Inverse selection*

Markus Brunnermeier Rohit Lamba Carlos Segura-Rodriguez

April 2020

Preliminary. Comments welcome.

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 that forcing the insurance company to reveal its statistical information can be welfare improving. Third, 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. We also discuss the significance our analysis through three stylized facts: the rise of data brokers, the importance of consumer activism and regulatory forbearance, and merits of a public data repository.

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

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, with big data, insurance companies 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 we restrict attention to a monopolistic screening setup.

Inverse selection does not only differ from the standard *adverse selection* but also from the more recent *advantageous selection* literature. Advantages selection stresses the importance of preference heterogeneity in order to overturn the standard theoretical, but empirically counterfactual, result that the high-risk agents get full insurance whereas low-risk agents opt for partial insurance. With

[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.

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

We model the inverse selection problem using a two-dimensional type space. Both dimensions determine the riskiness of the agent, but the agent only knows one (type of) characteristic, the first dimension of the type, in addition to the marginals along both dimensions. 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 novel *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; interestingly, complete separation along both dimensions is never optimal for the insurer. When the insurer pools certain correlation types she is giving up on price discrimination in order to maintain the statistical information advantage. We show that the insurer always offers a finite number of contracts. This result is based on a novel mechanism design problem that features "*ironing* almost everywhere" (in the sense of Myerson [1981]). To be best of our knowledge, this is the first paper to model beliefs as private information in an otherwise classical mechanism design setup, and deal with ironing on the entire type space. The number of contracts turns out to be small, highlighting that the belief gap-versus-price discrimination trade-off is firmly resolved in the favor of the former. Indeed, in many settings the

[2] Einav and Finkelstein [2011] provide an excellent 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.

contract space along the statistical information is partitioned into only two contracts.[4]

To further 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 contract 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. 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. 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 change/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.

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

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 to the standard monopolistic Rothschild and Stiglitz [1976] insurance problem. The key difference is following: While the benchmark case features either full or partial insurance, in each of the four cases studied in the paper, the optimal contract features some over insurance. This mirrors the findings in advantageous selection literature, but here following a novel mechanism of differential endowment of initial information between the principal and the agent, as opposed to multidimensional private information of risk and preferences on the side of the agent.

Finally, we look at welfare implications under the various settings considered in the paper. The following comparison are salient: First, the insurer's profit is extremely high and the insuree's surplus is uniformly negative in the special gutgläubig case. Second, when the optimal contract

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

(in the standard model) features at most two partitions, the insurer's profit and insuree's surplus are comparable to the benchmark model in the following sense: the insurer tries to replicate the price discrimination as in the benchmark model but while maintaining the maximal amount of belief gap permissible by feasibility constraints. As a consequence, optimal profits in the standard model correspond to a linear approximation around the benchmark case and beliefs are split to keep high differentiation among contracts. And, third, when the insurer is forced to reveal all of its private information to the insuree, then her profits are by a significant amount the lowest and the insuree's surplus is the highest, suggesting government intervention can significantly help buyers.

While our model is admittedly stylistic, it provides a conceptual 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]). Finally, the increase in consumer welfare from forcing the principals to make private statistical information public points towards the merits of a public data repository; see, for example, Rajan [2019].

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. In addition, the modeling of information in our paper is driven by the stylistic fact that nature of information asymmetry between the principal and agent itself can be classified into statistical and physical, soft and hard, with both components correlated to each other. Finally, the four different settings considered here under the monopolistic setup, are unique to our paper. While the setup and results are quite different, we view both of these other papers as being complimentary to our work in a push towards the aforementioned "third generation" of insurance models.[5]

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

2 Model

As discussed in the introduction, the model we present can be thought of as an informed principal problem (Myerson [1983]) where a risk-neutral principal can commit to a menu and the private information she holds is statistical. The agent is risk averse and informed about some underlying characteristic that influences 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, $x = 1$ means exact, and $x > 1$ means over insurance.

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

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

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

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

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

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

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

Information. The canonical CARA-Gaussian version of the insurance model would assume that the mean loss, μ, is the agent's private information. We depart from this crucial assumption on the "endowment" of information as follows. A relevant bidimensional state $\theta = (\theta_1, \theta_2)$ determines μ, where $\theta_i \in \{L, H\}$ for $i \in \{1, 2\}$. So, given 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 $\sigma = \sqrt{q_1(1-q_1)}\sqrt{q_2(1-q_2)}$. Then, 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 simply knows the joint distribution of θ. In terms of the primitives, we assume that q_1 and q_2 are common knowledge, the agent is privately informed about θ_1, and the principal privately knows ρ. Finally, to close the model, we assume that ρ is drawn from F on $[\underline{\rho}, \overline{\rho}]$, where F is differential and has a continuous density f, and is common knowledge.[6]

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

[6]The entire set of possible correlation is of course [-1,1]. However, once we fix the marginals to be q_1 and 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, technically, the distribution F is restricted by the marginals q_1 and q_2.

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 CARA-Gaussian one we write down, but is much harder to solve, because of the lack of structure on the agent's payoff.[7]

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 hard 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 $r : [\underline{\rho}, \overline{\rho}] \to \Delta(M)$ 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:

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

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

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

[7] One weakness of our model is that by assuming a Gaussian setup, we have to allow the possibility of losses to be positive, which is absent in this alternate model. This is a standard problem of interpretation in many asset pricing models such as Grossman and Stiglitz [1980]. We believe tractability trumps this limitation for representing the ideas we have in mind.

Stage 1

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

Stage 2

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

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

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

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

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

where $\mu_m(\theta_1)$ is the expected value of μ based on realized value of θ_1 and ρ which is drawn drawn from the posterior F_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 type of constraint are imposed on the optimization problem. First is the incentive constraint of the insurer, that the insurer wants to truthfully report her type to the mediator:

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

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

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

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, C) that satisfies these three (class of) constraints is said to be *incentive-feasible*. Finally, the optimization problem can then be written simply as:

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

4 Three "special" cases

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

4.1 ρ is common knowledge

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

Proposition 1. *$\exists \, \rho^*$ 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). Also, note that profit is maximized at ρ^*, because the agent's private information of θ_1 becomes statistically irrelevant: the principal offers a pooling contract and extracts all the surplus associated with it

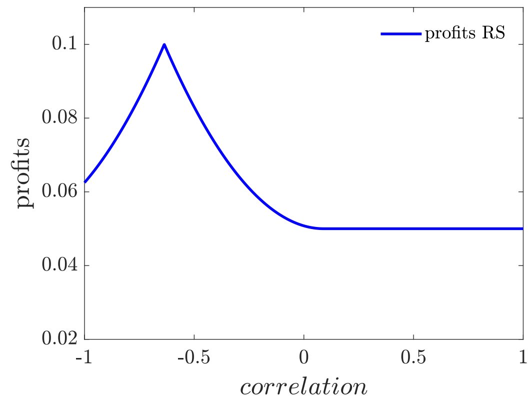

(a) optimal profit as a function of ρ

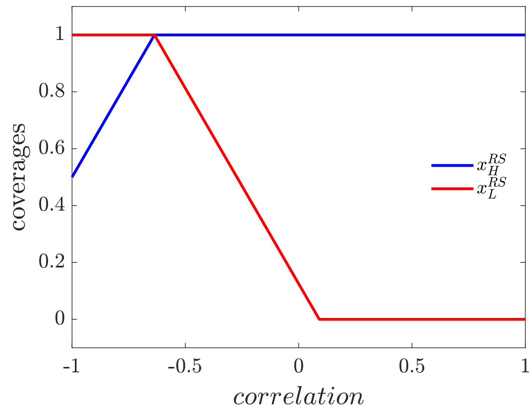

(b) optimal coverage as a function of ρ

Figure 1: Benchmark model when ρ is common knowledge

(see Figure 1a). We will refer to this as the *benchmark model*, and christen it RS, pointing towards the classical reference.[8]

The extent of distortion for the "low" risk type is further determined by the primitives of the problem. The economic force driving this result is typically known as the rent-versus-efficiency tradeoff. Since insuree is the residual claimant of the surplus, she wants to maximize efficiency 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 insurers to self select into the contract corresponding to his type. In fact if the proportion of "low" risk types is too low, the insurer will simply offer 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.

Definition 1. *A contract 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** $\hat{\rho}$.

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

[8] 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.

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 insurer about the way in which the two dimensions are correlated provides the insuree with great freedom in selecting contracts. She will choose r and C in tandem to create both the maximal belief gap and the maximal price discrimination.[9]

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

1. *binary messages are sent:* $M = \{\underline{m}, \overline{m}\}$ *s.t.* $r(\rho) = \overline{m}$ *for* $\rho < \tilde{\rho}$ *and* $m(\rho) = \underline{m}$ *for* $\rho > \tilde{\rho}$,
2. *posterior of the insuree is extreme:* $F_{\underline{m}} = \delta_{\underline{\rho}}$ *and* $F_{\overline{m}} = \delta_{\overline{\rho}}$,
3. *profits are uniformly higher than benchmark:* $\pi(\rho) > \pi^{RS}(\rho)\ \forall\ \rho$ *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, still 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, over insures $\theta_1 = L$ 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 it true for the case when the actual correlation is low: insuree reports large positive correlation, overinsures $\theta_1 = H$ and underinsures $\theta_1 = L$. In sum, the insurer sells a large amount of insurance to at 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.

[9] Since the Bayes' consistency condition is no longer valid, technically the class of contract is given by $C = (c_{m,\rho})$ because the contract offered for the actual realization of ρ has no relation to the reported value m.

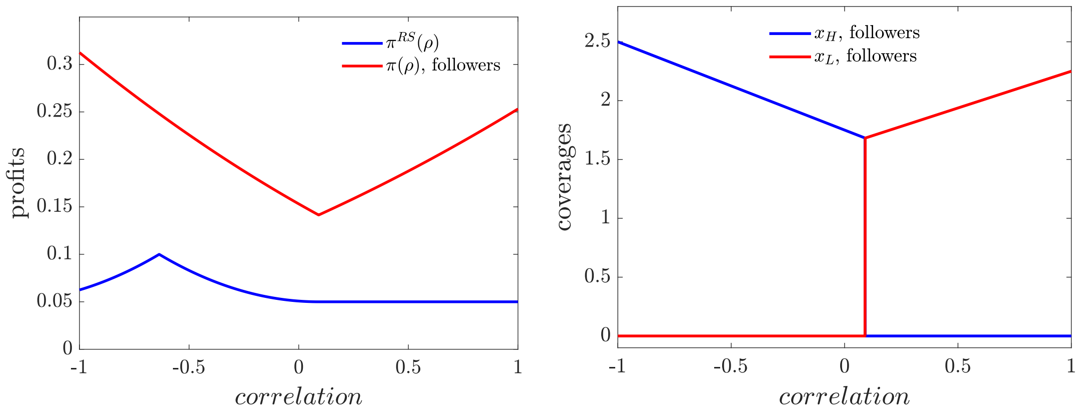

(a) optimal profit as a function of ρ

(b) optimal coverage as a function of ρ

Figure 2: Model with gutgläubig insuree

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). The last part of the proposition also points out that allocation are flipped in comparison to the benchmark model, signifying the role asymmetric information of correlation plays in the structure of contracts.

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 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.[10]

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 $F_m = F\ \forall m \in M$.[11] Thus, in this situation, the role of r is moot. The insuree designs the contract as a function of ρ with the knowledge that the insurer will evaluate his payoffs using the prior F.

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

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

[10] We are using the word "maximal" formally for belief gap but somewhat informally for price discrimination.

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

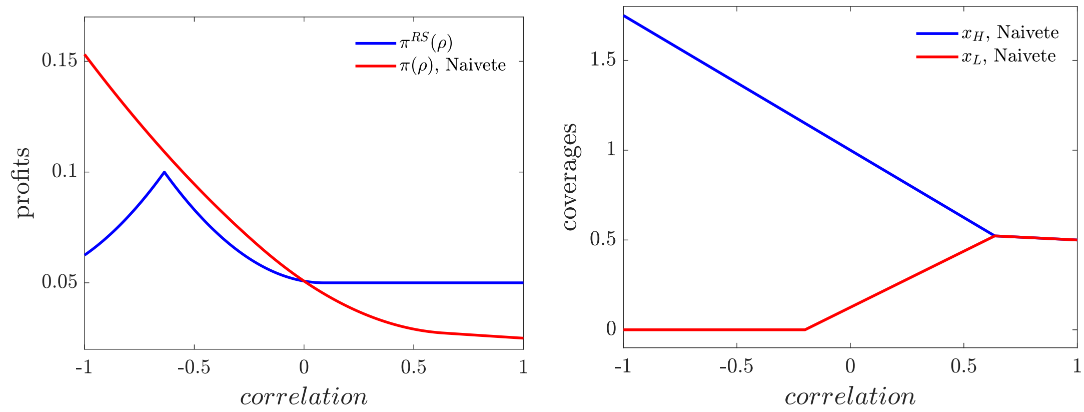

(a) optimal profit as a function of ρ (b) optimal coverage as a function of ρ

Figure 3: Model with naive insuree

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 correlation is close to $\underline{\rho}$, then the insurer wants to sell a lot of insurance to $\theta_1 = H$ and little insurance to the type $\theta_1 = L$, 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.[12]

4.4 Breakdown of the key forces

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

[12] 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.

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 insurer's belief, there is a no-longer a tradeoff, 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. Pice discrimination is then 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 and with the rent-versus-efficiency tradeoff to pin down the optimal contract.

5 Characterizing incentive compatibility and optimality

Our model differs from the standard screening problem in that it also features an incentive constraint for the principal, i.e. the insuree. In this section we analyze the incentive constraints of the insuree 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:*

$$\left.\frac{\partial \pi(\rho,\hat{\rho})}{\partial \rho}\right|_{\hat{\rho}=\rho} = \sigma x_{r(\rho)}(L)\cdot(\mu_{LH}-\mu_{LL}) - \sigma x_{r(\rho)}(H)\cdot(\mu_{HH}-\mu_{HL}) \equiv c(\rho) \tag{3}$$

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

Proof. Part two is standard property of value functions that satisfy incentive compatibility on a continuous type space (see for example Börgers [2015]). We show here the exact functional form of the envelope characterization stated in equation (3). 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:

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

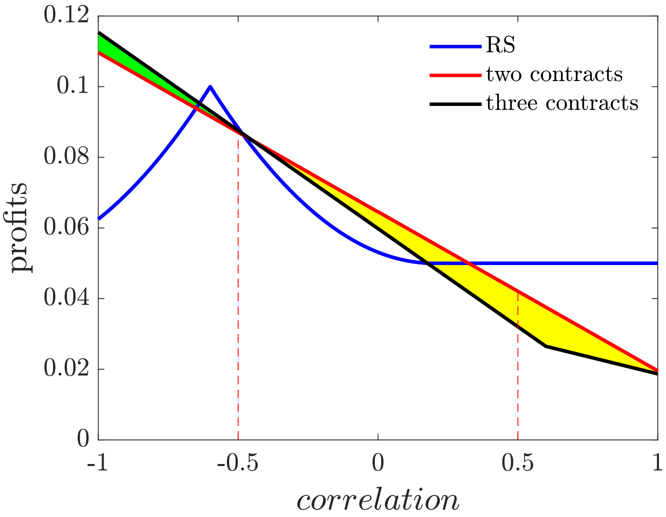

Figure 4: Costs and benefits of adding extra partitions.

and, substuiting $\hat{\rho} = \rho$ delivers equation (3). $\square$

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

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

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 choses for the low and high types corresponding to 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 can be feasibly chosen, specifically they limit the 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 C at the optimum.

All of the aforementioned contribute in reducing the richness of contracts considerably. In fact, we show that an optimal contract must be finite. Recollect that $C = \{c_m \mid m \in M\}$ where M is essentially a partition of $[\underline{\rho}, \overline{\rho}]$. For every additional element we introduce in M, there is a cost and benefit associated with it. Figure 4 presents precisely that for an example where the insurer moves from two to three partitions. Here the single-peaked blue curve is the profit associated with the benchmark model that we discussed in Section 4.1, the red straight line is the optimal profit with two partitions and the black step-function is the profit constructed for a three partition contract.

Take the straight red line representing the profit of optimal two partitions. Each partition has an 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, on each for 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. Alas, the incentive constraints of the insurer also put restrictions on the coverages that pushes the slope of the profit function (equation (3)) in the opposite direction. This culminates in a concave kink in the piecewise linear profit unction, which would violate the ignored convexity constraint no matter the primitives of the problem. Thus, the best highest convex profit function that the insurer can construct while satisfying incentive-feasibility is the one shown as the red line in Figure 4.

Finally, we increase the the number of partitions from two to three. 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 clearly, going from two to three partitions is sub-optimal.

It turns out that single-peakness of the benchmark profit function, along with the limitations that the insuree and insurer 's incentive compatibility constraints impose in the profits that can be reached in the subgame, make the costs of having an arbitrary number of partitions outweigh their benefits. Splitting the profit function while ensuring convexity and information rent for the agent is not very useful. The culmination of this reasoning is that at the optimum, the number of contracts offered is not just countable, it is also finite. Recollect $C = \{c^m \mid m \in \mathcal{M}\}$. The following definition and result document this point.

Definition 2. *A mechanism is f-partitional if C is a finite set.*

Proposition 4. *The optimal mechanism is f-partitional.*

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. Here we show further that at the optimum, under reasonable conditions, only two contracts are offered; that is, $|\mathcal{M}| = |C| \leqslant 2$. And, then we characterize the optimal one and two partition contracts.

Proposition 5. *Under (xx), the optimal contract features at most two partitions, i.e.* $|\mathcal{M}| = |C| \leqslant 2$.

At a conceptual level, Proposition 5 states in the duel between costs and benefits of more finely partitioning the already finite message space for reporting correlations, costs resoundingly win. In other words, the tradeoff between belief gap and price discrimination, maintaining belief gap strongly trumps greater price discrimination. At a technical level, the convexity constraint implied by incentive compatibility of the insurer and the monotonicity constraint on the allocation for fixed message implied by the incentive compatibility constraint leave such little wiggle room for optimum, that in culminates in the profit function being a straight line– this can only happen if the optimal contract has one or two partitions.

Ex ante, looking at the model and problem first defined in Sections 2 and 3 respectively, this result is rather striking. Put simply, it says that, Bayesian sophistication of the agent results in at most two contracts being offered at the optimum, instead of potentially uncountably many. Note that if only one contract is offered at the optimum, the insurer simply throws away her informational advantage by not using any price discrimination to maintain the maximal possible belief gap with the insuree.

Now, if the optimal number of 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 6. *When the optimal contract chooses* $|\mathcal{M}| = |C| = 1$*:*

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

Figure 5 plots the optimal profit and coverages for this case. The coverages are simply straight horizontal lines for the insuree is not using any of her private information about ρ and is instead offering a completely pooling contract along ρ. The profit function is a straight downward sloping line for the allocation are fixed, and π is linear in ρ. As is standard in the benchmark model, the (endogenously chosen) "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.[13]

[13]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 insuree, which actually separates our model form (most of) the literature on insurance markets.

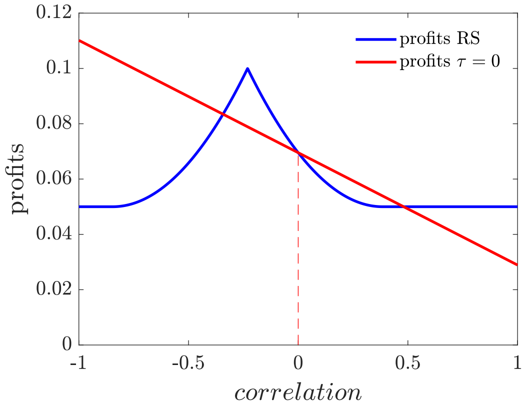

(a) optimal profit

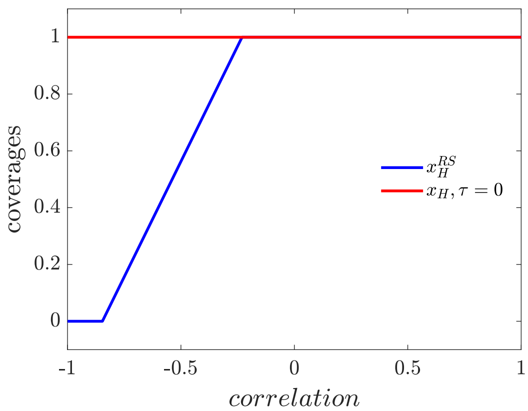

(b) optimal coverage for $\theta_1 = H$

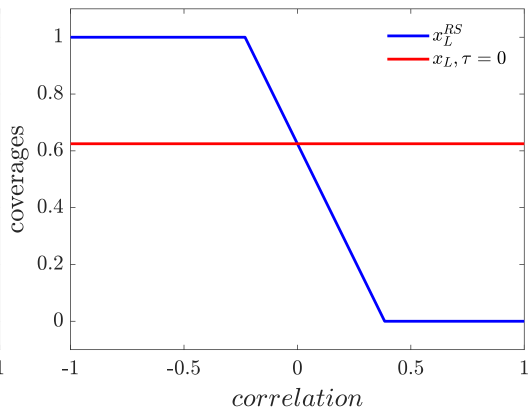

(c) optimal coverage for $\theta_1 = L$

Figure 5: Optimal contract features complete pooling

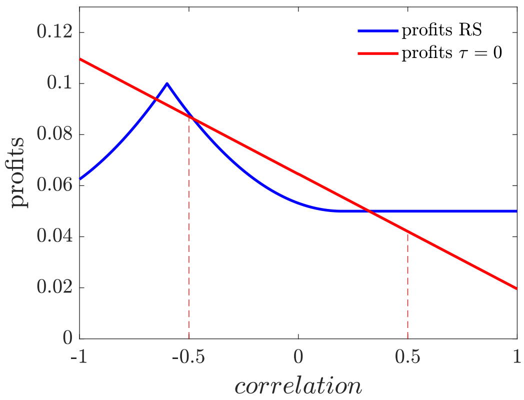

(a) optimal profit

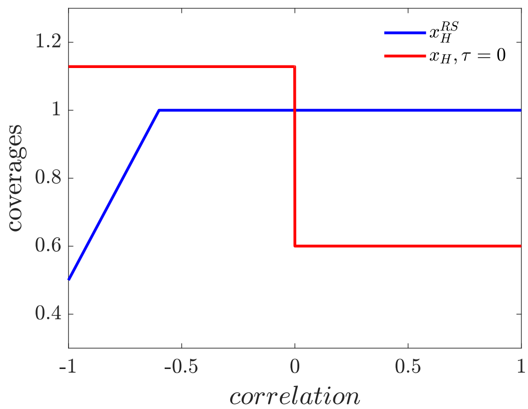

(b) optimal coverage for $\theta_1 = H$

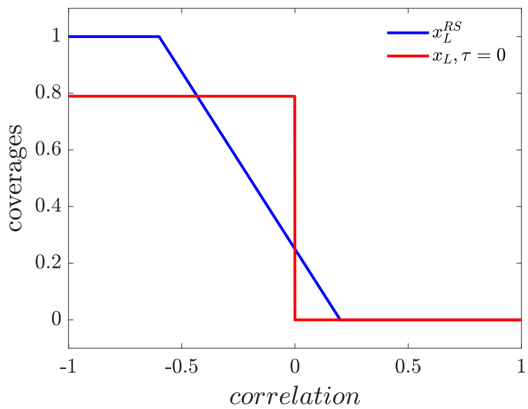

(c) optimal coverage for $\theta_1 = L$

Figure 6: Optimal contract features two partitions

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 7. *When the optimal contract chooses $|\mathcal{M}| = |C| = 2$, let $\mathbb{I}_1$ and $\mathbb{I}_2$ be the two intervals in the partition of M, $c_i = (x_i(H), x_i(L))$ be the two contracts offered in the and define $\rho^1 = \mathbb{E}[\rho \mid \rho \in \mathbb{I}_1]$, $\rho^2 = \mathbb{E}[\rho \mid \rho \in \mathbb{I}_2]$:*

1. *profits are linear in correlation: $\frac{d\pi(\rho)}{d\rho} = c$ f.s. c,*
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) $x_2(H) \leq x_{\rho^1}^{RS}(H)$ and $x_2(L) \geq x_{\rho^1}^{RS}(L)$, whenever $x_H^1 \neq x_L^1$.

Figure 6 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. In addition, the first result in Proposition 7 states that optimality forces both these slopes to be the same. The coverages variedly feature overinsurance for $\theta_1 = H$, partial insurance for both insuree types and no insurance for $\theta_1 = L$.

7 Optimal full revelation contract

So far we have analyzed the case where the insurer in sole proprietor of statistical information ρ and can decide, as part of the optimal contract, how much of it to reveal to the insurees. 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. In this case the model would become isomorphic to the one studied in Section 4.1, where ρ is common knowledge.

An alternate way is to assume that there are some shadow prices associated with publicly revealing the value of ρ. 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. We model this situation by exogenously fixing the message rule chosen by the mediator to be the identity mapping: $r(\rho) = \rho$, but then requiring the principal (the insuree) be compensated for this through her incentive constraint. This obviously has the downstream effect of influencing the contract C that are offered to the insurees. The entire optimization problem can be written in one piece as follows:

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

where $C^\star = \left\{c_\rho \mid \rho \in [\underline{\rho}, \overline{\rho}]\right\}, c_\rho = \left\{c_\rho(H), c_\rho(L)\right\}$ and $c_\rho(\theta_1) = \left(p_\rho(\theta_1), x_\rho(\theta_1)\right)$ 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 as opposed the expectation generated using the posterior by the message m is the earlier model. The optimal coverages will be denoted by $x_\rho(H)$ and $x_\rho(L)$. The following proposition summarizes the optimal full revelation contract.

Proposition 8. *When the insurer is forced to reveal ρ, that is, $r(\rho) = \rho$ is fixed exogenously:*

1. *profits are uniformly lower than the benchmark: $\pi(\rho) < \pi^{RS}(\rho)\ \forall \rho$,*
2. *generically inexact insurance: $x_i \neq 1$ for $i = H, L$.*
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 $\tilde{\rho}$.*

Figure 7 depicts the profit and coverages for the optimal full revelation contract. The profit of course lies uniformly below the benchmark. The dotted line shows the profit would also be uniformly higher if the insuree did not use her information at all and offered instead a pooling contract.

8 Welfare implications

The welfare implications of incorporating greater knowledge for the insurer in modeling insurance contracts is an important question. We take on these issues in at least three guises: How is consumer (or insuree) surplus impacted by equipping the insurer with some private statistical information? How much does it hurt the insuree to be unable to do Bayesian inference from contracts? What are the welfare consequence of forcing the insurer the publicly reveal the correlation? Since the profit performances of the insurer have already been documented, here we focus mostly on the insuree surplus under the various scenarios studied above.

Let $u_{m,\rho}(\theta_1)$ is the payoff of the insuree of type θ_1 when in the *gutgläubig* case where m is the correlation reported by the insurer and ρ is the actual realization correlation, and $u_\rho(\theta_!)$ be the

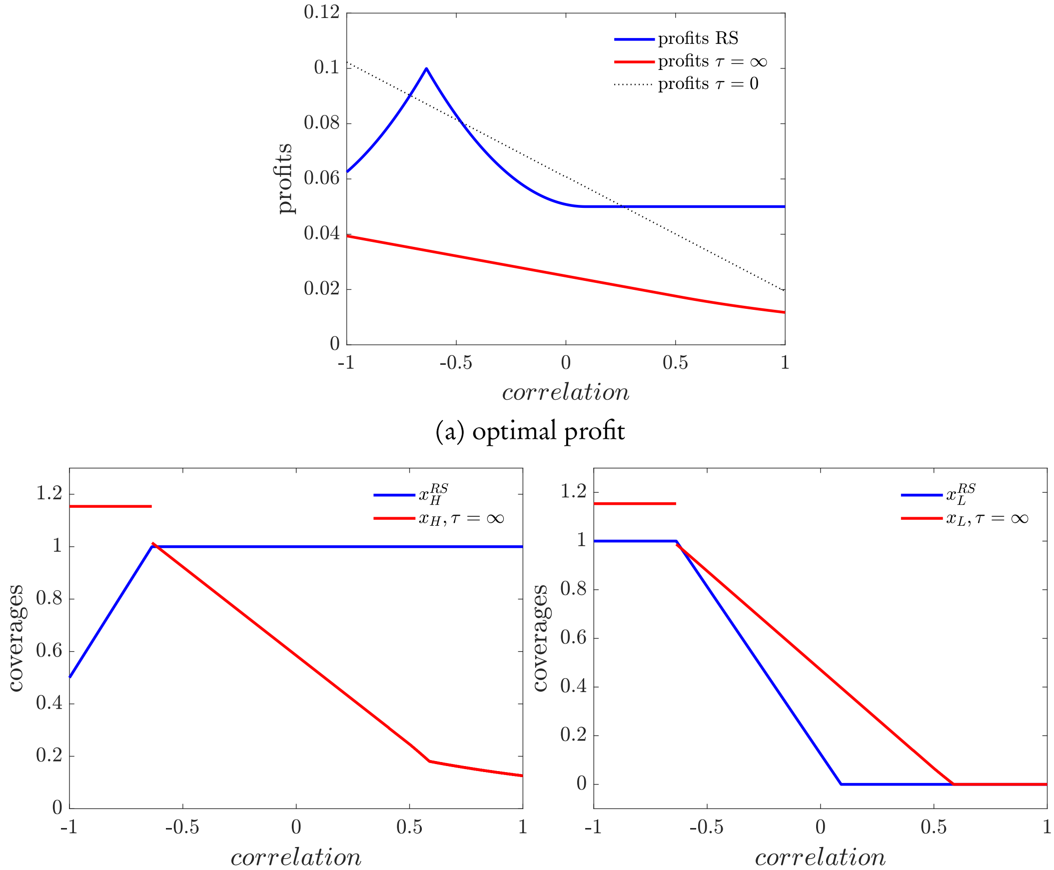

(a) optimal profit

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

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

Figure 7: Optimal contract with full information revelation

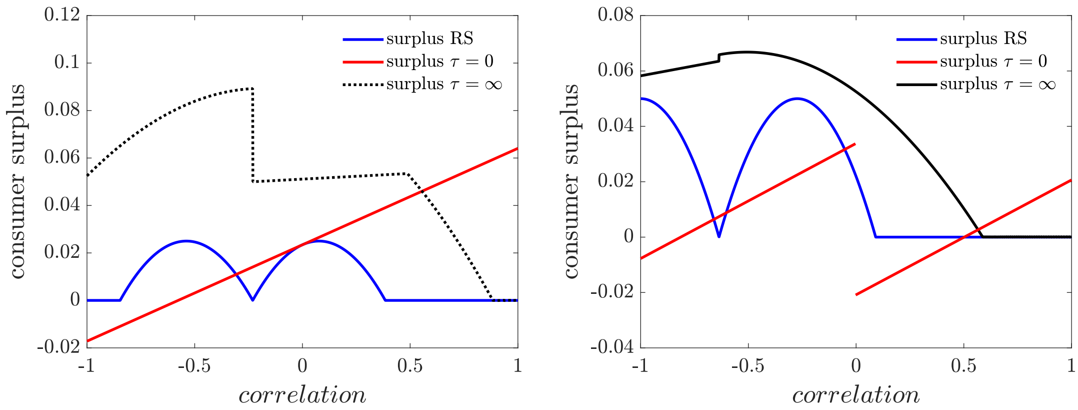

(a) When one partition is optimal in the standard case

(b) When two partitions are optimal in the standard case

Figure 8: Insuree surplus

payoff of the insuree of type θ_1 where ρ is the actual realized and reported correlation. It is easy to document the extreme cases.

Corollary 1. *The insuree's surplus (or utility) satisfies the following:*

1. *In the gutgläubig case it is uniformly negative, that is,* $\mathbb{E}\left[u_{m,\rho}(\theta_1)\right] < 0$ *for* $m = \underline{\rho}, \overline{\rho}$ *and* $\forall\ \rho$.

2. *For the full revelation contract, it is is uniformly positive:* $u_\rho(\theta_1) \geqslant 0$ *for all* ρ, θ_1.

So, in the case where the insuree creates the maximal belief gap and implements the maximal price discrimination, the insuree of course does badly in terms of payoff. Moreover, when the insuree can do Bayesian inference and the mediator or the government forces the insurer to reveal all private information about the correlation, never incurs a negative payoff.

The welfare consequences and ranking of the intermediate cases– the benchmark and standard model– depend more finely on the primitives of the model. At the outset is clear that in an ex ante sense, the expected payoff of the insuree is the highest in the model with full information revelation.

Corollary 2. *The insuree's ex ante surplus across in the full revelation model is strictly higher than that in the standard model and the benchmark models.*

This statement essentially means that the expected value of the black curve in Figure 8 is higher than the blue and red curves. Numerical results suggest that this ranking is not uniform in the ex post sense. There are some values of ρ for which the blue or red curves could lie above the black curve. At a high level, the revelation of information is good for the insuree in at least two ways: First. no matter the realization of ρ the insurer is guaranteed a non-negative ex post payoff. Second, in ex ante sense he always does better from this regulation than not. However, it is possible, owing to the incentive-feasibility constraints the it is not pointwise beneficial for the insuree to have the statistical information in the pubic domain.

9 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*.[14] One of the key issues 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 his 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 influences the probability of loss to be two dimensional, 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 agent too is privately informed.

The primary findings are as follows. Private statistical information on the side of the insurer introduces a novel belief gap versus price discrimination, in addition to the usual rent versus efficiency in standard screening contracts. The insurer wants to price discriminate along her private information dimension but is also wary that fine-tuning the contract too much to the details of the environment which will allow the insuree to infer her private 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. And, further in the case where the insurer is forced to reveal all the information, a larger variety of contracts are offered and the insuree (or consumer) surplus is higher, pointing to welfare improvements from making data public.

The approach invoked in the paper can potentially be pursed in various directions to better understand richer theoretical constructs in modeling insurance markets. We study insurance under the monopolistic set up. While there is considerable evidence that insurance companies do enjoy market power, it is not absolute. So, it is definitely useful to also think about the competitive version, and moreover, it is an important benchmark of the classical model (Rothschild and Stiglitz [1976]). Since our model has correlated types, it would be a non-trivial exercise to determine the off-path beliefs of both the insurers and insuree in an equilibrium setting.

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

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

statistical information about the credit worthiness of a client, in addition to the client knowing some hard information about his financial circumstances. In the capital markets, the venture capitalist may know more statistically about the probability success of various projects. Finally, aggregating across multiple principal-agent interactions. a greater information on the side of the principal may encourage more market concentration, of the kind we see in tech industry these days.

10 Appendix

Proof of Proposition 1. Let ρ^* be the correlation for which $\mu(H,\rho^*) = \mu(L,\rho^*)$. When $\rho = \rho^*$, the seller does not need to provide an information rent to any of the types and can maximize efficiency. Therefore, at this poin profits are maximized.

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

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

From the first and third conditions it is easy to conclude that $\lambda = (1 - q_1)$ and $\delta = 1$. Using these values it is straightforward to see that $x(H,\rho) = 1$ and $x(L,\rho) = 1 - \frac{1-q_1}{\eta q_1}(\mu(H,\rho) - \mu(L,\rho)) < 1$. In a corner solution it is trivial to show that $x(H,\rho) = 1$ and $x(L,\rho) = 0$.

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

Proof of Proposition 2. Let $\hat{\mu}_L$ and $\hat{\mu}_H$ the beliefs that are generated by using the seller's message. Then for each correlation ρ the problem of the seller is to choose those beliefs, coverage and prices that solve

$$\begin{aligned}
\max_{p_L,p_H,x_L,x_H,\hat{\mu}_L,\hat{\mu}_H} \quad & q_1(p_L - \mu(L,\rho)x_L) + (1 - q_1)(p_H - \mu(H,\rho)x_H)\\
s.t. \quad & \hat{\mu}_\theta x_\theta - \tfrac{\eta}{2}(1 - x_\theta)^2 - p_\theta \geq \hat{\mu}_\theta x_{\theta'} - \tfrac{\eta}{2}(1 - x_{\theta'})^2 - p_{\theta'} \;\; \forall \theta_1, \theta_1' \in \{L, H\} \quad \text{IC } \theta_1 - \theta_1'\\
& \hat{\mu}_\theta x_\theta - \tfrac{\eta}{2}(1 - x_\theta)^2 - p_\theta \geq -\tfrac{\eta}{2} \quad \forall \theta_1 \in \{L, H\} \quad \text{IR } \theta_1
\end{aligned}$$

First, notice that the problem is bang-bang in terms of $\hat{\mu}_\theta$. Then the solutions have to be on the extreme of the feasible set. Since the extreme beliefs are reached when the buyer believes that the correlation is either $\underline{\rho}$ or $\bar{\rho}$, the seller sends only send two signals $\underline{m}$ and $\bar{m}$ that generate buyer's posteriors $F_{\underline{m}} = \delta_{\underline{\rho}}$ and $F_{\bar{m}} = \delta_{\bar{\rho}}$.

Suppose that the seller sends the message $\bar{m}$. This message generates posterior beliefs $\bar{\mu}_H > \bar{\mu}_L$. Furthermore, for any $\rho < \bar{\rho}$, $\bar{\mu}_H > \mu(H,\rho)$ and $\bar{\mu}_L < \mu(L,\rho)$. Using the first order approach, it is straightforward to show that in an interior solution it has to be that $x_H = 1 + \frac{\bar{\mu}_H - \mu(H,\rho)}{\eta} > 1$, and $x_L = 1 - \frac{(1-q_1)\bar{\mu}_H + q_1\mu(L,\rho) - \bar{\mu}_L}{q_1\eta} < 1$, and 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 $x_L > 1$.

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

$$\frac{\partial\bar{\pi}(\rho)}{\partial\rho} - \frac{\partial\underline{\pi}(\rho)}{\partial\rho} = \frac{-\bar{\mu}_L + (1-q_1)\bar{\mu}_H + q_1\underline{\mu}_L}{\eta}\frac{\partial\mu(L,\rho)}{\partial\rho} + \frac{\underline{\mu}_H - (1-q_1)\bar{\mu}_H - q_1\underline{\mu}_L}{\eta}\frac{\partial\mu(H,\rho)}{\partial\rho} < 0,$$

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

Then if for a correlation ρ the seller send message $\underline{m}$ ($\bar{m}$), then for all $\rho' > \rho$ ($\rho' < \rho$) the seller sends the same message. This and our characterization above shows that there is a $\tilde{\rho} \in [\underline{\rho}, \bar{\rho}]$ and the contract flips around $\tilde{\rho}$. Finally, notice that when $\tilde{\rho} \in (\underline{\rho}, \bar{\rho})$ the argument above shows that the seller always offer contracts that over insure or under insure the insuree. □

Proof of Proposition 3. Denote by μ_H^e and μ_L^e the expected probabilities the buyers face a loss when he stick to his prior belief $\mathbb{E}$. Suppose that $\mu_H^e > \mu_L^e$

In an interior solution we obtain that $x_H(\rho) = 1 + \frac{\mu_H^e - \mu(H,\rho)}{\eta}$ and $x_L(\rho) = 1 - \frac{(1-q_1)\mu_H^e + q_1\mu(L,\rho) - \mu_L^e}{q_1\eta}$. Notice that $x_H(\rho)$ is decreasing in ρ and $x_L(\rho)$ is increasing in ρ. Then there are two corner solutions: one in which $x(L,\rho) = 0$ and $x(H,\rho) = 1 + \frac{\mu_H^e - \mu(H,\rho)}{\eta}$, and another one in which $x_L = x_H = 1 + \mu_L^e - \frac{q_1\mu(L,\rho) + (1-q_1)\mu(H,\rho)}{\eta}$.

Notice that coverage are flipped around $\mathbb{E}$. At correlation $\mathbb{E}(\rho)$ the contract is as in *RS*. Further, the slope of the coverage has the opposite sign than those of RS.

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

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

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

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

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

Proof of Proposition 4. Proposition ?? implies that the profit function's slope can take at most a countable number of different values. Suppose by contradiction that the profit function's slope takes an infinite number of different values. Then there exists two different messages m_1 and m_2 such that both preimages $r^{-1}(m_1)$ and $r^{-1}(m_2)$ are measure zero, they are contiguous, the contracts c^{m_1} and c^{m_2} generate different profits slopes $c_1 < c_2$, and there are not other contracts that are offered the generate the same slopes. Then the expected correlations $\rho_1 = E[\rho \mid r(\rho) = m_1] < \rho_2 = E[\rho \mid r(\rho) = m_1]$ are such that for any $\epsilon > 0$, $\rho_2 - \rho_1 < \epsilon$. Without loss assume that $\rho_1 > \rho^*$.

Suppose first that $c_1 > c(\rho_1)$ as defined in Lemma 2. The same lemma implies that $c_2 > c(\rho_2)$. Then Lemma 2 and Lemma 3 imply that $\hat{\pi}(\rho_2, c_1) - \hat{\pi}(\rho_2, c_2) = \delta > 0$. Furthermore, it is without loss to assume that $\pi(\rho_1) = \hat{\pi}(\rho_1, c_1)$ and $\pi(\rho_2) = \hat{\pi}(\rho_2, c_2)$; if not, the seller could slightly increase c_1 or c_2, respectively. Since the partitions containing either ρ_1 or ρ_2 have measure 0 this change does not affect what happens with the rest of the profit function.

Since the optimal profit function is convex imply, it is continuous. Therefore, $\hat{\pi}(\rho_1, c_1) + c_1(\dot{\rho} - \rho_1) = \hat{\pi}(\rho_2, c_2) + c_2(\rho_2 - \dot{\rho})$, where $\dot{\rho}$ is the correlation that defines the limit between the partitions that generate the expected correlations ρ_1 and ρ_2. By continuity of $\hat{\pi}(\rho, c)$ with respect to ρ there exists ϵ such that if $\rho_1 - \rho_2 < \epsilon$ then $| \hat{\pi}(\rho_1, c_1) - \hat{\pi}(\rho_2, c_1) | < \frac{\delta}{4}$. Take ϵ small enough such that if $c_1 < 0$ then $\epsilon < \frac{-\delta}{4c_1}$ and if $c_2 > 0$ then $\epsilon < \frac{\delta}{4c_2}$. Then we have that

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

a contradiction.

An analogous argument leads to a contradiction in the case in which $c_2 < c(\rho_2)$.

The case remaining to show is the one in which $c_1 < c(\rho_1)$ and $c_2 > c(\rho_2)$. In this case we can take $c \in (c(\rho_2), c(\rho_1))$ as the slope of the profit function in both partitions. Lemma 3 implies that with this profit the bound on the seller's profits increase, so that the firm's IC constraint is relaxed. Since the partitions containing either ρ_1 or ρ_2 have measure 0 the two partitions can be merged without loss, and these changes do not affect what happens with the rest of the profit function. Repeating this process we conclude that the seller's profits function can take only a finite number of different slopes.

In the optimal contract there can be at most two elements that generate a profit's slope of c. Lemma 2 implies that $\hat{\pi}(\cdot, c)$ is convex. Therefore, three different messages that generate the same slope is always dominated by a strategy in which there are only two different messages with the same slope. Therefore, it is without loss to consider partitions with a finite number of elements. $\square$

Proof of Proposition 6. 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 only one contract is going to be offered, 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, in this case the optimal contract coincides with the Rothschild-Stiglitz contract when $\rho = \mathbb{E}(\rho)$. Then it has to be that $\mathbb{E}(\pi(\rho)) = \pi^{RS}(\mathbb{E}(\rho))$, $x_H(\rho) = x_H^{RS}(\mathbb{E}(\rho))$, and $x_L(\rho) = x_L^{RS}(\mathbb{E}(\rho))$. $\square$

Proof of Proposition 7. We first argue that the profit function slope has to be constant. For this end we define the function $\hat{\pi}(\rho, c)$ as the maximum that the firm can obtain in the subgame in which both parties belief that the correlation is and the relation between the offered coverages in equation 3 is equal to c. This function has a nice behavior: it is single peak with respect to ρ, with a peak at ρ^* and it is convex with respect to ρ both to the right and the left of ρ^*; and it is strictly concave with respect to c.

Lemma 2. *Let ρ^* be the correlation such that $\mu(H, \rho^*) = \mu(L, \rho^*)$ and fix $\rho \geq \rho^*$. Then:*

1. *there exists correlations ρ_1 and ρ_2 with $\bar{\rho} \geq \rho_2 \geq \rho_1 \geq \rho^*$ and slopes $c_1 < c_2 < 0$ such that*
   (a) *for $\rho \in [\rho^*, \rho_1]$, $\hat{\pi}(\cdot, c)$ is linear and strictly decreasing with respect to ρ;*
   (b) *for $\rho \in [\rho_1, \rho_2]$, $\hat{\pi}(\cdot, c)$ is strictly convex and strictly decreasing with respect to ρ;*
   (c) *for $\rho > \rho_2$, $\hat{\pi}(\cdot, c)$ is constant with respect to ρ.*
2. *The function $\hat{\pi}(\rho, \cdot)$ is strictly concave with respect to c.*

An analogous characterization holds for $\rho < \rho^*$.

Proof. Suppose that $\rho \geq \rho^*$, that is, $\mu(H, \rho) \geq \mu(L, \rho)$. To simplify notation we let $\mu_L = \mu(L, \rho)$, $\mu_H = \mu(H, \rho)$, $x_L = x(L, \rho)$, $x_H = x(H, \rho)$, $p_L = p(L, \rho)$ and $p_H = p(H, \rho)$. Let $K_L = \sigma(\mu_{LH} - \mu_{LL})$ and $K_H = -\sigma(\mu_{HH} - \mu_{HL})$.

We prove 1. first. 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 $\beta = \frac{\eta q_1(1-q_1)(c-K_L-K_H)+(1-q_1)^2(\mu_H-\mu_L)K_L}{(1-q_1)K_L^2+q_1K_H^2}$, and up to β the profits in this solution are equal to

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

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

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

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

$$\frac{(1-q_1)^2 kH^2}{\eta((1-q_1)K_L^2+q_1K_H^2)}\left(\frac{\partial(\mu_H-\mu_L)}{\partial\rho}\right)^2>0.$$

Therefore, $\pi(\rho,c)$ is strictly decreasing and strictly convex with respect to ρ when the solution is interior.

A corner solution in which the buyer sells only to type H 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 μ 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 profits are constant with respect to ρ.

Finally, to satisfy both buyer's IC constraints it has to be that $x_H\geq x_L$, but this might no be the case in the interior solution we characterized above. In particular for $\rho<\rho_2$ the constraint $x_H\geq x_L$ 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 holds with equality is smaller than ρ_2, and the correlation that makes the second inequality to holds 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 seller 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 $x_L \geq 0$ 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 seller sells only to him.

To prove 2. we only need to take the second derivative of each of the possible profit functions with respect to c. 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. □

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 $[\underline{\rho}, \tilde{\rho}]$.

The problem that the seller wants to solve is

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

where ρ_i is the expected correlation after observing the respective message, and the first constraints are feasibility constraints; the profits that the firm is going to obtain at the expected value are possible to obtain in the subgame.

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. Lemma 3 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, by 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 increasing π_1 and/or π_2, contradicting that the initial contract is optimal.

In the second case increasing c_2 reduces the RHS of continuity constraint and by Lemma 2 increases the LHS of 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 $x_2(L) \geq x^{RS}_{\rho^1}(L)$. □

Lemma 3. *In the solution to problem $P-\rho$, if $\mu(H,\rho) > \mu(L,\rho)$, $x(H,\rho) = 1 > x(L,\rho)$ and $\frac{\partial x_L(\rho)}{\partial \rho}\big|_{x_L(\rho)>0} < 0$; and if $\mu(H,\rho) < \mu(L,\rho)$, $x(L,\rho) = 1 > x(H,\rho)$ and $\frac{\partial x_H(\rho)}{\partial \rho}\big|_{x_H(\rho)>0} > 0$. Therefore, $c(\rho)$ is strictly decreasing.*

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

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

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

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

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

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

Proof of Proposition 8. 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 those correlations, and they are guaranteed to exist by Lemma 2. Lemma 3 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 $\bar{c} < c^{RS}(\bar{\rho}) < c^{RS}(\underline{\rho})$. In the first case, by decreasing the slope of the profit function uniformly by Lemma 2 the profits that are feasible for each correlation increase uniformly. By increasing the optimal profit function by a constant, the seller's IC constraints are satisfied. Then the original contract was not optimal. Similarly, in the second case the objective function can be increased by increasing uniformly the slopes of the profit function.

Since convexity implies that $c(\rho)$ is weakly increasing and $c^{RS}(\rho)$ is decreasing by Lemma 3 they can coincide at most a one correlation, call it $\tilde{\rho}$.[15] 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 incresing in c and x_H is decreasing in c. Then when contracts are separating $x_H(\rho) \geq x_H^{RS}(\rho)$, and $x_L(\rho) \leq x_L^{RS}(\rho)$ for $\rho < \tilde{\rho}$, and $x_H(\rho) \leq x_H^{RS}(\rho)$, and $x_L(\rho) \geq x_L^{RS}(\rho)$ for $\rho > \tilde{\rho}$, with a strict inequality in each case. □

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.

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, 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.

[15] 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.

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.

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.

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.

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. 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.