Inverse selection*

Markus Brunnermeier Rohit Lamba Carlos Segura-Rodriguez

Febrauary 2020

Preliminary

Abstract

Big data, machine learning and AI inverts adverse selection problems. It allows insurers to infer statistical information and thereby reverses information advantage from the insuree to the insurer. In a setting with two-dimensional type space whose correlation can be inferred with big data we derive three results: Frist, a novel tradeoff between characteristics and statistical information emerges. The insurer tries to protect its statistical information by offering only a few screening contracts. Second, we show that forcing the insurance company to reveal its statistical information can be welfare reducing. Third, we show in a setting with naÃŕve agents that do not perfectly infer statistical information from the price of offered contracts, price discrimination significantly boosts insurersâĂŹ profits.

*Brunnermeier: Princeton, markus@princeton.edu; Lamba: Penn State, rlamba@psu.edu; Segura-Rodriguez: Banco Central de Costa Rica, segurarc@bccr.fi.cr.

1 Introduction

> Advances in big data analytics, artificial intelligence and the Internet of Things promise to fundamentally transform the insurance industry and the role data plays in insurance. New sources of digital data, for example in online media and the Internet of Things, reveal information about behaviours, habits and lifestyles that allows us to assess individual risks much better than before.
>
> *International "Geneva" Association for the Study of Insurance Economics*

The rise of big data, artificial intelligence (AI), and machine learning is one of the defining characteristics of the 21st century economy. Almost every action we take is recorded and correlates are constructed, to better predict our behavior.Âă Especially the insurance industry is undergoing a radical transformation. Price discrimination and contract structures will fundamentally change.Âă

Most models in information economics assume that customers have an informational advantage. Hence, the principal, e.g. the insurance company, faces an adverse selection problem, which it tries to mitigate by offering a menu of screening contracts to potential customers.[1] While customers might still have private information about some of their characteristics, with big data the insurance companies develop a informational statistical advantage. They can use new statistical tools to better infer correlates about the characteristics and the ultimate risk. In other words, the insurance companies can better "invert" the mapping from characteristics to risks: by screening a few of the client characteristics the insurance company can gain an information advantage. Thus, big data and AI transform many adverse selection problems to "inverse selection" problems.Âă

Our setting is close in spirit to the informed principal approach in mechanism design (Myerson [1983] and Maskin and Tirole [1990, 1992]), but it differs in two dimensions: first, while the agent has hard private information – family history, eating habits, zip code, etc; the principal has statistical private information – how all these characteristics interact and determine the agent's probability of say, getting cancer; and, second, as a regulatory constraint, it asks the principal to commit to a menu of contracts.

Inverse selection does not only differ from the standard *adverse selection* but also from the more recent *advantageous selection* literature. Advantages selection stresses the importance of preference heterogeneity in order to overturn the standard theoretical, but empirically counterfactual, result that the high-risk agents get full insurance whereas low-risk agents opt for partial insurance. With preference heterogeneity, highly risk-averse agents buy more insurance, despite the fact that they are less risky, since behave more cautiously.[2] In both settings, adverse and advantageous selection, the insurance provider suffers from an informational disadvantage, which is in sharp contrast to

[1] Akerlof [1970] pionnered the study of adverse selection and screening, The core idea has found applications is variety of settings: Rothschild and Stiglitz [1976] study the insurance problem, and Biais et al. [2000] and Tirole [2012] various aspects of financial markets, to name a few. See Laffont and Martimort [2009] for a geenral theoretical treatment of the principal agent screening problem

[2] Einav and Finkelstein [2011] provide an excellent overview of the key ideas. Important papers also include .

our “inverse selection" setting, which in the chronology of ideas may thus be regarded as pointing towards a third generation of models.

We model the inverse selection problem using a two-dimensional type space. Both dimensions determine the riskiness of the agent, but the agent only knows his own characteristic, the first dimension of the type. In contrast, the principal, e.g. the insurer, only knows the (statistical) correlation between the two dimensions of the type. In this sense, she has a better idea how agents' characteristics map into the risk exposure if she were to know the first dimension of the agents' type.

The basic tension the principal faces is the following: She can use a set of screening contracts, i.e. price discrimination, to elicit agent's private information, but she has to beware that by offering more fine-tuned screening contracts, she may partially reveal her iinformation advantage, the statistical correlation. In other words, the principal faces a novel *characteristics-versus-statistics trade-off*. By eliciting more information about agent's characteristics, the principal gives up some of its statistical informational advantage. Note that this trade-off is different from the rent-versus-efficiency trade-off prevalent in standard principle agent problems, where the principal worsens efficient risk-sharing in order to minimize the information rent that the agent can extract. Of course, the standard rent-versus-efficiency trade-off is also present in our setting (with respect to the agent's characteristics).

As in the classical setup, the optimal contract separates along the insuree's private characteristic information. However, along the private statistical information of the insurer, the optimal contract features either complete pooling or partial pooling; interestingly, complete separation along both dimensions is never optimal for the insurer. When the insurer pools certain correlation types she is giving up on price discrimination in order to maintain the statistical information advantage. We show that the insurer always offers a finite number of contracts. This result is based on a novel mechanism design problem that features “*ironing* almost everywhere” (in the sense of Myerson [1981]). To be best of our knowledge, this is the first paper to model beliefs as private information in an otherwise classical mechanism design setup, and deal with ironing on the entire type space. The number of contracts turns out to be small, highlighting that the characteristic versus statistics trade-off is firmly resolved in the favor of the latter. Indeed, in many setting the contract space along the statistical information is only partitioned into two contracts.

To further understand this trade-off it is instructive to consider two special cases: First, if the agent also knows the statistical correlation then the problem collapses to the standard monopolistic Rothschild and Stiglitz [1976] insurance problem. Second, if insurees are naive/gullible and do not try to infer the insurer's statistical information, then the insurer will only offer the two extreme contracts that correspond to lowest and highest possible statistical correlation. In the case with gullible insuree, the principal's returns to added statistical information is significantly higher than with rational insurees. This points towards a market for acquiring consumer information, which in reality has manifested in the rise of data brokers. On the other hand, the

revelation of this information by the insurer is often not beneficial to the insuree, pointing towards a more nuanced understanding of transparency in data management. The insuree's ability to do Bayesian inference is critical for his own welfare. There is a case to be made for limited probabilistic sophistication in the presence of information overload (cite xx). In such scenarios, our analysis points towards large returns from consumer activism in general, and subsidization of agencies that help insurees understand the implications of various choices offered by insurance companies in particular.

Another set of regulatory implications concerns the question of whether the insurer should be forced to reveal her private statistical information to the insuree prior to the posting of contracts. Such a regulatory or societal requirement would ensure that the insuree is not kept in the dark about his own risks. Formally, a mechanism design problem is solved as if the correlation is common knowledge in the extensive form of the interaction, for each possible report of correlation by the insurer. The conceptual innovation here is that the insurer has to be incentivized to reveal the information, and hence a family of shadow prices now constrained the size of the total pie.

2 Model

As discussed in the introduction, we present here what may be regarded as a third generation model of insurance, where the ex ante distribution of information tries to capture the fact that the insuree may also be informed about the underlying state of the world. The model we present can be thought of as an informed principal problem (Myerson [1983]) where the principal can commit to a menu and the private information she holds is statistical.

Preferences. A profit maximizing monopolist insurer (principal/seller) interacts with an insuree (agent/buyer) who wants to insure himself against some damage/loss. The insurer is risk neutral and offers a standard insurance contract (p, x), where p represents the price (or premium), and x represents the proportion of the insuree's loss that is covered by the contract. So, $x < 1$ means under, $x = 1$ means exact, and $x > 1$ means over insurance.

The insuree has an initial wealth w. The uncertain loss he faces is distributed according to a normal distribution $N(\mu, \nu)$. He has a CARA utility parameterized by γ. Let the realized loss be given by ℓ. Then under the contract (p, x), his ex post utility is given by:

$$u_0(w, l, x, p) = -exp(-\gamma(w - p - (1 - x)l)),$$

and his ex ante (expected) utility is given by:

$$u(w, \mu, x, p) = -exp\left(-\gamma\left(w - p - (1 - x)\mu - \frac{\eta}{2}(1 - x)^2\right)\right)$$

where $\eta = \gamma v > 0$. It is well known that maximizing expected utility in a CARA-Gaussian set up is equivalent to maximizing the certainty equilvalent. which is given by

$$\begin{aligned}\mathbb{CE}[u(w,\mu,x,p)] &= w - p - (1-x)\mu - \frac{\eta}{2}(1-x)^2 \\ &= \underbrace{w-\mu}_{a} + \underbrace{\left[x\mu_i - \frac{\eta}{2}(1-x)^2\right]}_{v(x)} - p \\ &= a + v(x) - p\end{aligned}$$

From hereon, we will use notation a bit in writing $u(w,\mu,p,x)$ to mean the certainty equivalent. Since a fully concave utility function makes the analysis intractable, we use the CARA-Gaussian setup to introduce risk aversion while maintaining linearity in money.

Information. The canonical CARA-Gaussian version of the insurance model would assume that the mean loss, μ, is the agent's private information. We depart from this crucial assumption on the "endowment" of information as follows. A relevant bidimensional state $\theta = (\theta_1, \theta_2)$ determines μ, where $\theta_i \in \{L, H\}$ for $i \in \{1, 2\}$. So, given the state θ, the mean loss of the agent is given by μ_θ. Without loss of generality, we assume that

$$\mu_{HH} > \mu_{HL} > \mu_{LL} \quad \text{and} \quad \mu_{HH} > \mu_{LH} > \mu_{LL}.$$

The joint distribution is of θ, given by $q = (q_{HH}, q_{HL}, q_{LH}, q_{LL})$, is depicted in Table 1. Here

		θ_2		
		L	H	
θ_1	L	q_{LL}	q_{LH}	q_1
	H	q_{HL}	q_{HH}	$1-q_1$
		q_2	$1-q_2$	

Table 1: Joint distribution of θ.

$q_1 = q_{LL} + q_{LH}$ and $q_2 = q_{LL} + q_{HL}$ are the marginal distributions of θ_1 and θ_2, respectively. The dimensionality of the problem can be reduced by rewriting the distribution in terms of marginals and the correlation coefficient. Let ρ be the correlation between θ_1 and θ_2, and define $\sigma = (q_1(1-q_1)q_2(1-q_2))^{1/2}$. The distribution can then be rewritten as in Table 2.

The insuree observes θ_1 and knows the marginal distribution of θ_2, and the insurer simply knows the joint distribution of θ. In terms of the primitives, we assume that q_1 and q_2 are common knowledge, the agent is privately informed about θ_1, and the principal privately knows ρ. Finally, to close the model, we assume that ρ is drawn from F on $[-1, 1]$, where F is differential

		θ_2		
		L	H	
θ_1	L	$q_1q_2 + \rho\sigma$	$q_1(1-q_2) - \rho\sigma$	q_1
	H	$(1-q_1)q_2 - \rho\sigma$	$(1-q_1)(1-q_2) + \rho\sigma$	$1-q_1$
		q_2	$1-q_2$	

Table 2: Joint distribution of θ in terms of correlation.

and has a continuous density f, and is common knowledge. lower bound upper bound..

The question we ask is: what is the principal optimal contract in this insurance problem?

Remarks on modeling. A few remarks on our modeling choices are in order. In a direct generalization of the monopolistic screening version of Rothschild and Stiglitz [1976], we could have written down the following model: the bidmensional state θ determined the probability of meeting an accident, say α_θ. The insurer is risk neutral as before, and the insuree has some general concave utility function over final wealth, which is $w - p$ in case of no-accident (with probability $1-\alpha_\theta$ and $w - p + x - l$ in case of an accident (with probability α_θ); and $0 \leq x \leq l$ here is the total coverage in monetary value. The information structure and initial endowment of information would be the same as above: $\mathbf{q}$ is be joint distribution of θ, etc. This model is similar in spirit to the CARA-Gaussian one we write down, but is much harder to solve, because of the lack of structure on the agent's payoff.

In addition we intentionally model the distribution of information between the insurer and insuree as the former knowing ρ and latter knowing θ_1 to capture the idea that [complete..]

3 The optimization problem

To write down the problem formally, we introduce the associated mechanism design lexicon in the spirit of Myerson [1983]. A message rule $r : [\underline{\rho}, \overline{\rho}] \to \Delta(M)$ represents how coarsely the insurer wants to communicate her information about the correlation coefficient to the insuree, as part of the optimal contract. Further, invoking the revelation principle, we simply look at a direct mechanism where the inusree reports her "type" θ_1 to the insurer, and in return the insurer offers her a contract:

$$c = (c_m)_{m \in M} \text{ where } c_m = \{c_m(H), c_m(L)\} \text{ and } c_m(\theta_1) = \big(p_m(\theta_1), x_m(\theta_1)\big) \text{ for } \theta_1 = H, L.$$

A direct mechanism is then completely captured by (r, c), which is chosen by a mediator with the objective of maximizing the profit of the insurer subject to incentive compatibility for the insurer, and incentive compatibility and individual rationality for the insuree.

Let q_ρ be the joint distribution of θ such that the correlation coefficient is restricted to be ρ. The exact timing of the (dynamic) mechanism is as follows.

Stage 1

- nature draws $\rho \sim F \wedge \theta \sim q_\rho$.
- seller learns ρ and reports it.
- r generates message m.
- buyer forms posterior F_m.

Stage 2

- menu $\{c_m(H), c_m(L)\}$ is offered.
- buyer learns θ_1 and reports it.
- contract $c_m(\theta_1)$ is implemented.
- payoffs π and u are realized.

The goal going forward is to characterize the optimal mechanism $(r, \mathbf{c})$. To that end, we now define the objective and constraints of the optimization problem. Let $\pi(\rho, \hat{\rho})$ be the (ex post) profit of the insurer if her type is ρ but she reports $\hat{\rho}$ to the mediator. So, under truthtelling, the optimal profit is given by $\pi(\rho; \rho)$ which we will simply refer to simply as $\pi(\rho)$. The (ex ante) objective of the mechanism design exercise is then given by:

$$\Pi = \int \pi(\rho) f(\rho) d\rho.$$

For a fixed menu c_m, the payoff of the insuree type $\theta_1 \in \{H, L\}$ from reporting $\hat{\theta}_1$ is:

$$\begin{aligned} u_m(\theta_1; \hat{\theta}_1) &= w - p_m(\hat{\theta}_1) - \left[1 - x_m(\hat{\theta}_1)\right] \mu_m(\theta_1) - \frac{\eta}{2}\left[1 - x_m(\hat{\theta}_1)\right]^2 \\ &= \underbrace{w - \mu_m(\theta_1)}_{a_m(\theta_1)} + \underbrace{\left[x_m(\hat{\theta}_1)\mu_m(\theta_1) - \frac{\eta}{2}\left\{1 - x_m(\hat{\theta}_1)\right\}^2\right]}_{v_m(\theta_1; \hat{\theta}_1)} - p_m(\hat{\theta}_1) \\ &= a_m(\theta_1) + v_m(\theta_1; \hat{\theta}_1) - p(\hat{\theta}_1) \end{aligned} \tag{1}$$

where $\mu_m(\theta_1)$ is the expected value of μ based on realized value of θ_1 and $\rho \sim F_m$. Assuming truthteling by the agent, the mathematical expression for the insurer's profit from reporting can then be written as:

$$\pi(\rho, \hat{\rho}) = q_1 \left[p_{r(\hat{\rho})}(L) - \mu_\rho(L) x_{r(\hat{\rho})}(L)\right] + (1 - q_1)\left[p_{r(\hat{\rho})}(H) - \mu_\rho(H) x_{r(\hat{\rho})}(H)\right] \tag{2}$$

where $\mu_\rho(\theta_1)$ is the expected value of μ based on realized value of ρ and (truthfully) reported value of θ_1.

Three type of constraint are imposed on the optimization problem. First is the incentive constraint of the insurer, that the insurer wants to truthfully report her type to the mediator:

$$IC_\rho : \pi(\rho; \rho) \geq \pi(\rho; \hat{\rho}) \; \forall \; \hat{\rho}.$$

Second is the incentive constraint for the insuree, that the insuree wants to truthfully report his type to the mediator:

$$IC_{\theta_1} : u_m(\theta_1;\theta_1) \geq u_m(\theta_1;\hat{\theta}_1) \; \forall \; \hat{\theta}_1.$$

And, third, is the individual rationality constraint of the insuree which guarantees him a minimum expected utility is any scenario:

$$IR_{\theta_1} : u_m(\theta_1;\theta_1) \geq .$$

Finally, the optimization problem can then be written simply as:

$$\max_{r,c} \; \Pi \text{ s.t. } IC_\rho, IC_{\theta_1}, IR_{\theta_1}.$$

4 Three "special" cases

Before we solve problem the main problem, we consider three related models that help identify the key economic forces at work. In what follows the mapping r is either superfluous or independent of the set of contracts c. In each case the insurer will find it optimal to offer contracts that change continuously as a function of ρ, so the set of contracts will be written as $c = (c_\rho)_{\rho\in[\underline{\rho},\overline{\rho}]}$.

4.1 ρ is common knowledge

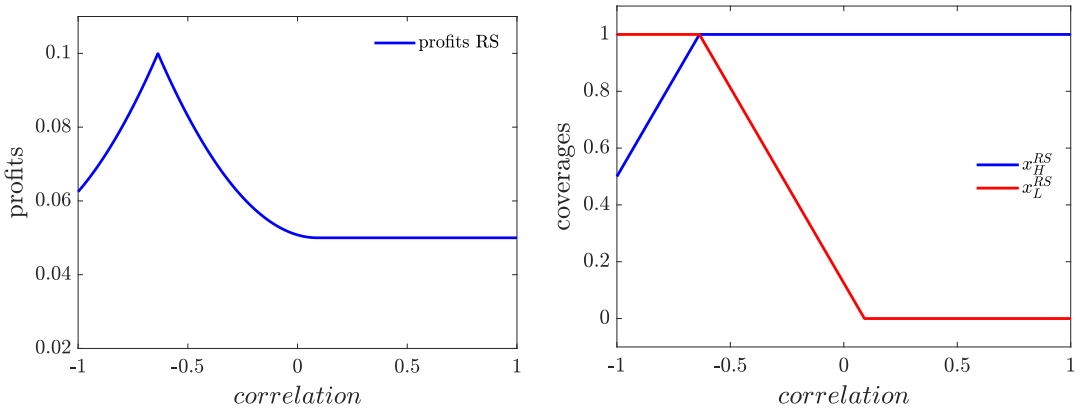
If ρ is common knowledge, our problem becomes isomorphic to the monopolistic version of the classical Rothschild and Stiglitz [1976] problem. Both parties take expectations over θ_2, and insuree is incentivized to reveal θ_1 truthfully. There exists a level of correlation, say ρ^*, above which the "high" risk type is $\theta_1 = H$, and below which the "high" risk type is in fact $\theta_1 = L$.

Proposition 1. *$\exists\, \rho^*$ s.t. $\pi^{RS}(\rho^*) = \max_\rho \pi^{RS}(\rho)$, and*

1. $\rho > \rho^* \Rightarrow 1 = x_\rho^{RS}(H) > x_\rho^{RS}(L)$,
2. $\rho < \rho^* \Rightarrow x_\rho^{RS}(H) < x_\rho^{RS}(L) = 1$.

Fix ρ^* to be the correlation where the expected value of mean loss is the same for both θ_1-types: that is ρ^* solves $\mu_\rho(H) = \mu_\rho(L)$. Then, for $\rho > \rho^*$, the $\theta_1 = H$ gets full insurance, and $\theta_1 = L$ gets partial insurance, while the opposite is true for $\rho < \rho^*$. As in the standard monopolistic screening model, the optimal contracting is always separating: "high" risk type is offered full coverage and "low" risk type is offered partial coverage.

The extent of distortion for the "low" risk type is further determined by the primitives of the problem. The economic force driving this result is typically known as the rent versus efficiency tradeoff. Since insuree is the residual claimant of the surplus, she wants to maximize efficiency which entails offering a separating contract, but subject to the payment of information rent that



(a) optimal profit as a function of ρ (b) optimal coverage as a function of ρ

Figure 1: Benchmark model when ρ is common knowledge

incentivizes insurers to self select into the contract corresponding to his type. In fact if the proportion of "low" risk types is too low, the insurer will simply offer one contract to the "high" risk types and exclude the "low" risk ones from the market.

Figure 1 plots the insurer's profits and optimal coverages as a function of the correlation. We will refer to this as the *benchmark model*, and christen it RS, pointing towards the classical reference. Figure 1a shows that the profit of the insurer is single peaked at ρ^*, and Figure 1b showa that coverages are generically separating with the "high" risk type pivoting around ρ^*.

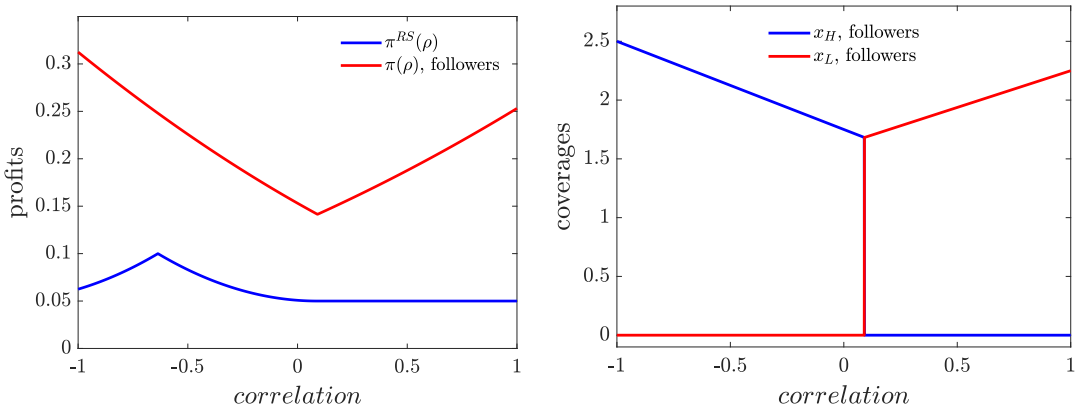
4.2 Gutgläubig insuree

Another useful, and rather non-standard, model to consider is one where the insuree tells the insurer the correlation coefficient in addition to offering a contract, and he simply believes it. This setting is different than (standard) naïveté model that we discuss in the next subsection. We will refer to such an insurer as *gutgläubig*, which is a German word that approximately translates to gullibly trusting.

Knowing that she can basically mislead the insurer about the way in which the two states correlate to produce the actual probability of an accident provides the insuree with great freedom in selecting contracts, formally it allows her to choose r and c independently.

Proposition 2. *If the buyer is a gutgläubig,* $\exists\ \tilde{\rho} \in (\underline{\rho}, \overline{\rho})$ *such that:*

1. *binary messages are sent:* $M = \{\underline{m}, \overline{m}\}$ *s.t.* $r(\rho) = \overline{m}$ *for* $\rho < \tilde{\rho}$ *and* $m(\rho) = \underline{m}$ *for* $\rho > \tilde{\rho}$,
2. *posterior of the inuree is extreme:* $F_{\underline{m}} = \delta_{\underline{\rho}}$ *and* $F_{\overline{m}} = \delta_{\overline{\rho}}$,
3. *profits are uniformly higher than benchmark:* $\pi(\rho) > \pi^{RS}(\rho)\ \forall\ \rho$,
4. *coverages are generically separating and inexact:* $x_\rho(H) \neq x_\rho(L)\ \forall \rho \neq \tilde{\rho}$, *and* $x_\rho \neq 1\ \forall\ \rho$ *a.s.*,
5. *coverages are flipped in comparison to the benchmark:*



(a) optimal profit as a function of ρ (b) optimal coverage as a function of ρ

Figure 2: Model with gutgläubig insuree

(a) $x_\rho(H) > x_\rho^{RS}(H)$ *and* $x_\rho(L) < x_\rho^{RS}(L)$ *for* $\rho < \tilde{\rho}$,

(b) $x_\rho(H) < x_\rho^{RS}(H)$ *and* $x_\rho(L) > x_\rho^{RS}(L)$ *for* $\rho > \tilde{\rho}$.

Figure 2 depicts the profit and coverages when the insuree is gutgläubig. That the profits are uniformly higher is intuitive: the model allows the insurer to input any value of the correlation in the insuree's incentive compatibility condition, which in turn allows her to manipulate which type is perceived to be high risk type the insuree, $\theta_1 = H$ or $\theta_1 = L$, and then decide what coverages to offer each θ_1 type.

Part 1 of the result states that the ...

4.3 Naive insuree

A more standard "behavioral" way of modeling limitations on information processing is to assume that the agent ignores the signals offered by the contract about the correlation coefficient, so that $F_m = F \ \forall m \in M$.[3]

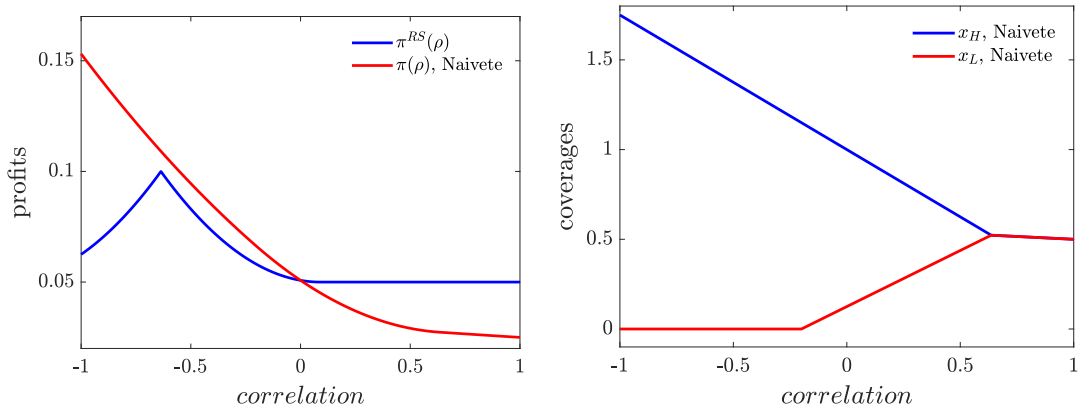
Proposition 3. *If the buyer is naive (and thus sticks to the prior):*

1. *profits are higher in expectation:* $\mathbb{E}(\pi(\rho)) > \mathbb{E}(\pi^{rs}(\rho))$,
2. *coverages features both pooling and separation,*
3. *coverages are generically inexact:* $x_\rho(\theta_1) \neq 1 \ \forall \ \rho$ *a.s.,*
4. *coverages are flipped in comparison to the benchmark:*

   (a) $x_\rho(H) > x_\rho^{RS}(H)$ *and* $x_\rho(L) < x_\rho^{RS}(L)$ *for* $\rho < \mathbb{E}(\rho)$,

   (b) $x_\rho(H) < x_\rho^{RS}(H)$ *and* $x_\rho(L) > x_\rho^{RS}(L)$ *for* $\rho > \mathbb{E}(\rho)$.

[3] This is for example the chosen way in benanou and tirole xx



(a) optimal profit as a function of ρ (b) optimal coverage as a function of ρ

Figure 3: Model with naive insuree

Figure 3 plots depicts the profit and coverages for the naïveté model. While the profits are higher in expectation they are not uniformly so. This might seem a bit counterintuitive at first. The coverages are feature both separation and pooling for different values of ρ.

5 Characterizing incentive compatibility

Our model differs from the standard screening problem in that it also features an incentive constraint for the principal. It also differs somewhat from the standard informed principal problem in that the principal can commit to post a menu of contracts. In this section we carefully analyze the incentive constraints of the principal for any fixed reporting strategy $r : [\underline{\rho}, \overline{\rho}] \to \Delta(M)$ by the mediator.

The standard envelope condition gives us the Myersonian characterization of the insurer's incentive compatibility.

Lemma 1. *IC_ρ holds if and only if*

$$\left.\frac{\partial \pi(\rho, \rho')}{\partial \rho}\right|_{\rho'=\rho} = \sigma x(L, \rho)(\mu_{LH} - \mu_{LL}) - \sigma x(H, \rho)(\mu_{HH} - \mu_{HL}) \equiv c(\rho)$$

and, $\pi(.)$ is non-increasing in ρ.

But, we can say more: the slope of the profit function, i.e. $c(\rho)$ is also required to be non-decreasing in the correlation.

Lemma 2. *$IC_\rho \implies c(.)$ is non-increasing.*

Proposition 4. *$IC_\rho \implies c(.)$ is a step function with a countable number of steps, and $\pi(\rho)$ is piecewise linear with the non-decreasing slopes.*

Recollect $C = \{c^m \mid m \in \mathcal{M}\}$.

Definition 1. *A contract is partitional if C is a countable set and f-partitional if it is finite.*

Proposition 4 immediately implies that C is C-parititonal. In the next section we show that is in fact f-partitional.

6 Optimal contract

In the previous section we showed that ironing almost everywhere will lead to an optimal contract that features a countable number of contracts. Here we show that the cardinality is finite, and in fact, under reasonable restrictions on the primitives, at most two contracts are offered. Then, we provide a precise characterization of the optimal contract that features complete pooling, and one that features two elements.

Lemma 3. *The optimal contract is f-partitional.*

But, we can say more...

Theorem 1. *The optimal contract features at most two partitions.*

One partition..

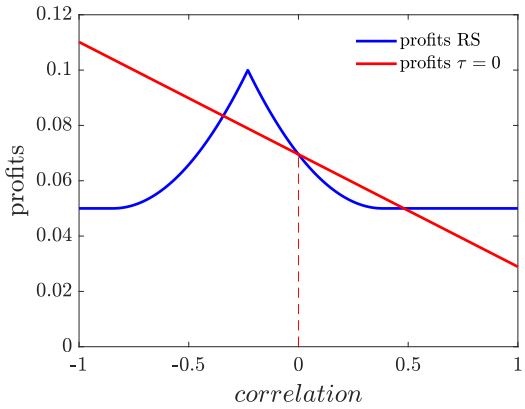
Proposition 5. *When the optimal contract chooses $|\mathcal{M}| = |C| = 1$:*

1. *optimal expected profits are the same as in benchmark: $\mathbb{E}[\pi(\rho)] = \pi^{RS}(\mathbb{E}[\rho])$,*
2. *optimal coverages are the ones offered for the expected correlation in the benchmark: $x_H(\rho) = x_H^{RS}(\mathbb{E}[\rho])$ and $x_L(\rho) = x_L^{RS}(\mathbb{E}[\rho])$.*

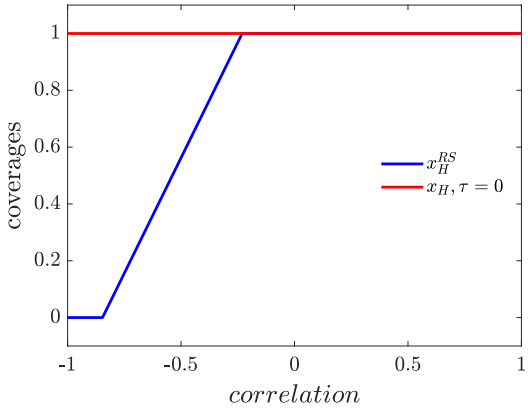
Two partitions..

Proposition 6. *When the optimal contract chooses $|\mathcal{M}| = |C| = 2$, let $\mathbb{I}_1$ and $\mathbb{I}_2$ be the two intervals in the partition of $\mathcal{M}$. In the optimal contract:*

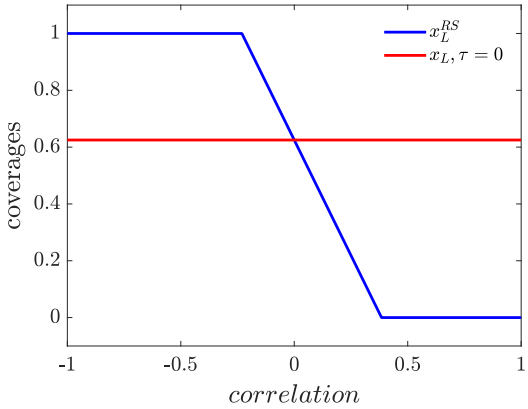
1. *profits are linear in correlation: $\frac{d\pi(\rho)}{d\rho} = c$ f.s. c,*
2. *both separation and pooling in coverages can arise,*
3. *coverages are flipped in comparison to the benchmark:*
   (a) $x_H^1 \geq x_H^{rs}(\rho^1)$ and $x_L^1 \leq x_L^{rs}(\rho^1)$, whenever $x_H^1 \neq x_L^1$,
   (b) $x_H^2 \leq x_H^{rs}(\rho^2)$ and $x_L^2 \geq x_L^{rs}(\rho^2)$, whenever $x_H^2 \neq x_L^2$.



(a) optimal profit

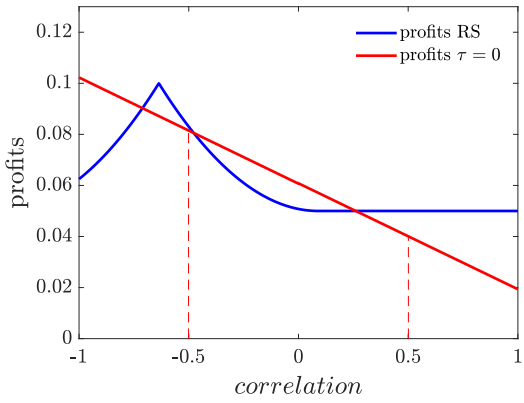


(b) optimal coverage for $\theta_1 = H$

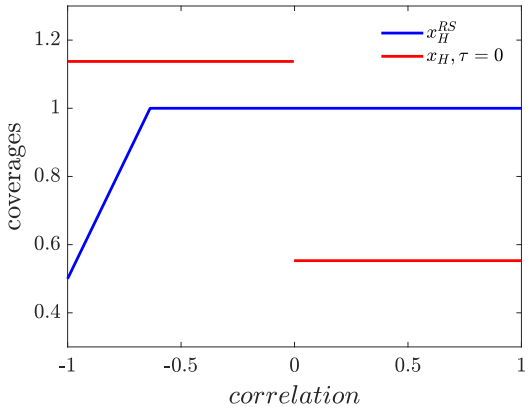


(c) optimal coverage for $\theta_1 = L$

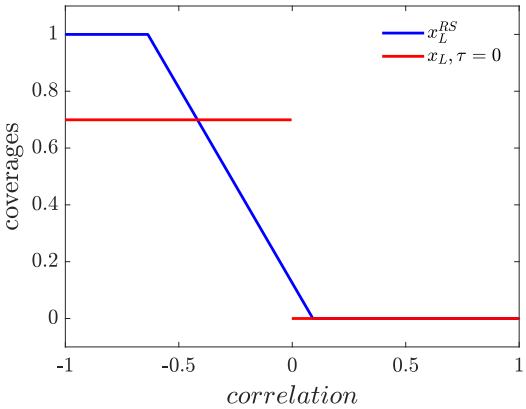
Figure 4: Optimal contract features complete pooling



(a) optimal profit



(b) optimal coverage for $\theta_1 = H$



(c) optimal coverage for $\theta_1 = L$

Figure 5: Optimal contract features two partitions

7 Appendix

Proof of Proposition 4. Let us focus first on the second stage game. Suppose that the firm's IC constraint does not bind, i.e., for a fixed ρ the mediator wants to solve problem

$$
\begin{aligned}
(P-\rho) \qquad & \max_{p(L,\rho),p(H,\rho),x(L,\rho),x(H,\rho)} && \pi(\rho,\rho) && \\
& s.t. && U(\theta_1,\theta_1,\rho) \geq U(\theta_1,\theta_1',\rho) \quad \forall \theta_1,\theta_1' \in \{L,H\} && \text{IC } \theta_1-\theta_1' \\
& && U(\theta_1,\theta_1,\rho) \geq \bar{U}(\theta_1,\rho) \quad \forall \theta_1 \in \{L,H\} && \text{IR } \theta_1
\end{aligned}
$$

This problem is the classical insurance problem after we have fixed the expected means $\mu(L,\rho)$ and $\mu(H,\rho)$. Lemma 4 implies that the solution to this relaxed problem is not IC compatible for the firm, since it does not satisfy the condition in equation **??**.

Lemma 4. *In the solution to problem $P-\rho$, if $\mu(H,\rho) > \mu(L,\rho)$, $x(H,\rho) = 1 > x(L,\rho)$ and $\frac{\partial x_L(\rho)}{\partial \rho}\big|_{x_L(\rho)>0} < 0$; and if $\mu(H,\rho) < \mu(L,\rho)$, $x(L,\rho) = 1 > x(H,\rho)$ and $\frac{\partial x_H(\rho)}{\partial \rho}\big|_{x_H(\rho)>0} > 0$. Therefore, $c(\rho)$ is strictly decreasing.*

Proof. Suppose that $\mu(H,\rho) > \mu(L,\rho)$. Then the only constraints that bind are *IC H-L* and *IR L*. Letting λ be the multiplier in the first constraint and δ the multiplier in the second constraint, we have that the FOCs that characterize an interior solution are given by:

$$
\begin{aligned}
& q_1 - \delta + \lambda = 0 \\
& -\mu(L,\rho)q_1 + \delta\mu(L,\rho)q + \sigma(1-x(L,\rho))\delta - \lambda\mu(H,\rho) - \lambda\eta(1-x(L,\rho)) = 0 \\
& (1-q_1) - \lambda = 0 \\
& -(1-q_1)\mu(H,\rho) + \lambda\mu(H,\rho) + \eta(1-x(H,\rho))\lambda = 0.
\end{aligned}
$$

From the first and third conditios it is easy to conclude that $\lambda = (1-q_1)$ and $\delta = 1$. Using these values it is straighforward to see that $x(H,\rho) = 1$ and $x(L,\rho) = 1 - \frac{1-q_1}{\eta q_1}(\mu(H,\rho) - \mu(L,\rho))$.

Since $\frac{\partial \mu(L,\rho)}{\partial \rho} = \frac{\sigma}{q_1}(\mu_{LL} - \mu_{LH}) < 0$ and $\frac{\partial \mu(H,\rho)}{\partial \rho} = \frac{\sigma}{1-q_1}(\mu_{HH} - \mu_{HL}) > 0$, we have that $\frac{\partial x_L(\rho)}{\partial \rho}\big|_{x_L(\rho)>0} < 0$.

Finally, in a corner solution it is trivial to see that $x(H,\rho) = 1$ and $x(L,\rho) = 0$.

An analogous argument shows the result for the case in which $\mu(H,\rho) < \mu(L,\rho)$. □

In other words the solution to the relaxed version of the mediator's problem without the non-decreasing condition does not satisfy this constraint whenever there is an interior solution. Furthermore, this is true even if the non-decreasing condition is not impose for some arbitrary interval. This implies that the non-decreasing condition binds almost everywhere, meaning that in the solution to the mediator's problem the function $c(\rho)$ has to be a step function, which implies that the function $\pi(\rho,\rho')$ is a piecewise linear function, a result stated in Proposition 4.

□

References

G. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

B. Biais, D. Martimort, and J. Rochet. Competing mechanisms in a common value environment. *Econometrica*, 68(4):799–837, 2000.

L. Einav and A. Finkelstein. Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, 25(1):115–138, 2011.

J.-J. Laffont and D. Martimort. *The theory of incentives: the principal-agent model*. Princeton university press, 2009.

E. Maskin and J. Tirole. The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, 58(2):379–409, 1990.

E. Maskin and J. Tirole. The principal-agent relationship with an informed principal ii: Common values. *Econometrica*, 60(1):1–42, 1992.

R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

R. B. Myerson. Mechanism design by an informed principal. *Econometrica*, 51(6):1767–1797, 1983.

M. Rothschild and J. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 90(4):629–649, 1976.

J. Tirole. Overcoming adverse selection: How public intervention can restore market functioning. *American Economic Review*, 102(1):29–59, 2012.