly) optimal contract e^*. We will show that the principal has a profitable deviation: Change contract 1 by decreasing $\underline{u}_1$ by $\varepsilon > 0$ and increasing $\bar{u}_1$ by $\varepsilon' > 0$.

Define the effort level e' as the optimal effort level in the changed menu; that is, e' solves

$$\alpha\beta_h'(e')(\Delta_2 - \bar{u}_1 - \varepsilon' + \underline{u}_1 - \varepsilon) = c'(e').$$

Note that $e' < e^*$. Choose ε' such that the expected utility of the agent is not affected by the contract change:

$$\alpha(\beta_h(e^*)\underline{u}_1 + (1-\beta_h(e^*))\bar{u}_1) + (1-\alpha)(\beta_l(e^*)\underline{u}_2 + (1-\beta_l(e^*))\bar{u}_2) - c(e^*)$$
$$= \alpha\left(\beta_h(e')(\underline{u}_1 - \varepsilon) + (1-\beta_h(e'))(\bar{u}_1+\varepsilon')\right) + (1-\alpha)(\beta_l(e')\underline{u}_2 + (1-\beta_l(e'))\bar{u}_2) - c(e').$$

Note that $\beta_h(e^*)\underline{u}_1 + (1-\beta_h(e^*))\bar{u}_1 > \beta_h(e^*)(\underline{u}_1 - \varepsilon) + (1-\beta_h(e^*))(\bar{u}_1+\varepsilon')$: If this was not the case, the agent would—under the modified menu—get an ex ante rent at least as high as in the original menu by exerting effort e^*. Choosing the optimal effort e' would then result in strictly higher ex ante utility, which contradicts the definition of ε'. Using this insight, we show that profits are higher in the modified menu:

$$\begin{aligned}
\pi^{old} &= \alpha\left[-\beta_h(e^*)D + w - \beta_h(e^*)h(\underline{u}_1) - (1-\beta_h(e^*))h(\bar{u}_1)\right] \\
&\quad +(1-\alpha)\left[-\beta_l(e^*)D + w - \beta_l(e^*)h(\underline{u}_2) - (1-\beta_l(e^*))h(\bar{u}_2)\right] \\
&< \alpha\left[-\beta_h(e^*)D + w - \beta_h(e^*)h(\underline{u}_1-\varepsilon) - (1-\beta_h(e^*))h(\bar{u}_1+\varepsilon')\right] \\
&\quad +(1-\alpha)\left[-\beta_l(e^*)D + w - \beta_l(e^*)h(\underline{u}_2) - (1-\beta_l(e^*))h(\bar{u}_2)\right] \\
&\leq \alpha\left[-\beta_h(e')D + w - \beta_h(e')h(\underline{u}_1-\varepsilon) - (1-\beta_h(e'))h(\bar{u}_1+\varepsilon')\right] \\
&\quad +(1-\alpha)\left[-\beta_l(e')D + w - \beta_l(e')h(\underline{u}_2) - (1-\beta_l(e'))h(\bar{u}_2)\right] \\
&= \pi^{new}
\end{aligned}$$

The first inequality follows from the convexity of h as well as $\beta_h(e^*)\underline{u}_1 + (1-\beta_h(e^*))\bar{u}_1 > \beta_h(e^*)(\underline{u}_1-\varepsilon) + (1-\beta_h(e^*))(\bar{u}_1+\varepsilon')$ and the monotonicity of h. The second inequality follows from $\alpha\beta_h'(e) = -(1-\alpha)\beta_l'(e)$ (which implies that the expression is decreasing in e) and the fact that $e^* > e'$.

Next we must show that no constraint is violated under the modified menu. As the ex ante expected utility and also the low coverage contract did not change, (IGl) is not affected by the modification of the menu. As (IGh) is slack under the optimal menu (see Lemma 2), this constraint is not violated for small changes ε. To check (EH), define the function $z(\Delta_1)$ as

$$\begin{aligned}
z(\Delta_1) &= \alpha[\beta_h(e^h) - \beta_h(e)\Delta_1] + (1-\alpha)[\beta_l(e)\underline{u}_2 + (1-\beta_l(e))\bar{u}_2] - c(e) \\
&\quad -(1-\alpha)[\beta_l(e^h)\underline{u}_0 - (1-\beta_l(e^h))\bar{u}_0] + c(e^h),
\end{aligned}$$

where e and e^h are also functions of Δ_1 defined in the obvious manner through (7). (EH) is satisfied if $z(\Delta_1) \geq 0$. Using the envelope theorem, we can derive $z'(\Delta_1) = \alpha(\beta_h(e^h) - \beta_h(e)) > 0$. Hence, increasing Δ_1 relaxes (EH). As our modification increased Δ_1, the modification relaxes (EH). All other constraints are, by Lemma 2, implied by (EH) and (IGl).

This shows that the modification of the menu increased profits while relaxing (or not affecting) the relevant constraints. Consequently, overinsurance cannot occur in a separating or exclusion equilibrium.

Last we consider a pooling equilibrium. Note that there cannot be overinsurance in a pooling equilibrium. The same argument as in the previous step (where the changes ε and ε' apply to the pooling contract instead of contract 1) shows this immediately.

It remains to show that full insurance is not optimal in a pooling equilibrium. There is only one binding constraint in a pooling equilibrium: The agent is indifferent between (i) zero effort and buying the contract and (ii) exerting positive effort and buying the contract only if he receives a high signal. The slope of the indifference curve of an agent exerting zero effort and buying a contract $(\bar{u}, \underline{u})$ is

$$\left.\frac{d\bar{u}}{d\underline{u}}\right|_{EU=const} = -\frac{\beta}{1-\beta},$$

where $\beta = \alpha\theta^h + (1-\alpha)\theta^l$. The slope of the indifference curve of an agent exerting positive effort $e^h > 0$ and buying the contract only when receiving a high signal is (when deriving this slope we use the fact that e^h is chosen optimally and, therefore, the effect through e^h is zero by the envelope theorem):

$$\left.\frac{d\bar{u}}{d\underline{u}}\right|_{EU=const} = -\frac{\beta_h(e^h)}{1-\beta_h(e^h)} < -\frac{\beta}{1-\beta},$$

where the inequality follows from $\beta_h(e^h) > \beta$. Therefore, the indifference curve of the effort exerting agent is steeper than the indifference curve of an agent with zero effort. This implies that—starting from a full coverage contract—there are partial coverage contracts which strictly relax the binding constraint. Furthermore, the slope of the isoprofit curve of the principal at a full coverage contract with zero effort is $\beta/(1-\beta)$. Therefore, for $\varepsilon > 0$ small enough, a partial coverage contract $(\tilde{u} + \kappa\varepsilon, \tilde{u} - \varepsilon)$ with $\kappa = \beta/(1-\beta)$ will (i) keep the utility of an agent who does not exert effort and buys the contract constant, (ii) will keep profits at the same level and (iii) strictly relax the binding (EH) constraint. Therefore, there exists a partial coverage contract $(\tilde{u}+\kappa\varepsilon, \tilde{u}-\varepsilon)$ with κ slightly below $\beta/(1-\beta)$ such that the binding constraint is not violated while profits are higher than under the full coverage contract $(\tilde{u}, \tilde{u})$. This shows that the contract in a pooling equilibrium will have partial coverage only. □

Proof of Proposition 3: First, we show that the game has a unique optimal contract for $\gamma = 0$ which is separating and induces infinite effort; i.e. the Stiglitz contract menu $(\bar{u}_1^S, \underline{u}_1^S, \bar{u}_2^S, \underline{u}_2^S)$ results if $\gamma = 0$.

If the contract menu is separating (i.e. if strictly positive effort is better than zero effort), then $e = \infty$. To see this, note that the derivative of (5), stated in (6), is strictly positive if the menu is separating and costs are zero. This implies that the Stiglitz contract menu is the best separating contract menu if $\gamma = 0$. It remains to check that no pooling contract leads to higher profits.

With $\gamma = 0$, a pooling contract must satisfy (IRl) for $e = \infty$. Otherwise, the agent has an incentive to deviate to effort $e' = \infty$ and buy the contract only in case he receives a high signal. Hence, the most profitable pooling contract has full coverage and (IRl) binds. But then the standard proof (Stiglitz, 1977, see Property 3) showing that this pooling is not profit maximizing applies. This establishes that the Stiglitz contract menu is the unique optimal menu if $\gamma = 0$.

Second, we show that the set of optimal contract menus as a function of γ is closed. More precisely, let $(\gamma^m)_{m=1}^{\infty}$ be a sequence in $\mathbb{R}_+$ converging to γ. Let $u^m = (\bar{u}_1^m, \underline{u}_1^m, \bar{u}_2^m, \underline{u}_2^m)$ be an optimal contract under γ^m. Then the limit $u = \lim_{m\to\infty} u^m$ is an optimal contract menu under γ (whenever the limit exists). This follows from the continuity of profits (9) in utilities and the fact that the weak inequality constraints (EH) and (IRl) hold in the limit if they hold for every m.

Third, we show that the first two properties imply that the optimal contract menu converges to the Stiglitz contract menu as $\gamma \to 0$. The proof is by contradiction. Suppose, to the contrary, that there was a sequence of γ^m and u^m with $\lim_{m\to\infty} \gamma^m = 0$ such that $\max\{|\bar{u}_1^m - \bar{u}_1^S|, |\underline{u}_1^m - \underline{u}_1^S|, |\bar{u}_2^m - \bar{u}_2^S|, |\underline{u}_2^m - \underline{u}_2^S|\} \geq \varepsilon$ for all m and some $\varepsilon > 0$. Since $u^m \in [\underline{u}_0, \bar{u}_0]^4$, the Bolzano-Weierstrass theorem implies that u^m has a convergent subsequence. By the previous two properties, the limit of this subsequence must be the Stiglitz contract menu. But then $\max\{|\bar{u}_1^m - \bar{u}_1^S|, |\underline{u}_1^m - \underline{u}_1^S|, |\bar{u}_2^m - \bar{u}_2^S|, |\underline{u}_2^m - \underline{u}_2^S|\} \geq \varepsilon$ cannot hold for all m: It will be violated for elements of the convergent subsequence with m high enough. This is the desired contradiction.

Hence, we have shown that equilibrium menus are close to the Stiglitz menu for γ close enough to 0. As the Stiglitz menu is separating, the equilibrium menu for small enough γ cannot be pooling.

Now we turn to the case $\gamma \to \infty$.

First, it is shown that equilibrium effort e is less than some level $\tilde{e}(\gamma)$ and that $\lim_{\gamma\to\infty} \tilde{e}(\gamma) = 0$. As equilibrium contracts (by Proposition 1) have partial coverage, (7) implies $e < \tilde{e}(\gamma)$, where $\tilde{e}(\gamma)$ is the effort level solving $\alpha\beta_h'(\tilde{e})\Delta_0 = \gamma c'(\tilde{e})$. By the implicit function theorem, $\tilde{e}'(\gamma) = \frac{\gamma c'(\tilde{e})}{\alpha\beta_h''(\tilde{e})\Delta_0 - \gamma c''(e)} < 0$ and therefore $\tilde{e}$ is uniquely determined by the defining equation. $\tilde{e}$ exists for all $\gamma \in \mathbb{R}_+$ by the assumption $c'(0) = 0$. From $c'' > 0$, it follows that $c'(e) > 0$ for every $e > 0$. Therefore, $\lim_{\gamma\to\infty} \tilde{e}(\gamma) = 0$ (otherwise, the defining equation cannot be satisfied for γ high enough as β_h' is bounded).

Second, we show that $\lim_{\gamma\to\infty} \tilde{e}(\gamma) = 0$ implies $\lim_{\gamma\to\infty} \Delta_2 = 0$. We divide this step into two cases. In case 1, we concentrate on separating (including exclusion) equilibria. Case 2 covers pooling equilibria.

In case 1, (17) holds with equality. Note that the first two terms in (17) are negative and the third is positive. Therefore, (17) can only hold with equality if

$$(1-\alpha)\beta_l(e)(1-\beta_l(e))(h'(\bar{u}_2)-h'(\underline{u}_2)) < \alpha\left(\beta_h(e)h'(\underline{u}_1)+(1-\beta_h(e))h'(\bar{u}_1)\right)(\beta_h(e)-\beta_l(e)).$$

Note $\lim_{\gamma\to\infty}\beta_h(e)-\beta_l(e)=0$ because $\lim_{\gamma\to\infty}\tilde{e}(\gamma)=0$ and $e<\tilde{e}(\gamma)$. Consequently, the right hand side of the inequality above tends to zero as $\gamma\to\infty$. Therefore, the left hand side also has to converge to 0. As $\lim_{\gamma\to\infty}\beta_l(e(\gamma))=\beta$, this implies that $\lim_{\gamma\to\infty}h'(\bar{u}_2)-h'\underline{u}_2)=0$ which implies $\lim_{\gamma\to\infty}\Delta_2=0$. As $0\le\Delta_1\le\Delta_2$, $\lim_{\gamma\to\infty}\Delta_1=0$ has to hold as well. This completes the proof for case 1.

In case 2, we look at pooling contracts. This implies $0=e<e^h$ and therefore (18) holds with equality. (18) can then be rearranged to (assuming pooling, i.e. $e=0$)

$$h'(\bar{u})=\frac{\beta}{1-\beta}\frac{1-\beta_l(e^h)}{\beta_l(e^h)}h'(\bar{u})$$

where $(\bar{u},\underline{u})$ denotes the equilibrium pooling contract. Note that also $e^h\le\tilde{e}(\gamma)$ and therefore $\lim_{\gamma\to\infty}\beta_l(e^h(\gamma))=\beta$. Therefore, the equation above can hold for high γ only if $\lim_{\gamma\to\infty}\bar{u}(\gamma)-\underline{u}(\gamma)=0$, i.e. the contract converges to full coverage as γ goes to infinity. □

Proof of Proposition 4: Consider first the claim about the profits in the pooling case (where $e=0$). Note that the only binding constraint (EH) is relaxed if γ is increased: Using the envelope theorem, the derivative with respect to γ of the left hand side is $-c(e)$ and the derivative of the right hand side is $-c(e^h)$. This means that a contract satisfying (EH) for γ' will also satisfy (EH) for all $\gamma''>\gamma'$. Hence, profits are non-decreasing in γ as the set of feasible pooling contracts expands as γ increases. Since (EH) is binding, profits are in fact increasing.

In an equilibrium with exclusion, the ex ante individual rationality constraint of a consumer is binding. As the consumer's outside option utility does not depend on γ, the consumer surplus must be constant in γ. □

C. Numerical example

This section shows that both zero effort and positive effort occur in equilibrium in the most straightforward example. For this example, we use the linear signal technology in Appendix A and a square root utility function: $u(x)=\sqrt{x}$.[25] The cost function is

[25] As the function β^h is non-differentiable at $e=1$ and flat thereafter, our assumptions are strictly speaking only satisfied if $e<1$. From the first order condition (7), we can—with the here assumed functions—conclude that $e^3=\frac{\alpha(1-\alpha)(\theta^h-\theta^l)(\Delta_2-\Delta_1)}{4\gamma}$. As $\Delta_2-\Delta_1\le\Delta_0=1$ with our numbers, $e^3\le\frac{0.007875}{\gamma}$ and $e<1$ if $\gamma>0.007875$.

assumed to be $c(e) = \gamma e^4$.[26] As the program is not necessarily quasiconcave, we use a grid maximization where we search for the highest profits on a grid of 500 equally spaced effort levels between 0 and 1 (and an equally spaced grid for e^h). The result of this grid search is then given as starting value to a maximization algorithm.[27] We take $w = 4$ and $D = 3$ which leads to $\bar{u}_0 = 2$ and $\underline{u}_0 = 1$. For $\theta^h = 0.35$, $\theta^l = 0.2$ and $\alpha = 0.7$, the optimal contract is summarized for several values of γ in the following table:

	$\bar{u}_1$	$\underline{u}_1$	$\bar{u}_2$	$\underline{u}_2$	e	e^h	π	EU
$\gamma = 0.05$	1.6929	1.6439	2	1	0.5311	0.5311	0.1408	1.695
$\gamma = 0.1$	1.6946	1.6512	2	1	0.4223	0.4223	0.1424	1.695
$\gamma = 0.2$	1.7052	1.6576	1.9646	1.0954	0.3186	0.3347	0.1442	1.6995
$\gamma = 0.5$	1.7316	1.6415	1.8556	1.3671	0.1844	0.2429	0.1546	1.7066
$\gamma = 0.7$	1.7352	1.6378	1.8263	1.4348	0.1490	0.2165	0.1593	1.7069
$\gamma = 1.0$	1.7559	1.5970	1.7559	1.5970	0	0.1878	0.1643	1.7074
$\gamma = 1.3$	1.7512	1.6050	1.7512	1.6050	0	0.1729	0.1679	1.7067

Table C1: optimal contracts with the parameter values above (rounded on fourth digit)

[26] Also with a quadratic cost function equilibria with positive equilibrium effort exist. However, we did not find parameter values where only varying γ allowed us to move through all three equilibrium types.

[27] We use the "ralg" algorithm of the openopt package (`http://openopt.org`); see our websites for the Python code.

References

Bardey, D. and P. De Donder (2012). Genetic testing with primary prevention and moral hazard. Technical Report 2012-10, Documento CEDE.

Bergemann, D. and J. Välimäki (2002). Information acquisition and efficient mechanism design. *Econometrica 70*(3), 1007–1033.

Bergemann, D. and J. Välimäki (2006). Information in mechanism design. In R. Blundell, W. Newey, and T. Persson (Eds.), *Proceedings of the 9th World Congress of the Econometric Society*. Cambridge University Press.

Crémer, J. and F. Khalil (1992). Gathering information before signing a contract. *American Economic Review 82*(3), 566–578.

Crémer, J., F. Khalil, and J. Rochet (1998). Contracts and productive information gathering. *Games and Economic Behavior 25*(2), 174–193.

Dafny, L. (2010). Are health insurance markets competitive? *American Economic Review 100*(4), 1399–1431.

Doherty, N. and P. Thistle (1996). Adverse selection with endogenous information in insurance markets. *Journal of Public Economics 63*(1), 83–102.

European Commission and others (2007). EU Consumer Policy strategy 2007-2013. Empowering consumers, enhancing their welfare, effectively protecting them. *Communication from the commission to the council, the European Parliament, and the European Economic and Social Committee. COM (2007) 99.*

Lagerlöf, J. N. M. and C. Schottmüller (2013, November). Facilitating consumer learning in insurance markets—what are the welfare effects? Discussion Paper No. 9753, Centre for Economic Policy Analysis.

Ligon, J. A. and P. D. Thistle (1996). Information asymmetries and informational incentives in monopolistic insurance markets. *Journal of Risk and Insurance 63*(3), 434–459.

Persico, N. (2000). Information acquisition in auctions. *Econometrica 68*(1), 135–148.

Peter, R., A. Richter, and P. Thistle (2012). Endogenous information and adverse selection under loss prevention. Mimeo.

Shi, X. (2012). Optimal auctions with information acquisition. *Games and Economic Behavior 74*(2), 666–686.

Stiglitz, J. E. (1977). Monopoly, non-linear pricing and imperfect information: the insurance market. *Review of Economic Studies 44*(3), 407–430.

Szalay, D. (2009). Contracts with endogenous information. *Games and Economic Behavior 65*(2), 586–625.