Inverse selection[*]

Markus Brunnermeier Rohit Lamba Carlos Segura-Rodriguez

December 2021

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: First, a novel tradeoff between a belief gap and price discrimination emerges. The insurer tries to protect its statistical information by offering only a few screening contracts. Second, we show in a setting with naïve agents that do not perfectly infer statistical information from the price of offered contracts, price discrimination significantly boosts insurer's profits. Third, introducing competition sucks out the informational advantage of the insurer, and forcing the monopolistic insurer to reveal its statistical information can be helpful to the insuree even though it may reduce total surplus. We also discuss the significance of our analysis through four stylized facts: the rise of data brokers, the perils of market concentration with advent of big data, the importance of consumer activism and regulatory forbearance, and the merits of a public data repository.

Keywords: Insurance, Big Data, Informed Principal, Belief Gap, Price Discrimination.

JEL codes: G22, D82, D86, C55.

[*]Brunnermeier: markus@princeton.edu, Princeton University; Lamba: rlamba@psu.edu, Pennsylvania State University; Segura-Rodriguez: segurarc@bccr.fi.cr Banco Central de Costa Rica. We are indebted to Hanming Fang (discussant at NBER-Summer Institute 2021), Nemanja Antić, Ben Handel, Rakesh Vohra, Motohiro Yogo, and participants at the NBER-Insurance Conference 2020, Decentralization Conference 2021 and NASM-Econometric Society 2021 for useful comments.

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, Keller et al. [2018]*

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. The direct effects of these developments are being felt in the insurance industry, which is undergoing a radical transformation— price discrimination and contract structures will fundamentally change.

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, big data allows insurance companies to develop superior aggregate information, using new statistical tools to better infer correlates about the characteristics and the ultimate risk. In other words, the principal here can "invert" the mapping from characteristics to risks through an informational and technical advantage. Thus, big data and AI transform many adverse selection problems to what we call "inverse selection" problems.

Our setting is close in spirit to the informed principal approach in mechanism design (Myerson [1983] and Maskin and Tirole [1990, 1992]). It departs from the canonical structure in two ways: 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. Also, the basic structure of our model is inspired from the classical insurance problem studied by Rothschild and Stiglitz [1976] with two key differences: we consider a richer information structure, and for the most part restrict attention to a monopolistic screening setup.

Inverse selection does not only differs 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

[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, Mailath and Postlewaite [1990] study public goods provision, and Biais, Martimort, and Rochet [2000] and Tirole [2012] study various aspects of financial markets, to name a few. See Green and Laffont [1979] and Laffont and Martimort [2009] for general theoretical treatments of the principal agent screening problem.

agents opt for partial insurance. With preference heterogeneity, highly risk-averse agents buy more insurance, despite the fact that they are less risky, since they behave more cautiously.[2] In both settings, adverse and advantageous selection, the insurance provider suffers from an informational disadvantage. This is in contrast to our inverse selection setting, which in the chronology of ideas may thus be regarded as 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, and the marginal distribution along both dimensions is common knowledge. The agent perfectly knows one (type of) characteristic, the first dimension of the type. In contrast, the principal, e.g. the insurer, knows the entire joint distribution, her statistical advantage manifests in private information about the correlation between the two dimensions. At a high level, we equip the agent with greater hard or physical information and the principal with greater soft or statistical information. This marks a departure from most standard principal-agent models of asymmetric information.[3]

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 informational advantage, the statistical correlation. In other words, the principal faces a *belief gap-versus-price discrimination* trade-off. By offering a richer set of contracts, the principal can discriminate more but will also end up giving up some of its statistical informational advantage. Note that this trade-off is different from the *rent-versus-efficiency trade-off* prevalent in standard principal 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 private information.

As in the classical setup, the optimal contract separates along the insuree's private information. However, along the private statistical information of the insurer, the optimal contract features either complete pooling or partial pooling; 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, i.e. the belied gap. We show that revealing small bits of information is too costly, so that the insurer prefers to offer only a finite number of contracts. In the language of Myerson [1981], the optimum features *ironing* almost everywhere. In fact, we further show that the number of contracts turns out to be small, highlighting that the belief gap-versus-price discrimination trade-off is firmly resolved in the favor of the former. In most cases the contract space along the statistical information is partitioned into one or two contracts.[4]

[2] Einav and Finkelstein [2011] provide an overview of the key ideas. See Finkelstein and McGarry [2006] and Fang, Keane, and Silverman [2008] for empirical evidence on adverse and advantageous selection.

[3] The model can be equivalently interpreted as the first dimension being the set of all characteristics and the second dimension being the riskiness of the agent. Then the agent has private information about personal characteristics, and the principal understands the mapping between characteristics and risks.

[4] Eilat, Eliaz, and Mu [2020] study a standard quasi-linear monopolistic screening where the information

For a large class of parameters, the optimal contract features one partition, i.e. complete pooling along the private information of the insurer. So the insurer prefers to not use her information in contract design at all—the benefit of greater of price discrimination is offset by cost imposed by larger number of incentive constraints to satisfy for the insuree.

To better understand this trade-off it is instructive to consider a few "special cases": First, we say the insuree is *gutgläubig* if he does not infer any statistical information from the menu of contracts and in addition believes whatever the insurer tells him about the correlation coefficient. For such an insuree, only two correlations are ever reported— the lowest and highest possible values, and a distinct contract is chosen for each possible actual realization of the correlation. The low risk agent is overinsured and the high risk agent is often excluded from the market. This model, although theoretically non-standard, clarifies the direction in which the insurer would like to push the contract if she could create the maximal belief gap and implement the maximal price discrimination. The profits of the insurer are uniformly higher in comparison to the standard model. Second, we say that insuree is *naïve* if he again does not infer any statistical information from the menu of contract, but unlike gutgläubig, sticks to the prior. Here too the insurer gains on average, but ex post the ranking is not uniform: dictated by feasibility constraints, the insurer would like the insuree to update his belief (even correctly) in certain situations. For the naïve case, the belief gap is exogenously fixed by the prior and the insurer maximizes on the price discrimination channel, given this constraint.[5]

A reasonable regulatory question to ask is 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 innovation here is that the insurer has to be incentivized to reveal the information, and hence a family of shadow prices now constrain the size of the pie. The profit of insurer is uniformly reduced (and sometimes the total size of the pie too), but the hope is that it can still increase the insurance coverage towards the efficient value. We show that the insuree's surplus is uniformly positive, the variance in insurance coverage decreases (so less discrimination), and the total coverage may go up or down as a function of the primitives of the environment.

Now, in each of the four cases, the standard model, *gutgläubig*, naïve, and optimal full revelation, we compare the insurance premiums to the benchmark model where the statistical correlation is common knowledge at the outset— in this latter case, the problem collapses

change of the principal is exogenously restricted by a cap on KL-divergence between the prior and posterior. They too find that the number of contracts offered at the optimum is finite. Their model, mechanism and the application are however quite different than ours.

[5] In recent work, Fang and Wu [2020] also study a model of belief based behavioral biases and how these can be exploited systematically by firms in the market for insurance contracts.

to the standard monopolistic Rothschild and Stiglitz [1976] insurance problem, but with a twist. What is the "high risk" or "low risk" type is determined by the realized correlation between the two dimensions. The same piece of information would signify a high risk type if correlation is high and it would signify a low risk type if the correlation is low.

Finally, we introduce competition into the model in the following tractable way: There are other 'regular' insurers who do not have the big data technology available and are thus uninformed about correlation structure between the two dimensions. They offer the benchmark (or Rothschild-Stiglitz) contract averaging over all possible correlations in order to screen the insuree along the first dimension. How does this outside competition impact the insurer with big data at its disposal? If the insuree is *gutgläubig*, it reduces the extent to which the insurer can be misled and reduces price discrimination. If the insuree is can do full Bayesian inference, it reduces the number of contracts offered at the optimum; in fact for most parameters, the insurer completely pools along its private insurance offering the same contract as other 'regular' insurers. The introduction of competition benefits the insuree by reducing the extent to which statistical information can be used in an adversarial fashion.

While our model is admittedly stylistic, it provides a framework to think about the role of big data and AI in the design of screening contracts. The contrast between our standard model and the gutgläubig case shows that the returns to statistical information for the principal can be quite large, especially when the agents are not sophisticated. This points towards a market for acquiring consumer information, which in reality has manifested in the rise of data brokers such as Oracle, Nielsen and Salesforce; see, for example, Financial Times [2019]. On the other hand, the limits to exploitation of consumer data when consumers are completely sophisticated points towards the returns to consumer activism and greater regulatory forbearance; see, for example, the call for transparency by the Federal Trade Commission (Ramirez et al. [2014]) and the framework for a general data protection regulation issued by the European Parliament (Council of the European Union [2016]). The positive surplus guarantee and greater equality of insurance provision from forcing the principals to make private statistical information public points towards the merits of a public data repository; see, for example, Rajan [2019]. Finally, competition leading to a reduction in extent of price discrimination along the dimension of statistical information generated from big data nudges towards the regulation of data monopolies; see, for example, Khan [2017].

The informed principal problem seems to us a likely candidate to capture the essence of inverse selection. To the best our knowledge, Villeneuve [2005] is the first paper to think systemically about insurance markets in the realm of the informed principal model. This has been followed up by Abrardi, Colombo, and Tedesch [2020], simultaneously, with our work. Both these papers though focus on competing principals, in contrast to our monopolistic setup. Moreover, Villeneuve [2005], and for the most part, Abrardi et al. [2020] focus on one-dimensional private information on the side of the principal, whereas, we look at a two-dimensional state, part of which is known to the principal and part is known to the agent.[6] In recent work, Luz, Gottardi, and Moreira [2020] and Bhaskar, McClellan, and Sadler [2021] also look at a two-dimensional type space for insurance contracts; the former considers heterogeneity in preferences, specifically risk-aversion, as the second-dimension and the latter assumes that the first dimension is commonly known and can used by a third party such as a regulator to offer a large number of contracts to implement the efficient outcome. While the setups and results of all these papers are quite different to ours, we view these papers as being complimentary to our work in a push towards the aforementioned "third generation" of insurance models.

2 Model

As discussed in the introduction, the model we present can be thought of as an informed principal problem (Myerson [1983]): A risk-neutral principal holds some statistical private information about an underlying state and can commit to a menu of insurance contracts to screen the agent. The agent is risk averse and holds some hard private information about the risk he faces.

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 insurance, $x = 1$ means exact insurance, and $x > 1$ means over insurance.

The insuree has an initial wealth w. The uncertain loss he faces is a random variable with a well defined mean μ and variance ν. So, given a contract (p, x) and realized loss ℓ, his final wealth is given by $z = w - p - (1-x)\ell$. The insuree is assumed to have a standard mean-variance preference. Thus, his utility is given by:

$$u(x, p) = \mathbb{E}[z] - \frac{\gamma}{2}\mathbb{V}[z] = w - p - (1-x)\mu - \frac{\eta}{2}(1-x)^2$$

where γ measures the extent of risk aversion, $\mu = \mathbb{E}[\ell]$ and $\eta = \gamma \times \mathbb{V}[\ell] = \gamma \times \nu$ captures the level of risk faced by the insuree. This expression can be simplified further as follows:

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

which lends a tractable structure to the insuree's preferences so that his utility is linear in

[6] Beyond insurance markets, see also Mylovanov and Tröger [2014] and Koessler and Skreta [2019] for related theoretical models of the informed principal.

money and concave in the extent of loss.[7]

Information. A standard approach to 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 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	θ_2: 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. Let ρ be the correlation between θ_1 and θ_2, and define $\sigma = \sqrt{q_1(1-q_1)}\sqrt{q_2(1-q_2)}$. As shown in Table 2, the distribution can then be rewritten using three parameters: ρ, q_1, q_2.

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

The insuree observes θ_1 and knows the marginal distribution of θ_2, and the insurer 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 $[\underline{\rho}, \overline{\rho}]$, where F is differentiable, has a continuous density f, and is common knowledge.[8]

[7] A standard behavioral foundation for the mean-variance preference is the CARA-Gaussian model, which has been used in many seminal papers, including Grossman and Stiglitz [1980].

[8] The entire set of possible correlation is of course [-1,1]. However, once we fix the marginals to be q_1 and

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

Remarks on modeling. A few remarks on 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 determines 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 $x \geq 0$ 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 the joint distribution of θ, etc. This model is similar in spirit to the 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 the insurer has some statistical knowledge and the insurer has some concrete knowledge about the underlying state. After the endowment of initial information, the insurer knows more about the general environment in the form of the correlation coefficient between the two dimensions, and the insuree knows something specific about his situation in the form of θ_1. Once the insurer incentivizes the insuree to reveal θ_1, the insurer can make better inference about the state than the insuree, this inverts the selection problem.

3 The optimization problem

To write down the problem formally, we introduce the associated mechanism design lexicon in the spirit of Myerson [1982, 1983]. A message rule $r : [\underline{\rho}, \overline{\rho}] \to \Delta(\mathcal{M})$ represents how coarsely (or finely) 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 insurer reports her "type" ρ, the insuree reports his "type" θ_1, and a contract is selected from the menu:

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

A direct mechanism is then completely captured by $(r, \mathcal{C})$, 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.

The exact timing of the (dynamic) mechanism is as follows.

q_2, it can be easily checked that the set of feasible correlations is $[\underline{\rho}, \overline{\rho}]$, where $\overline{\rho} = \min\left\{\frac{q_1(1-q_2)}{\sigma}, \frac{q_2(1-q_1)}{\sigma}\right\}$ and $\underline{\rho} = \max\left\{-\frac{q_1 q_2}{\sigma}, -\frac{(1-q_1)(1-q_2)}{\sigma}\right\}$. Thus, the support of F is restricted by the marginals q_1 and q_2.

Stage 1	**Stage 2**
• nature draws $\rho \sim F \wedge \theta \sim \mathbf{q}$. • insurer learns ρ and reports it. • r generates message m. • insuree forms posterior F_m.	• menu $\{c_m(H), c_m(L)\}$ is offered. • insuree 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 choice of $(r, \mathcal{C})$. 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 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:

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

where $\mu_m(\theta_1)$ is the expected value of μ based on the realized value of θ_1 and the insuree's beliefs about ρ after observing the message m. Assuming truthteling by the agent, the mathematical expression for the insurer's profit is:

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

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 types 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:

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

As pointed out in the description of the dynamic mechanism above, insuree's incentive constraint incorporates the report of the insurer by conditioning the (expected) utility on the message m, and hence the posterior F_m. Third is the individual rationality constraint of the insuree which guarantees him a minimum expected utility:

$$IR_{\theta_1} : u_m(\theta_1; \theta_1) \geq 0.$$

Any contract $(r, \mathcal{C})$ that satisfies these three (class of) constraints is said to be *incentive-feasible*. Finally, the optimization problem can be written simply as:

$$\boxed{\max_{r,\mathcal{C}} \ \Pi \text{ s.t. } IC_\rho, IC_{\theta_1}, IR_{\theta_1}.}$$

4 Three "special" cases

Before we solve the main problem, we consider three related models that help identify the key economic forces at work.

4.1 ρ is common knowledge

If ρ is common knowledge, the problem becomes isomorphic to the monopolistic version of the classical Rothschild and Stiglitz [1976] problem, studied first by Stiglitz [1977]. Both parties take expectations over θ_2, and insuree is incentivized to reveal θ_1 truthfully. Since there is no need of communication from the insurer, r here is irrelevant. The optimal contract is as follows.

Proposition 1. *$\exists \ \rho^* \in [\underline{\rho}, \overline{\rho}]$ s.t. $\pi^{RS}(\rho^*) = \max_\rho \pi^{RS}(\rho)$ and coverages are generically separating:*

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

As in the standard monopolistic screening model, the optimal contract is always separating: "high" risk type is offered exact coverage and "low" risk type is offered partial coverage, though which type is high risk pivots around ρ^*. 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 $\rho > \rho^*$, high risk type is $\theta_1 = H$ and for $\rho < \rho^*$, the high risk type is $\theta_1 = L$ (see Figure 1b). The profit is maximized at ρ^*, because the agent's private information of θ_1

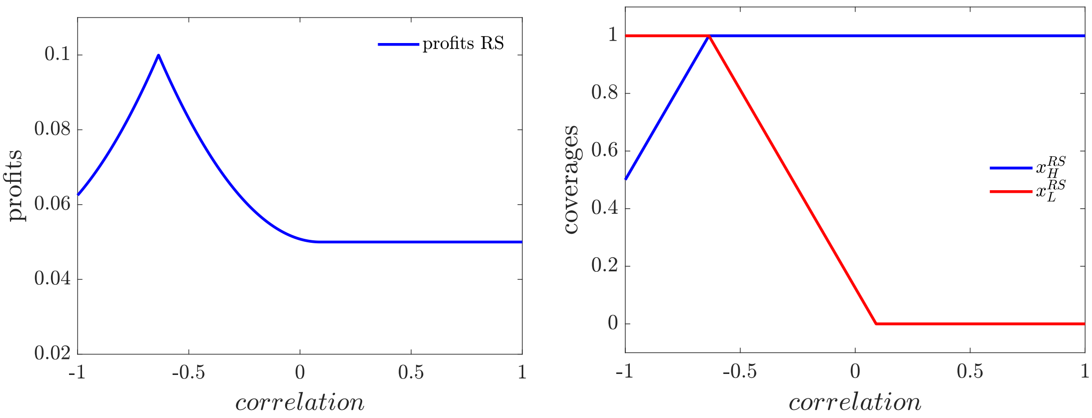

(a) profit as a function of ρ

(b) coverage as a function of ρ

Figure 1: Benchmark model when ρ is common knowledge

becomes statistically irrelevant: the principal offers a pooling contract and extracts all the surplus associated with it (see Figure 1a). We will refer to this as the *benchmark model*, and christen it RS, pointing to the classical reference.[9]

For the case $\rho > \rho^*$, for an interior solution, the optimal coverages are given by:

$$x_\rho(H) = 1 \text{ and } x_\rho(L) = 1 - \frac{1-q_1}{\eta q_1}(\mu_\rho(H) - \mu_\rho(L)) < 1 \tag{3}$$

As can be seen transparently, the extent of distortion for the "low" risk type is further determined by the primitives of the problem. In particular, the distortion is decreasing in η and q_1, and increasing in $\Delta_\rho = \mu_\rho(H) - \mu_\rho(L)$. Analogous comparative statics emerge for the case $\rho < \rho^*$.

The economic force driving this result is typically known as the rent-versus-efficiency tradeoff. Since insurer is the residual claimant of the surplus, she wants to maximize efficiency by offering full (or exact) insurance to both types with different premia (or prices) that hold each of them at their reservation utility. But due to asymmetric information she provides two different coverages, one full and another partial, and chooses premia in way that incentivizes insurees to self select into the contract corresponding to their type. In fact if the proportion of low risk types is too small, the insurer will simply offer full insurance to the high risk types and exclude the low risk ones from the market (see Figure 1b for high values of ρ).

In all the models that follow, ρ is *not* common knowledge, rather it is the insurer's private information. These will feature an inversion of adverse selection: by designing an incentive compatible mechanism, once the insurer learns θ_1, she knows more than the agent about the probability of the state θ. The insurer will exploit, to varying degrees, this *belief gap*, and one of the tools we will use to capture this intuition will be termed flipped allocation.

[9]Technically speaking for $\rho > \rho^*$, IC_H binds at the optimum, and for $\rho < \rho^*$, IC_L binds at the optimum. This determines which type is offered the efficient contract and which one is distorted.

Definition 1. *A contract $\mathcal{C}$ is said to feature flipped coverages if there exists $\hat{\rho}$ such that*

$$x_{r(\rho)}(H) > x_{\rho}^{RS}(H) \text{ and } x_{r(\rho)}(L) < x_{\rho}^{RS}(L) \text{ for } \rho < \hat{\rho};$$

$$x_{r(\rho)}(H) < x_{\rho}^{RS}(H) \text{ and } x_{r(\rho)}(L) > x_{\rho}^{RS}(L) \text{ for } \rho > \hat{\rho}.$$

In that case, will say that the coverages are **flipped around** $\hat{\rho}$.

This represents the idea that with private information on the joint distribution of the underlying state, insurer will want to sell more insurance to the "low" risk type and less insurance to the "high" risk type. The final allocation is then *flipped* in comparison to the standard case where ρ is common knowledge. In the inequalities above, $\theta_1 = H$ is the "low" risk type and $\theta_1 = L$ is the "high" risk type for $\rho < \hat{\rho}$. So the flipped allocation assigns more coverage to the $\theta_1 = H$ and less coverage to $\theta_1 = L$ in comparison to the benchmark. The analogous statement is true for $\rho > \hat{\rho}$.

Further, the insurer limits the extent of *price discrimination* in the contract on the basis of ρ through the function r because if the insuree is sophisticated, she can extract the private information of the insurer through Bayesian inference. In what follows we first discuss two cases in which the insuree cannot perfectly infer the value of ρ from the set of contracts offered by the insurer.

A final thought on the appropriate benchmark: It is also possible to let the benchmark to be the case where both parties are perfectly uninformed about the correlation coefficient and take expectations over it. In this case the optimal profits and coverages will be given by their counterparts in Proposition 1 evaluated at the expected correlation: $\pi_e^{RS} = \pi^{RS}(\mathbb{E}[\rho])$, $x_e^{RS}(H) = x_{\mathbb{E}[\rho]}^{RS}(H)$, and $x_e^{RS}(L) = x_{\mathbb{E}[\rho]}^{RS}(L)$.

4.2 Gutgläubig insuree

Another useful, and rather non-standard model to consider is one where in addition to offering a contract, the insuree tells the insurer the correlation coefficient and the latter simply believes it. This setting is different than the (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 insuree about the way in which the two dimensions are correlated provides the insuree with great freedom in selecting contracts. She will choose r and $\mathcal{C}$ in tandem to create both the maximal belief gap and the maximal price discrimination.[10]

Proposition 2. *If the insuree is a gutgläubig, $\exists\ \tilde{\rho} \in [\underline{\rho}, \overline{\rho}]$ such that:*

[10]Since the Bayes' consistency condition is not valid, technically the class of contract is given by $\mathcal{C} = (c_{m,\rho})$ because the contract offered for the actual realization of ρ has no relation to the reported value m.

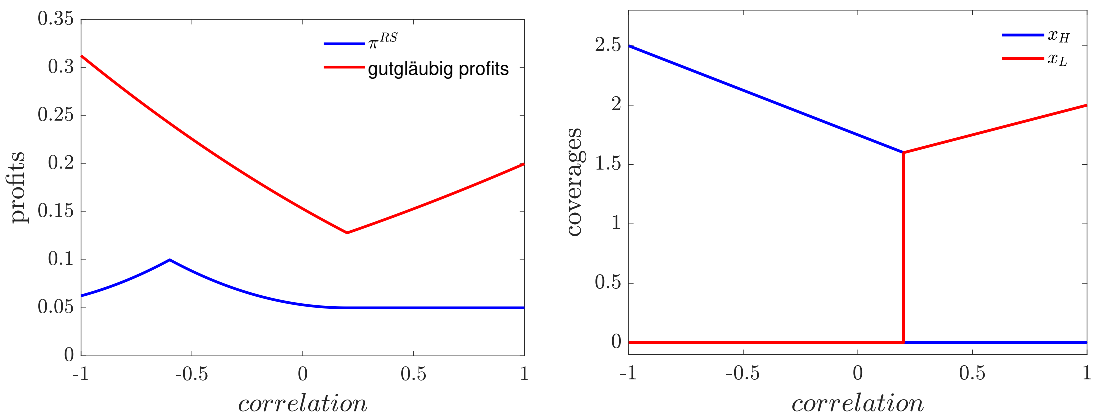

(a) profit as a function of ρ (b) coverage as a function of ρ

Figure 2: Model with gutgläubig insuree

1. *binary messages are sent:* $\mathcal{M} = \{\underline{m}, \overline{m}\}$ *s.t.* $r(\rho) = \overline{m}$ *for* $\rho < \tilde{\rho}$ *and* $m(\rho) = \underline{m}$ *for* $\rho > \tilde{\rho}$*,*

2. *posterior of the insuree is extreme:* $F_{\underline{m}} = \delta_{\underline{\rho}}$ *and* $F_{\overline{m}} = \delta_{\overline{\rho}}$ *where* δ_ρ *is Dirac measure on* ρ*,*

3. *profits are uniformly higher than benchmark:* $\pi(\rho) > \pi^{RS}(\rho)\ \forall\ \rho$ *almost surely,*

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 around* $\tilde{\rho}$*.*

There exists a threshold value of ρ, to the right of which the insuree reports the extreme negative correlation, $\underline{\rho}$, and to the left of which she reports the extreme positive correlation, $\overline{\rho}$. Even though the cardinality of the message space is just 2, a distinct contract is offered for each value of ρ, since the insuree does not infer anything about ρ from the menu of contracts.

When the *actual* correlation is high, it means that the type $\theta_1 = H$ is likely to suffer a large loss and $\theta_1 = L$ is likely to suffer a small loss. In this scenario, the insurer *reports* a large negative correlation, in fact the largest possible negative value, and overinsures $\theta_1 = L$ and underinsures $\theta_1 = H$. In the process, she is able to achieve dramatic price discrimination while maintaining an extreme belief gap. The exact opposite is true for the case when the actual correlation is low: insuree reports large positive correlation, and overinsures $\theta_1 = H$ and underinsures $\theta_1 = L$. In sum, the insurer sells a large amount of insurance at a high price to the type who actually has a low probability of loss, and a small amount of insurance to the type who actually has a high probability of loss.

Figure 2 depicts the profit and coverages when the insuree is gutgläubig. That the profits are uniformly higher (Figure 2a) 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 and then decide what coverages to offer each θ_1 type (Figure 2b).[11]

The last part of the proposition points out that allocations are flipped in comparison to the benchmark model, capturing the role of private information of the insuree. The extent of flip will be maximal in the gutgläubig case because the competing force limiting the extent of price discrimination, i.e. Bayesian sophistication of the agent, is shut down and further through gullibility, made to work against the agent.

This analysis has two take away messages: First, private information on the side of the insuree, especially statistical information, fundamentally changes the incentives of the insurer and hence the nature of contracts that are observed in the market for insurance. Second, an inability on part of the insuree to infer information and further be misled by the insurer results in a maximal belief gap and maximal price discrimination at the optimum, leading to large increase in profits for the insurer in comparison to the benchmark.

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 $F_m = F \ \forall m \in M$.[12] Thus, in this situation, the role of r is moot. The insurer designs the contract as a function of ρ with the knowledge that the insuree will evaluate his payoffs using the prior F.

Proposition 3. *If the insuree 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 around* $\mathbb{E}(\rho)$.

The salient difference between the naive model and *gutgläubig* case (and also the general model we will present next) is that here the belief gap is determined exogenously by the fixed prior and the realization of ρ, and the insurer cannot influence it. This works in the insurer's favor sometimes and other times it works against her. As a consequence, when the insuree is naive, the insurer is better off on average in comparison to the benchmark, however, unlike the *gutgläubig* case, this ranking is not uniform (see Figure 3a).

Here is a simple intuition for the result: Suppose the expected correlation according to F is high enough, so that according to insuree, the "high" risk type is $\theta_1 = H$. If the realized

[11] The large overinsurance offered at the extremes brings out the message starkly. One can limit the coverage exogneosuly to be below one, that is, $x \leq 1$. Whenever it is optimal to set $x > 1$, the bound will be hit and will get full but not over insurance. All other results will continue to hold qualitatively.

[12] See Benjamin [2019] for an overview of the literature.

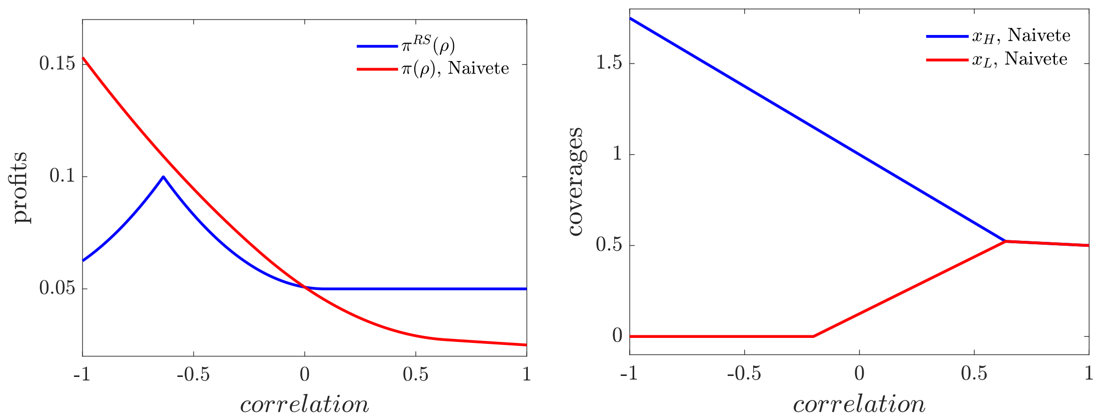

(a) profit as a function of ρ (b) coverage as a function of ρ

Figure 3: Model with naive insuree

correlation is close to $\underline{\rho}$, the insurer wants to sell a lot of insurance to $\theta_1 = H$ and little insurance to $\theta_1 = L$, because $\theta_1 = H$ is actually the "low" risk type but believes his risk to be at a higher level, according to F, and $\theta_1 = L$ is actually "high" risk (see left part of Figure 3b). On the other hand, when the realized correlation is close to $\overline{\rho}$, the insurer cannot sell a lot of insurance to $\theta_1 = H$ because he does not internalize the extent of risk he faces, and moreover, she cannot sell a lot of insurance to $\theta_1 = L$, because the nature of binding incentive constraints demands $x_\rho(H) \geq x_\rho(L)$; thus, for extremely high correlations, the insuree is forced to pool the coverages.[13]

4.4 Breakdown of the key forces

The key take away message from these special cases is this. The coverages vary as a function of ρ, the insurer's private information, and θ_1, the insuree's private information. The latter due to is the classical rent-versus-efficiency tradeoff which runs through each of the cases since the insuree's incentive constrain needs to be satisfied. The former generates a distinct tension of belief gap versus price discrimination.

In the first case, when correlation is common knowledge, belief gap is zero, and price discrimination is determined exogenously through the realized value of ρ. In the gutgläubig case, both belief gap and price discrimination are endogenously determined. Since the insurer can choose the contract independently from the insuree's belief, there is a no-longer a tradeoff between belief gap and price discrimination, and both are selected to maximize the insurer's profit. In the naïve case the belief gap exists but is determined exogenously for the insuree sticks to the prior no matter what contract is offered. Price discrimination is endogenously chosen to maximize the insurer's profit given the exogenous belief gap constraint. In what follows, both these forces will be determined endogenously and will interact with each other

[13] If expected correlation according to F is low enough, then in a symmetric contrast to Figure 3, the profit curve would intersect benchmark profits from below, and pooling in coverages will happen for high negative correlations.

and with the rent-versus-efficiency tradeoff to pin down the optimal contract.

5 Characterizing incentive compatibility

Our model differs from the standard screening problem in that it also features an incentive constraint for the principal, i.e. the insurer. In this section, we analyze the incentive constraints of the insurer for any fixed reporting strategy $r : [\underline{\rho}, \overline{\rho}] \to \Delta(M)$ by the mediator. The standard Myersonian characterization of the insurer's incentive compatibility is first stated.

Lemma 1. *IC_ρ holds if and only if π satisfies the following*

1. *envelope characterization of local incentives:*

$$\left.\frac{\partial \pi(\rho;\hat{\rho})}{\partial \rho}\right|_{\hat{\rho}=\rho} = \sigma x_{r(\rho)}(L) \cdot (\mu_{LH} - \mu_{LL}) - \sigma x_{r(\rho)}(H) \cdot (\mu_{HH} - \mu_{HL}) \equiv c(\rho), \qquad (4)$$

   and,

2. *convexity: $\pi(\rho)$ is convex in ρ.*

Proof. Part two is a standard property of value functions that satisfy incentive compatibility on a continuous type space (see, for example, Börgers [2015], Chapter 3). We show here the exact functional form of the envelope characterization stated in Equation (4). Start with Equation (2), i.e. assuming truthteling by the insuree, the profit function from (mis)reporting $\hat{\rho}$ is given by

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

where the only terms that are a function of ρ are

$$\mu_\rho(L) = (q_2 + \rho\sigma/q_1)\mu_{LL} + ((1-q_2) - \rho\sigma/q_1)\mu_{LH}, \text{ and}$$
$$\mu_\rho(H) = (q_2 - \rho\sigma/(1-q_1))\mu_{HL} + (1 - q_2 + \rho\sigma/(1-q_1))\mu_{HH}$$

Taking a derivative with respect to ρ, then gives us:

$$\frac{\partial \pi(\rho;\hat{\rho})}{\partial \rho} = -\sigma x_{r(\hat{\rho})}(L)(\mu_{LL} - \mu_{LH}) - \sigma x_{r(\hat{\rho})}(H)(-\mu_{HL} + \mu_{HH})$$

and, substuiting $\hat{\rho} = \rho$ delivers Equation (4). □

By fixing r, we fix M, which partitions type space of possible correlations, $[\underline{\rho}, \overline{\rho}]$. Hence we also fix the number of contracts offered at the optimum, $|\mathcal{C}| = |M|$. Now, for a given r, Lemma 1 tells us two things. First, the slope of the profit function can be written as

$$c(\rho) = k_L \phi_L(\rho) - k_H \phi_H(\rho),$$

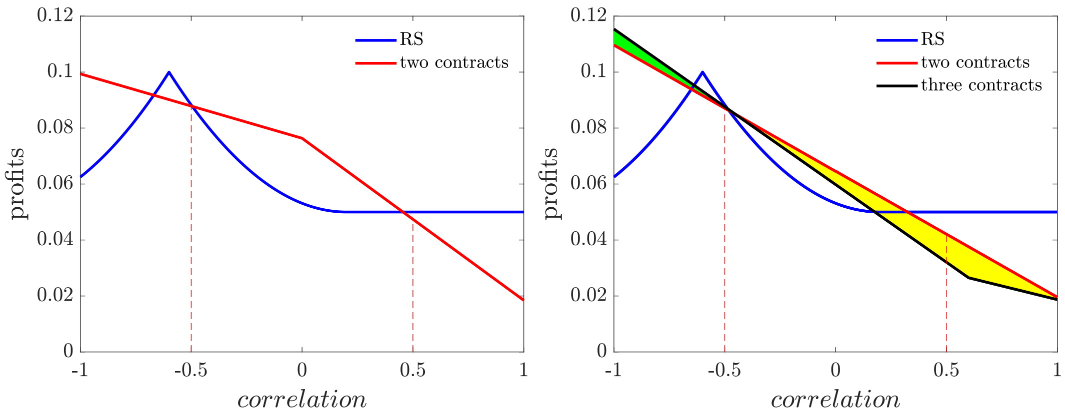

(a) optimal profit ignoring convexity (b) costs/benefits of adding an extra partition

Figure 4: Structure of optimal partitions

where k_L and k_H are positive constants, and $\phi_L(\rho) = x_{r(\rho)}(L)$ and $\phi_H(\rho) = x_{r(\rho)}(H)$ are the coverages chosen for $\theta_1 = L$ and $\theta_1 = H$, as a function of the partition of M in which ρ falls. And, second, by convexity of π, that $c(\rho)$ must be non-decreasing. These two together put restrictions on what coverages/allocations are feasible, specifically they limit the extent of price discrimination that the insurer can employ even for a fixed number of contracts.

The typical approach taken in mechanism design is to ignore the convexity constraints, solve the relaxed problem using only the envelope condition, and invoke a regularity condition such as the monotone hazard rate. But this problem is not standard in at least three ways: (i) the "policy function" is multidimensional, there are two allocation rules in the envelope condition, ϕ_L and ϕ_H, (ii) these functions in turn solve another downstream screening problem for the agent, and (iii) the mechanism still has to jointly choose r and $\mathcal{C}$ at the optimum.

All of the aforementioned constraint the contract space in non-trivial ways. At the first pass, we show that an optimal contract must in fact be finite:

Proposition 4. *The optimal mechanism has a finite number of messages and contracts:* $|\mathcal{M}| = |\mathcal{C}| = \kappa$ *for some* $\kappa \in \mathbb{N}$.

Recollect that $\mathcal{C} = \{c_m \mid m \in \mathcal{M}\}$ where $\mathcal{M}$ is essentially a partition of $[\underline{\rho}, \overline{\rho}]$. For every additional element we introduce in $\mathcal{M}$, there is a cost and benefit associated with it. Figure 4 presents an example where the optimal number of contracts offered at the optimum is two, that is, $|\mathcal{M}| = |\mathcal{C}| = 2$. Figure 4a delineates the role of the convexity constraint and Figure 4b shows why going from two to three partitions is not profitable.

In both figures, the single peaked blue curve is the optimal profit in the benchmark model discussed in Section 4.1. The red line depicts the optimal profit line with two partitions—in Figure 4a while ignoring the convexity constraint, and in Figure 4b while imposing it. Each partition has its expected correlation marked at the two vertical dotted lines. Feasibility demands that red profit line must not be above the blue benchmark profit at each of those two points. This is because in the subgame in which correlation is common knowledge, the best the insurer can do is to achieve a profit of $\pi^{RS}(\rho)$. In the two subgames, one for each of the two partitions, it is as if the insurer is in the benchmark model with the correlation being the expectation of correlation in those partitions.

Now, in the relaxed problem in which we ignore the convexity constraint from Lemma 1, we could choose the highest piecewise linear curve that crosses the blue curve at those expected correlations. This would culminates in a concave kink in the piecewise linear profit function, as shown in Figure 4a. The ignored convexity constraint is obviously violated, and this observation is not limited to the parameters chosen here—*any* two (or more) partition contract which solves the relaxed problem will generate such a concave kink. Thus, to make the contract incentive compatible, it has to be *ironed*. The highest convex profit function that the insurer can construct while satisfying incentive-feasibility is the one shown as the straight red line in Figure 4b.

Finally, we increase the the number of partitions from two to three, as shown in Figure 4b. The new profit curve is the piecewise liner black line. This transition still needs to satisfy all the incentive-feasibility restrictions we imposed before. Following those similar logics, we draw the best piecewise linear function that is convex and weakly below the benchmark profit at each of the three expected correlations corresponding to the three partitions. In doing this, the insurer incurs some costs and some benefits. The cost is shown in the lower yellow triangle in what constitutes the loss in profit, and the benefit is shown in the upper green triangle in what constitutes the gain in profit. In this case going from two to three partitions is clearly sub-optimal.

This intuition holds more generally. At the optimum the principal does not want to have an arbitrary number of contracts for the costs of doing so in terms of the restrictions imposed on the slope of the profit function, i.e. distortions introduced to satisfy incentive compatibility across partitions, outweigh the benefits accrued from greater price discrimination. Thus, the number of contracts offered is not just countable, it is also finite. This is documented in Proposition 4.

6 Optimal contract

In the previous section, we showed that incentive compatibility restricts the shape of the profit function and further evaluated the cost and benefit of having partitions of the correlation type space to conclude that set of the contract offered at the optimum must be finite. This dramatically simplifies the search for the optimal contract. The problem though is still quite hard to pin down in general. This is because both the number of partitions and their placement is endogenously determined. Here we argue that the number of contracts offered at the optimum is actually quite small, often the number of optimal partitions is one or two, that is, $|\mathcal{M}| = |\mathcal{C}| \leqslant 2$, and then characterize the contract for one and two partitions.

Limit result on number of partitions: Recollect that $\eta = \gamma\nu$, where γ is the risk aversion parameter, and ν is the variance of loss. Thus, η is a sufficient statistic of risk in our setup. We show that in both limits of arbitrarily large and arbitrarily small risk, the number of contracts offered at the optimum is one.

Proposition 5. *Fix $\eta = \gamma\nu$ and $\varepsilon = q_1\frac{\eta}{2}$. Then:*

1. *As $\eta \to \infty$, insurer finds it optimal to offer one message and contract, $|\mathcal{M}| = |\mathcal{C}| = 1$.*

2. *For any small $\eta > 0$, the difference between the optimal profit and the profit generated by the optimal one-partition contract is smaller than ε.*

When the insurer sets $|\mathcal{M}| = |\mathcal{C}| = 1$, there is complete pooling across ρ. Thus, she simply does not use her informational advantage towards price discrimination, opting rather to maintain the ex ante belief gap. Thus, the power of Bayesian inference essentially compels the seller to resolve the trade-off between belief gap and price discrimination completely in favor of the former force, at least in the two limits specified in Proposition 5.

As η becomes large, for the case when ρ is common knowledge, we can note from Equation (3) that the optimal contract approximates full insurance for both $\theta_1 = H$ and $\theta_1 = L$. Since the insurer is risk neutral, fixing the average correlation in a partition, the ideal outcome is to offer the benchmark contract corresponding to the average correlation in that partition. Since this coverage is converging to 1 irrespective of the value of ρ, the only price discrimination the insurer can introduce is to charge different prices (or premia) as a function of ρ for this approximately full insurance contract. But, given that there is not much room to discriminate along the quantity dimension, discriminating only along the price dimension is not beneficial enough to outweigh the costs associated with revealing information about ρ to the insuree. On the other hand, when η becomes small, since the surplus to be accrued from insurance is so small, the insuree again does not find it worthwhile to temper the belief gap in favor of price discrimination.

Intermediate numerical results on number of partitions: What about other parameters? Recollect that the number of contracts offered at the optimum depends on the slope of the profit function: $c(\rho) = k_L\phi_L(\rho) - k_H\phi_H(\rho)$, which in turn depends on the primitives k_L and k_H, and the allocations $\phi_L(\rho)$ and $\phi_H(\rho)$. The allocations are of course driven by the extend of risk and uncertainty in the environment, viz η. Hence, to understand the structure of optimal contracts, we split parametric space along these two dimensions.

Figure 5 documents when the optimal contract features complete pooling, $|\mathcal{M}| = |\mathcal{C}| = 1$, and when it features some price discrimination along ρ, that is $|\mathcal{M}| = |\mathcal{C}| > 1$. In fact, in the numerical simulation we allow the program to accept up to three partitions and it selects either one (in the dark/purple region) or two (in the light/yellow region) but never three. This further provides credence to the claim that the *number of contracts offered at the*

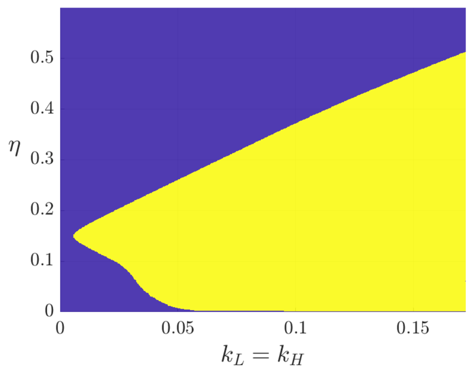

(a) case $k_L = k_H$

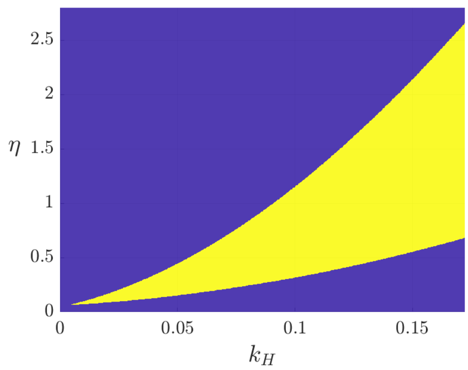

(b) fixing k_L, varying k_H

Figure 5: Splitting the parameters space into two regions: $|\mathcal{M}| = |\mathcal{C}| = 1$ and $|\mathcal{M}| = |\mathcal{C}| > 1$.

optimum with Bayesian sophisticated agents is small.[14]

Restricting the possible correlations to two: To gain further intuition on the optimal contract, we ask what if type space of the possible correlations is also restricted to be two, say $\rho \in \{\rho_1, \rho_2\} \subset [\underline{\rho}, \overline{\rho}]$. When does the insurer choose to offer a pooling contract in ρ and when does she choose to separate the two types? Loosely speaking, this question is "equivalent" to when the insurer offers one or two partitions in the original continuous type space model. The next result offers a characterization of pooling versus separation in the two types model. Recollect that ρ is defined as the solution to $\mu_\rho(H) = \mu_\rho(L)$.

Proposition 6. *Suppose $\rho \in \{\rho_1, \rho_2\}$ where $\rho_1 < \rho_2$.*

1. *If $\rho_1 < \rho^* < \rho_2$, the insurer offers one pooling contract across ρ: $c_{\rho_1} = c_{\rho_2}$.*
2. *If $\rho_1 < \rho_2 \leqslant \rho^*$ or $\rho^* \leqslant \rho_1 < \rho_2$, depending upon the slope of the profit function the insurer may pool or separate across ρ.*

When offering a separating contract, it is as if the insurer is creating two partitions in the continuous types model where the expected correlations of the partitions are ρ_1 and ρ_2 respectively. One way to interpret Proposition 6 is that when the primitives of the model push the two partitions to be such that the expected correlations are on different sides of ρ^*, it is always better to offer a completely pooling contract along ρ. Maintaining the belief gap is more valuable for such parameters than the price discrimination afforded by separating ρ_1

[14]In fact in all our numerical simulations across a range of parameters, the optimal number of partitions is capped at two. This we conjecture is a global result: $|\mathcal{M}| = |\mathcal{C}| \leqslant 2$. If we don't impose the convexity constraint on the optimization problem then the allocation rule generates a $c(\rho)$ which turns out to be a decreasing function. However, the convexity constraint demands that $c(.)$ be non-decreasing. Thus, the contract must be ironed everywhere. In a typical Myerson problem, ironing everywhere would generate a constant slope of the value function, see for example, Hartline and Roughgarden [2008]. Due to the complexity of our mechanism design problem with an informed principal, we cannot at the outset rule out that ironing could lead to the slope increasing locally at some point. Though we conjecture this case to not be optimal so that $c(.)$ is always constant at the optimum and the total number of partitions is either one or two.

and ρ_2. Moreover, $\{\rho_1, \rho_2\}$ on the same side of ρ^* is necessary but not a sufficient conditions for separation. The result, as show in the appendix, depends further on the slope of the profit function.

Characterization of the optimal one and two partition contracts: Now back to the model $\rho \in [\underline{\rho}, \overline{\rho}]$. If the optimal number of partitions turns out to be one, it is fairly intuitive to conclude that the coverages offered would be same as those offered in the benchmark model at the ex ante expected correlation, and the optimal profit too will be equal to the optimal profit at that correlation. This result is summarized in the next proposition. Recollect that $\pi_e^{RS} = \pi^{RS}(\mathbb{E}[\rho])$, $x_e^{RS}(H) = x_{\mathbb{E}[\rho]}^{RS}(H)$, and $x_e^{RS}(L) = x_{\mathbb{E}[\rho]}^{RS}(L)$.

Proposition 7. *When the optimal contract chooses $|\mathcal{M}| = |\mathcal{C}| = 1$:*

1. *expected profits are the same as in benchmark at the expected correlation: $\mathbb{E}[\pi(\rho)] = \pi_e^{RS}$,*
2. *coverages are the ones offered for the expected correlation in the benchmark: $x_{r(\rho)}(H) = x_e^{RS}(H)$ and $x_{r(\rho)}(L) = x_e^{RS}(L)$ $\forall$ ρ.*

Figure 6 plots the optimal profit and coverages for this case. The coverages are simply straight horizontal lines for the insurer is not using any of her private information about ρ and instead offers a completely pooling contract along ρ. As in the benchmark model, the "high" risk insuree (which is type $\theta_1 = H$ in the figure) is given full insurance and the "low" risk insuree is given partial insurance. The profit function is a straight downward sloping line since the allocations are fixed, and π is linear in ρ. The dotted vertical line captures the expected correlation at which point the red straight line and benchmark blue curve intersect.[15]

Next, we consider the case where the optimal number of partitions is two. In this case, the type space of correlations is split into two intervals, say $\mathbb{I}_1$ and $\mathbb{I}_2$. The coverages in each interval are evaluated using the expected correlation in those intervals while ensuring that the insuree's incentive constraint is satisfied between reporting interval $\mathbb{I}_1$ or $\mathbb{I}_2$ and within each interval, the insurer's incentive constraint is satisfied in reporting $\theta_1 = H$ or L. The following result summarizes the key aspects of the optimal contract.

Proposition 8. *When the optimal contract chooses $|\mathcal{M}| = |\mathcal{C}| = 2$, let $\mathbb{I}_1$ and $\mathbb{I}_2$ be the two intervals in the partition of M, let $c_i = (x_i(H), x_i(L))$ be the two contracts offered, and define $\rho_1 = \mathbb{E}[\rho \mid \rho \in \mathbb{I}_1]$, $\rho_2 = \mathbb{E}[\rho \mid \rho \in \mathbb{I}_2]$. Then:*

1. *profits are linear in correlation: $\frac{d\pi(\rho)}{d\rho} = c$ for some constant c,*

[15] For all the optimal contracts, we plot the profit and the coverages of the benchmark model simultaneously to help motivate the impact of the privacy of statistical information on the side of the insurer, which separates our model form (most of) the literature on insurance markets.

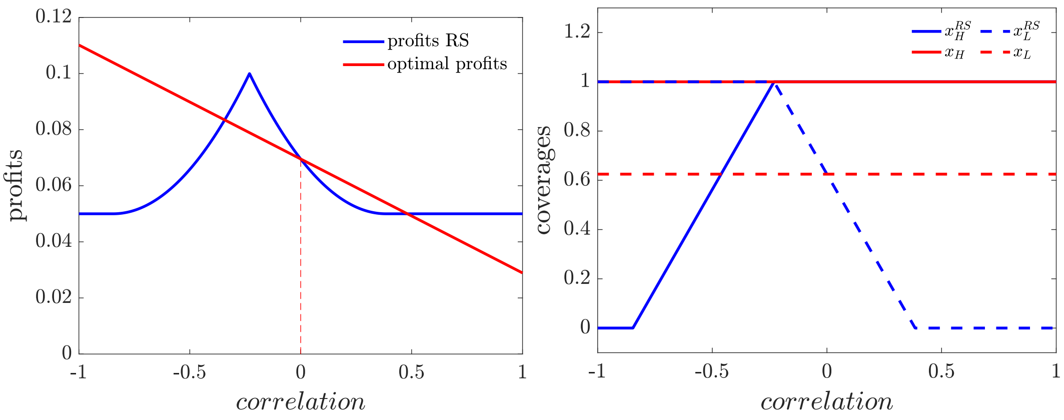

(a) profit (b) coverage and its benchmark counterpart

Figure 6: Optimal contract features complete pooling

2. *expected profits are larger than the benchmark:* $\mathbb{E}[\pi(\rho)] > \pi_e^{RS}$

3. *coverages are flipped in comparison to the benchmark in the following sense:*

   (a) $x_1(H) \geq x_{\rho_1}^{RS}(H)$ *and* $x_1(L) \leq x_{\rho_1}^{RS}(L)$*, whenever* $x_H^1 \neq x_L^1$*,*

   (b) $x_2(H) \leq x_{\rho_2}^{RS}(H)$ *and* $x_2(L) \geq x_{\rho_2}^{RS}(L)$*, whenever* $x_H^1 \neq x_L^1$*.*

Figure 7 plots the optimal profit and coverages when the optimal number of partitions is two. Each partition corresponds to two coverages, one for each insuree type, which gives the profit function its slope. The first result in Proposition 8 states that optimality forces both these slopes to be the same. This follows from the intuition given in Section 5 that without imposing the convexity constraint the optimal profit line has a concave link. So the highest profit line that satisfies convexity is then simply the straight line which equates the slope of the profit function along the two partitions, as shown in Figure 7a. The second result in Proposition 8 simply states that if the insuree is employing two partitions at the optimum than the profit must be greater than expected profit in the benchmark model since the latter can always be attained by offering a completely pooling contract as in Proposition 7.

The third result in Proposition 8 documents that the coverages are flipped in each partition. In the left partition where $\theta_1 = L$ is the "high" risk type, $\theta_1 = H$ is overinsured. But, unlike the *gutgläubig* case, incentive constraints force the allocation to always satisfy $x_1(H) \geqslant x_1(L)$. So, in the right partition, $\theta_1 = H$, which is now the "high" risk type, is offered under insurance, and $\theta_1 = L$ is forced out of the market with no coverage. See Figure 7b.

Summary: To summarize, when the insuree is Bayesian sophisticated, the total number of contracts offered at the optimum is small, often at most two. This illustrates the fact the trade-off between belief gap and price discrimination is resolved mostly in favor of the former. When the optimal contract features complete pooling across ρ, the contract corresponding

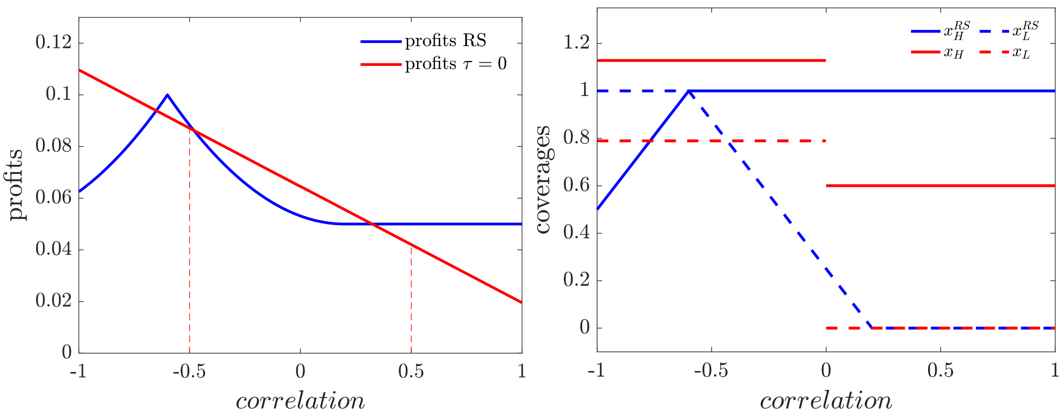

(a) profit (b) coverage and its benchmark counterpart

Figure 7: Optimal contract features two partitions

to the benchmark model at the expected correlation is offered and the insurer does not use her informational advantage at all. When the optimal number of contracts offered is two, we see both over and under insurance at the optimum, owing to forces similar to the ones seen in the *gutgläubig* case, but significantly tempered by the fact the agent cannot be misled and perfectly infers all information from the offered contracts.

7 Two interventions

In this section we look at two extensions of the basic model. First, we allow for competition from firms that do not posses the big data advantage, they average over θ_2 using the prior. Second, we explore the implications of forcing the insurer to publicly reveal her informational advantage, that is report ρ to the insuree. In each case we make a comparison with the *gutgläubig* case and the case where the agent can do perfect Bayesian inference.

7.1 Introducing competition from 'regular' insurers

Suppose there is an insurer that knows the value of ρ and there are other insurers in the market that do not know this value and work with the prior F. Borrowing from the benchmark (in Section 4.1), the latter group of insurers are assumed to offer the Rothschild-Stiglitz contract evaluated at the expected value of ρ. As before, the big data insurer can send a message disclosing some information about the correlation and offer a contract. The idea here is to introduce competition in a tractable way from regular firms who do not have the in-house expertise of big data.

The insuree can choose to buy insurance from any one seller. However, independent of which insurer he buys from, the belief is updated as before from the contract offered by the big data insuree and as a function of his level of Bayesian sophistication. Therefore, the insurer will buy from the big data seller only if he obtains a utility higher than the one offered by the Rothschild-Stiglitz contract evaluated at the expected correlation. So, in effect, the introduction of regular insurers simply modifies the individual rationality constraint for the optimization problem of the big data insurer. This problem can be stated as follows:

$$\boxed{\max_{r,\mathcal{C}} \Pi \text{ s.t. } IC_\rho, IC_{\theta_1}, IR^e_{\theta_1}}$$

where r, $\mathcal{C}$, IC_ρ, IC_{θ_1} are as defined before, and $IR^e_{\theta_1}$ is modified so that the insuree's outside option is evaluated at the contract $c^e = \{p^{RS}_e(H), x^{RS}_e(H), p^{RS}_e(L), x^{RS}_e(L)\}$. Here $x^{RS}_e(H)$ and $x^{RS}_e(L)$ are defined in Section 4.1, and $p^{RS}_e(H)$ and $p^{RS}_e(H)$ are respectively the prices that maximize the insurer's expected utility when the correlation is common knowledge and equal to $\mathbb{E}[\rho]$.[16]

We document below the change in the structure of the optimal contract from the introduction of competition from regular insurers.

Proposition 9. *Suppose the big data insurer faces competition from regular insurers that offer c^e.*

1. *Suppose the insuree is gutgläubig. Then, all features of the optimal contract stated in Proposition 2 continue to hold. The main departure is that the coverage of the "low type" insuree goes up.*

2. *Suppose the insuree is Bayesian sophisticated. Then the number of partitions at the optimum is weakly lower the standard model.*

In summary, the introduction of competition forces the insurer not to exploit price discrimination even further. In the gutgläubig case, it increases the coverage of the "low" risk type. In the Bayesian sophisticated case, it sometimes forces the insurer to reduce the number of partitions from two to one, and if it is was one in benchmark case, then it stays at one with competition as well. Competition thus weakly increases the surplus of the insuree.

7.2 Optimal full revelation contract

So far we have analyzed the case where the insurer is the sole proprietor of statistical information ρ and can decide, as part of the optimal contract, how much of it to reveal to the insuree. For a variety of regulatory and (presumably) welfare concerns, the insurer can be asked to reveal the information about ρ publicly. One obvious way to model this is to assume that we "nationalize" the system by taking over the insurance company and putting this information in the public domain. Abstracting from the costs of this nationalization, the model would then become equivalent to the one studied in Section 4.1, where ρ is common knowledge.

An alternate way is to assume the insurer develops her private information at an arm's length distance from the government, and is then incentivized to put this in the public

[16]Here, we assume that if indifferent, the insuree buys from the informed (or big data) insurer.

domain. The underlying assumption here is that it is difficult to construct a statistical model that increases the precision in predicting θ and then use it effectively in designing insurance contracts. At a high level, what we have in mind is that information is dispersed in a society and collecting it and making it public is a non-trivial exercise (Hayek [1945]). There are shadow prices associated with incentivizing the insurer to publicly reveal ρ.

We model this situation by exogenously fixing the message rule chosen by the mediator to be the identity mapping: $r(\rho) = \rho$, and requiring the principal (i.e., the insurer) be compensated for this through her incentive constraint. This obviously has the downstream effect of influencing the contract C that is offered to the insuree. The entire optimization problem can be written in one piece as follows:

$$\boxed{\max_{\mathcal{C}^\star} \Pi \text{ s.t. } IC_\rho, IC^\star_{\theta_1}, IR^\star_{\theta_1}}$$

where $\mathcal{C}^\star = \{c_\rho \mid \rho \in [\underline{\rho}, \overline{\rho}]\}, c_\rho = \{c_\rho(H), c_\rho(L)\}$ and $c_\rho(\theta_1) = (p_\rho(\theta_1), x_\rho(\theta_1))$ for $\theta_1 = H, L$, and $IC^*_{\theta_1}$ $IR^*_{\theta_1}$ are evaluated using u_m which plugs in the actual realization of ρ since m is pre-fixed to be ρ here. The optimal coverages will be denoted by $x_\rho(H)$ and $x_\rho(L)$. The following proposition summarizes the optimal full revelation contract.

Proposition 10. *Suppose the insurer is required to reveal all information, that is $r(\rho) = \rho$ is fixed exogenously. Then:*

1. *profits are uniformly lower than the benchmark: $\pi(\rho) < \pi^{RS}(\rho)\ \forall \rho$,*
2. *coverages are generically inexact: $x_\rho(\theta_1) \neq 1\ \forall\ \rho$ a.s.,*
3. *there is pooling and separation at the optimum:*
   (a) *$\rho > \rho^* \Rightarrow x_\rho(H) \geq x_\rho(L)$,*
   (b) *$\rho < \rho^* \Rightarrow x_\rho(H) \leq x_\rho(L)$.*
   (c) *one of these may hold as an equality.*
4. *$\exists\ \tilde{\rho}$ such that the contract is flipped around $\tilde{\rho}$.*

The first observation is that the insurer's profits are uniformly lower in the full revelation contract than the benchmark, see also Figure 8a. This makes intuitive sense for the subgame following the reporting of ρ is exactly the same as the benchmark model, but the mandatory revelation of ρ is subjected to the binding constraint, IC_ρ. The second result documents the fact that there is both under and over provision of insurance at the optimum; Figures 8b and 8c document that partial insurance or under-provision dominates. The third result states that similar to the benchmark model, the "high" types that is offered larger insurance changes around ρ^*, but unlike the benchmark model this ranking is weak for there can also

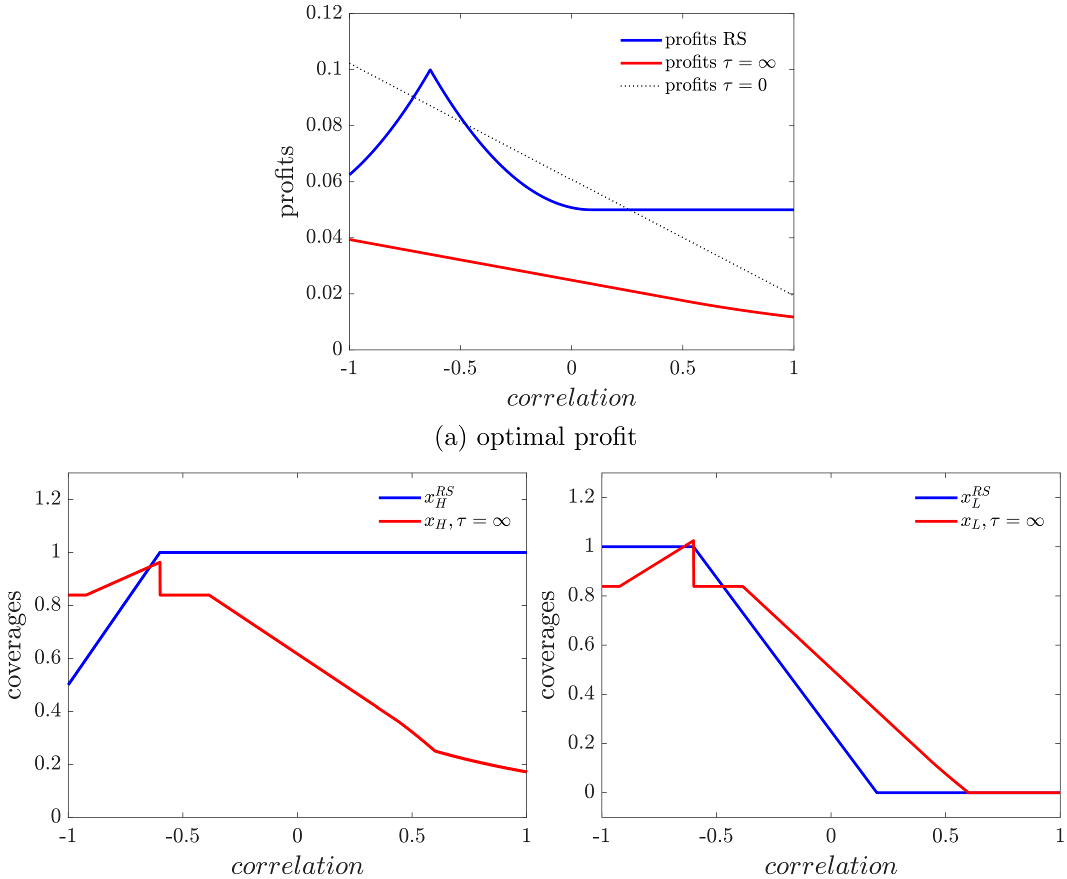

(a) optimal profit

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

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

Figure 8: Optimal contract with full information revelation

be pooling at the optimum. Pooling across the insuree's type occurs because of the binding incentive constraints of the insurer. Finally, the coverage is flipped around $\tilde{\rho}$ in comparison to the benchmark. As can be seen in the figures, typically $\tilde{\rho}$ is very close to ρ^*.

It can be noted (in Figures 8b and 8c) that a continuum of contracts are offered at the optimum in the full revelation model in comparison to the optimal contract where the insurer controls the release of information (in Section 6). This is because the the insuree is being forced to relinquish the belief gap, so it tries to maximize on the price discrimination part to the extent feasible.

What are welfare implications of forcing the insurer to reveal her private information? A global result is elusive due to the complexity of the mechanism design problem with an informed principal. However, for most parameters numerical results suggest that the total surplus often reduces in comparison to the benchmark and standard models and expected utility of the insuree goes up. Here we often two qualitative results on the payoffs and coverage.

We have seen thus far that the insurer's profits are uniformly higher in the *gutgläubig* model than the benchmark and uniformly lower in the full revelation model than the benchmark. In the next result we document the surplus guarantees for the insuree under the two specifications. Let $u_{m,\rho}(\theta_1)$ is the payoff of the insuree of type θ_1 in the *gutgläubig* case where m is the correlation reported by the insurer and ρ is the actual realization correlation. In addition, let $u_\rho(\theta_!)$ be the payoff of the insuree of type θ_1 in the full revelation case where ρ is the actual realized and reported correlation and correspondingly $\underline{u}^\rho(\theta_1)$ be the payoff if the insuree does not accept any insurance. Then we have the following simple result.

Corollary 1. *The insuree's surplus (or utility)*

1. *is uniformly negative in the gutgläubig case, i.e.,* $\mathbb{E}_\theta\left[u_{m,\rho}(\theta_1)\right] < 0$ *for* $m = \underline{\rho}, \overline{\rho}$, $\forall\ \rho$*;*

2. *is uniformly positive for the full revelation contract.i i.e.,* $u_\rho(\theta_1) \geqslant \underline{u}_\rho(\theta_1)$ *for all* ρ*,* θ_1*.*

In the first case the insurer is able to maximally mislead the insurer about the true probability of loss. The high price charged to the "low" risk insurer culminates in a negative surplus. On the other hand, by construction, since the individual rationality constraint for the insurer holds pointwise–for each θ_1 and ρ– we get that the insurer's payoff is non-negative.

In the appendix we offer another result that characterizes the total extent discrimination and total coverage offered at the full information revelation model. Since the insurer has to incentivized to reveal ρ the extra shadow prices constraint the amount of price discrimination that can be sustained at the optimum. Moreover, as a function of the parameters the total coverage offered at the optimum can go up or down in comparison to the benchmark model.

8 Final remarks

A big debate in ensuing right now on the merits of technological advancements in data documentation and processing. Foregrounding these issues, in the summer of 2019, the New York Times carried a series of articles under the rubric of *The Privacy Project.*[17] One of the key topics of discussion therein was the impact of big data and AI on the insurance industry. This paper is an attempt to mainstream these discussions in the modeling choices made by classical economic theory in formalizing the key ideas in insurance markets.

Traditionally mechanism design models of insurance assume that the agent (or insuree) has some private information about the probability of incurring a loss or meeting with an accident. This results in the proverbial rent-versus-efficiency trade-off wherein the principal (or insurer) gives up on efficiency and provides information rents in order to separate the high risk from the low risk agents. We depart from this standard model in one crucial way—we make the state of world that parametrizes the loss to be two dimensional, and allow the agent to posses information about one of these dimensions and the principal to know the statistical correlation between the two dimensions. This creates an informed principal problem where the principal too has private information.

[17] See www.nytimes.com/interactive/2019/opinion/internet-privacy-project.html.

Private statistical information on the side of the insurer introduces a novel trade-off between belief gap and price discrimination, in addition to the usual rent versus efficiency in standard screening contracts. The insurer wants to price discriminate using her private information dimension but is also wary that fine-tuning the contract too much to the details of the environment will allow the insuree to infer that information. This latter desire to maintain a belief gap pulls against the desire to price discriminate.

In the standard framework in which the agent is Bayesian sophisticated, the insuree resolve this tradeoff by offering very few contracts (at most two in most cases) in order to maintain the belief gap. In the case where the insurer is gullible, this tradeoff disappears, the the insurer is able to maximize price discrimination while maintaining the maximal belief gap. Forcing the insurer to reveal the statistical information leads to less price discrimination and often greater total coverage. Moreover, introducing competition reduces the informational advantage of big data.

The result on fewness of contracts under Bayesian sophistication should be viewed as a theoretical benchmark— when the consumers can do proper inference there are limits to deployment of big data in extracting surplus. However, at the other extreme the gutgläubig case shows that there are significant gains to be made from big data when the consumer has limited inference capacities. This provides a foundation of sorts for both the rise of data markets and the returns to consumer activism whereby implications of data disclosure and its deployment by sellers can be better understood. Finally, putting this data in public domain along with an understanding of how to interpret this information can benefit consumers, and so will competition by endogenously limiting the extent to which the big data can be used against the consumers.

The ideas developed here can potentially be applied to contexts other than insurance. For example, in credit markets, owing to big data and AI, the credit issuing agency may also have some statistical information about the credit worthiness of a client, in addition to the client knowing some hard information about his financial circumstances.[18] Finally, aggregating across multiple principal-agent interactions. greater statistical information on the side of the principal may encourage more market concentration, of the kind we see in the tech-industry these days. Endogenizing data collection and market size is a promising question for future work.

9 Appendix

Proofs of results stated in the main text are presented in this section.

[18] See Vives and Ye [2021] for some recent work in this direction.

9.1 Proofs for Section 4

Proof of Proposition 1. Since ρ is common knowledge, the mapping r is redundant and the seller's profit for a specific ρ is given by

$$\pi(\rho) = q_1(p_\rho(L) - \mu_\rho(L)x_\rho(L)) + (1-q_1)(p_\rho(H) - \mu_\rho(H)x_\rho(H)) \tag{5}$$

The insurer's problem is to choose a contract $c_\rho = \{p_\rho(\theta_1), x_\rho(\theta_1)\}_{\theta_1=H,L}$ to maximize $\pi(\rho)$ to subject to incentive feasibility. Using Equation (1), the constraints can be written as:

$$\mu_\rho(\theta_1)x_\rho(\theta_1) - \tfrac{\eta}{2}(1-x_\rho(\theta_1))^2 - p_\rho(\theta_1) \geqslant \mu_\rho(\theta_1)x_\rho(\theta_1') - \tfrac{\eta}{2}(1-x_\rho(\theta_1'))^2 - p_\rho(\theta') \;\; \forall\theta_1,\theta_1' \in \{L,H\} \quad IC_{\theta_1-\theta_1'}$$
$$\mu_\rho(\theta_1)x_\rho(\theta_1) - \tfrac{\eta}{2}(1-x_\rho(\theta_1))^2 - p_\rho(\theta_1) \geqslant -\tfrac{\eta}{2} \quad \forall\theta_1 \in \{L,H\} \quad IR_{\theta_1}$$

Let ρ^* be the correlation for which $\mu_\rho(H) = \mu_\rho(L)$. Thus, when $\rho = \rho^*$, there is no asymmetric information, and the principal/insurer does not need to provide an information rent to any of the types. She can simply maximize efficiency by offering a unique pooling contract of full insurance $x_{\rho^*}(H) = x_{\rho^*}(L) = 1$ and bind the IR constraints to extract expected surplus.

Suppose that $\mu_\rho(H) > \mu_\rho(L)$, that is, $\rho > \rho^*$. Then, we are in the standard Rothschild and Stiglitz [1976] setup where $\theta_1 = H$ is the "high" type and $\theta_1 = L$ is the "low" type. It is standard practice (see for example Laffont and Martimort [2009]) to show that in this case IC_H and IR_L bind, and IC_L and IR_H are slack. Let λ be the multiplier on IC_H and δ the multiplier on IR_L. Th following FOCs that characterize an interior solution:

$$\begin{aligned}
&[p_\rho(L)]: \quad q_1 - \delta + \lambda = 0\\
&[x_\rho(L)]: \quad -\mu_\rho(L)q_1 + \delta\mu_\rho(L)q + \sigma(1-x_\rho(L))\delta - \lambda\mu_\rho(H) - \lambda\eta(1-x_\rho(L)) = 0\\
&[p_\rho(H)]: \quad (1-q_1) - \lambda = 0\\
&[x_\rho(L)]: \quad -(1-q_1)\mu_\rho(H) + \lambda\mu_\rho(H) + \eta(1-x_\rho(H))\lambda = 0.
\end{aligned}$$

From the first and third conditions it can be concluded that $\lambda = (1-q_1)$ and $\delta = 1$. Using these values it is straightforward to see that

$$x_\rho(H) = 1 \text{ and } x_\rho(L) = 1 - \frac{1-q_1}{\eta q_1}(\mu_\rho(H) - \mu_\rho(L)) < 1.$$

In case of a corner solution, $x_\rho(H) = 1$ and $x_\rho(L) = 0$.

An analogous argument shows the result for the case in which $\mu_\rho(H) < \mu_\rho(L)$, that is, when $\rho < \rho^*$. □

Proof of Proposition 2. Since the insuree doesn't do Bayesian inference, the mapping r is independent of the contract offered at the optimum. The insurer's profit for a specific ρ is $\pi(\rho)$ given by Equation (5). And due to the disentangling of the inference problem from the contractual problem, the insurer can solve the optimization independently for each value of ρ:

$$\max_{c,\hat{\mu}} \pi(\rho) \text{ subject to}$$

$$\begin{array}{lr}\hat{\mu}(\theta_1)x(\theta_1) - \frac{\eta}{2}(1-x(\theta_1))^2 - p(\theta) \geqslant \hat{\mu}(\theta)x(\theta_1') - \frac{\eta}{2}(1-x(\theta'))^2 - p(\theta_1') \;\; \forall \theta_1, \theta_1' \in \{L, H\} & IC_{\theta_1 - \theta_1'} \\ \hat{\mu}(\theta_1)x(\theta_1) - \frac{\eta}{2}(1-x(\theta))^2 - p(\theta_1) \geqslant -\frac{\eta}{2} \quad \forall \theta_1 \in \{L, H\} & IR_{\theta_1}\end{array}$$

where $c = \{p(\theta_1), x(\theta_1)\}_{\theta_1=H,L}$ is the contract and $\hat{\mu}$ is the belief that the insurer generates for the insuree. Note that both the contract and the belief chosen by insurer depend on ρ, but since ρ is fixed for the optimization problem, we have simplified notation here for the rest of the calculations.

The objective is independent of $\hat{\mu}(\theta_1)$, and the constraint set is linear. Thus, it is straightforward to conclude that the solution in bang-bang in $\hat{\mu}(\theta_1)$. So, the insurer will report either extremes of the feasible set of correlations, $\underline{\rho}$ or $\overline{\rho}$. This implies only two messages are sent at the optimum, say $\underline{m}$ and $\overline{m}$, that generate buyer's posteriors $F_{\underline{m}} = \delta_{\underline{\rho}}$ and $F_{\overline{m}} = \delta_{\overline{\rho}}$.

Suppose that the seller sends the message $\overline{m}$. This message generates posterior beliefs $\overline{\mu}(H) = \mu_{\overline{\rho}}(H)$ and $\overline{\mu}(L) = \mu_{\overline{\rho}}(L)$, where $\overline{\mu}(H) > \overline{\mu}(L)$. Furthermore, for any ρ, we have $\overline{\mu}(H) > \mu_\rho(H)$ and $\overline{\mu}(L) < \mu_\rho(L)$. Now, using the first-order approach, following steps from the proof of Proposition 1, it is straightforward to show that in an interior solution

$$x(H) = 1 + \frac{\overline{\mu}(H) - \mu_\rho(H)}{\eta} > 1 \text{ and } x(L) = 1 - \frac{(1-q_1)\overline{\mu}(H) + q_1\mu_\rho(L) - \overline{\mu}(L)}{q_1\eta} < 1,$$

and at a corner solution $x(L) = 0$ and $x(H)$ takes the same value. An analogous argument shows that when sending the message $\underline{m}$, we obtain $x(H) < 1$ and $x(L) > 1$.

As a final step, we need to argue that for low correlations the seller will send messages $\overline{m}$ and for high correlations the seller will send the message $\underline{m}$. Let $\overline{\pi}(\rho)$ be the profits the seller obtains after sending message $\overline{m}$ when the actual correlation is ρ, and analogously define $\underline{\pi}(\rho)$. Plugging in we obtain that when the optimal contract is interior:

$$\begin{aligned}\frac{\partial \overline{\pi}(\rho)}{\partial \rho} - \frac{\partial \underline{\pi}(\rho)}{\partial \rho} &= \frac{-\overline{\mu}(L) + (1-q_1)\overline{\mu}(H) + q_1\underline{\mu}(L)}{\eta}\frac{\partial \mu_\rho(L)}{\partial \rho} + \frac{\underline{\mu}(H) - (1-q_1)\overline{\mu}(H) - q_1\underline{\mu}(L)}{\eta}\frac{\partial \mu_\rho(H)}{\partial \rho} \\ &< 0,\end{aligned}$$

since $\frac{\partial \mu_\rho(L)}{\partial \rho} < 0$, $\frac{\partial \mu_\rho(H)}{\partial \rho} > 0$, $\overline{\mu}(L) < \underline{\mu}(L)$, $\overline{\mu}(L) < \overline{\mu}(H)$, $\overline{\mu}(H) > \underline{\mu}(H)$, and $\underline{\mu}(H) < \underline{\mu}(L)$.

Thus, if for a correlation $r(\rho) = \underline{m}$, then for all $\rho' > \rho$ the seller sends the same message, $r(\rho') = \underline{m}$. Analogously if $r(\rho) = \overline{m}$, then for all $\rho' < \rho$, $r(\rho') = \overline{m}$. This and our characterization above shows that there is a $\tilde{\rho} \in [\underline{\rho}, \overline{\rho}]$ such that $r(\rho) = \overline{m}$ for $\rho < \tilde{\rho}$ and $r(\rho) = \underline{m}$ for $\rho > \tilde{\rho}$. Moreover, it is easy to see that the contract flips around $\tilde{\rho}$ (in the sense of Definition 1). Finally, notice that when $\tilde{\rho} \in (\underline{\rho}, \overline{\rho})$ the argument above shows that the seller always offer contracts that over insure or under insure the insuree. □

Proof of Proposition 3. Here the mapping r is redundant since the insuree does not update his prior. Since the IC and IR constraints are linear in beliefs, what matters for the insuree is the expected probability of loss as evaluated through the prior F, call these $\mu^e(H)$ and $\mu^e(L)$. The insurer offers a contract $c_\rho = \{p_\rho(\theta_1), x_\rho(\theta_1)\}_{\theta_1=H,L}$ to maximize $\pi(\rho)$ subject to appropriate notions of IC and IR—same as in the proof of Proposition 1 with $\mu_\rho(\theta_1)$ replaced with $\mu^e(\theta_1)$. Suppose $\mu^e_H > \mu^e_L$.

Following steps as before, we obtain that at an interior solution:

$$x_\rho(H) = 1 + \frac{\mu^e(H) - \mu_\rho(H)}{\eta} \text{ and } x_\rho(L) = 1 - \frac{(1-q_1)\mu^e(H) + q_1\mu_\rho(L) - \mu^e(L)}{q_1\eta}.$$

Notice that $x_\rho(H)$ is decreasing in ρ and $x_\rho(L)$ is increasing in ρ. Then there are two "corner" solutions: one in which $x_\rho(L) = 0$ and $x_\rho(H) = 1 + \frac{\mu^e(H)-\mu_\rho(H)}{\eta}$, and another one in which $x_\rho(L) = x_\rho(H) = 1 + \mu^e(L) - \frac{q_1\mu_\rho(L)+(1-q_1)\mu_\rho(H)}{\eta}$. This happens when the two allocation listed above violate $x_\rho(L) = 0$ and $x_\rho(H) \geqslant x_\rho(L)$, respectively.

This proves part 2 that coverages can be pooling and separating and part 3 that coverages are generically not equal to 1. Moreover, comparing the allocations with Proposition 1 it is easy to see that coverages are the same as the benchmark at the expected correlation $\mathbb{E}(\rho)$ and flipped around it otherwise (part 4). So, we are now only left to show part 1, that is, $\mathbb{E}(\pi(\rho)) > \mathbb{E}(\pi^{rs}(\rho))$.

We show that the profit generated by contract above is concave in ρ. Since at correlation $\mathbb{E}(\rho)$ we have $\pi(\mathbb{E}(\rho)) = \pi^{RS}(\mathbb{E}(\rho))$, Jensen's inequality implies that $\mathbb{E}(\pi(\rho)) > \pi^{RS}(\mathbb{E}(\rho))$.

In an interior contract, we have that

$$\frac{\partial^2\pi}{\partial\rho^2} = \frac{q_1}{\eta}\left(\frac{\partial\mu_L}{\partial\rho}\right)^2 + \frac{1-q_1}{\eta}\left(\frac{\partial\mu_H}{\partial\rho}\right)^2 > 0,$$

for the corner solution in which $x_L(\rho) = 0$ we have that $\frac{\partial^2\pi}{\partial\rho^2} = 0$ and for the corner solution in which $x(L,\rho) = x(H,\rho)$ we have

$$\frac{\partial^2\pi}{\partial\rho^2} = \frac{1}{\eta}\left(q_1\frac{\partial\mu_L}{\partial\rho} + (1-q_1)\frac{\partial\mu_H}{\partial\rho}\right)^2 > 0.$$

Therefore, the profit function is concave and we obtain the inequality as desired.

An analogous argument shows the proposition for the case $\mu^e_H < \mu^e_L$. □

9.2 Interim profit function

To prove the next set of results, we will introduce the concept of interim profit for the insurer, and state and prove a lemma characterizing it. Define the function $\hat{\pi}(\rho, c)$ to be the maximum profit that an insurer can obtain in the subgame in which both parties believe that the correlation is ρ, and the slope of the profit function as defined by Equation (4) is equal to c, that is $c(\rho) = c$.

The analysis here is stated for the case when ρ is common knowledge in the subgame, that is, r is the identity mapping, $r(\rho) = \rho \; \forall \; \rho$. Since profits ($\pi$) and utility ($u$) are both linear in ρ, we will later replace ρ with the expectation of the partition in which it lies and all results stated here would carry through.

We reproduce Equation (4) here for completeness:

$$\left.\frac{\partial \pi(\rho;\hat{\rho})}{\partial \rho}\right|_{\hat{\rho}=\rho} = \sigma x_{r(\rho)}(L) \cdot (\mu_{LH} - \mu_{LL}) - \sigma x_{r(\rho)}(H) \cdot (\mu_{HH} - \mu_{HL}) \equiv c(\rho).$$

The following lemma establishes that the interim profit function, $\hat{\pi}(\rho, c)$ is (i) single peaked with respect to ρ, with a peak at ρ^*, (ii) it is convex with respect to ρ both to the right and the left of ρ^*; and (iii) it is strictly concave with respect to c. The next result states some other key properties of the interim profit function.

Lemma 2. *Let ρ^* solve $\mu_\rho(H) = \mu_\rho(L)$ and fix $\rho \geq \rho^*$. Then:*

1. *$\exists$ ρ_1 and ρ_2 with $\bar{\rho} \geq \rho_2 \geq \rho_1 \geq \rho^*$ such that*
   - *(a) for $\rho \in [\rho^*, \rho_1]$, $\hat{\pi}(\cdot, c)$ is linear and strictly decreasing in ρ;*
   - *(b) for $\rho \in [\rho_1, \rho_2]$, $\hat{\pi}(\cdot, c)$ is strictly convex and strictly decreasing in ρ;*
   - *(c) for $\rho > \rho_2$, $\hat{\pi}(\cdot, c)$ is constant in ρ.*
2. *The function $\hat{\pi}(\rho, \cdot)$ is strictly concave in c.*

Analogous characterization holds for $\rho < \rho^*$.

Proof. We prove the result in xx steps.

Step 1. Fix $\rho \geq \rho^*$, so that $\mu_\rho(H) \geqslant \mu_\rho(L)$. To simplify notation, let $\mu_L = \mu_\rho(L)$, $\mu_H = \mu_\rho(H)$, $x_L = x_\rho(L)$, $x_H = x_\rho(H)$, $p_L = p_\rho(L)$ and $p_H = p_\rho(H)$. As in the main text, let $k_L = \sigma(\mu_{LH} - \mu_{LL})$ and $k_H = \sigma(\mu_{HH} - \mu_{HL})$. We prove part 1 first.

Step 2. In an interior solution we obtain that:

$$\begin{aligned} x_L &= 1 - \tfrac{1-q_1}{\eta q_1}(\mu_H - \mu_L) + \tfrac{\beta}{\eta q_1} k_L \\ x_H &= 1 - \tfrac{\beta}{\eta(1-q_1)} k_H \end{aligned}$$

where $\beta = \frac{\eta q_1(1-q_1)(c-k_L+k_H)+(1-q_1)^2(\mu_H-\mu_L)k_L}{(1-q_1)k_L^2+q_1k_H^2}$ is the Lagrange multiplier for the convexity constraint. Further, substituting for x_L and x_H, the optimal profit is

$$\hat{\pi} = \frac{\eta}{2} - (1-q_1)(\mu_H - \mu_L) + \frac{(1-q_1)^2}{2\eta q_1}(\mu_H - \mu_L)^2 - \frac{(1-q_1)k_L^2 + q_1 k_H^2}{2\eta q_1(1-q_1)}\beta^2.$$

Since β depends on ρ only through $\mu_H - \mu_L$, π is quadratic in ρ. Its first derivative with respect to ρ is

$$(1-q)\left(\frac{\partial(\mu_H - \mu_L)}{\partial \rho}\right)\left(-1 + \frac{1-q_1}{\eta q_1}(\mu_H - \mu_L) - \frac{\beta k_L}{\eta q_1}\right) < 0,$$

since the partial derivative is positive and the last term has to be negative to guarantee that x_L is positive. Its second derivative with respect to ρ is given by

$$\frac{(1-q_1)^2 k_H^2}{\eta((1-q_1)k_L^2+q_1k_H^2)}\left(\frac{\partial(\mu_H-\mu_L)}{\partial\rho}\right)^2>0.$$

Therefore, $\pi(\rho, c)$ is strictly decreasing and strictly convex with respect to ρ when the solution is interior. Next we will construct the bounds for interiority: ρ_1 and ρ_2.

Step 3. A corner solution in which the insurer sells only to $\theta_1 = H$ insuree occurs when x_L is negative, that is, when $\frac{\eta(q_1k_H^2+(1-q_1)(c+k_H)k_L)}{(1-q_1)k_H^2} < \mu_H-\mu_L$. Since $\mu_H-\mu_L$ is increasing in ρ this condition may hold only for large correlations. Let ρ_2 to be equal to the correlation that makes this condition to hold with equality if it is smaller than $\bar{\rho}$ and equal to $\bar{\rho}$, otherwise.

In such a corner it has to be that $x_H = \max\{\frac{-c}{k_H}, 0\}$. Using the constraint IR_H we obtain that $(1-q_1)(p_H-\mu_Hx_H) = (1-q_1)\max\{\frac{-\eta c}{2k_H}\left(2+\frac{c}{k_H}\right), 0\}$, so that the profit function is constant in ρ.

Step 4. Finally, to satisfy both IC constraints of the isnuree it has to be that $x_H \geqslant x_L$, but this might not be the case in the interior solution we characterized above. In particular for $\rho < \rho_2$ the constraint $x_H \geqslant x_L$ is not satisfied if

$$\begin{cases} \mu_H-\mu_L < \frac{\eta(c-k_L+k_H)((1-q_1)k_L+q_1k_H)}{-(1-q_1)k_H(-k_H+k_L)} & \text{if } -k_H+k_L>0 \\ \mu_H-\mu_L > \frac{\eta(c-k_L+k_H)((1-q_1)k_L+q_1k_H)}{-(1-q_1)k_H(-k_H+k_L)} & \text{if } -k_H+k_L<0. \end{cases}$$

Notice that in the first case the inequality is never true if $c-k_L+k_H<0$ and in the second case it is never true if $c-k_L+k_H>0$. In the domain in which the inequalities can be true, the correlation that makes the first inequality to hold with equality is smaller than ρ_2, and the correlation that makes the second inequality to hold with equality is larger than ρ_2. Then we define ρ_1 in the first case as the maximum of the correlation that makes the inequality to hold with equality and ρ^*, and in the second case we just define it as ρ^*.

For correlations in $[\rho^*, \rho_1]$ the insurer offers a unique package $x_H = x_L = \frac{c}{-k_H+k_L} > 0$ at the price that makes the constraint IR_L to hold with equality. This generates profits equal to

$$\hat{\pi} = \frac{\eta}{2} - \frac{\eta}{2}\left(1-\frac{c}{k_L-k_H}\right)^2 - (1-q_1)(\mu_H-\mu_L)\frac{c}{k_L-k_H}.$$

Since the profits depend on ρ only trough $\mu_H-\mu_L$ and this dependence is linear we conclude that the profits are linear with a slope $s_1 = -\frac{((1-q_1)k_L+q_1k_H)c}{q_1(k_L-k_H)} < 0$.

The second inequality is true only for correlations for which the constraint $x_L \geqslant 0$ binds as well, that is, in the solution to the problem without these constraints both x_L and x_H are negative. Therefore, since the H type is willing to pay more for insurance, the insurer sells only to him and we are back to the case in step 3.

Step 5. To prove 2, we only need to take the second derivative of the profit functions with respect to c for all possible cases. In the corner solution with $x_L = 0$ and $x_H > 0$ we have

$$\frac{\partial^2 \hat{\pi}}{\partial c^2} = \frac{-(1-q_1)\eta}{k_H^2} < 0,$$

in the corner solution with $x_L = x_H > 0$ we have

$$\frac{\partial^2 \hat{\pi}}{\partial c^2} = \frac{-\eta}{(k_L - k_H)^2} < 0,$$

and in the interior solution with $x_H > x_L > 0$ we have that

$$\frac{\partial^2 \hat{\pi}}{\partial c^2} = \frac{-\eta q_1(1-q_1)}{(1-q_1)k_L^2 + q_1 k_H^2} < 0.$$

Therefore, the function $\hat{\pi}(\rho, \cdot)$ is strictly concave with respect to c. □

As a final thought, the interim profit function is useful not only for the case in which both players hold the same ex-post belief. It can be used as well to characterize the expected or ex-ante profits of the insurer by replacing the first argument of $\hat{\pi}$ with the expected correlation of the various partitions the insurer intends to communicate to the insuree. Therefore, we will use this function in multiple cases by replacing the observe correlation by the expected correlation that is generated by a message function in more complicated selling mechanisms, i.e. where r is not the identity mapping.

9.3 Proof of finiteness: Proposition 4

Proof. We do the proof in five steps. First we argue that the optimal mechanism has at most countable partitions and then use the next four steps to gradually build towards the finiteness of the number of partitions.

Step 1. From Lemma 1 we know that incentive compatibility of the insuree is equivalent to the envelope and a convexity condition. Start by considering the relaxed version of the optimization problem without the convexity constraint, i.e, without the condition that $c(\rho)$ is non-decreasing. Let $c^{RS}(\rho)$ the implied profits' slope obtained from the solution of the Rothschild-Stiglitz problem. From Proposition 1 it is clear that whenever there is an interior solution the resulting $c^{RS}(\rho)$ is strictly decreasing. Furthermore, this is true even if the non-decreasing condition is not imposed for some arbitrary interval.

Now, go back to the original problem which imposes the convexity condition. The analysis of the relaxed problem implies that the non-decreasing (or convexity) condition binds almost everywhere. So in the solution to the original problem, $c(\rho)$ must be a step function, which implies that $\pi(\rho, \rho')$ is piecewise linear. Therefore, the profit function's slope can take at most a countable number of different values. We will now show that the profit function's slope actually can take only a finite number of different values.

Step 2. Suppose by contradiction that the profit function's slope takes an infinite number of different values. Fix $\epsilon_1 > 0$. As the slope take an infinite number of different values, there exists two different messages m_1 and m_2 with the following properties:

1. the preimages $r^{-1}(m_1)$ and $r^{-1}(m_2)$ are contiguous,

2. the contracts c_{m_1} and c_{m_2} generate different profits slopes $c_1 < c_2$ such that there are not other posted contracts that generate the same slopes, and

3. the insuree's expected correlations after observing those messages, $\rho_1 = E[\rho \mid r(\rho) = m_1] < \rho_2 = E[\rho \mid r(\rho) = m_2]$, are such that $\rho_2 - \rho_1 < \epsilon$.

Without loss assume that $\rho_1 > \rho^*$. By Proposition 1 we have that $c^{RS}(\rho_1) > c^{RS}(\rho_2)$. Further, Lemma 2 shows that the functions $\hat{\pi}(\rho_1, \cdot)$ and $\hat{\pi}(\rho_2, \cdot)$ are concave in c. Therefore, $c^{RS}(\rho_1)$ and $c^{RS}(\rho_2)$ correspond to the unique maximizers of these functions, respectively, because Rothschild-Stiglitz problem does not impose constraints on c. Therefore, for any $c < c^{RS}(\rho)$, $\hat{\pi}(\rho, \cdot)$ is increasing, and for any $c > c^{RS}(\rho)$, $\hat{\pi}(\rho, \cdot)$ is decreasing.

Step 3. Before delving into impossibility to infinite slopes, we show as an intermediate step that it cannot be that $c_1 < c^{RS}(\rho_1)$ and $c_2 > c^{RS}(\rho_2)$. If they were, c_1 would be in the increasing part of $\hat{\pi}(\rho_1, \cdot)$ and c_2 would be in the decreasing part of $\hat{\pi}(\rho_2, \cdot)$. Then by slightly increasing c_1 and slightly decreasing c_2 the firm relaxes the feasibility constraints and is able to increase its profits (see Figure 9). Therefore, the original c_1 and c_2 cannot be optimal, a contradiction. Thus, we conclude that $c_1 \geq c^{RS}(\rho_1)$ or $c_2 \leq c^{RS}(\rho_2)$.

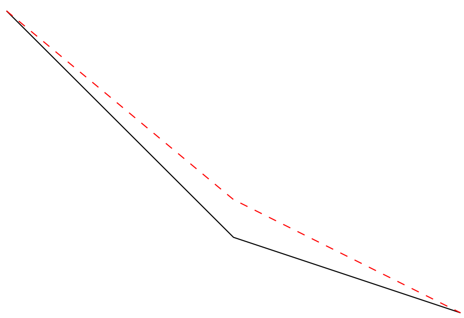

Figure 9: Improvement arguments for Step 3 of the proof.

Step 4. Suppose first that $c_1 \geq c^{RS}(\rho_1)$. We show by contradiction that in this case whenever the slope takes an infinite number of different values, the resulting profit function cannot be continuous, contradicting the convexity constraint.

As $c_2 > c_1$ and $c^{RS}(\rho_1) > c^{RS}(\rho_2)$, the previous discussion implies that both c_1 and c_2 are located in the decreasing part of $\hat{\pi}(\rho_2, \cdot)$. Therefore, $\hat{\pi}(\rho_2, c_1) - \hat{\pi}(\rho_2, c_2) = \delta > 0$.

Let $\pi(.)$ denote the optimal profit function. By feasibility it must be that $\pi(\rho_1) < \hat{\pi}(\rho_1, c_1)$ and $\pi(\rho_2) < \hat{\pi}(\rho_2, c_2)$. We argue that it has to be that $\pi(\rho_1) = \hat{\pi}(\rho_1, c_1)$ by showing perturbations that increase the seller's profits. We pick perturbations that do not change the function π outside the two partitions we are considering. There are two cases.

If both $\pi(\rho_1) < \hat{\pi}(\rho_1, c_1)$ and $\pi(\rho_2) < \hat{\pi}(\rho_2, c_2)$ then the insurer/seller can slightly increase c_1 and move ρ_2 to the right. Figure 10a shows that this allows the seller to increase profits and it is feasible because both inequalities are strict and all functions are continuous.

Instead, if $\pi(\rho_1) < \hat{\pi}(\rho_1, c_1)$ and $\pi(\rho_2) = \hat{\pi}(\rho_2, c_2)$ then the seller can slightly decrease c_2 and move ρ_1 to the left. As the initial c_2 is in the decreasing part of $\hat{\pi}(\rho_2, \cdot)$ this relaxes the constraint at the second partition and makes the perturbation feasible. Figure 10b depicts why this change benefits the firm.

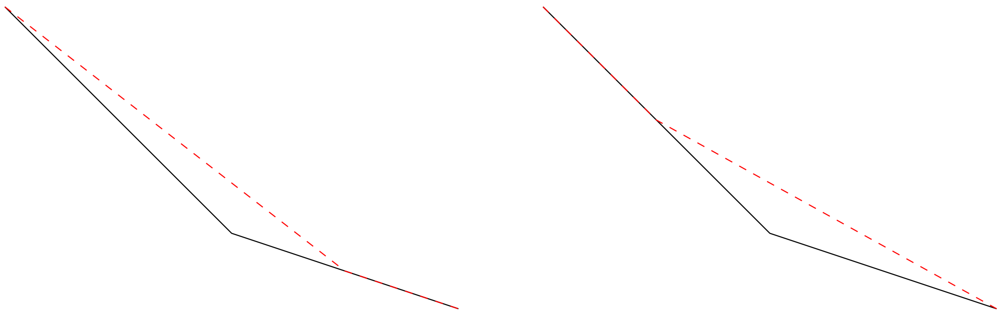

(a) Increase c_1 and move ρ_2 to the right. (b) Decrease c_2 and move ρ_1 to the right.

Figure 10: Improvement arguments for Step 4 of the proof.

Let $\dot{\rho}$ be the threshold correlation between the two partitions. Fixed $\delta > 0$. By continuity of $\hat{\pi}(\rho, c)$ with respect to ρ there exists ϵ_2 such that if $\mid \rho_1 - \rho_2 \mid < \epsilon_2$ then $\mid \hat{\pi}(\rho_1, c_1) - \hat{\pi}(\rho_2, c_1) \mid < \frac{\delta}{4}$. Take ϵ_2 small enough such that if $c_1 < 0$ then $\epsilon_2 < \frac{-\delta}{4c_1}$ and if $c_1 > 0$ then $\epsilon_2 < \frac{\delta}{4c_1}$.

Let $\epsilon = \min\{\epsilon_1, \epsilon_2\}$. Putting all our calculations together we obtain that

$$\begin{aligned}
\pi(\rho_1) + c_1(\dot{\rho} - \rho_1) = \quad & \hat{\pi}(\rho_1, c_1) + c_1(\dot{\rho} - \rho_1) \\
& > (\hat{\pi}(\rho_2, c_1) - \tfrac{\delta}{4}) - \tfrac{\delta}{4} \\
& = \hat{\pi}(\rho_2, c_2) + \delta - \tfrac{\delta}{2} \\
& > \hat{\pi}(\rho_2, c_2) + \tfrac{\delta}{4} + c_2(\rho_2 - \dot{\rho}) \\
& > \hat{\pi}(\rho_2, c_2) + c_2(\rho_2 - \dot{\rho}) \\
& \geq \pi(\rho_2) + c_2(\rho_2 - \dot{\rho}),
\end{aligned}$$

but this implies that the pasting conditions is not satisfied. This is a contradiction because the insurer's optimal profit function π is convex and therefore it is continuous. Therefore, there cannot exist such elements of the partition.

An analogous argument leads to a contradiction in the case in which $c_2 \leq c^{RS}(\rho_2)$.

Step 5. Finally, we argue that in the optimal contract there can be at most two elements of the partition that generate the same profit's slope c. Lemma 2 implies that $\hat{\pi}(\cdot, c)$ is convex. Therefore, a contract with three different messages that generate the same slope is trivially dominated by a contract with only two different messages that lead to the same slope.

So, the partition in the optimal contract can contain only a finite number of elements. □

9.4 Proofs for Section 6

Proof of Proposition 5. We prove part 1 first. Notice that $\mathbb{E}(\rho) > \rho^*$. From Proposition 1 we obtain that as $\eta \to \infty$, $x_H^{RS} = 1$ and $x_L^{RS} \to 1$ if $\rho > \rho^*$, and $x_L^{RS} = 1$ and $x_H^{RS} \to 1$ otherwise. That is, the contract that is offered in Rothschild-Stiglitz becomes independent of the correlation. Therefore, the seller can approximate the optimal profits at each point by offering only one contract that consists of the Rothschild-Stiglitz menu at the ex-ante expected correlation. In the limit, when the seller offers full insurance to both types, can obtain profits equal to $\frac{\eta}{2}$.

If instead the seller decides to offer more than one contract, he has to face a non-vanishing cost that is imposed by the convexity constraint. Actually, from the proof of Lemma 2 can be observed that $\hat{\pi}(\rho, c)$ is always smaller and away from $\frac{\eta}{2}$ as long as $c \neq k_L - k_H$, the contract's slope that is generated by the full insurance contract. Therefore, in the limit the seller cannot improve upon having one contract.

Now to the second part. At the expected correlation the seller can always attain a profit at least equal to $(1 - q_1)\frac{\eta}{2}$ by offering only full insurance to the H type at the monopolist price, and this is exactly the optimal profit when $\eta \leq \frac{1-q_1}{q_1}(\mu_H(E(\rho)) - \mu_L(E(\rho)))$. Therefore, $(1 - q_1)\frac{\eta}{2}$ is a lower bound of the seller profits when offering a one-partition contract and it coincides with its profit for η small.

The profits that the seller can obtain at ρ^* are equal to $\pi(\rho^*) = \frac{\eta}{2}$. This value is an upper bound of the profits that the seller can obtain: for any belief the buyer has, the seller can never obtain this profit unless the buyer's belief is exactly ρ^*.

Let $\epsilon = q_1\frac{\eta}{2}$, an upper bound of the difference between the profits in the optimal contract and the one-partition contract. As η converges to 0, ϵ converges to 0. □

Proof of Proposition 6. Let the prior be given by $\mathbb{P}(\rho = \rho_1) = f$ and $\mathbb{P}(\rho = \rho_2) = 1 - f$, where $\rho_1 < \rho_2$. Also, denote by $\bar{\rho}$ the expected ex-ante value of ρ, that is, $\bar{\rho} = f\rho_1 + (1-f)\rho_2$. And, recollect ρ^* solves the equation $\mu_\rho(H) = \mu_\rho(L)$.

We first show that for $\rho_1 < \rho^* < \rho_2$, the insurer offers a pooling contract along ρ. In this case, Lemma 2 implies that for any c, $\hat{\pi}(\bar{\rho}, c) > p\hat{\pi}(\rho_1, c) + (1-p)\hat{\pi}(\rho_2, c)$. Suppose the firm chooses to offer different contracts for correlations ρ_1 and ρ_2 with profit slopes c_1 and c_2, respectively. To satisfy the IC constraints it has to be that $c_2 \geq c_1$. Then, Lemma 2 implies that either $\hat{\pi}(\rho_1, c_1) < \hat{\pi}(\rho_1, c_2)$ or $\hat{\pi}(\rho_2, c_2) < \hat{\pi}(\rho_2, c_1)$ (since the optimal c for ρ_2 is smaller than the one that is optimal for ρ_1). Both pieces imply that in the first case the firm would be better off by offering the same contracts for both correlations with an implied slope c_2, and in the second case with implied slope c_1.

Now, suppose $\rho_1 < \rho_2 \leqslant \rho^*$ <u>or</u> $\rho^* \leqslant \rho_1 < \rho_2$. We offer parametric examples to show that both pooling or separation is possible at the optimum. Assume $f = \frac{1}{2}$ and $k_L = k_H = k$ (say).

We first compute the profits that the insurer can get by pooling along ρ. When pooling, the insurer can do no better than offering the Rothschild-Stiglitz contract at the expected correlation $\bar{\rho}$. Therefore, the insurer's profit in this case is given by

$$\pi_1 = \frac{\eta}{2} - \frac{\gamma + \bar{\rho}K}{2} + \frac{\gamma^2 + 2\gamma\bar{\rho}K + \bar{\rho}^2 K^2}{4\eta}.$$

On the other hand, if the insurer wants to offer different contracts for different correlations, then it has to solve the problem

$$\max_c \frac{1}{2}\hat{\pi}(\rho_1, c) + \frac{1}{2}\hat{\pi}(\rho_2, c)$$
$$\text{such that } \hat{\pi}(\rho_1, c) + c(\rho_2 - \rho_1) = \hat{\pi}(\rho_2, c).$$

The equality constraint follows from the two IC constrains that need to be satisfied: The insurer wants to offer contracts with a unique slope because it wants to offer a large slope for correlation ρ_1 and a small one for ρ_2, but the IC constraint requires that the second one is at least as large as the first one, hence the equality.

The only value of c that satisfies the constraint is $c = \frac{k(-2\eta+\gamma+\bar{\rho})}{3\eta}$. With a slope c the firm will obtain the following ex-ante profits:

$$\pi_2 = \frac{\eta}{2} - \frac{\gamma + \bar{\rho}k}{2} + \frac{\gamma^2 + 2\gamma\bar{\rho}k}{8\eta} + \frac{\rho_1^2 k^2 + \rho_2^2 k^2}{16\eta} - \frac{\eta c^2}{8k^2} + \frac{c(\gamma + \bar{\rho}k)}{4k}$$

Therefore, with the optimal slope it is better for the insurer to offer the pooling contract for both correlations rather than separating one iff

$$\frac{\gamma^2 + 2\gamma\bar{\rho}k + \rho_1\rho_2 k^2}{2} > -\frac{(-2\eta + \gamma + \bar{\rho})^2}{18} + \frac{(-2\eta + \gamma + \bar{\rho})(\gamma + \bar{\rho}k)}{3}$$
$$\Leftrightarrow 4\gamma^2 + 6\gamma\bar{\rho}k + 9\rho_1\rho_2 k^2 + 4\eta^2 + \bar{\rho}^2 + 8\eta\gamma + 12\eta\bar{\rho}k > 4\eta\bar{\rho} + 4\gamma\bar{\rho} + 6\bar{\rho}^2 k.$$

When $\rho_1 < 0$, $\bar{\rho} > 0$, η and γ are small and k is large the condition is violated and the insurer will offer a separating contract. For the same correlations, when η and γ are large and k is small, the condition is satisfied and the firm would prefer to offer pool across the two correlations.

For example, when $\gamma = 0.05$, $\eta = 0.01$, $k = 2$, $\rho_1 = -0.1$ and $\rho_2 = 0.3$, the insurer prefers to offer a separating contract for the two correlations, but when $\gamma = 0.1$, $\eta = 0.05$, $k = 0.1$, the insurer prefers to offer the same contract for both correlations. □

Proof of Proposition 7. When the seller sends only one message, the buyer's belief is equal to the prior. Since the profits are linear on the correlation and the same contract is offered for all correlations, the expected value of the profits is equal to the profits generated by the contract when the correlation is equal to the expected correlation.

Therefore, the optimal contract coincides with benchmark at the correlation $\mathbb{E}(\rho)$. Then we have:

$$\mathbb{E}[\pi(\rho)] = \pi_e^{RS} \quad \text{and} \quad x_{r(\rho)}(H) = x_e^{RS}(H), x_{r(\rho)}(L) = x_e^{RS}(L) \;\forall\; \rho.$$

$\square$

Proof of Proposition 8. We first argue that the profit function slope has to be constant. Suppose by contradiction that optimally the seller offers two contracts that generate slopes $c_1 < c_2$. To satisfy convexity of the profit function it has to be that the first contract is targeted to small correlations and the second one to large correlations. Fix the optimal message function r which sends message $\underline{m}$ when the realized correlation is in $[\underline{\rho}, \tilde{\rho}]$ and message $\bar{m}$ when the realized correlation is in $(\tilde{\rho}, \overline{\rho}]$.

In terms of the interim profit function $\hat{\pi}(\rho, c)$ defined in Section 9.2, the insurer solves the following problem:

$$\begin{aligned} \max_{c_1, c_2, \tilde{\rho}, \pi_1, \pi_2} \quad & Pr(\rho < \tilde{\rho})\pi_1 + Pr(\rho > \tilde{\rho})\pi_2 & \\ s.t. \quad & \hat{\pi}(\rho_i, c_i) \geq \pi_i \quad \forall i \in \{1, 2\} & \text{feasibility } i \\ & \pi_1 + c_1(\tilde{\rho} - \rho_1) = \pi_2 + c_2(\tilde{\rho} - \rho_2) & \text{continuity} \end{aligned}$$

where ρ_i is the expected correlation after observing the respective message; the first constraints are feasibility constraints, i.e., they guarantee that the insurer does not accrue more profits than those that the insurer can obtain in the subgame when the buyer beliefs correspond to the partition's expected value.

First, notice that the feasibility constraint has to bind. If not we can decrease c_1 (increase c_2) which relaxes the continuity constraint, and allows to increase π_1 (π_2).

Denote by $c(\rho_i)$ the unique value the maximizes $\hat{\pi}(\rho_i, c)$, which exists by Lemma 2 and coincides with Rothschild-Stiglitz implied profits'slope by Proposition 1. Proposition 1 implies that $c(\rho_1) > c(\rho_2)$. Then since $c_2 > c_1$ there are three cases: $c(\rho_2) \leq c_2$ and $c(\rho_1) \geq c_1$, $c_1 < c_2 < c(\rho_2) < c(\rho_1)$, and $c(\rho_2) < c(\rho_1) < c_1 < c_2$.

In the first case, when increasing c_1 and decreasing c_2 simultaneously, Lemma 2 implies that the LHS of the feasibility constraints increases (at least one of them and the other one stays constant). Then it is possible to increase π_1 and/or π_2, contradicting that the initial contract is optimal.

In the second case, increasing c_2 reduces the RHS of the continuity constraint and increases the LHS of the feasibility constraint 2. Then it is possible to increase π_2, contradicting that the initial contract is optimal. A similar argument shows that in the third case the initial contract cannot be optimal.

Then, it has to be that $c_1 = c_2 = c$, that is, the slope of the profit function is constant. Further, we have shown that $c(\rho_1) \geq c \geq c(\rho_2)$.

The characterization of equilibrium in the proof of Lemma 2 shows that when contracts are separating, x_L is increasing in c and x_H is decreasing in c. Since Rothschild-Stiglitz contracts use constants $c(\rho_1)$ and $c(\rho_2)$, we have that when contracts are separating $x_1(H) \geq x^{RS}_{\rho^1}(H)$, $x_1(L) \leq x^{RS}_{\rho^1}(L)$, $x_2(H) \leq x^{RS}_{\rho^1}(H)$, and $x_2(L) \geq x^{RS}_{\rho^1}(L)$. $\square$

9.5 Proofs for Section 7

Proof of Proposition 9. First note that the insuree can choose from which seller to buy. However, independently of which seller he buys from, he has the same belief and this belief corresponds to the one that is implied by the message sent by the informed insurer and insuree's updating rule (which depends on her sophistication).

Therefore, the insuree will buy from the informed seller only if he can obtain a (weakly) higher utility from buying an insurance contract from this insurer than buying the benchmark contract at the expected correlation, c^e .

Fix $\mu^e_H = \mu_{\bar{\rho}}(H)$ and $\mu^e_L = \mu_{\bar{\rho}}(L)$ be the expected probabilities at the expected correlation $\bar{\rho} = \mathbb{E}[\rho]$. Let $\tilde{\mu}_H$ and $\tilde{\mu}_L$ be the expected beliefs of types $\theta_1 = L$ and H after observing the message send by the informed insurer, respectively, and μ_L and μ_H the actual loss probabilities which are only known to the informed insurer.

To simplify notation, since the optimization is pointwise, we will write x^e_H for $x^{RS}_e(H)$ and x_H for $x_{r(\rho)}(H)$, etc. The informed insurer solves the following problem:

$$
\begin{aligned}
&\max\ q_1(p_L - \mu_L x_L) + (1-q_1)(p_H - \mu_H x_H) \\
&s.t.\ \tilde{\mu}_H x_H - \frac{\eta}{2}(1-x_H)^2 - p_H \geq \tilde{\mu}_H x_L - \frac{\eta}{2}(1-x_L)^2 - p_L \\
&\quad \tilde{\mu}_L x_L - \frac{\eta}{2}(1-x_L)^2 - p_L \geq \tilde{\mu}_L x_H - \frac{\eta}{2}(1-x_H)^2 - p_H \\
&\quad \tilde{\mu}_H x_H - \frac{\eta}{2}(1-x_H)^2 - p_H \geq \max\{\tilde{\mu}_H x^e_H - \frac{\eta}{2}(1-x^e_H)^2 - \bar{p}_H, \tilde{\mu}_H x^e_L - \frac{\eta}{2}(1-x^e_L)^2 - \bar{p}_L, -\frac{\eta}{2}\} \\
&\quad \tilde{\mu}_L x_L - \frac{\eta}{2}(1-x_L)^2 - p_L \geq \max\{\tilde{\mu}_L x^e_H - \frac{\eta}{2}(1-x^e_H)^2 - \bar{p}_H, \tilde{\mu}_L x^e_L - \frac{\eta}{2}(1-x^e_L)^2 - \bar{p}_L, -\frac{\eta}{2}\}
\end{aligned}
$$

Suppose that $\tilde{\mu}_H > \tilde{\mu}_L$. Then $\theta_1 = H$ is the high risk type and as is standard the IC_H constraint binds and IC_L is slack, which implies that x_H is offered the efficient contract since x_H does appear on the RHS on any binding constraint. Analogously, when $\tilde{\mu}_H < \tilde{\mu}_L$, then $\theta_1 = L$ is the low risk type and it gets the efficient contract.

So, when the insuree is gutgläubig, as before $r(\rho) = \overline{\rho}$ or $r(\rho) = \underline{\rho}$, the seller always sells to both types due to the binding IR constraints at the expected correlation and rest of the analysis carried through as is from Proposition xx.

Now, suppose that the insuree can do Bayesian inference. If the informed insurer only offers one partition in the original model without competition then it is obvious that she continues to to do the same here and the optimal contract coincides with Rothschild-Stiglitz contract at the expected correlation.

Next, suppose the optimal contract without competition features more than two partitions. The new RHS of the IC constraints impose extra restrictions in the problem : for $\rho > \rho^e$ x^e_H generates to type H a surplus larger than $-\frac{\eta}{2}$ and for $\rho < \rho^e$, x^e_L to type L a surplus larger than $-\frac{\eta}{2}$. Therefore, the maximum profits that can be reached have decreased for any possible partition that the firm can offer.

Finally, if the insurer offers two partitions in the model without competition, then it is obvious that in the model with competition ether one or two partitions will be offered which completes the proof. □

Proof of Proposition 10. Let $[\underline{c}, \bar{c}]$ be the smallest interval that contains all the slopes of the profit function in the optimal contract. Let $c^{RS}(\underline{\rho})$ and $c^{RS}(\bar{\rho})$ be the maximizers of the function $\hat{\pi}(\rho, c)$ for correlations $\underline{\rho}$ and $\bar{\rho}$, respectively. Proposition 1 implies that $c^{RS}(\bar{\rho}) < c^{RS}(\underline{\rho})$.

We argue that $c^{RS}(\underline{\rho}) \geq \underline{c}$ or $c^{RS}(\bar{\rho}) \leq \bar{c}$. Suppose none of the two is true in the optimal contract. Then $c^{RS}(\bar{\rho}) < c^{RS}(\underline{\rho}) < \underline{c}$ or $\bar{c} < c^{RS}(\bar{\rho}) < c^{RS}(\underline{\rho})$. Suppose the first case is true. Then, by decreasing the slope of the profit function uniformly, Lemma 2 implies that the profits that are feasible for each correlation increase uniformly. Therefore, since increasing the optimal profit function by a constant, does not affect the seller's IC constraints, the original contract was not optimal. The argument for the second case is analogous.

Since convexity requires that the profits' slope $c(\rho)$ is weakly increasing and $c^{RS}(\rho)$ is decreasing by Proposition 1, they can coincide at most a one correlation, call it $\tilde{\rho}$.[19] Therefore, only at $\tilde{\rho}$ is possible that the optimal profits are equal to Rothschild-Stiglitz' profits.

By definition of $\tilde{\rho}$ we have that for $\rho < \tilde{\rho}$, $c(\rho) < c^{RS}(\rho)$ and for $\rho > \tilde{\rho}$, $c(\rho) > c^{RS}(\rho)$. The characterization of equilibrium in the proof of Lemma 2 shows that when contracts are separating, x_L is increasing in c and x_H is decreasing in c. Then when contracts are separating $x_H(\rho) \geq x^{RS}_H(\rho)$, and $x_L(\rho) \leq x^{RS}_L(\rho)$ for $\rho < \tilde{\rho}$, and $x_H(\rho) \leq x^{RS}_H(\rho)$, and $x_L(\rho) \geq x^{RS}_L(\rho)$ for $\rho > \tilde{\rho}$, with a strict inequality in each case. □

Lastly, we state and prove a result which is described in words in Section 7.

Proposition 11. *1. For any $\rho \in [\underline{\rho}, \bar{\rho}] \setminus \mathbb{I}$, discrimination along θ_1 in the full revelation model is less than in the benchmark when ρ is common knowledge: $\|x_\rho(H) - x_\rho(L)\| < \|x^{RS}_\rho(H) - x^{RS}_\rho(L)\|$.*

2. If $\Gamma < 0$, then

(a) total coverage is higher for the full information revelation model when $\rho > \tilde{\rho}$: $x_H(\rho) + x_L(\rho) < x^{RS}_H(\rho) + x^{RS}_L(\rho)$, and

[19] If they do not cross it means that there is a discontinuity at some correlation in $c(\rho)$. In this case call $\tilde{\rho}$ the correlation at which the discontinuity occurs.

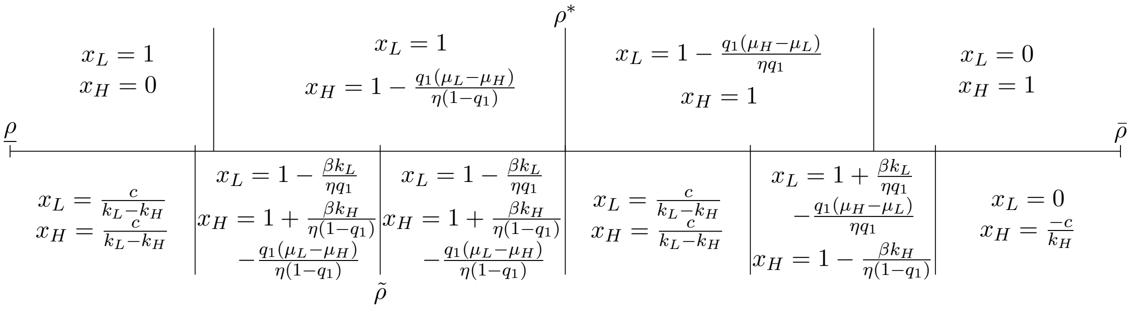

Figure 11: Configuration of contracts for the case $\tilde{\rho} < \rho^*$.

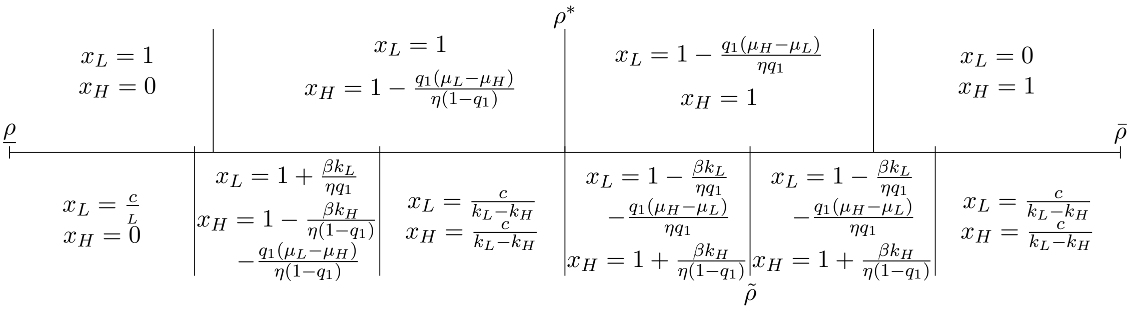

Figure 12: Configuration of contracts for the case $\tilde{\rho} > \rho^*$

.

(b) *total coverage is higher for the benchmark model when* $\rho < \tilde{\rho}$*:* $x_H(\rho) + x_L(\rho) < x_H^{RS}(\rho) + x_L^{RS}(\rho)$.

And, the opposite is true if $\Gamma > 0$.

Proof of Proposition 11. From Lemma 2 we have that the multiplier on the convexity constraint, $\beta(\rho)$, is positive for each correlation $\rho < \tilde{\rho}$ and negative for each correlation $\rho > \tilde{\rho}$.

There are two main cases: $\tilde{\rho} < \rho^*$ and $\tilde{\rho} > \rho^*$. Suppose first that $\tilde{\rho} < \rho^*$. The configuration of contracts according to the value of the correlation is depicted in Figure 11. In the figure the optimal full information revelation contracts are shown on the bottom line and the Rothschild-Stiglitz contracts in the top one. To simplify the expressions, we do not specify how they depend on ρ.

Let us show the first property on the extent of discrimination. Suppose the solution is interior for both sets of assumptions. In that case, it is easy to check by using the sign of β that for $\rho < \rho^*$ $x_L < x_L^{RS}$ and $x_H > x_H^{RS}$ and for $\rho > \rho^*$ $x_L > x_L^{RS}$ and $x_H < x_H^{RS}$. Next, suppose the correlations in such that we have a corner solution, the same properties are concluded from continuity since the multipliers are larger in absolute value as ρ is farther away from $\tilde{\rho}$, and $\mu_H - \mu_L$ is increasing in ρ.

Now we prove the second property on total coverage. Suppose the solution is interior. We calculate the difference between $x_H + x_L$ and $x_H^{RS} + x_L^{RS}$ which equals $\beta\left(\frac{k_L}{\eta q_1} - \frac{k_H}{\eta(1-q_1)}\right)$. Then the second property follows for these cases follows from the sign of β and rearranging terms.

For correlations in which there is no interior solution, the result follows from continuity and that the multipliers are larger as the absolute value of ρ is farther away from $\tilde{\rho}$.

If $\tilde{\rho} > \rho^*$ the configuration of contracts slightly changes, and it is depicted in Figure 12. The same argument as before implies that both properties are satisfied. $\square$

References

L. Abrardi, L. Colombo, and P. Tedesch. The value of ignoring risk: Competition between better informed insurers. Politecnico di Torino andUniversit'a Cattolica del Sacro Cuore, 2020.

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

D. J. Benjamin. Errors in probabilistic reasoning and judgment biases. In B. D. Bernheim, S. DellaVigna, and D. Laibson, editors, *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, chapter 2, pages 69–186. Elsevier, 2019.

D. Bhaskar, A. McClellan, and E. Sadler. Regulation design in insurance markets. Queen Mary University, University of Chicago and Columbia University, 2021.

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

T. Börgers. *An Introduction to the Theory of Mechanism Design*. Oxford University Press, 2015.

Council of the European Union. Regulation 2016/679 of the European Parliament and of the Council: General Data Protection Regulation. eur-lex.europa.eu/eli/reg/2016/679/oj, *Official Journal of the European Union*, 2016.

R. Eilat, K. Eliaz, and X. Mu. Bayesian privacy. Ben Gurion University and Tel Aviv University and Columbia University, 2020.

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

H. Fang and Z. Wu. Consumer vulnerability and behavioral biases. NBER, 2020.

H. Fang, M. P. Keane, and D. Silverman. Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy*, 116(2):303–350, 2008.

Financial Times. Data brokers: regulators try to rein in the 'privacy deathstars'. www.ft.com/content/f1590694-fe68-11e8-aebf-99e208d3e521, January 7, 2019.

A. Finkelstein and K. McGarry. Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review*, 96(4):938–958, 2006.

J. R. Green and J.-J. Laffont. *An Introduction to the Theory of Mechanism Design*. North-Holland Publishing Company, 1979.

S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. *STOC 2008 Conference Proceedings*, 2008.

F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.

B. Keller, M. Eling, H. Schmeise, M. Christen, and M. Loi. Big data and insurance: Implications for innovation, competition and privacy. Technical report, International "Geneva" Association for the Study of Insurance Economics, www.genevaassociation.org/research-topics/cyber-and-innovation-digitalization/big-data-and-insurance-implications-innovation, 2018.

L. M. Khan. Amazon's antitrust paradox. *The Yale Law Jounral*, 126(3):710–805, 2017.

F. Koessler and V. Skreta. Selling with evidence. *Theoretical Economics*, 14(2):345–371, 2019.

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

V. F. Luz, P. Gottardi, and H. Moreira. Risk classification in insurance markets with risk and preference heterogeneity. University of British Columbia, University of Essex and FGV EPGE Brazilian School of Economics and Finance, 2020.

G. J. Mailath and A. Postlewaite. Asymmetric information bargaining problems with many agents. *Review of Economic Studies*, 57(3):351–367, 1990.

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. Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, 10(1):67–81, 1982.

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

T. Mylovanov and T. Tröger. Mechanism design by an informed principal: Private values with transferable utility. *Review of Economic Studies*, 81(4):1668–1707, 2014.

R. Rajan. Big technology companies are modern monopolies, but don't break them up. www.barrons.com/articles/big-tech-companies-modern-monopolies-google-facebook-51552076382, March 09, 2019.

E. Ramirez, J. Brill, M. K. Ohlhausen, J. D. Wright, and T. McSweeny. Data brokers: A call for transparency and accountability. Technical report, Federal Trade Commission, www.ftc.gov/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014, 2014.

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

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

B. Villeneuve. Competition between insurers with superior information. *European Economic Review*, 49(2):321–340, 2005.

X. Vives and Z. Ye. Information technology and bank competition. IESE Business School, 2021.