Market Power and Insurance Coverage*

Daniel Gottlieb†and Humberto Moreira‡

This version: February 29, 2024

Abstract

This paper examines how market power affects coverage in markets with adverse selection. We show that market power decreases coverage for individuals who are less willing to pay for insurance but increases coverage for those with a higher willingness to pay. Under weak conditions, a monopolist always excludes a positive mass of customers, whereas competitive firms do not. However, to avoid cream skimming, competitive firms provide less coverage than a monopolist for consumers who are willing to pay more. The welfare comparison between competitive and monopolistic markets depends on whether the distortion at the bottom (higher under monopoly) exceeds the distortion at the top (higher under competition). Using simulations based on an empirical model of preferences and costs, we find that both effects are quantitatively important although the effect at the bottom dominates. Therefore, the market power distortion exceeds the cream skimming distortion from competition in our calibrated model.

*We thank Mark Armstrong, Eduardo Azevedo, Piero Gottardi, Jon Kolstad, Camille Landais, Stephen Ryan, and Johannes Spinnewijn for helpful comments. Diogo Wolff Surdi provided excellent research assistance. Gottlieb gratefully acknowledges the financial support of the Engineering and Physical Sciences Research Council Grant EP/X021912/1. Moreira acknowledges the financial support of FAPERJ and CNPq. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

†London School of Economics, d.gottlieb@lse.ac.uk.

‡FGV EPGE, humberto.moreira@fgv.br.

1 Introduction

Access to insurance coverage is crucial for both efficiency and equity reasons. For instance, in the past century, all high-income countries implemented measures to ensure access to health insurance. While most of them opted for public provision, the United States chose to keep health insurance privately provided. Several policies were introduced to increase insurance coverage, including tax benefits for employer-provided plans, free health insurance for the poor, elderly, and disabled, risk adjustments, penalties for lacking coverage through individual and employer mandates, and setting up health insurance exchanges. While these policies have considerably reduced the number of uninsured individuals, over 30 million Americans remain without coverage. Understanding why so many people do not have health insurance is a key policy issue.

In this paper, we show that a general prediction of market power in insurance is that a positive fraction of the population does not purchase any coverage. In contrast, under mild assumptions, all consumers buy some coverage in competitive markets. These theoretical predictions suggest that market power may be an important feature of health insurance markets in the United States. This is consistent with a large body of empirical work, which finds that health insurance markets are highly concentrated and insurers have substantial market power.[1]

A monopolist faces a trade-off between efficiency and rent extraction. Reducing coverage is costly because risk-averse customers would pay more than the actuarially fair price to increase their coverage. However, this reduction in coverage allows the firm to charge higher premiums to those willing to pay more. We show that the rent-extraction effect always dominates for customers at the lower end of the willingness-to-pay spectrum. Thus, a monopolistic insurer profits from excluding a positive mass of customers. In contrast, in a competitive market with endogenous contracts, there is no exclusion of risk-averse consumers. Therefore, consumers with lower willingness to pay buy less coverage under monopoly than in a competitive market, with some not purchasing coverage at all.

However, the effect of market power on coverage is heterogeneous across consumers. Those with a high willingness to pay obtain higher coverage under monopoly than in competitive markets (see Figure 1). This is because competitive firms have an incentive to cream skim, stealing safer consumers from other firms by offering less coverage. A monopolist does not face the risk of losing its safest customers to other firms, so it can provide more coverage to those at the higher end of the willingness-to-pay spectrum.

We illustrate the quantitative importance of these effects in a calibrated health insurance model based on Einav et al. (2013). As depicted in Figure 2, there is substantial exclusion

[1] See Dafny (2010); Dafny et al. (2012); Starc (2014); Ho and Lee (2017); Cabral et al. (2018); Cicala et al. (2019); Polyakova and Ryan (2021; 2023); Saltzman et al. (2021); and Tebaldi (Forthcoming).

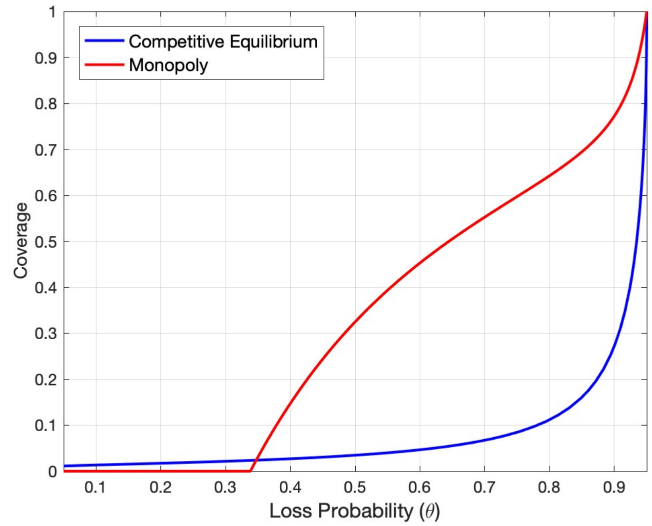

Figure 1: Coverage under Monopoly and Perfect Competition (Rothschild & Stiglitz)

Notes: Coverage in example 2. Consumers have CARA utility with risk aversion $A = \frac{1}{2}$ and face a loss of $L = 1$. The horizontal axis depicts each consumer's loss probability, which is uniformly distributed between 5 and 95 percent. The blue line depicts each consumer's coverage with perfect competition. The red line depicts coverage with a monopoly. A monopolist provides less coverage for types with a low loss probability and excludes those with low enough loss probabilities. Competitive firms provide more coverage for those types and do not exclude any types. However, consumers with high loss probabilities receive less coverage under perfect competition than with a monopolist.

with monopoly, with 70 percent of consumers not purchasing any coverage. In contrast, with perfect competition, all consumers buy coverage. At the top, 12 percent of consumers purchase a higher coverage under monopoly than under perfect competition. In our simulations, both consumer and total surplus are higher under perfect competition than under monopoly. Specifically, the average annual consumer surplus equals \$3,047 in the competitive equilibrium and \$1,308 under monopoly. The average annual total surplus is again \$3,047 in the competitive equilibrium (insurers make zero profits) and \$2,223 with a monopolist.

While our simulation considers a model of health insurance, our theoretical results apply to many other markets with adverse selection, including other insurance and credit markets. See Einav et al. (2021) for a survey of the literature.

Our paper is structured as follows. In Section 2, we introduce the general model, examples, and main definitions. Section 3 presents the results for one-dimensional types. We then generalize the results to multi-dimensional types in Section 4. Section 5 illustrates the quantitative implications of our results using a calibrated health insurance model. We review the related literature in Section 6. Then, Section 7 concludes. Intuitive proofs are presented in the text, whereas more technical proofs are in the appendix.

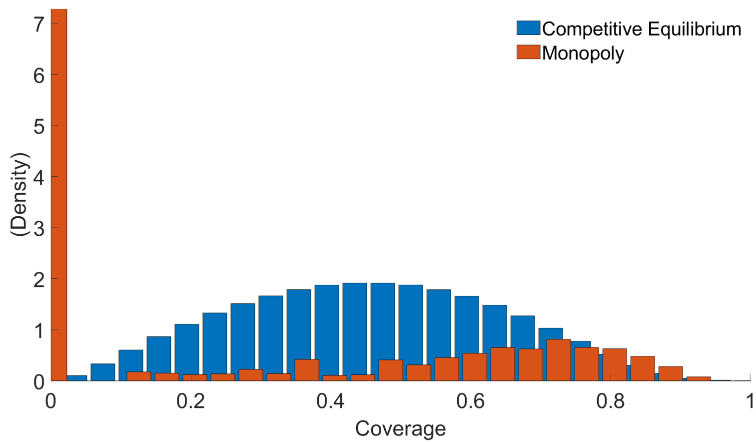

Figure 2: Coverage under Monopoly and Perfect Competition (Calibrated Model)
Notes: Distribution of coverage choices in the numerical example from Section 5. The horizontal axis depicts the contracts chosen by consumers, with coverage ranging between 0 to 100 percent of expenses. Blue bars represent the distribution of coverage in the competitive equilibrium. Orange bars represent coverage with a monopolist. With monopoly pricing, approximately 70 percent of consumers are uninsured. With perfect competition, all consumers purchase some coverage. However, more consumers purchase policies with coverage levels above 70 percent with monopoly than with perfect competition.

2 Model

2.1 The Model

Following a large applied literature, we assume that consumers have quasi-linear utility. This assumption is consistent with insurance models in which consumers have constant absolute risk aversion.[2] Consumer private information is represented by a K-dimensional type θ drawn from $\Theta \equiv [\underline{\theta}_1, \overline{\theta}_1] \times ... \times [\underline{\theta}_K, \overline{\theta}_K] \subset \mathbb{R}_+^K$. Types are distributed according to an absolutely continuous measure μ on Θ with a continuous probability density function f with full support.

A type-θ consumer's utility from buying a policy with coverage level $x \in [0, 1]$ and premium $p \in \mathbb{R}$ is:

$$u(\theta, x) - p. \tag{1}$$

The firm's expected profit from selling a policy with coverage x to a type-θ consumer is:

$$p - c(\theta, x). \tag{2}$$

[2]Constant absolute risk aversion is a common assumption in the applied literature (see, for example, Cardon and Hendel (2001), Handel (2013), Einav et al. (2013), Handel et al. (2015), and Hackmann et al. (2015)). All of our results hold with arbitrary utility in the standard setting of Rothschild and Stiglitz (1976).

The surplus from providing coverage x to type θ is:

$$S(\theta, x) \equiv u(\theta, x) - c(\theta, x).$$

The willingness to pay and the cost of providing zero coverage are both zero: $u(\theta, 0) = c(\theta, 0) = 0$ for all θ.

We maintain the following assumptions throughout the paper:[3]

Assumption 1. *The utility function $u : \mathbb{R}_+^K \times [0,1] \to \mathbb{R}_+$ is twice continuously differentiable and is strictly increasing in θ for each $x > 0$. The cost function $c : \mathbb{R}_+^K \times [0,1] \to \mathbb{R}_+$ is continuously differentiable.*

This framework is general enough to allow for multidimensional heterogeneity, as formulated in many empirical models. Ex-post moral hazard can be incorporated through the definitions of the utility and cost functions. By specifying the utility function appropriately, it also allows for other types of consumer behavior, such as overconfidence, inertia to abandon a default choice, or misunderstanding the benefits from being insured.[4]

2.2 Examples

The following examples clarify how our general framework can be applied to different settings. The first example is the model we use in our simulations (Figure 2 and Section 5):

Example 1. (Einav et al. (2013); Azevedo and Gottlieb (2017)) Consumers face a stochastic health shock l, which is normally distributed with mean M and variance S^2.[5] After the shock, they decide how much to spend on health services e. Consumers are heterogeneous in their distribution of health shocks (parametrized by M and S^2), risk aversion parameter A, moral hazard parameter H, and initial wealth W. Utility after the shock equals

$$CE(e, l; x, p, H) = [W - p - (1-x)e] + [(e-l) - \frac{1}{2H}(e-l)^2].$$

Substituting the privately optimal health expenditure, $e = l + Hx$, we can write the utility after the shock as

$$CE^*(l; x, p, H) = W - p - l + l \cdot x + \frac{H}{2} \cdot x^2.$$

[3]Note that u and c are also defined for types outside of the type space Θ. This allows us to obtain conditions on the type space Θ, while ensuring that preferences and costs are well defined.

[4]See Handel (2013) and Polyakova (2016) for inertia and Handel and Kolstad (2015) and Handel et al. (2019) for other frictions.

[5]Appendix A generalizes the model for arbitrary distributions and presents simulations for the truncated normal.

Consumers have constant absolute risk aversion (CARA), so their expected ex-ante utility equals

$$\mathbb{E}[-e^{-A\cdot CE^*(l;x,p,H)}|l \sim N(M,S^2)].$$

Using the expression for the normal distribution, the model can be described as in (1) and (2) with

$$\begin{aligned} u(\theta,x) &= xM + x(2-x)\frac{S^2A}{2} + x^2\frac{H}{2}, \text{ and} \\ c(\theta,x) &= xM + x^2H. \end{aligned} \tag{3}$$

The consumer's willingness to pay for coverage depends on three terms: average covered expenses xM, utility from risk-sharing $x(2-x)\cdot S^2A/2$, and utility from overconsumption of health services $x^2H/2$. Since the firm has to pay the covered expenses, the first term is subtracted from the firm's profits. Over-consuming health services (moral hazard) costs firms twice as much as consumers are willing to pay for it. Risk neutral firms have no cost of absorbing risk. The value of risk-sharing is increasing in coverage, in the consumer's risk aversion, and in the variance of health shocks.

Farinha Luz et al. (2023) consider a special case of example 1 without moral hazard ($H = 0$ for all consumers). We now turn to simpler, one-dimensional models.

Example 2. (Rothschild and Stiglitz, 1976; Chade and Schlee, 2021) Consumers have initial wealth W and face a potential loss of $L \in (0, W)$. They have heterogeneous loss probabilities $\theta \in [\underline{\theta}, \overline{\theta}]$, which is their private information, where $0 \leq \underline{\theta} < \overline{\theta} < 1$. Risk-neutral firms sell insurance policies. An insurance policy with coverage x repays $x \cdot L$ if the consumer experiences a loss.

Consumers have constant absolute risk aversion, so their preferences can be represented as in equation (1) with

$$u(\theta,x) = \frac{\ln\left[1-\theta+\theta e^{AL}\right] - \ln\left[1-\theta+\theta e^{AL(1-x)}\right]}{A}, \tag{4}$$

where $A > 0$ is the coefficient of risk aversion.[6] Firm profits can be written as in (2) with $c(\theta, x) = x\theta L$.

The next example follows Rothschild-Stiglitz in assuming that types are one-dimensional, but does not restrict losses to be binary:

[6] All of our results hold in this model with arbitrary concave utility functions.

Example 3. (Levy and Veiga, 2022) Consider the following willingness to pay and cost of providing coverage functions

$$\begin{aligned} u(M,x) &= xM + g(x), \text{ and} \\ c(M,x) &= xM, \end{aligned} \tag{5}$$

where $g : [0,1] \to \mathbb{R}_+$ is a strictly concave function satisfying $g(0) = 0$, $g'(x) > 0$ for $x < 1$, and $g'(1) = 0$. In this model, types are one-dimensional and correspond to the consumer's expected loss: $\theta = M$.

The two examples below help illustrate results that hold beyond insurance.

Example 4. (Lemons Market with a Divisible Asset) Buyers are privately informed about the quality of an asset, represented by the type θ.[7] The valuations of buyers and sellers are:

$$\begin{aligned} u(\theta,x) &= \alpha\theta x, \text{ and} \\ c(\theta,x) &= \theta x, \end{aligned} \tag{6}$$

where $\alpha > 1$ (so there are positive gains from trade).

Example 5. (Non-Linear Pricing) In models of non-linear pricing, the seller's cost does not depend on the buyer's type: $c(\theta,x) = c(\widetilde{\theta},x)$ for all $\theta, \widetilde{\theta}$.

2.3 Definitions

An **allocation** is a measure α over $\Theta \times [0,1]$ such that the marginal distribution satisfies $\alpha|\Theta = \mu$. That is, $\alpha(\{\theta,x\})$ is the measure of θ types who purchase contract x. A **price** is a measurable function $p : [0,1] \to \mathbb{R}$ with $p(0) = 0$, with $p(x)$ denoting the premium charged for coverage x.[8] A **mechanism** (p,α) consists of a price p and an allocation α. A mechanism (p,α) is **incentive compatible** if for almost every (θ,x) with respect to α, consumers pick their contracts optimally:

$$u(\theta,x) - p(x) = \sup_{x' \in X} u(\theta,x') - p(x').$$

An allocation α is **deterministic** if for each θ, there exists $x(\theta)$ such that $\alpha(\theta,[0,1]) = \alpha(\theta,x(\theta))$. That is, a deterministic allocation assigns the same contract to each type.

[7] This example assumes that buyers are the informed party. It is straightforward to relabel players and renormalize prices and quantities to allow buyers to be uninformed and sellers to be informed.

[8] The requirement that zero coverage has a premium of zero corresponds to the standard participation constraint.

The monopolist's expected profits from a mechanism (p, α) are given by the expectations over the firm's profits (2) with respect to the measure α. A mechanism maximizes the monopolist's profits if: (a) it is incentive compatible, and (b) no other incentive-compatible mechanism gives higher expected profits. We use the following shorthand notation for conditional moments:

$$\mathbb{E}_x[c|\alpha] = \mathbb{E}[c(\widetilde{\theta}, \widetilde{x})|\alpha, \widetilde{x} = x].$$

That is, $\mathbb{E}_x[c|\alpha]$ is the expectation of $c(\widetilde{\theta}, \widetilde{x})$ according to the measure α and conditional on coverage $\widetilde{x} = x$. Our competitive equilibrium concept is based on Azevedo and Gottlieb (2017):

Definition 1. The pair (p^*, α^*) is a **competitive equilibrium** if

1. For each contract x, firms make zero profits: $p^*(x) = \mathbb{E}_x[c|\alpha^*]$ almost everywhere according to α^*.

2. Consumers select contracts optimally: for almost every (θ, x) with respect to α^*, we have
$$u(\theta, x) - p^*(x) = \sup_{x' \in X} u(\theta, x') - p^*(x').$$

3. For every contract $x' \in X$ with strictly positive price, there exists (θ, x) in the support of α^* such that
$$u(\theta, x) - p^*(x) = u(\theta, x') - p^*(x') \text{ and } c(\theta, x') \geq p^*(x').$$
That is, every contract that is not traded in equilibrium has a low enough price for some consumer to be indifferent between buying it or not, and the cost of this consumer is at least as high as the price.

Conditions 1 and 2 state that consumers and firms to optimize taking prices as given. Since condition 1 only requires prices to be equal to the average cost of consumers almost everywhere, it does not place restrictions on the prices of contracts that are not traded. Therefore, these two conditions alone do not rule out unreasonable equilibria where firms do not offer a contract because they fear that it will attract overly risky consumers when, in fact, if they offered this contract, they would attract consumers with much lower risk. Condition 3 is a refinement that rules out this type of "unreasonably pessimistic" beliefs. It requires contracts that are not traded to be such that, if their prices were slightly reduced, some consumers would choose to purchase them and the firm would not make money.[9]

[9] As Azevedo and Gottlieb (2017) show, condition 3 can be obtained from requiring the equilibrium to be robust to small perturbations in the spirit of proper equilibrium (Myerson (1978)). Moreover, competitive

Existence of competitive equilibrium follows from Theorem 1 and Proposition 1 in Azevedo and Gottlieb (2017). Existence of a mechanism that maximizes the monopolist's profits follows from Theorem 5.11 in Kadan et al. (2017). The following lemma will be useful in our characterization of competitive equilibria.

Lemma 1. *Suppose Assumption 1 holds and let (p^*, α^*) be a competitive equilibrium. Then, p^* is Lipschitz continuous and, therefore, Lebesgue almost everywhere differentiable.*

3 One-Dimensional Types

To illustrate our results, we start with the case of one-dimensional types: $\Theta = [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}_+$. We establish three results: (i) there is no exclusion with perfect competition, (ii) a monopolist excludes a positive mass of consumers if $\underline{\theta}$ is low enough, and (iii) high enough types get more coverage with a monopolist than with perfect competition. These results are generalized in Section 4.

3.1 No Exclusion with Perfect Competition

We impose the following assumption:

Assumption 2. *There exist $\varepsilon > 0$ and bounded functions $v, \kappa : [0,1] \to \mathbb{R}$ such that $\frac{\partial u}{\partial x}(\theta, x) \geq \theta v(x)$ and $c(\theta, x) \leq \theta \kappa(x)$ for all $x < \varepsilon$ and all $\theta > \underline{\theta}$. Moreover, κ is differentiable at 0 and $v(0) > \kappa'(0) \geq 0$.*

The first part of the assumption is technical. It requires marginal utility to be bounded from below and cost to be bounded from above for coverages close to zero. The second part is more economically substantial. It states that a small amount of coverage increases surplus, implying that exclusion is not a property of the first best.[10]

Proposition 1. *Suppose Assumptions 1 and 2 hold and let (p^*, α^*) be a competitive equilibrium. Then there is no exclusion: $\alpha^*(\Theta, 0) = 0$.*

Proof. Fix a competitive equilibrium in which a positive mass of types are excluded. Suppose that for every $\varepsilon > 0$ there exists $x \in (0, \varepsilon)$ such that no type is indifferent between their equilibrium contract and x. Then, by condition 3 of Definition 1, the price of that contract

equilibrium corresponds to the limit of a game in which horizontally differentiated firms with small scales compete through prices, as the amount of horizontal differentiation vanishes. For one-dimensional settings with single crossing, competitive equilibrium gives the same allocation as in Riley (1979) (see Subsection 3.3).

[10] It is perhaps not surprising that if there is exclusion in the first-best allocation, some types may be excluded in the competitive equilibrium. However, in most applications, it is not first-best optimal to exclude a positive mass of types (for example, because consumers are risk averse and firms are risk neutral).

must be zero, which implies that no type can be excluded (since any such type prefers to purchase $x > 0$ small enough at price zero over zero coverage). Thus, for any contract in a neighborhood of zero, there must exist some type for whom picking that contract maximizes utility. Moreover, again by condition 3, the cost of selling x to that type must be weakly higher than the price. Letting $m(x) \geq \underline{\theta}$ denote such a type, we must have:

$$(p^*)'(x) = \frac{\partial u}{\partial x}(m(x), x) \geq m(x)v(x) \tag{7}$$

at all points of differentiability of x (where the equality follows from the first-order condition and the inequality uses the bound in Assumption 2), and

$$p^*(x) \leq c(m(x), x) \leq m(x)\kappa(x) \tag{8}$$

where the first inequality follows from condition 3 of Definition 1 and the second uses the bound in Assumption 2.

Integrate (7) and use $p^*(0) = 0$ to obtain:

$$\int_0^x m(\widetilde{x})v(\widetilde{x})d\widetilde{x} \leq p^*(x) \leq m(x)\kappa(x),$$

where the last inequality uses (8). Divide both sides by $x > 0$ and rearrange:

$$m(x)\frac{\kappa(x)}{x} - \frac{\int_0^x m(\widetilde{x})v(\widetilde{x})d\widetilde{x}}{x} \geq 0.$$

Note that this condition must hold for all $x > 0$. We consider the limit of this expression as $x \searrow 0$. By the Fundamental Theorem of Calculus, $\lim_{x\searrow 0} \frac{\int_0^x m(\widetilde{x})v(\widetilde{x})d\widetilde{x}}{x} = m(0_+)v(0)$, where $m(0_+) \equiv \lim_{x\searrow 0} m(x)$. Therefore, the condition above for x in a neighborhood of $x = 0$ requires:

$$m(0_+)\left[\kappa'(0) - v(0)\right] \geq 0.$$

Since $v(0) > \kappa'(0)$ (Assumption 2), this condition requires $m(0_+) \leq 0$. Since $m(x) \geq \underline{\theta}$ for all x, this is a contradiction if $\underline{\theta} > 0$. If instead $\underline{\theta} = 0$, it implies $m(0_+) = 0$, so at most the lowest type (if $\underline{\theta} = 0$), which has zero measure, can be excluded. Therefore, there is no equilibrium in which a positive mass of types is excluded. $\square$

Note that allowing contracts to be endogenously determined is key for the no-exclusion result. With a single exogenous contract, the competitive equilibrium often excludes some types, and it may even exclude almost all types, as in Akerlof (1970).[11]

[11]Other restrictions to the contract space, such as a minimum coverage, may also generate exclusion (see Azevedo and Gottlieb (2017) and Levy and Veiga (2022)). However, maximum coverage regulations, which set

3.2 Exclusion with Monopoly

In this subsection, we show that exclusion is optimal whenever the type space includes types with low enough surplus from coverage. To make this statement formal, we need to define what it means for a type to have low enough surplus. This is done by considering a "safe type" ($\theta = 0$), which may not be in the type space (since we may have $\underline{\theta} > 0$):

Assumption 3. *(Safe type)* $\frac{\partial S}{\partial x}(0, x) \leq 0$ *for all* $x > 0$ *and* $\frac{\partial^2 u}{\partial \theta \partial x}(0, 0) > 0$*.*

Assumption 3 states that providing coverage does not increase the safe type's surplus ($\frac{\partial S}{\partial x}(0, x) \leq 0$), and types close enough to the safe type have a higher willingness to pay for coverage than the safe type ($\frac{\partial^2 u}{\partial \theta \partial x}(0, 0) > 0$).[12]

To understand Assumption 3, consider the Rothschild and Stiglitz (1976) model, where a type corresponds to the loss probability. An individual with $\theta = 0$ has probability zero of having a loss and therefore has both a zero willingness to pay and zero cost of coverage. If $\underline{\theta} > 0$, all types have a positive probability of experiencing a loss, so they all have a positive surplus from purchasing coverage. In fact, the surplus-maximizing contract for all types is full insurance. Nevertheless, as we show next, a monopolist excludes a positive mass of types as long as $\underline{\theta}$ is close enough to zero.

Proposition 2. *Suppose Assumptions 1 and 3 hold. There exists* $\theta^* > 0$ *such that if* $\underline{\theta} < \theta^*$*, then any mechanism that maximizes the monopolist's profits excludes a set of types with positive measure:* $\alpha^*(\Theta, 0) > 0$*.*

Proof. For simplicity, we restrict attention to deterministic allocations here (see the Supplementary Appendix D for the general proof, which also allows for stochastic allocations). Fix an incentive-compatible allocation $x(\cdot)$ and let $U(\theta) \equiv u(\theta, x(\theta)) - p(x(\theta))$ denote the indirect utility of type θ. By the envelope theorem, incentive compatibility implies:

$$\dot{U}(\theta) = \frac{\partial u}{\partial \theta}(\theta, x(\theta)) > 0. \tag{9}$$

Since $U(\cdot)$ is increasing in θ, the exclusion region is an interval: $[\underline{\theta}, \theta^*]$. If all types participate ($\theta^* = \underline{\theta}$), any allocation that maximizes the firm's profits must give zero utility to the lowest type. If there is exclusion ($\theta^* > \underline{\theta}$), all types who do not participate get zero utility. So in either case we have $U(\theta^*) = 0$.

the contract space as $[0, \bar{x}]$ for $\bar{x} < 1$, can be accommodated by renormalizing the coverage level x. Therefore, the no-exclusion result from Proposition 1 generalizes to settings in which regulators ban sufficiently generous plans ("cadillac plans"). Note that the no-exclusion result is immediate if there is no adverse selection (such as in example 5), since the competitive equilibrium has each good supplied at its marginal cost, so all consumers with a positive surplus buy positive amounts.

[12] This second requirement is related to but substantially weaker than the single-crossing property (see Assumption 4 below).

Integrate equation (9) to obtain:

$$U(\theta) - \underbrace{U(\theta^*)}_{0} = \int_{\theta^*}^{\theta} \frac{\partial u}{\partial \tilde{\theta}}(\tilde{\theta}, x(\tilde{\theta}))d\tilde{\theta}.$$

Substituting in the firm's expected profits and integrating by parts, we obtain the virtual surplus expression:

$$\int_{\theta^*}^{\bar{\theta}} \left[S(\theta, x(\theta)) - \frac{\partial u}{\partial \theta}(\theta, x(\theta)) \cdot \frac{1 - F(\theta)}{f(\theta)} \right] f(\theta)d\theta.$$

For the moment, suppose $\underline{\theta} = 0$. To show that any optimal allocation excludes some types, we verify that the integrand of the expression above is negative for θ close enough to zero. The integrand evaluated at $\theta = 0$ equals

$$S(0, x(0)) - \frac{\partial u}{\partial \theta}(0, x(0)) \cdot \frac{1}{f(0)}.$$

Note that $S(0, x) \leq 0$ for all x (Assumption 3) and $\frac{\partial u}{\partial \theta}(0, x) > 0$ for all $x > 0$ (Assumption 1). Therefore, the expression above is strictly negative if $x(0) > 0$, implying that it is optimal to set $\theta^* > 0$. If $x(0) = 0$, the expression above equals zero. Its derivative with respect to $x(0)$ equals:

$$\frac{\partial S}{\partial x}(0, 0) - \frac{\partial^2 u}{\partial x \partial \theta}(0, 0) \cdot \frac{1}{f(0)} < 0,$$

where the inequality follows from $\frac{\partial S}{\partial x}(0, 0) \leq 0$ and $\frac{\partial^2 u}{\partial x \partial \theta}(0, 0) > 0$ (Assumption 3). Since $x(\theta) > 0$ for all $\theta > \theta^* = 0$ (types above θ^* are not excluded), by the continuity of the expression on the LHS, it follows that the integrand is strictly negative for all θ in a neighborhood of $\theta^* = 0$. Since S is continuously differentiable and u is twice continuously differentiable, and the integrand remains negative if $\underline{\theta}$ is sufficiently small. □

3.3 Monopoly provides more coverage at the top

Propositions 1 and 2 imply that low types buy more coverage in a competitive market than with a monopolist. We now consider the coverage purchased by high types. We impose the following conditions:

Assumption 4. *(Single Crossing) The utility function satisfies* $\frac{\partial^2 u}{\partial \theta \partial x}(\theta, x) > 0$ *for all* θ, x.

Assumption 4 is standard in one-dimensional models. It states that higher types to have a higher marginal utility from coverage and implies that incentive-compatible mechanisms are non-decreasing.

Assumption 5. *(Monotonicity) The utility function u is strictly increasing in x for all $\theta > \underline{\theta}$. The cost function c is strictly increasing in x for all $\theta > \underline{\theta}$, and strictly increasing in θ for all $x > 0$.*

In addition to stating that customers like coverage but providing coverage is costly, Assumption 5 states that there is adverse selection, so that it costs more to provide coverage to types who are willing to pay more for coverage.

Assumption 6. $\frac{\partial S}{\partial x}(\overline{\theta}, \overline{x}) = 0$ *where* $\overline{x} \equiv \arg\max_x S(\overline{\theta}, x) > 0$ *and* $\frac{\partial^2 S}{\partial x^2}(\theta, x) < 0$ *for all* θ.

Assumption 6 states that the efficient contract for the highest type solves a first-order condition. Intuitively, this assumption precludes exogenous upper bounds on the space of contracts. It requires that either the efficient coverage of the highest type is interior ($0 < \overline{x} < 1$) or that it would be inefficient to offer a coverage greater than 100 percent (so the consumer makes money by incurring a loss). Our last assumption imposes technical smoothness conditions:

Assumption 7. *The PDF f is continuously differentiable and u is three times continuously differentiable.*

The **least-costly separating allocation** is the deterministic allocation that solves the following ordinary differential equation:

$$\dot{x}(\theta) = \frac{\frac{\partial c}{\partial \theta}(\theta, x(\theta))}{\frac{\partial S}{\partial x}(\theta, x(\theta))}, \tag{10}$$

with boundary condition $x(\bar{\theta}) = \bar{x}$.[13] To understand this equation, note that incentive compatibility requires type $x(\theta)$ to maximize type the utility of type θ, giving the necessary first-order condition:

$$\frac{\partial u}{\partial x}(\theta, x(\theta)) = p'(x(\theta)) \tag{11}$$

at all points in which p is differentiable. Since types are separated, the zero profits condition gives:

$$p(x(\theta)) = c(\theta, x(\theta)).$$

Differentiating and substituting back in equation (11), gives (10). Note that the least costly separating allocation has $x(\overline{\theta}) = \overline{x}$ and $\lim_{\theta \nearrow \overline{\theta}} \dot{x}(\theta) = +\infty$. Therefore, coverage is very steep close to the top.

The lemma below shows that this is the allocation in the unique competitive equilibrium:

[13] Formally, equation (10) is not well defined at $\theta = \bar{\theta}$. However, one can define the least-costly separating allocation in terms of its inverse, the "type-assignment function," which specifies the type that picks each coverage. Existence and uniqueness then follow from the Picard-Lindelöf Theorem.

Lemma 2. *Suppose Assumptions 1 and 4-7 hold. The competitive equilibrium is unique and its allocation $x^c(\cdot)$ is the least-costly separating allocation.*

To illustrate the reason for screening in a competitive market, suppose the surplus-maximizing coverage does not depend on types: $\frac{\partial S}{\partial x}(\theta, \bar{x}) = 0$ for all θ. This is the case, for example, in the Rothschild and Stiglitz model, where full insurance maximizes the surplus of each type. The problem is that competitive firms cannot prevent cream skimming by other firms. If multiple types purchase the same coverage level, a firm can always profit by offering slightly less coverage at a lower premium, attracting only the less risky types. The competitive equilibrium allocation distorts each type's coverage downwards by the exact amount needed to prevent cream skimming while preserving zero profits.

To understand the amount of distortion needed to prevent cream skimming, note that each type equates the marginal utility from coverage with the price of an additional coverage, as in equation (11). The price of an additional unit of coverage $p'(x(\theta))$ has two components: the marginal cost of coverage holding selection constant, $\frac{\partial c}{\partial x}(\theta, x(\theta))$, and the increase in cost due to adverse selection, $\frac{\frac{\partial c}{\partial \theta}(\theta, x(\theta))}{\dot{x}(\theta)}$. Since the first best equates marginal utility with the marginal cost of coverage, $\frac{\partial u}{\partial x}(\theta, x(\theta)) = \frac{\partial c}{\partial x}(\theta, x(\theta))$, adverse selection distorts the equilibrium allocation downwards. Moreover, as types increase, we approach the efficient allocation. Close to the efficient allocation, the loss in surplus from reducing coverage becomes arbitrarily small. Therefore, in order to prevent cream skimming, firms must substantially reduce coverage: $\lim_{\theta \nearrow \bar{\theta}} \dot{x}(\theta) = +\infty$.

We now turn to the monopolist's program. For simplicity, we restrict ourselves to deterministic allocations (that is, we focus on mechanisms that maximize the monopolist's profits among those with deterministic allocations). In the appendix, we generalize the analysis to allow for stochastic allocations.

Lemma 3. *Suppose Assumptions 1 and 4-7 hold and suppose $x^m(\cdot)$ maximizes the monopolist's profits among deterministic allocations. Then, $x^m(\overline{\theta}) \geq \overline{x}$ and $\dot{x}^m(\overline{\theta}) < \infty$.*

The proof solves the monopolist's program and shows that there are two cases. If the monotonicity constraint associated with incentive compatibility does not bind, the solution has the same boundary condition $x^m(\overline{\theta}) = \overline{x}$ (no distortion at the top) but has a flatter slope $\dot{x}^m(\overline{\theta}) < +\infty$. If the monotonicity constraint does not bind, there is bunching at the top at a point above the efficient coverage: $x^m(\overline{\theta}) \geq \overline{x}$ and $\dot{x}^m(\overline{\theta}) = 0$.[14]

A monopolist offers partial insurance for a different reason than competitive firms. A monopolist does not face the risk of cream skimming. However, providing the efficient coverage to all types would require the monopolist to charge the willingness to pay of the lowest

[14]In the Rothschild-Stiglitz model (example 2), the allocation that maximizes profits is deterministic and has no bunching at the top, so the solution entails $x^m(\overline{\theta}) = \overline{x}$ and $\dot{x}^m(\overline{\theta}) < +\infty$.

type, leaving excessive informational rents. The monopolist balances the efficiency gain from increasing coverage against informational rents that are left to higher types. As long as the distribution of types is continuously differentiable around $\bar{\theta}$, the optimal policy has a finite slope.

From Lemmas 2 and 3, it follows that the monopolist always provides more coverage for high enough types:

Proposition 3. *Suppose Assumptions 1 and 4-7 hold. There exists $\theta^* < \overline{\theta}$ such that $x^m(\theta) > x^c(\theta)$ for all $\theta \in (\theta^*, \overline{\theta})$.*

Recall that Assumption 5 requires costs to be increasing in types, ruling out non-linear pricing models (example 5). If instead costs were constant in consumer types, the competitive equilibrium would feature each product being sold at cost and the allocation would be efficient. Under standard regularity conditions on the distribution of types, the monopolist would distort quantity downwards. Therefore, Proposition 3 does not generalize to private values settings.

4 Multidimensional Model

This section generalizes the results from Section 3 to settings with multidimensional types.

4.1 No Exclusion with Perfect Competition

The following assumption generalizes Assumption 2:

Assumption 2′. *There exist $\epsilon > 0$ and bounded functions $v_0, v_k, \kappa_k : [0,1] \to \mathbb{R}$ such that for all $x < \epsilon$ and all $\theta \in \Theta$:*

a. (Bounded marginal utility) $\frac{\partial u}{\partial x}(\theta, x) \geq v_0(x) + \sum_{k=1}^{K} \theta_k v_k(x) > 0$;

b. (Bounded cost) $c(\theta, x) \leq \sum_{k=1}^{K} \theta_k \kappa_k(x)$;

c. (Positive surplus) $v_0(0) + \sum_{k=1}^{K} \theta_k \left[v_k(0) - \kappa_k'(0)\right] > 0$.

Part (a) is a technical condition requiring the marginal utility of a small amount of coverage to be bounded from below by functions that are linear in types. Part (b) is also a technical condition requiring the cost of a small amount of coverage to be bounded above by a function that is linear in types. Part (c) states that the marginal utility of a small amount of coverage exceeds the marginal cost for each type, so exclusion is not a property of the first best. Note that Assumption 2' does not impose an order on marginal utility of coverage (such as the single-crossing property), so incentive-compatible allocations may be non-monotonic.

It is straightforward to verify Assumption 2' in all the examples from Subsection 2.2.[15] The theorem below establishes that there is no exclusion in competitive equilibrium:

Theorem 1. *Suppose Assumptions 1 and 2' hold and let* (p^*, α^*) *be a competitive equilibrium. Then there is no exclusion:* $\alpha^*(\Theta, 0) = 0$.

The proof follows similar steps as in the one-dimensional case and is presented in the appendix.

4.2 Exclusion with Monopoly

Next, we show that a monopolist always excludes a positive mass of consumers if some types have low enough surplus from coverage. Let $\|\cdot\|$ denote the Euclidean norm and let $\nabla_\theta u$ denote the gradient of u with respect to θ.

We impose the following technical conditions on the marginal utility of coverage for types close enough to zero:

Assumption 8. *There exist* $\kappa > 0$ *and* $\varepsilon > 0$ *such that* $\theta \cdot \nabla_\theta u(\theta, x) \geq \kappa u(\theta, x)$ *for all* $x \geq 0$ *and all* θ *with* $\|\theta - \underline{\theta}\| < \varepsilon$.

It is straightforward to verify Assumption 8 in examples 1-4. If types are one-dimensional, the assumption states that the utility function increases faster than some power function.[16] Assumption 8 is useful as it allows us to obtain a lower bound on the consumer's informational rent using a homogeneous function.

Our second assumption generalizes the requirement that there exist some types with low enough surplus in Assumption 3. Before stating the assumption, we need to introduce some notation. For each $\theta \in \Theta$, let $\ell(\theta) \equiv \max\{t : \ t \geq 1 \text{ and } t\theta \in \Theta\}$ denote the length of the ray connecting θ to the boundary of Θ. For each $\epsilon > 0$, let

$$\Theta_\epsilon \equiv \{\theta \in \Theta : \ ||\theta - \underline{\theta}|| < \epsilon\}$$

denote the set of types distanced less than ϵ from $\underline{\theta}$. Let $\underline{\ell} \equiv \inf\{\ell(\theta) : \ \theta \in \Theta_\epsilon\} > 1$ denote the shortest ray connecting a type in Θ_ϵ and the boundary of Θ and let $\underline{f} \equiv \min\{f(\theta) : \ \theta \in \Theta\} > 0$ denote the lowest density of types. Finally define the following constant:

$$\eta \equiv \frac{1}{1 - \kappa \frac{\underline{\ell}^K - 1}{K} \frac{\underline{f}}{f(\underline{\theta})}} > 1, \tag{12}$$

[15] In example 1, we have $\theta = (M, S^2 A, H)$, and the conditions are satisfied with $v_0(x) = 0$, $v_1(x) = 1$, $v_2(x) = 1 - x$, $v_3(x) = x$, $\kappa_1(x) = x$, $\kappa_2(x) = 0$, and $\kappa_3(x) = x^2$.

[16] That is, there exists $\kappa > 0$ such that $\frac{u(\theta,x)}{\theta^\kappa}$ is increasing in θ for all x and all $\theta < \epsilon$.

where $\underline{\theta} \equiv (\underline{\theta}_1, ..., \underline{\theta}_K)$. Note that η depends on the geometry of the type space (through $\underline{\ell}$ and K), the distribution of types (through $\underline{f}$ and $f(\underline{\theta})$), and the utility function (through κ and ϵ).

Assumption 3′. *(Safe type) $\eta c(\underline{\theta}, x) \geq u(\underline{\theta}, x)$ for all $x > 0$ and $\eta \frac{\partial c}{\partial x}(\underline{\theta}, 0) \geq \frac{\partial u}{\partial x}(\underline{\theta}, 0)$.*

Note that Assumption 3' always holds if the lowest type has non-positive surplus and marginal surplus: $c(\underline{\theta}, x) \geq u(\underline{\theta}, x)$ and $\frac{\partial c}{\partial x}(\underline{\theta}, 0) \geq \frac{\partial u}{\partial x}(\underline{\theta}, 0)$. With the terminology of Assumption 3, this condition always holds if the "safe type" belongs to the type space. But just as with Assumption 3, we do not require that the safe type belongs to the type space. Since $\eta > 0$, Assumption 3' allows for all types to have a positive surplus and marginal surplus from coverage, as long as they are "not too high" for the lowest type.

Theorem 2. *Suppose Assumptions 1, 3', and 8 hold. Then, any mechanism that maximizes the monopolist's profits excludes a set of types with positive measure: $\alpha^*(\Theta, 0) > 0$.*

As in the one-dimensional model, a monopolist balances the efficiency gain in providing additional coverage to some consumers against the ability to change more from those with a higher willingness to pay for coverage. When consumers with the lowest willingness to pay have low enough surplus, the efficiency gain from excluding them is lower than the rent extraction. Equation (12) provides an upper bound on the surplus of those consumers. It is increasing in the range of types for which informational rents can be bounded by a power function (since $\underline{\ell}$ is increasing in ϵ) and is decreasing in the density of lowest types $f(\underline{\theta})$ (which makes exclusion more costly to the firm). Finally, note that Theorem 2 also applies to economies without adverse selection since it does not require any assumptions about how costs depend on θ.

4.3 Distortions at the Top

Theorems 1 and 2 imply that a monopolist provides less coverage to individuals with a low willingness to pay than competitive firms. We now consider the coverage offered to those with a high willingness to pay with multidimensional types.

In this subsection, we assume that utility and costs are separable:

$$u(\theta, x) = \sum_{i=1}^{K} \theta_i \phi_i(x) \text{ and } c(\theta, x) = \sum_{i=1}^{K} \theta_i \kappa_i(x), \tag{13}$$

where $\phi_i(\cdot)$ and $\kappa_i(\cdot)$ are continuously differentiable with $\kappa_i(x) > 0$ for all $x > 0$, for all i.

With this specification, the surplus from providing coverage x to type θ becomes

$$S(\theta, x) = \sum_{i=1}^{K} \theta_i \left[\phi_i(x) - \kappa_i(x)\right].$$

We assume that the total surplus $S(\theta, \cdot) : [0,1] \to \mathbb{R}$ is strictly concave in coverage with a second derivative uniformly bounded away from zero:

$$\gamma \equiv \inf \left\{ \left| \frac{\partial^2 S}{\partial x^2}(\theta, x) \right| : \ (\theta, x) \in \Theta \times [0,1] \right\} > 0. \tag{14}$$

It is straightforward to verify that this condition holds in all the insurance models in Subsection 2.2 (Examples Einav et al. (2013), 2, and 3).

In this section, we focus on deterministic allocations (see the online appendix for a generalization to stochastic allocations). Recall that an allocation α is deterministic if there exists $\widetilde{x} : \Theta \to [0,1]$ such that $\alpha(\theta, [0,1]) = \alpha(\theta, \widetilde{x}(\theta))$. In this case, we abuse of notation and refer to the function $\widetilde{x}(\cdot)$ as a deterministic allocation.[17]

In the one-dimensional model, the restriction on prices of non-traded contracts imposed by Condition 3 of Definition 1 rules out any pooling. If an multiple types buy the same contract, a firm can profit by offering a slightly lower coverage at a discounted price, stealing only the safest types ("cream skimming"). With multidimensional types, some pooling is unavoidable. For instance, a consumer with high risk and low risk aversion may end up buying the same coverage as a low-risk, high-risk-aversion consumer. We therefore need to generalize the concept of cream skimming for settings in which necessarily involve some pooling.

Fix an incentive-compatible, deterministic mechanism. At any point of differentiability of p, if θ chooses to purchase an interior coverage x, the following first-order condition must hold

$$\frac{\partial u}{\partial x}(\theta, x) = \sum \theta_i \phi_i'(x) = p'(x),$$

where we used separable utility function from equation (13). Therefore, at points of differentiability of p, types choosing the same contract will have the same marginal utility from coverage.

Formally, let

$$\Theta_x^{\nu} \equiv \left\{ \widetilde{\theta} \in \Theta : \ \sum \widetilde{\theta}_i \phi_i'(x) = \nu \right\}$$

denote the set of types with the same marginal utility ν from coverage x. The set Θ_x^{ν} includes

[17] We are abusing notation here because the allocation α also specifies distributions conditional on contracts that are not chosen by any type. Intuitively, these distributions correspond to "beliefs" conditional on contracts that off the equilibrium path. Such beliefs are not included in $\widetilde{x}(\cdot)$, which only specifies contracts on the equilibrium path, but need to satisfy the equilibrium refinement (Condition 3 from Definition 1).

all types that choose contract x when the marginal price is ν. In the model of Farinha Luz et al. (2023), this set is a hyperplane that combines types with different levels of risk and risk aversion. In Example 1, this hyperplane combines individuals with different levels of risk, risk aversion, mean and variance of losses, and moral hazard. In the one-dimensional model with single crossing (Subsection 3.3), Θ_x^ν is a singleton, since there is at most one type with a given marginal utility from coverage x (all types above this type have a higher marginal utility and all types below have a lower marginal utility).

Since pooling is generally unavoidable with multidimensional types, we need to consider which type among those buying the same coverage has the highest incentive to deviate from a contract. At all points of differentiability of p, all pooled types have the same marginal utility from coverage. Therefore, their incentive to deviate is given by the convexity of their utility function. As we show in the appendix, the type with the highest incentive to increase coverage is the one with the most convex utility function:

$$\arg\max_{\widetilde{\theta}\in\Theta_x^\nu}\left\{\sum \widetilde{\theta}_i \phi_i''(x)\right\}.$$

Reciprocally, the types with the highest incentive to decrease coverage is the one with the most concave utility function.

The following condition states that by offering a slightly higher coverage, a firm attracts the type with the highest cost among those who were originally choosing the same coverage.

Assumption 9. *(Strong Adverse Selection) For all ν, x,*

$$\sum \theta_i^\nu \kappa_i(x) \geq \sum \theta_i \kappa_i(x)$$

for each $\theta \in \Theta_x^\nu$ and $\theta^\nu \in \arg\max_{\widetilde{\theta}\in\Theta_x^\nu}\left\{\sum \widetilde{\theta}_i \phi_i''(x)\right\}$.

For example, in Farinha Luz et al. (2023), each pool combines types with different levels of risk and risk aversion. Among those, the type with the highest risk and lowest risk aversion is the most willing to increase coverage. This is also the type with the highest cost. In Example 1, the type who is most willing to increase coverage among those in the same pool is the one with the lowest risk aversion and the highest moral hazard. This is again the type with the highest cost.

An implication of strong adverse selection is that a firm that offers a slightly lower coverage at a discounted price steals only the cheapest types. Therefore, as we show in the appendix, with strong adverse selection all types who purchase the same contract have the same marginal utility from coverage. In particular, we cannot have a positive mass of types picking the same contract.

While it is often straightforward to check for strong adverse selection directly, the following proposition gives an intuitive sufficient condition:

Proposition 4. *Suppose that for all $i, j \in \{1, ..., K\}$ and all $x \in [0, 1]$,*

$$\frac{\phi_j''(x)}{\phi_j'(x)} \geq \frac{\phi_i''(x)}{\phi_i'(x)} \iff \frac{\kappa_j(x)}{\phi_j'(x)} \geq \frac{\kappa_i(x)}{\phi_i'(x)}.$$

Then, Assumption 9 holds.

The condition in Proposition 4 relates the curvature of the utility function with the cost per marginal willingness to pay. The dimension in which consumers have the highest incentive to increase coverage is the one with the most convex willingness to pay. Strong adverse selection requires an increase in that direction to be the one that raises cost the most.

Fix a deterministic allocation $\widetilde{x}(\cdot)$. Let $\widetilde{x}_+ \equiv \sup\{\widetilde{x}(\theta) : \theta \in \Theta\}$ denote the highest contract chosen by any type in this allocation. For each $\varepsilon > 0$, let $N_{\widetilde{x},\varepsilon} \equiv \{\theta \in \Theta : \widetilde{x}(\theta) > \widetilde{x}_+ - \varepsilon\}$ denote the set of types who pick contracts in an ε-neighborhood of this highest contract. We can now state the main result from this section:

Theorem 3. *Suppose Assumption 9 holds. Let $\widetilde{x}(\cdot)$ be a deterministic equilibrium allocation and let x be an incentive-compatible deterministic allocation. Suppose there exists $\varepsilon > 0$ such that:*

- *$x(\theta) \leq \widetilde{x}(\theta)$ for all $\theta \in N_{\widetilde{x},\varepsilon}$, and*
- *$x(\theta) < \widetilde{x}(\theta)$ in a subset of $N_{\widetilde{x},\varepsilon}$ with positive measure.*

Then, there exists a contract assignment that coincides with $\widetilde{x}(\cdot)$ in $N_{\widetilde{x},\varepsilon}$ and gives the principal a strictly higher profit.

Theorem 3 shows that starting from a competitive allocation, a firm with market power profits by providing greater coverage for those with a high willingness to pay.

5 Illustrative Calibration

To illustrate the quantitative implications of our results, we calibrated the multidimensional health insurance model from example 1 based on Einav et al. (2013)'s estimates from employees in a large US corporation.[18] In the appendix, we present results for when losses are

[18] Our simulations are not aimed at matching any specific market, which would take the estimates from Einav et al. (2013) too far from the range of contracts in their data. Our goal is to illustrate that the theoretical effects obtained previously can be quantitatively important in a realistic setting.

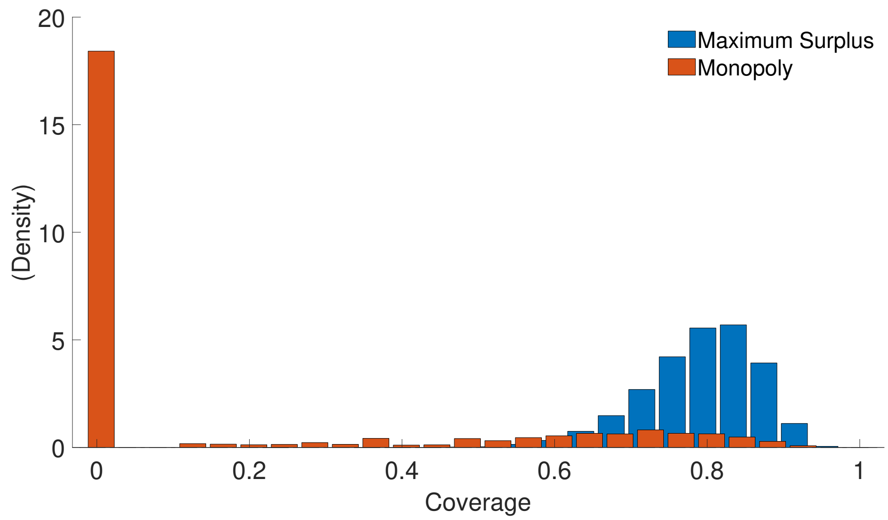

Figure 3: Monopoly and Efficient Coverage

Notes: The figure depicts the distribution of coverage choices in the numerical example. The horizontal axis depicts the contracts chosen by consumers, with coverage ranging between 0% (uninsured) to 100% of expenses. The blue bars represent the distribution of coverage in the competitive equilibrium. The orange bars represent coverage under monopoly pricing. With monopoly pricing, over 70 percent of consumers remain uninsured. In the surplus-maximizing allocation, all consumers purchase coverage.

distributed according to a truncated normal, which are similar to the ones discussed here. Consumers are heterogeneous along four dimensions: expected value of the health shock, standard deviation of the health shock, moral hazard, and risk aversion. We assumed that the distribution of parameters in the population is log-normal and picked the same parameters as in Azevedo and Gottlieb (2017), which were chosen to match the central estimates of Einav et al. (2013).

Figure 2 (in the introduction) contrasts the coverage under monopoly and perfect competition. Under monopoly, there is substantial exclusion, with 70 percent of the population choosing not to purchase any coverage. In contrast, with perfect competition, all customers purchase some coverage. This pattern reverses when considering higher levels of coverage, with 12 percent of customers opting for more coverage under monopoly compared to perfect competition. Nevertheless, as Figure 3 illustrates, coverage under monopoly is still lower than the level that maximizes total surplus (second-best).

Due to these offsetting effects, it is theoretically possible to have a higher surplus either under a monopoly or perfect competition. In our simulations, we consistently observed a higher total surplus under perfect competition (see Table 1). With parameter values considered here, the total surplus amounts to \$3,047 per consumer per year in the competitive equilibrium and \$2,223 under a monopoly. There is a substantial difference in consumer surplus (\$1,308 with monopoly versus \$3,047 with perfect competition) because a monopolist charges much higher prices. With perfect competition, there is approximately 8 percent loss in total surplus relative to the second-best optimum. A monopoly results in a much larger loss, approximately 32 percent of the surplus.

	No coverage	$\geq 70\%$ coverage	$\geq 80\%$ coverage	CS	Profit
Competition	0.0	11.1	4.1	3,047	0
Monopoly	70.8	11.4	5.7	1,308	915
Efficient	0.0	89.4	62.8	5,769	-2476

Table 1: Coverage with Perfect Competition, Monopoly, and Maximum Surplus

Notes: The table presents the coverage, consumer surplus (CS), and profits in the numerical example with perfect competition, monopoly, and surplus maximization. The first column corresponds to the percentage consumers excluded in each market structure. The second and third columns describe the proportion of consumer who purchase at least 70% and 80% coverage. The two last columns correspond to consumer surplus and profits.

6 Related Literature

A large theoretical literature studies competitive markets with adverse selection. Most of this literature follows two distinct approaches, building on either Rothschild and Stiglitz (1976) or Akerlof (1970). Models that build on Rothschild and Stiglitz allow contracts to be endogenously determined, while restricting consumers to be heterogeneous along a single dimension. This is an important limitation when considering applications, since there is abundant evidence that multiple dimensions of private information are important.[19]

Models that build on Akerlof consider a market with a single insurance contract with exogenous characteristics.[20] This setting allows for rich consumer heterogeneity, which is important to capture realistic insurance demand patterns. However, the assumption of a single exogenous contract makes it impossible to use the model to study which policies are offered in equilibrium. In particular, the model cannot distinguish between effects at the extensive and intensive margin (such as in Figure 1). More recently, Azevedo and Gottlieb (2017), Levy and Veiga (2022), and Farinha Luz et al. (2023) consider competitive models with multiple dimensions of consumer heterogeneity and endogenous contracts. Our paper

[19]Since pure strategy Nash equilibria often fails to exist in their model, a large literature has studied alternative equilibrium concepts to capture competitive markets (Wilson (1977), Miyazaki (1977), Riley (1979), Bisin and Gottardi (1999), Gale (1992), Dubey and Geanakoplos (2002), Azevedo and Gottlieb (2017)). Farinha Luz (2017) characterizes the mixed strategy equilibria of the original Rothschild and Stiglitz game.

[20]See Einav and Finkelstein (2011), Hackmann et al. (2015), Spinnewijn (2017), Scheuer and Smetters (2018), Fang and Wu (2018), and Handel et al. (2019). Handel et al. (2015) and Landais et al. (2021) consider a choice between two exogenous contracts, while allowing for rich consumer heterogeneity. Einav et al. (2021) survey the recent industrial organization literature on selection markets.

builds on this literature by studying the effect of market power in these settings.[21]

There is also some work on models with market power. Stiglitz (1977) considers optimal pricing for a monopolist facing consumers with two risk types. Chade and Schlee (2021) consider a continuum of consumers.[22] Mahoney and Weyl (2017) generalize the Einav et al. (2010) model to allow for imperfect competition among symmetric firms, retaining the assumption that firms offer a single exogenous contract. They study how changes in the degree of selection affects consumer surplus and profits and show that some intuitions from competitive markets do not carry over to settings with market power.

Veiga and Weyl (2016) allow for endogenous quality, while retaining the assumption that firms cannot offer multiple contracts. This single-contract assumption leads to a tractable model in which exclusion is determined by an intuitive first-order condition. They find that market power has an ambiguous impact on coverage, due to two conflicting effects. On the one hand, as in settings without selection, a monopolist that cannot price discriminate would like to cut quantity to increase prices. On the other hand, since market power limits the firms' incentives to engage in cream skimming, firms with market power can offer more coverage. Our results show these two effects persist when firms can offer multiple policies, and each of them dominates at each end of the distribution of coverage. At the bottom, the rent extraction motive dominates, making coverage lower under monopoly than under competition. At the top, the cream skimming effect dominates, making coverage higher under monopoly.

Chade et al. (2022) consider a multidimensional model with a finite set of contracts. They show that the unrestricted optimum can be approximated with a finite number of contracts and contrast policies that maximize social welfare with those that maximize monopoly profits. They also consider a simplified program, show that a monopolist has higher incentives to exclude than a utilitarian social planner, and obtain conditions for exclusion. Our numerical results suggest that our model and Chade et al. (2022) agree on many qualitative predictions. For example, both our simulations find substantial exclusion (39 percent in their simulations versus 70 in ours).

As illustrated in example 5, our framework is also related to the literature on non-linear pricing with a single instrument.[23] The main difference between non-linear pricing and

[21]Working in the Rothschild and Stiglitz setting, Hendren (2013, 2014) obtains conditions for consumers not to obtain any coverage in incentive-compatible mechanisms that break even. The conditions require the type space to include a type that incurs the loss with certainty. When no type incurs a loss with certainty, the competitive equilibrium features the least costly separating allocation. Shutdown does not happen in Proposition 1 because Assumption 2 rules out models in which it is efficient for the highest type to purchase zero coverage.

[22]Castro-Pires et al. (Forthcoming) propose a decoupling method to study the monopolist's problem with both adverse selection and moral hazard.

[23]Examples of such models include Mussa and Rosen (1978) and Maskin and Riley (1984) for one-dimensional types and Laffont et al. (1987) for two-dimensional types. See Rochet and Stole (2003) for

insurance is that values are private in non-linear pricing, whereas common values play an important role in insurance due to adverse selection. Since the exclusion results do not rely on common values, they still apply in non-linear pricing settings. Without adverse selection, a monopolist still excludes a positive mass of consumers if some of them have low enough willingness to pay (Proposition 2 and Theorem 2).[24] On the other hand, since competitive firms set prices equal to marginal cost and sell the efficient quantity, there is no exclusion with perfect competition. However, because there are no incentives to cream skim, monopolists do not provide greater coverage at the top if there is no adverse selection (Proposition 3 and Theorem 3 do not hold).

7 Conclusion

This paper studies how market power affects insurance policies in a general class of models. We show that market power creates different distortions compared to perfect competition, with each distortion dominating at different ends of the willingness-to-pay spectrum.

A monopolist faces a trade-off between efficiency and rent extraction. Reducing coverage results in an efficiency loss, as customers would be willing to pay more than the actuarially fair price to increase their coverage. However, this reduction in coverage allows the firm to extract higher rents by increasing the premiums charged from those willing to pay more. At the lower end of the willingness-to-pay spectrum, the rent extraction effect dominates, so a monopolist prefers not to sell to a positive mass of customers.

In a competitive market, firms are concerned about cream skimming by other firms. With cream skimming, a competitor steals the less risky customers, leaving the original firm only with the riskier, less profitable ones. The competitive equilibrium allocation distorts coverage downwards to prevent cream skimming. At the higher end of the willingness-to-pay spectrum, the incentives to cream skim are very high, so the competitive allocation provides less coverage than a monopolist.

The welfare effect of market power depends on whether the distortion at the bottom exceeds the one at the top. Using simulations based on an empirical model of preferences, we find that those both effects are quantitatively important, although the effect at the bottom

a survey.

[24] Armstrong (1996) shows the optimality of exclusion in a non-linear pricing model with multidimensional types. His result is distinct from our Proposition 2 and Theorem 2 in multiple ways. First, he considers a monopolist selling multiple goods, so the instrument has as many dimensions as the type, whereas we consider the allocation of a one-dimensional object (coverage). Second, he assumes that the utility function is homogeneous of degree one, while we consider more general utility functions. Third, he considers private values. And fourth, he considers types that are bounded away from zero, so there may not be exclusion in the one-dimensional version of his model. However, we build on his approach of integrating along rays in our proof of Theorem 2.

usually dominates. Consequently, in our simulations, total surplus is higher under perfect competition than under monopoly.

References

Akerlof, G. A. (1970): "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 83, 488–500.

Armstrong, M. (1996): "Multiproduct Nonlinear Pricing," *Econometrica*, 64, 51–75.

Azevedo, E. and D. Gottlieb (2017): "Perfect Competition in Markets with Adverse Selection," *Econometrica*, 85, 67–105.

Bisin, A. and P. Gottardi (1999): "Competitive Equilibria with Asymmetric Information," *Journal of Economic Theory*, 87, 1–48.

Cabral, M., M. Geruso, and N. Mahoney (2018): "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage," *American Economic Review*, 108, 2048–87.

Cardon, J. H. and I. Hendel (2001): "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *The RAND Journal of Economics*, 32, 408–427.

Castro-Pires, H., H. Chade, and J. Swinkels (Forthcoming): "Disentangling Moral Hazard and Adverse Selection," *American Economic Review*.

Chade, H., V. Marone, A. Starc, and J. Swinkels (2022): "Multidimensional Screening and Menu Design inHealth Insurance Markets," *Working paper*.

Chade, H. and E. Schlee (2021): "Optimal insurance with adverse selection," *Theoretical Economics*, 7, 571–607.

Cicala, S., E. M. J. Lieber, and V. Marone (2019): "Regulating Markups in US Health Insurance," *American Economic Journal: Applied Economics*, 11, 71–104.

Dafny, L., M. Duggan, and S. Ramanarayanan (2012): "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry," *American Economic Review*, 102, 1161–85.

Dafny, L. S. (2010): "Are Health Insurance Markets Competitive?" *American Economic Review*, 100, 1399–1431.

Dubey, P. and J. Geanakoplos (2002): "Competitive Pooling: Rothschild-Stiglitz Reconsidered," *Quarterly Journal of Economics*, 117, 1529–1570.

Einav, L. and A. Finkelstein (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 25, 115–138.

Einav, L., A. Finkelstein, and M. R. Cullen (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 125, 977–921.

Einav, L., A. Finkelstein, and N. Mahoney (2021): "The IO of Selection Markets," *NBER Working Paper 29039.*

Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013): "Selection on Moral Hazard in Health Insurance," *Amercian Economic Review*, 103, 178–219.

Fang, H. and Z. Wu (2018): "Multidimensional private information, market structure, and insurance markets," *The Rand Journal of Economics*, 49, 751–787.

Farinha Luz, V. (2017): "Characterization and Uniqueness of Equilibrium in Competitive Insurance," *Theoretical Economics*, 12.

Farinha Luz, V., P. Gottardi, and H. Moreira (2023): "Risk Classification in Insurance Markets with Risk and Preference Heterogeneity," *Review of Economic Studies, forthcoming.*

Gale, D. (1992): "A Walrasian Theory of Markets with Adverse Selection," *The Review of Economic Studies*, 59, 229–255.

Hackmann, M. B., J. T. Kolstad, and A. E. Kowalski (2015): "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *Amercian Economic Review*, 105, 1030–66.

Handel, B., I. Hendel, and M. D. Whinston (2015): "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 83, 1261–1313.

Handel, B. R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *The American Economic Review*, 103, 21643–2682.

Handel, B. R. and J. T. Kolstad (2015): "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare," *The American Economic Review*, 105, 2449–2500.

Handel, B. R., J. T. Kolstad, and J. Spinnewijn (2019): "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets," *The Review of Economics and Statistics*, 101, 326–340.

Hendren, N. (2013): "Private Information and Insurance Rejections," *Econometrica*, 81, 1713–62.

——— (2014): "Unravelling vs Unravelling: A Memo on Competitive Equilibriums and Trade inInsurance Markets," *The Geneva Risk and Insurance Review*, 2, 176–183.

Ho, K. and R. S. Lee (2017): "Insurer Competition in Health Care Markets," *Econometrica*, 85, 379–417.

Kadan, O., P. J. Reny, and J. M. Swinkels (2017): "Existence of Optimal Mechanisms in Principal-Agent Problems," *Econometrica*, 85, 769–823.

Laffont, J.-J., E. Maskin, and J.-C. Rochet (1987): *Optimal Nonlinear Pricing with Two-Dimensional Characteristics* -, Minneapolis: University of Minnesota Press, 256–66.

Landais, C., A. Nekoei, P. Nilsson, and J. S. David Seim (2021): "Risk-Based Selection in Unemployment Insurance: Evidence and Implications," *Amercian Economic Review*, 111, 1315–55.

Levy, Y. J. and A. Veiga (2022): "Optimal Contract Regulation in Selection Markets," *Woriking paper.*

Mahoney, N. and E. G. Weyl (2017): "Imperfect Competition in Selection Markets," *The Review of Economics and Statistics*, 99, 637–651.

Maskin, E. and J. Riley (1984): "Monopoly with Incomplete Information," *The RAND Journal of Economics*, 15, 171–96.

Miyazaki, H. (1977): "The Rat Race and Internal Labor Markets," *The Bell Journal of Economics*, 8, 394–418.

Mussa, M. and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–17.

Myerson, R. (1978): "Refinements of the Nash equilibrium concept," *International Journal of Game Theory*, 7, 73–80.

Polyakova, M. (2016): "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D," *American Economic Journal: Applied Economics*, 8, 165–195.

Polyakova, M. and S. P. Ryan (2021): "Subsidy Targeting with Market Power," *NBER Working Paper 26367*.

——— (2023): "Market Power and Redistribution: Evidence from the Affordable Care Act," Tech. rep., Stanford and Washington University in St Louis.

Riley, J. G. (1979): "Informational Equilibrium," *Econometrica*, 47, 331–359.

Rochet, J.-C. and L. A. Stole (2003): *The Economics of Multidimensional Screening*, Cambridge University Press, Econometric Society Monographs, advances in economics and econometrics: theory and applications ed.

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

Saltzman, E., A. Swanson, and D. Polsky (2021): "Inertia, Market Power, and Adverse Selection in Health Insurance: Evidence from the ACA Exchanges," *NBER Working Paper No. w29097*.

Scheuer, F. and K. Smetters (2018): "How Initial Conditions Can Have Permanent Effects: The Case of the Affordable Care Act," *Amercian Economic Journal: Economic Policy*, 10, 302–43.

Spinnewijn, J. (2017): "Heterogeneity, Demand for Insurance, and Adverse Selection," *American Economic Journal: Policy*, 9, 308–43.

Starc, A. (2014): "Insurer pricing and consumer welfare: evidence from Medigap," *The RAND Journal of Economics*, 45, 198–220.

Stiglitz, J. E. (1977): "Monopoly, Non-linear Pricing and Imperfect Information: the Insurance Market," *Review of Economic Studies*, 44, 407–30.

Tebaldi, P. (Forthcoming): "Estimating Equilibrium in Health Insurance Exchanges:Price Competition and Subsidy Design under the ACA," *Review of Economic Studies*.

Veiga, A. and E. G. Weyl (2016): "Product Design in Selection Markets," *The Quarterly Journal of Economics*, 131, 1007–1056.

Wilson, R. (1977): "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16, 167–207.

Appendix: Proofs

Proof of Lemma 1

Since u is twice continuously differentiable and x lies in the compact space $[0,1]$, $u(\theta,\cdot)$ has a bounded derivative and must therefore be Lipschitz continuous. So, there exists $K>0$ such that

$$|u(\theta,x)-u(\theta,x')| \leq K\left\|x-x'\right\|$$

for all θ, x, x'.

Consider two contracts x and x'. Without loss of generality, suppose that $p^*(x)>p^*(x')$. Since prices are non-negative, we must have $p^*(x)>0$ so there must be some type who prefers x to x':

$$u(\theta,x)-p^*(x) \geq u(\theta,x')-p^*(x').$$

By monotonicity of $u(\theta,\cdot)$, we must have $x>x'$. Rearrange this expression and use the Lipschitz continuity of u to obtain:

$$p^*(x)-p^*(x') \leq u(\theta,x)-u(\theta,x') = |u(\theta,x)-u(\theta,x')| \leq K\left\|x-x'\right\|.$$

The fact that p^* is Lebesgue almost everywhere differentiable follows from Rademacher's Theorem. □

Proof of Lemmas 2 and 3

See Appendix B.

Proof of Theorem 1

Fix a competitive equilibrium in which a positive mass of types are excluded. By consumer optimization (condition 2 of Definition 1), whenever (θ,x) is in the support of α^*, we must have:

$$p'(x) = \frac{\partial u}{\partial x}(\theta,x) \geq v_0(x)+\sum_{k=1}^{K}\theta_k v_k(x) \tag{15}$$

for all points of differentiability of p. Since the inequality holds for each θ picking x, it must also hold for the expected θ among those choosing x:

$$p'(x) \geq v_0(x)+\sum_{k=1}^{K} m_k(x) v_k(x), \tag{16}$$

where $m_k(x) \equiv \mathbb{E}_x[\theta_k|\alpha^*]$.

Suppose that for every $\varepsilon > 0$ there exists $x \in (0, \varepsilon)$ such that no type is indifferent between their equilibrium contract and x. Then, by condition 3 of Definition 1, the price of that contract must be zero, which implies that no type can be excluded (since any such type prefers to purchase $x > 0$ at price zero over zero coverage). Thus, for any contract x in a neighborhood of zero, there must exist some type for whom picking that contract maximizes utility. Therefore, x must also maximize some type θ's utility, satisfying condition (15). Moreover, the cost of selling x to θ must be weakly higher than the price $p(x)$. For each such x, let $\mathbf{m}(x) \equiv (m_1(x), ..., m_K(x))$ denote that type, so that (16) must again hold.

Integrate and use $p(0) = 0$ to obtain:

$$p(x) \geq \int_0^x \left[v_0(\widetilde{x}) + \sum_{k=1}^{K} m_k(\widetilde{x}) \cdot v_k(\widetilde{x}) \right] d\widetilde{x}.$$

By Conditions 1 and 3 of Definition 1,

$$p(x) \leq \mathbb{E}_x[c|\alpha^*] \leq \sum_{k=1}^{K} \mathbb{E}_x[\theta_k|\alpha^*]\kappa_k(x) = \sum_{k=1}^{K} m_k(x)\kappa_k(x)$$

(where the first inequality holds as an equality if (θ, x) is in the support of α^* and we extend the conditional expectation to assign full mass to some selection $\mathbf{m}(x)$ if there is no (θ, x) is in the support of α^*). Combine both of these inequalities to obtain:

$$\sum_{k=1}^{K} m_k(x)\kappa_k(x) \geq \int_0^x \left[v_0(\widetilde{x}) + \sum_{k=1}^{K} m_k(\widetilde{x}) \cdot v_k(\widetilde{x}) \right] d\widetilde{x}.$$

Divide both sides by $x > 0$ and rearrange:

$$\sum_{k=1}^{K} m_k(x) \frac{\kappa_k(x)}{x} - \frac{\int_0^x v_0(\widetilde{x}) d\widetilde{x}}{x} - \sum_{k=1}^{K} \frac{\int_0^x m_k(\widetilde{x}) \cdot v_k(\widetilde{x}) d\widetilde{x}}{x} \geq 0.$$

Note that this condition must hold for all $x > 0$ in a neighborhood of 0.

We consider the limit of the expression on the LHS as $x \searrow 0$. By the Fundamental Theorem of Calculus, we have $\lim_{x \searrow 0} \frac{\int_0^x v_0(\widetilde{x}) d\widetilde{x}}{x} = v_0(0)$ and $\lim_{x \searrow 0} \frac{\int_0^x m_k(\widetilde{x}) \cdot v_k(\widetilde{x}) d\widetilde{x}}{x} = m_k(0_+) v_k(0)$, where $m_k(0_+) \equiv \lim_{x \searrow 0} m_k(x)$. Therefore, the condition above for x in a neighborhood of $x = 0$ requires:

$$0 \geq v_0(0) + \sum_{k=1}^{K} m_k(0_+) \left[v_k(0) - \kappa_k'(0) \right],$$

which contradicts part 3 from Assumption 2'. Therefore, there is no equilibrium with a positive mass of excluded types.

□

Proof of Theorem 2:

The proof of Theorem 2 has three steps. First, we construct a family of subsets of the type space $\Theta^\xi \subset \Theta \equiv [\underline{\theta}_1, \overline{\theta}_1] \times ... \times [\underline{\theta}_K, \overline{\theta}_K]$ that converges to the whole type space Θ as $\xi \searrow 0$. Second, we show that, for any mechanism that excludes a zero measure set of types in the original type space Θ, a small enough increase in prices raises the firm's profits in the economy with type space Θ^ξ uniformly in ξ. Lastly, since this increase is uniform in ξ, it also applies to Θ.

Without loss of generality, we can consider an economy with type space equal to $\mathbb{R}^K_+$ by assigning zero mass to all types outside Θ.[25] Let α be an incentive compatible allocation and let $\mathcal{U}(\theta)$ be the informational rent of type θ associated to α. Since $\mathcal{U}$ is differentiable a.e. along rays, we can write

$$\mathcal{U}(\theta) = \mathcal{U}(0) + \int_0^1 \frac{d}{dr}\mathcal{U}(r\theta)dr. \tag{17}$$

By the envelope condition,

$$\frac{d}{dr}\mathcal{U}(r\theta) = \theta \cdot \mathbb{E}\left[\nabla_\theta u(\widetilde{\theta}, \widetilde{x})|\alpha, \widetilde{\theta} = r\theta\right], \tag{18}$$

where the Leibniz integral rule is justified by the Dominated Convergence Theorem.

By the monotonicity of $u(\cdot, x)$, it follows that $\mathcal{U}$ is non-decreasing, so the lowest value of $\mathcal{U}$ is attained at $\theta = 0$. Thus, any mechanism that maximizes the monopolist's profits must have $\mathcal{U}(0) = 0$. Also by monotonicity, if $\mathcal{U}(\theta) = 0$ for any $\theta >> \underline{\theta}$, then $\mathcal{U}(\widetilde{\theta}) = 0$ for all $\widetilde{\theta} \in \{\widetilde{\theta}; \underline{\theta} \leq \widetilde{\theta} \leq \theta\}$, which is a positive-measure set. Therefore, in any incentive-compatible mechanism in which the set of excluded types has measure zero, any excluded type $(\theta_1, ..., \theta_K)$ must have at least one coordinate equal to the lowest value $\underline{\theta}_k$.

To introduce the family of subsets of the type space mentioned previously, it is helpful to introduce some terminology. For each set of vectors $\{a_1, ..., a_K\} \subset \mathbb{R}^K$, let

$$C(a_1, ..., a_K) \equiv \left\{ \sum_{k=1}^K \lambda_k a_k : \ \lambda_k \geq 0, \ k = 1, ..., K \right\}$$

denote the convex cone generated by $\{a_1, ..., a_K\}$. For each $\xi > 0$, let $e_k^\xi = (\xi, ..., 1, ..., \xi)$ denote the vector with $\xi \in \mathbb{R}$ in all but the k-th coordinate, where it equals 1. Finally, for

[25] Every incentive-compatible mechanism with type space Θ can be mapped into an incentive compatible mechanism with type space $\mathbb{R}^K_+$ such that its restriction to Θ coincides with the original mechanism. For instance, we can assign for each type $\theta \notin \Theta$ a contract that maximizes this type's payoff among those offered to types in Θ. By construction, this new mechanism is also incentive compatible. Since all such types have zero mass, this extension does not affect the firm's profits.

each $\xi > 0$, let

$$\Theta^{\xi} \equiv \Theta \cap \left(\underline{\theta} + C(e_1^{\xi}, ..., e_K^{\xi})\right)$$

denote the ξ-perturbed type space. Note that $\Theta^{\xi} \subseteq \Theta$, and Θ^{ξ} converges to $\Theta^0 = \Theta$ as $\xi \searrow 0$ in the Hausdorff distance.

An **economy** is defined by a pair (Θ, μ). A ξ-perturbed economy consists of an economy with the ξ-perturbed type space Θ^{ξ} and a distribution of types given by the restriction of μ to Θ^{ξ}. We first consider mechanisms in ξ-perturbed economies. The first lemma shows that any incentive-compatible mechanism that excludes a set of measure zero of types can only exclude at most the lowest type $\underline{\theta}$.

Lemma 4. *Fix an incentive-compatible mechanism for the $\xi-$perturbed economy, where $\xi \in (0,1)$. Let E be the set of excluded types and suppose $\mu(E) = 0$. Then, $E \subseteq \{\underline{\theta}\}$.*

Recall that, by monotonicity, if θ is excluded then any $\widetilde{\theta} << \theta$ must also be excluded. The proof below shows that, in any $\xi-$perturbed economy, for any $\theta \neq \underline{\theta}$, the set of types $\widetilde{\theta} \in \Theta^{\xi}$ such that $\widetilde{\theta} << \theta$ has positive measure. Therefore, the firm cannot exclude a measure zero of types while excluding any type other than $\underline{\theta}$.

Proof. By the definition of Θ^{ξ}, if $\theta \in \Theta^{\xi} \backslash \{\underline{\theta}\}$, then $\theta = \underline{\theta} + \sum_{k=1}^{K} \lambda_k e_k^{\xi}$, for some $\lambda_1^0, ..., \lambda_K^0 \geq 0$. Moreover, since $\theta \neq \underline{\theta}$, we must have $\lambda_{k'}^0 > 0$ for some k'. Let $\widehat{\theta} \equiv \underline{\theta} + \sum_{k \neq k'} \lambda_k^0 e_k^{\xi} + \frac{\lambda_{k'}^0}{2} e_{k'}^{\xi} \in \Theta^{\xi} \backslash \{\underline{\theta}\}$. Since $\xi > 0$, each coordinate of the vector $\sum_{k=1}^{K} \lambda_k e_k^{\xi}$ is strictly positive. Therefore,

$$\widehat{\theta} = \underline{\theta} + \sum_{k \neq k'} \lambda_k^0 e_k^{\xi} + \frac{\lambda_{k'}^0}{2} e_{k'}^{\xi} << \underline{\theta} + \sum_{\forall k} \lambda_k^0 e_k^{\xi} = \theta.$$

Since the expression on the LHS is linear (and therefore continuous) in $(\lambda_1^0, ..., \lambda_K^0)$, there exists $r > 0$ small enough such that

$$\underline{\theta} + \sum_{k=1}^{K} \lambda_k e_k^{\xi} << \theta, \text{ if } \lambda_k < \lambda_k^0 + r, \text{ for all } k \neq k' \text{ and } \lambda_{k'} < \frac{\lambda_{k'}^0}{2} + r. \tag{19}$$

Let $\Lambda \equiv \{(\lambda_1, ..., \lambda_K) \in \mathbb{R}_{++}^K$ such that (19) holds$\}$. It is straightforward to see that Λ is a non-empty open subset of $\mathbb{R}_{+}^K$. Finally, if $\xi \in (0,1)$, then the affine transformation $T : \mathbb{R}^K \to \mathbb{R}^K$ given by $T(\lambda_1, ..., \lambda_K) = \underline{\theta} + \sum_{k=1}^{K} \lambda_k e_k^{\xi}$ is a bijection (i.e., an affine isomorphism). Hence, $T(\Lambda) \subset \Theta^{\xi}$ is a non-empty open set such that $\widehat{\theta} << \theta$ for all $\widehat{\theta} \in T(\Lambda)$. Therefore, if type θ is excluded, every type in $T(\Lambda)$ is also excluded, which is a set with positive measure. $\square$

Let $\mathbb{E}_{\theta}[S|\alpha] \equiv \mathbb{E}\left[S(\tilde{\theta}, \tilde{x})|\alpha, \tilde{\theta} = \theta\right]$ denote the expected surplus according to measure α conditional on type θ. We use analogous notation for the conditional expectation of other functions. For each $\widetilde{\Theta} \subset \mathbb{R}^K_+$, let $\pi_\alpha(\widetilde{\Theta})$ denote the firm's expected profit under allocation α among types in $\widetilde{\Theta}$:

$$\pi_\alpha(\widetilde{\Theta}) \equiv \int_{\widetilde{\Theta}} \left\{ \mathbb{E}_\theta \left[S|\alpha \right] - \mathcal{U}(\theta) \right\} f(\theta) d\theta. \tag{20}$$

For each $\epsilon > 0$, let $\Theta^\xi_\epsilon \equiv \{\theta \in \Theta^\xi : \ ||\theta - \underline{\theta}|| < \epsilon\}$ denote the set of types in the perturbed type space that are close to the lowest type. The lemma below provides an upper bound on the firm's profits restricted to Θ^ξ_ϵ:

Lemma 5. *For every $\epsilon > 0$ and $\xi > 0$, we have*

$$\pi_\alpha(\Theta^\xi_\epsilon) \leq \int_{\Theta^\xi_\epsilon} \left\{ \mathbb{E}_\theta \left[S|\alpha \right] - \kappa \frac{g(\theta)}{f(\theta)} \mathbb{E}_\theta \left[u|\alpha \right] \right\} f(\theta) d\theta, \tag{21}$$

where $g(\theta) \equiv \int_1^\infty t^{K-1} f(t\theta) dt$.

Proof. By the envelope condition (18) and the Assumption 8 for type $r\theta$, we have:

$$\frac{d}{dr}\mathcal{U}(r\theta) = \theta \cdot \mathbb{E}_\theta \left[\nabla_\theta u(r\theta, x)|\alpha \right] \geq \frac{\kappa}{r} \mathbb{E}_\theta \left[u(r\theta, x)|\alpha \right],$$

for all $\theta \in \Theta_\epsilon$. Substitute in (17) to obtain:

$$\mathcal{U}(\theta) \geq \int_0^1 \frac{\kappa}{r} \mathbb{E} \left[u(\widetilde{\theta}, \tilde{x}) \mid \alpha, \widetilde{\theta} = r\theta \right] dr.$$

Taking expectations gives:

$$\begin{aligned} \textstyle\int_{\Theta^\xi_\epsilon} \mathcal{U}(\theta) f(\theta) d\theta \ &\geq \textstyle\int_{\Theta^\xi_\epsilon} \left\{ \int_0^1 \frac{\kappa}{r} \mathbb{E} \left[u(\widetilde{\theta}, \widetilde{x}) \mid \alpha, \widetilde{\theta} = r\theta \right] dr \right\} f(\theta) d\theta \\ &= \textstyle\int_0^1 \left[\int_{\Theta^\xi_\epsilon} \frac{\kappa}{r} \mathbb{E} \left[u(\widetilde{\theta}, \widetilde{x}) \mid \alpha, \widetilde{\theta} = r\theta \right] f(\theta) d\theta \right] dr. \end{aligned} \tag{22}$$

Apply the change of variables $\widehat{\theta} = r\theta$ and use the fact that $d\widehat{\theta} = r^K d\theta$ to get

$$\int_{\Theta^\xi_\epsilon} \frac{1}{r} \mathbb{E} \left[u(\widetilde{\theta}, \widetilde{x}) \mid \alpha, \widetilde{\theta} = r\theta \right] f(\theta) d\theta = \int_{\Theta^\xi_\epsilon} \frac{1}{r^{K+1}} \mathbb{E}_{\widehat{\theta}} \left[u|\alpha \right] f\left(\frac{\widehat{\theta}}{r} \right) d\widehat{\theta}$$

where, as defined previously, $\mathbb{E}_{\widehat{\theta}} \left[u|\alpha \right] \equiv \mathbb{E} \left[u(\widetilde{\theta}, \widetilde{x}) \mid \alpha, \widetilde{\theta} = \widehat{\theta} \right]$. Substituting in (22), gives:

$$\begin{aligned} \textstyle\int_{\Theta^\xi_\epsilon} \mathcal{U}(\theta) f(\theta) d\theta \ &\geq \textstyle\int_0^1 \left[\int_{\Theta^\xi_\epsilon} \frac{\kappa}{r^{K+1}} \mathbb{E}_{\widehat{\theta}} \left[u|\alpha \right] f\left(\frac{\widehat{\theta}}{r} \right) d\widehat{\theta} \right] dr \\ &= \textstyle\int_{\Theta^\xi_\epsilon} \frac{\kappa}{r^{K+1}} \mathbb{E}_{\widehat{\theta}} \left[u|\alpha \right] \left[\int_0^1 \frac{1}{r^{K+1}} f\left(\frac{\widehat{\theta}}{r} \right) dr \right] d\widehat{\theta}. \end{aligned} \tag{23}$$

Substitute $t = 1/r$ to obtain:

$$\int_0^1 \frac{1}{r^{K+1}} f\left(\frac{\widehat{\theta}}{r}\right) = \int_1^\infty t^{K-1} f(t\widehat{\theta}) dt.$$

Plugging this expression back into (23) gives:

$$\int_{\Theta_\epsilon^\xi} \mathcal{U}(\theta) f(\theta) d\theta \geq \kappa \int_{\Theta_\epsilon^\xi} E_\theta \left[u|\alpha\right] g(\theta) d\theta, \tag{24}$$

where $g(\theta)$ is as defined in the statement of the lemma. Substitute (24) in (20) to obtain (21). □

Lemma 6. *Fix an incentive-compatible mechanism (p^ξ, α^ξ) for the $\xi-$perturbed economy, where $\xi \in (0,1)$. If the set of excluded types has measure zero ($\alpha(\Theta, 0) = 0$) then the mechanism does not maximize the firm's profits.*

Proof. Let $\mathcal{U}$ denote the informational rent function associated with the mechanism. By Lemma 4, this mechanism can only exclude type $\underline{\theta}$. Now consider a small uniform price increase of $\delta > 0$ for all $x \neq 0$. There are two effects. First, types who were getting surplus below the price increase choose not to participate, i.e., types in

$$A_\delta^\xi \equiv \{\theta \in \Theta^\xi; \mathcal{U}(\theta) < \delta\}.$$

Second, the firm increases profits by δ from all types who remain, i.e., types in $\Theta^\xi \backslash A_\delta$. So the increase in the firm's profits from those who remain is equal to

$$\delta. \int_{\Theta^\xi \backslash A_\delta^\xi} f(\theta) d\theta. \tag{25}$$

For the first effect, we will show that it also leads to a positive gain for the firm when δ is small enough. For this, we claim that there exists $\epsilon > 0$, uniform in ξ, such that

$$\left[1 - \kappa \frac{g(\theta)}{f(\theta)}\right] \mathbb{E}_\theta \left[u|\alpha\right] - \mathbb{E}_\theta \left[c|\alpha\right] < 0, \tag{26}$$

for all $\theta \in \Theta_\epsilon^\xi$. Indeed, let $\underline{x} \equiv \inf\{x : x \in \text{ support } \alpha(\theta, \cdot);\ \theta \in \Theta^\xi\}$ denote the lowest coverage among the coverages in the support of the mechanism distributions. Recall that by the definition of η in the text, we have $\eta < \left[1 - \kappa \frac{\underline{g}}{f(\underline{\theta})}\right]^{-1}$, where $\underline{g} := \inf\{g(\theta) :\ \theta \in \Theta_\epsilon\} > 0$.

If $\underline{x} > 0$, Assumption 3' implies that

$$\left[1 - \kappa \frac{\underline{g}}{f(\underline{\theta})}\right] u(\underline{\theta}, \underline{x}) - c(\underline{\theta}, \underline{x}) < 0.$$

By the continuity of the above expression at $(\underline{\theta}, \underline{x})$ and the definition of $\underline{x}$, we can find ϵ sufficiently small so that inequality (26) holds for all $\theta \in \Theta_\epsilon$. Since $\Theta_\epsilon^\xi \subset \Theta_\epsilon$, the choice of ϵ is uniform in ξ.

If $\underline{x} = 0$, Assumption 3' again implies that

$$\left[1 - \kappa \frac{g}{f(\underline{\theta})}\right] \frac{\partial u}{\partial x}(\underline{\theta}, 0) - \frac{\partial c}{\partial x}(\underline{\theta}, 0) < 0.$$

By the continuity of the above expression at $(\underline{\theta}, 0)$, we can find $\epsilon > 0$ sufficiently small such that (26) holds for all $\theta \in \Theta_\epsilon$. Thus, as before, the choice of ϵ is uniform in ξ.

Integrating (26) on Θ_ϵ^ξ, we conclude that the right hand side of (21) (see Lemma 5) restricted to Θ_ϵ^ξ is negative, where ϵ is determined in the two cases considered above ($\underline{x} > 0$ and $\underline{x} = 0$) and is uniform on ξ. Finally, since the only excluded type is $\underline{\theta}$, for such $\epsilon > 0$, there exists $\delta > 0$ sufficiently small such that $A_\delta^\xi \subset \Theta_\epsilon^\xi$. By increasing the price by δ, the firm excludes types that lead to ex-post losses (i.e., types in $A_\delta^\xi \subset \Theta_\epsilon^\xi$) and gain extra profit with one that are still participating (i.e., types in $\left(A_\delta^\xi\right)^c$). Therefore, the firm can ensure a positive gain of at least (25), which is uniformly (in ξ) bounded away from zero, which implies the result. □

Proof of Theorem 2. Fix an incentive compatible mechanism that excludes a zero measure set of types in Θ. For each $\xi \in (0,1)$, consider the restriction of this mechanism to Θ^ξ. By Lemma 6, there exists a uniform increase in price $\delta > 0$ that ensures that the firm increases profits by at least (25) in each $\xi-$perturbed economy. Since this gain is uniform in $\xi > 0$, taking $\xi \searrow 0$, this uniform increase in prices also increases profits by at least (25) for $\xi = 0$, which concludes the proof. □

Distortions at the Top and Proof of Theorem 3:

Recall that in the one-dimensional model, the competitive equilibrium is characterized by zero profits, separation, and no distortion at the top.

The proof of Proposition 3 used these three conditions to show that the allocation becomes infinitely steep as we approach the highest type: $\lim_{\theta \nearrow \bar{\theta}} \dot{x}(\theta) = +\infty$. This condition can alternatively be written in terms of the type assignment function $\widetilde{\theta} : x(\Theta) \to [0,1]$, which indicates the type that chooses each contract $\widetilde{\theta} = x^{-1}$. The equivalent statement is that the type assignment function becomes flat as we approach the efficient contract: $\lim_{x \nearrow \bar{x}} \dot{\widetilde{\theta}}(x) = 0$.

With multiple dimensions, we cannot have full separation (as the dimension of types exceeds the dimension of instruments), so the allocation function is no longer an injection and its inverse (the type assignment function) is not well defined. Instead, we work with the expectation of types choosing each contract.

For a fixed allocation α, let $m_i(x) \equiv \mathbb{E}_x[\theta_i|\alpha]$ denote the expected i-th dimension of the types choosing contract x and let $\widetilde{x}_+ \equiv \sup\{x : (\theta, x) \in \text{supp } \alpha\}$ denote the highest contract chosen by some type.

Definition 2. An allocation α has **separation at the top** if $\lim_{x \nearrow \widetilde{x}_+} m_i(x) = \overline{\theta}_i$ for all $i = 1, ..., K$.

Separation at the top means that, as we approach the highest contract chosen by some type, the expected type picking that contract converges to the highest type. It is immediately satisfied in the one-dimensional model with single crossing, in which case the expected type coincides with the type assignment function $\tilde{\theta}$.

Definition 3. An allocation α has **no insufficient coverage at the top** if

$$\mathbb{E}_x\left[\frac{\partial S}{\partial x}(\theta, \widetilde{x}_+) \mid \alpha\right] = \sum_{i=1}^{K} m_i(\widetilde{x}_+)\left[\phi_i'(\widetilde{x}_+) - \kappa_i'(\widetilde{x}_+)\right] \leq 0.$$

It has **no distortion at the top** if the inequality above holds as an equality.

Recall that the surplus is a concave function of coverage. No insufficient coverage at the top means that the highest coverage $\widetilde{x}_+$ is weakly higher than the level that maximizes the average surplus (among types who pick that coverage). No distortion at the top means that the highest coverage maximizes the average surplus for those who pick the highest coverage. No distortion at the top is also trivially satisfied in the one-dimensional model with single crossing.

Definition 4. An allocation α has **high incentives to cream skim at the top** if $\lim_{x \nearrow \widetilde{x}_+} \frac{m_i(x) - \overline{\theta}_i}{x - \widetilde{x}_+} = 0$ for all $i = 1, ..., K$.

High incentives to cream skim at the top is related to the slope of the type assignment function at the top. Recall that in the one-dimensional model, the type assignment function becomes flat as we approach the highest contract: $\lim_{x \nearrow \widetilde{x}_+} \dot{\widetilde{x}}(x) = 0$. This condition is equivalent to the allocation function becoming infinitely steep as we approach the highest type. Definition 4 is the multidimensional counterpart of this condition, using expected types rather than the allocation function, which is not well defined in multidimensional settings.

As discussed in the text, fact that the type assignment function becomes arbitrarily flat as we approach the efficient contract arises from the incentives to cream skim. It is very easy to convince a type who is close to his efficient contract to deviate by offering a lower coverage at a discounted price. So, to prevent cream skimming close to the top, firms must substantially reduce coverage, making the allocation function very steep or, equivalently, the type allocation function very flat. Definition 4 is analogous, except that it replaces the unique type picking each contract in the one-dimensional model by the average type.

The lemma below shows that in any deterministic equilibrium, those three conditions are equivalent.

Lemma 7. *Let (p, α) be a competitive equilibrium. Suppose α is deterministic and let $\widetilde{x}$ denote its on-path coverage. The following statements are equivalent:*

(a) $\widetilde{x}$ has separation at the top;

(b) $\widetilde{x}$ has no insufficient coverage at the top;

(c) $\widetilde{x}$ has high incentives to cream skim at the top;

(d) $\widetilde{x}$ has no distortion at the top.

Proof. (a $\implies$ b) By single crossing and Condition 3 from Definition 1, the highest type $\overline{\theta}$ must be indifferent between his equilibrium contract $\widetilde{x}_+$ and each off-path contract $x > \widetilde{x}_+$:

$$p(x) - p(\tilde{x}_+) = \sum_{i=1}^{K} \overline{\theta}_i \left[\phi_i(x) - \phi_i(\tilde{x}_+)\right].$$

Moreover, the profits from each such contract cannot be positive:

$$p(x) \leq \sum_{i=1}^{K} \overline{\theta}_i \kappa_i(x). \tag{27}$$

By the zero profits condition of on-path contract $\widetilde{x}_+$ and separation at the top, we have $p(\widetilde{x}_+) = \sum_{i=1}^{K} \overline{\theta}_i \kappa_i(\widetilde{x}_+)$. Therefore, the previous equation becomes:

$$p(x) = \sum_{i=1}^{K} \overline{\theta}_i \left[\phi_i(x) - \phi_i(\widetilde{x}_+) + \kappa_i(\widetilde{x}_+)\right].$$

Using (27), we obtain:

$$\sum_{i=1}^{K} \overline{\theta}_i \left[\phi_i(x) - \phi_i(\widetilde{x}_+) + \kappa_i(\widetilde{x}_+)\right] \leq \sum_{i=1}^{K} \overline{\theta}_i \kappa_i(x)$$

$$\therefore \underbrace{\sum_{i=1}^{K} \overline{\theta}_i \left[\phi_i(x) - \kappa_i(x)\right]}_{S(\overline{\theta}, x)} \leq \underbrace{\sum_{i=1}^{K} \overline{\theta}_i \left[\phi_i(\widetilde{x}_+) - \kappa_i(\widetilde{x}_+)\right]}_{S(\overline{\theta}, \widetilde{x}_+)}.$$

Therefore, the highest contract on the equilibrium path cannot have less coverage than the surplus-maximizing contract, i.e., $\widetilde{x}$ has excess coverage at top.

(b $\Longrightarrow$ c) Suppose that $\widetilde{x}$ has excess coverage at the top. Let p be the equilibrium price. Consumer optimization requires that at all points of differentiability of p, we have

$$p'(x) = u_x(\theta, x) = \sum_{i=1}^{K} \theta_i \phi_i'(x) = \sum_{i=1}^{K} m_i(x)\phi_i'(x), \tag{28}$$

where the first equality uses the necessary FOC, the second substitutes the expression for the utility function, and the third takes expectations among all types who pick contract x.

Recall the zero profits condition:

$$p(x) = E_x[c(\theta, x)|\alpha] = \sum_{i=1}^{K} m_i(x)\kappa_i(x). \tag{29}$$

Differentiate this condition with respect to x, take the limit as $x \nearrow \widetilde{x}_+$, and use the previous equation to obtain:

$$\sum_{i=1}^{K} m_i(\widetilde{x}_+)\phi_i'(\widetilde{x}_+) = p'(\widetilde{x}_+) = \sum_{i=1}^{K} m_i(\widetilde{x}_+)\kappa_i'(\widetilde{x}_+) + \sum_{i=1}^{K} \lim_{x \to \widetilde{x}_+} \frac{m_i(x) - m_i(\widetilde{x}_+)}{x - \widetilde{x}_+}\kappa_i(x). \tag{30}$$

Since there is excess of coverage at the top, the expression in (30) cannot exceed $\sum_{i=1}^{K} m_i(x)\kappa_i'(x)$, so that:

$$\sum_{i=1}^{K} \lim_{x \nearrow \widetilde{x}_+} \frac{m_i(x) - m_i(\widetilde{x}_+)}{x - \widetilde{x}_+}\kappa_i(x) \leq 0.$$

By the single-crossing property in each dimension, $\widetilde{x}$ must be weakly increasing in each dimension. Hence, $m_i(x) \leq m_i(\widetilde{x}_+)$, for all $x < \widetilde{x}_+$ and, consequently, $\frac{m_i(x)-m_i(\widetilde{x}_+)}{x-\widetilde{x}_+} \geq 0$ and $\kappa_i(\widetilde{x}_+) > 0$ by assumption, then

$$\lim_{x \nearrow \widetilde{x}_+} \frac{m_i(x) - m_i(\widetilde{x}_+)}{x - \widetilde{x}_+} = 0,$$

for all i.

(c $\Longrightarrow$ d) Using the same argument of (a), we will get equality (30). If $\widetilde{x}$ has high incentive to cream skim at top, the last term on the right hand side of (30) is zero. Hence, it is immediate that $\widetilde{x}$ has no distortion at the top.

(d $\Longrightarrow$ a) From Definition 3 it follows that (d) implies (b). We have already proved that (b) implies (c), which, from Definition 4, obviously implies (a). □

Let (p, α) be an allocation satisfying condition 3 from Definition 1. Then, for any contract

$x' \in X$ with strictly positive price, there exists (θ, x) in the support of α such that

$$u(\theta, x) - p(x) = u(\theta, x') - p(x') \text{ and } c(\theta, x') \geq p(x').$$

This condition trivially holds if x' is on the equilibrium path. For off-path contracts (i.e. those that are not chosen by any type according to α), prices must be: (i) low enough to make some type θ indifferent between deviating to x' or keeping the equilibrium contract x, and (ii) a deviation by this type would make firms (weakly) lose money.

Suppose p is not differentiable at some contract x that is on the equilibrium path, so that $p'(x_-) \neq p'(x_+)$. Since x is chosen by some type, it must be the case that $p'(x_-) < p'(x_+)$, so there is a small neighborhood of contracts around x that are not chosen by any type. The lemma below identifies the types who are most willing to deviate from any such contract:

Lemma 8. *Let (p, α) be a competitive equilibrium and suppose p is not differentiable at some on-path contract x. Then, there exists $\varepsilon > 0$ such that for each $x' \in (x - \varepsilon, x)$, there exists θ with*

$$\theta \in \arg \min_{\widetilde{\theta} \in \Theta_x^{p'(x_+)}} \sum \widetilde{\theta}_i \phi_i''(x)$$

and

$$u(\theta, x) - p(x) = u(\theta, x') - p(x') \text{ and } p(x') \leq c(\theta, x'). \tag{31}$$

Next, we establish that strong adverse selection rules out pooling by types with different marginal utility of coverage:

Lemma 9. *Suppose Assumption 9 holds and let $\widetilde{x}(\cdot)$ be the on-path coverage in a deterministic competitive equilibrium. If $\widetilde{x}(\theta) = \widetilde{x}(\widetilde{\theta}) = x$ then $\sum \theta_i \phi_i'(x) = \sum \widetilde{\theta}_i \phi_i'(x)$.*

Proof. At any point of differentiability of p, the result follows from the necessary first-order condition. We claim that p cannot have kinks.

To see this, consider a point x in which p is not differentiable. Lemma 8 determines the the consumer types who are most willing to deviate. By condition 3 of Definition 1, we need to verify that the price obtained by the indifference condition (31) for any such type exceeds the cost of selling to that type. By strong adverse selection (Assumption 9), this type is also the one with the lowest cost among those purchasing x. By the zero profits condition for (on-path) contract x, it follows that, for x' close enough to x, the price exceeds the cost. □

A direct implication of Lemma 9 is that any competitive equilibrium with a deterministic allocation, prices are differentiable at all points in which they are positive.

Lemma 10. *Suppose Assumption 9 holds and let (p, α) be a deterministic competitive equilibrium. Then, α has separation at the top.*

Proof. We first show that $p'(\cdot)$ does not have a kink at the top. Suppose first that $\lim_{x \searrow \widetilde{x}_+} p'(x) > \lim_{x \nearrow \widetilde{x}_+} p'(x)$, so that no type chooses contracts in a small neighborhood around $\widetilde{x}_+$. Suppose, to obtain a contradiction, that a positive mass of types different from $\overline{\theta}$ picks $\widetilde{x}_+$. By single crossing, the contracts $x < \widetilde{x}_+$ in this neighborhood of $\widetilde{x}_+$ must be priced at the indifference curve of some type $\theta < \overline{\theta}$ in that pool, which contradicts Lemma 9.

So we must have $\lim_{x \searrow \widetilde{x}_+} p'(x) \leq \lim_{x \nearrow \widetilde{x}_+} p'(x)$. Of course, if this holds as a strict inequality, no one would pick $\widetilde{x}_+$ and some types would instead prefer $x > \widetilde{x}_+$, contradicting the fact that it is the highest contract on path. So the solution cannot have a kink at $\widetilde{x}_+$. Since there is no kink at $\widetilde{x}_+$, all types θ choosing this contract must satisfy:

$$p'(\widetilde{x}_+) = \sum_{i=1}^{K} \theta_i \phi_i'(\widetilde{x}_+),$$

which cannot simultaneously hold for $\overline{\theta}$ and $\theta < \overline{\theta}$.

From the consumer's optimality condition calculated at the top type $\overline{\theta}$ and since $p'(\cdot)$ is continuous at $\widetilde{x}_+$ we get

$$\sum_{i=1}^{K} \lim_{x \nearrow \widetilde{x}_+} m_i(x)\phi_i'(x) = \sum_{i=1}^{K} \overline{\theta}_i \lim_{x \nearrow \widetilde{x}_+} \phi_i'(x).$$

Since $m_i(\widetilde{x}_+) \leq \overline{\theta}_i$, for all i, we must have that $\lim_{x \to \widetilde{x}_+} m_i(x) = \overline{\theta}_i$. That is, $m_i(x)$ is continuous at $\widetilde{x}_+$ and $m_i(\widetilde{x}_+) = \overline{\theta}_i$. □

Lemma 11. *Consider a deterministic competitive equilibrium with on-path coverage $\widetilde{x}$.*

(a) Suppose $m_i(x) = \overline{\theta}_i$ for all i for x in a neighborhood of $\widetilde{x}_+$. Then, there exists $\xi > 0$ such that $\widetilde{x}(\theta) < \widetilde{x}_+ - \xi$, for almost all $\theta \in \Theta \backslash \{\overline{\theta}\}$.

(b) Suppose there exists i_0 such that $m_{i_0}(\cdot)$ is not equal to $\overline{\theta}_{i_0}$ (constant) in a neighborhood of $\widetilde{x}_+$. Then, there exists i such that the ratio $R_{j,i}(x) := \frac{\overline{\theta}_j - m_j(x)}{\overline{\theta}_i - m_i(x)}$ is uniformly bounded for all $j \neq i$.

Proof. (a) Let $\xi > 0$ be such that $m_i(x) = \overline{\theta}_i$, for all $x \in [\widetilde{x}_+ - \xi, \widetilde{x}_+]$. Suppose, in order to obtain a contradiction, that $\Pr[\widetilde{x}(\theta) \in [\widetilde{x}_+ - \xi, \widetilde{x}_+)]$ is positive (where the probability is with respect to α). This implies that

$$\sum_{i=1}^{K} \overline{\theta}_i > E\left[\sum_{i=1}^{K} \theta_i | \widetilde{x}(\theta) \in [\widetilde{x}_+ - \xi, \widetilde{x}_+]; \alpha\right].$$

However, by the law of interacted expectation we have that the conditional expectation on right hand side is

$$\frac{\int_{\tilde{x}_+-\xi}^{\tilde{x}_+}\sum_{i=1}^K \mathbb{E}_x[\theta_i|\alpha]d\overline{\alpha}(x)}{\int_{\tilde{x}_+-\xi}^{\tilde{x}_+} d\overline{\alpha}(x)},$$

where α is the equilibrium allocation, whose realization is $\tilde{x}(\cdot)$, and $\overline{\alpha}(x) = \int \alpha(x,\theta)f(\theta)d\theta$ is the marginal distribution. Notice that $m_i(x) = \mathbb{E}_x[\theta_i|\alpha] = \overline{\theta}_i$, for all i, which gives us a contradiction.

(b) If $K = 2$, then this condition is trivially satisfied because either $R_{1,2}(x)$ or $R_{2,1}(x)$ is bounded since we cannot have both the numerator and denominator identical to zero in a neighborhood of $\tilde{x}_+$. By induction, suppose that the result is true for some K. Take a model with $K+1$ dimensions. By induction, without loss of generality we can assume that the result holds for the first K dimensions. If $R_{K+1,i}(x)$ is bounded, then the condition holds for the model of $K+1$ dimensions. If not, we have that $R_{K+1,i}(x)$ is unbounded as $x \to \tilde{x}_+$. Since $R_{j,K+1}(x) = R_{j,i}(x)/R_{K+1,i}(x)$, then by the induction hypothesis $R_{j,K+1}(x)$ must be bounded for all $j \neq K+1$,which concludes the proof. □

Next, we adapt the bounds from the proof of Theorem 2 to obtain an expression for the firm's profits restricted to contracts close to the top.

Proof of Theorem 3:

By the same argument as in the proof of Theorem 2 (and the homogeneity of degree one of $u(\cdot, x)$), the monopolistic profit restricted to $N_{\tilde{x},\varepsilon}$ is

$$\pi(x) = \int_{N_{\tilde{x},\varepsilon}} \left[S(\theta, x(\theta)) - \frac{g(\theta)}{f(\theta)}u(\theta, x(\theta))\right] f(\theta)d\theta,$$

where $g(\theta) \equiv \int_1^\infty t^{K-1}f(t\theta)dt$. The Gateaux derivative in the direction $h > 0$ equals:

$$\delta\pi(x,h) = \int_{N_{\tilde{x},\varepsilon}} \left[\frac{\partial S}{\partial x}(\theta, x(\theta)) - \frac{g(\theta)}{f(\theta)}\frac{\partial u}{\partial x}(\theta, x(\theta))\right] h(\theta)f(\theta)d\theta. \tag{32}$$

By Proposition 10 and Lemma 7 we have that $\frac{\partial S}{\partial x}(\overline{\theta}, \tilde{x}_+) = 0$. By the mean value theorem, there exists $z(\theta) \in [x(\theta), \tilde{x}_+]$ such that

$$\frac{\partial S}{\partial x}(\theta, x(\theta)) = \frac{\partial S}{\partial x}(\theta, \tilde{x}_+) + \frac{\partial^2 S}{\partial x^2}(\theta, z(\theta))\,[x(\theta) - \tilde{x}_+]$$

and there exists $\tilde{\theta}_i \in [\theta_i, \overline{\theta}_i]$ such that

$$\frac{\partial S}{\partial x}(\theta, \tilde{x}_+) = \frac{\partial^2 S}{\partial x \partial\theta_i}(\tilde{\theta}_i, \theta_{-i}, \tilde{x}_+)(\theta_i - \overline{\theta}_i).$$

By the assumption (14), we have

$$\delta\pi(x,h) \geq \int_{N_{\tilde{x},\varepsilon}} \left[\frac{\partial S}{\partial x}(\theta,\tilde{x}_+) - \gamma(x(\theta)-\tilde{x}_+) - \frac{g(\theta)}{f(\theta)}\frac{\partial u}{\partial x}(\theta,x(\theta))\right] h(\theta)f(\theta)d\theta.$$

By Lemma 11, we have two possibilities: (i) there exists i such that $m_i(x) < \overline{\theta}_i$ for all $x < \tilde{x}_+$ and

$$m_i'(\tilde{x}_+) = \lim_{x\to\tilde{x}_+} \frac{m_i(x)-\overline{\theta}_i}{x-\tilde{x}_+} = 0;$$

or (ii) $m_i(x) = \overline{\theta}_i$ in a neighborhood of $\tilde{x}_+$, for all i. Let us consider each case:

(i) Since $x(\theta) < \tilde{x}(\theta) < \tilde{x}_+$we have

$$\begin{aligned}\delta\pi(x,h) > \;& \textstyle\int_{N_{\tilde{x},\varepsilon}} \frac{\partial^2 S}{\partial x\partial\theta_i}(\tilde{\theta}_i,\theta_{-i},\tilde{x}_+)\left(\frac{\theta_i - m_i(\tilde{x}(\theta))}{\overline{\theta}_i - m_i(\tilde{x}(\theta))} - 1\right)[\overline{\theta}_i - m_i(\tilde{x}(\theta))]h(\theta)f(\theta)d\theta \\ & + \textstyle\int_{N_{\tilde{x},\varepsilon}}\left[\gamma\frac{\tilde{x}_+ - \tilde{x}(\theta)}{\overline{\theta}_i - m_i(\tilde{x}(\theta))} - l\frac{g(\theta)}{\overline{\theta}_i - m_i(\tilde{x}(\theta))}\right][\overline{\theta}_i - m_i(\tilde{x}(\theta))]h(\theta)f(\theta)d\theta,\end{aligned} \tag{33}$$

where i is the type that has high incentive to cream skim according to Lemma 11 (b) and l is a uniform bound on $\frac{\partial u}{\partial x}(\theta,x)/f(\theta)$.

Let us prove that the first term on the right hand side of (33) is bounded. From the first-order condition of the consumer's problem, let θ_i^m be the minimum value of θ_i which is achieved at $\theta_{-i} = \overline{\theta}_{-i}$, i.e,

$$\theta_i^m \phi_i'(x) + \sum_{j\neq i}\overline{\theta}_j\phi_j'(x) = \sum_j m_j(x)\phi_j'(x)$$

or

$$\left(m_i(x) - \theta_i^m\right)\phi_i'(x) = \sum_{j\neq i}\left(\overline{\theta}_j - m_j(x)\right)\phi_j'(x).$$

Now dividing both sides by $\overline{\theta}_i - m_i(x)$, by Lemma 11 (b) and the choice of i we have that the right hand side is uniformly bounded in x. Therefore,

$$\frac{|\theta_i - m_i(x)|}{\overline{\theta}_i - m_i(x)} \leq \max\left\{\frac{m_i(x)-\theta_i^m}{\overline{\theta}_i - m_i(x)}, 1\right\} \tag{34}$$

is uniformly bounded.

For the second term on the right hand side of (33), notice that

$$\lim_{\theta\to\overline{\theta}}\frac{\tilde{x}_+ - \tilde{x}(\theta)}{\overline{\theta}_i - m_i(\tilde{x}(\theta))} = \lim_{\theta\to\overline{\theta}}\frac{\tilde{x}_+ - \tilde{x}(\theta)}{\overline{\theta}_i - m_i(\tilde{x}(\theta))} = \lim_{x\to\tilde{x}_+}\frac{1}{(\overline{\theta}_i - m_i(x))/(\tilde{x}_+ - x)} = \frac{1}{m_i'(\tilde{x}_+)} = \infty.$$

By the mean value theorem, there exists $\eta_{\theta,\tilde{x}(\theta)}$ between θ_i and $m_i(\tilde{x}(\theta))$ such that $g(\theta) =$

$g(m_i(\widetilde{x}(\theta), \theta_{-i}) + \partial_{\theta_i} g(\eta_{\theta,\widetilde{x}(\theta)}, \theta_{-i})(\theta_i - m(\widetilde{x}(\theta)))$. Hence,

$$\lim_{\theta_i \to \overline{\theta}_i} \frac{g(\theta)}{\overline{\theta}_i - m_i(\widetilde{x}(\theta))} \leq \lim_{x \to \widetilde{x}_+} \sup \left[\frac{g(m_i(x), \theta_{-i})}{\overline{\theta}_i - m_i(x)} + \partial_{\theta_i} g(\eta_{\theta,x}, \theta_{-i}) \frac{|\theta_i - m_i(x)|}{\overline{\theta}_i - m_i(x)} \right],$$

and the first term of the right hand is clearly bounded because g is differentiable at dimension i and $g(\overline{\theta}_i, \theta_{-i}) = 0$. The second term is also bounded because of the argument that leads to the conclusion of (34).

Therefore, for $\epsilon > 0$ sufficiently small, the second term on the right hand side of (33) is strictly positive on a positive measure subset of $N_{\widetilde{x},\varepsilon}$ where $h(\theta) > 0$.

(ii) Lemma 11 (a) gives $\widetilde{x}(\theta) < \widetilde{x}_+ - \xi$, for almost all $\theta \in \Theta \backslash \{\overline{\theta}\}$. From the strict concavity of $S(\theta, \cdot)$, we have that $\frac{\partial S}{\partial x}(\theta, x(\theta)) \geq \frac{\partial S}{\partial x}(\theta, \widetilde{x}(\theta)) \geq \frac{\partial S}{\partial x}(\theta, \widetilde{x}_+ - \xi) > 0$, for $\epsilon < \xi$, since $\frac{\partial S}{\partial x}(\theta, \widetilde{x}_+ - \xi) \geq \frac{\partial S}{\partial x}(\theta, \widetilde{x}(\theta)) \geq 0$ and $x(\theta) \leq \widetilde{x}(\theta)$ for all $\theta \in N_{\widetilde{x},\varepsilon}$. From the expression (32), this implies that

$$\delta\pi(x, h) \geq \int_{N_{\widetilde{x},\varepsilon}} \left[\frac{\partial S}{\partial x}(\theta, \widetilde{x}_+ - \xi) - \frac{g(\theta)}{f(\theta)} \frac{\partial u}{\partial x}(\theta, x(\theta)) \right] h(\theta) f(\theta) d\theta.$$

Since $\frac{\partial u}{\partial x}(\theta, x)/f(\theta)$ is bounded and $g(\theta) \to 0$ when $\theta \to \overline{\theta}$, we have the term in the bracket will strictly positive on a positive measure subset of $N_{\widetilde{x},\varepsilon}$ where $h(\theta) > 0$, if we take $\epsilon > 0$ sufficiently small.

Therefore, for sufficiently small $\epsilon > 0$, increasing from the allocation $x(\cdot)$ to $\widetilde{x}(\cdot)$ at top generates profit gains for the monopolist whenever $x(\theta) < \widetilde{x}(\theta)$ in a positive measure set at the top. We have to show that there exists an incentive-compatible allocation that implements this improvement. Let U denote the informational rent associated with mechanism x and let V denote the rent associated with $\widetilde{x}$. Notice that the utility function considered here (separable and with non-decreasing $\phi_i(\cdot)$, for each i) imply that $\nabla V(\theta) \geq \nabla U(\theta)$ if and only if $\widetilde{x}(\theta) \geq x(\theta)$. Applying Lemma 14 below restricted to the domain $N_{\widetilde{x},\varepsilon}$, we find that $\overline{U} = \max\{U, V^\tau\}$ for some $\tau \in \mathbb{R}$ is the informational rent associated with an incentive-compatible mechanism that increases the monopolist's profit in the domain $N_{\widetilde{x},\varepsilon}$. Finally, we can extend the informational rent (and the allocation) outside of $N_{\widetilde{x},\varepsilon}$ to coincide with U (coincide with x). By construction, this mechanism satisfies the participation constraint because its informational rent is greater or equal to U.

Lemma 12. *U is an informational rent function from an incentive compatible mechanism if and only if U is a convex function and the consistency condition is satisfied (i.e., there exists a mechanism (p, x) that leads to the informational rent U).*

Proof. (Necessity) Fix an incentive compatible mechanism (p, x). The informational rent associated to this mechanism is

$$U(\theta) = \max_{\widehat{\theta} \in \Theta} u(\theta, x(\widehat{\theta})) - p(\widehat{\theta}).$$

Since for each $\widehat{\theta} \in \Theta$, the function $\theta \to u(\theta, x(\widehat{\theta})) - p(\widehat{\theta})$ is convex by assumption, $U(\theta)$ is the upper envelope of convex functions. Therefore, $U(\theta)$ is a convex function.

(Sufficiency) By the subgradient inequality, we have that

$$U(\theta) - U(\widehat{\theta}) \geq \nabla U(\widehat{\theta}) \cdot (\theta - \widehat{\theta}) = u_\theta(\widehat{\theta}, x(\widehat{\theta})) \cdot (\theta - \widehat{\theta}),$$

where the equality is a consequence of the envelope theorem. By separability of u, $u_\theta(\widehat{\theta}, x(\widehat{\theta})) \cdot (\theta - \widehat{\theta}) = u(\theta, x(\widehat{\theta}) - u(\widehat{\theta}, x(\widehat{\theta}))$. Hence, using the definition of U and consistency condition, the previous inequality is equivalent to

$$U(\theta) \geq u(\theta, x(\widehat{\theta})) - p(\widehat{\theta}),$$

which shows that the incentive compatibility holds for the mechanism (p, x). □

In the next lemma, we denote by ∇V and ∇U the gradients of functions V and U.

Lemma 13. *Let $U, V : \Theta \to \mathbb{R}$ be functions such that $\nabla V \geq \nabla U$. If $V(\theta) \geq U(\theta)$, then $V(\widehat{\theta}) \geq U(\widehat{\theta})$, for all $\widehat{\theta} \geq \theta$.*

Proof. The result follows by considering the restriction of functions V and U to the line segment connecting θ and $\widehat{\theta}$. □

Lemma 14. *Suppose that Θ is a convex set with a maximal type $\overline{\theta}$. Let $U, V : \Theta \to \mathbb{R}$ be a convex functions such that $\nabla V \geq \nabla U$. Denote $V^\tau := V + \tau$ the τ-translation of V, for each $\tau \in \mathbb{R}$. Then:*

(i) there exists $\tau_0 \in \mathbb{R}$ such that $V^{\tau_0}(\overline{\theta}) = U(\overline{\theta})$;

(ii) for every $\epsilon > 0$, there exists $\delta > 0$ such that $V^\tau \geq U$ in an open neighborhood contained in the ball $B(\overline{\theta}, \epsilon)$, for all $\tau \in (\tau_0, \tau_0 + \delta)$;

(iii) if u is a separable function and U and V are informational rent functions associated with incentive-compatible mechanisms, then $\overline{U} = \max\{U, V^\tau\}$ is an informational rent function associated with an incentive-compatible mechanism.

Proof. (i) Define τ_0 as the supremum of τ such that $V^\tau \leq U$. The result follows from Lemma 13.

(ii) Since U and V are continuous functions in the interior of Θ, the result follows immediately from item (i) and Lemma 13.

(iii) Notice that $\overline{U}$ is a convex function. By Lemma 12, it is an informational rent function of an incentive-compatible mechanism because the consistency condition is satisfied in each region where $\overline{U} = U$ or $\overline{U} = V^{\tau}$. Moreover, $\overline{U} \geq U$ implies that it also satisfies the participation constraint. □

Proof of Proposition 4:

Before presenting the proof, we start with a helpful result. Let $a \in \mathbb{R}^K$, $\nu \in \mathbb{R}$ and $P = \times_{i=1}^{K}[\underline{\theta}_i, \overline{\theta}_i] \subset \mathbb{R}_+^K$. Consider the following minimization problem $M(a, \nu)$:

$$\begin{array}{l} \min_{\theta \in P} a \cdot \theta \\ \text{s.t. } \mathbf{1} \cdot \theta = \nu \end{array}, \tag{35}$$

where $\mathbf{1} = (1, ..., 1)$ and $a_1 \leq ... \leq a_K$. If $\mathbf{1} \cdot \underline{\theta} \leq \nu \leq \mathbf{1} \cdot \overline{\theta}$, let us define the lowest $\widetilde{k}$ such $\mathbf{1} \cdot (\overline{\theta}_1, ..., \overline{\theta}_{\widetilde{k}}, \underline{\theta}_{\widetilde{k}+1}, ..., \underline{\theta}_K) \geq \nu$.

Definition 5. We say that vector $\widetilde{a}$ preserves the order of a with respect to $\widetilde{k}$ when $a_i \leq a_{\widetilde{k}}$ if and only if $\widetilde{a}_i \leq \widetilde{a}_{\widetilde{k}}$, for all $i \in \{1, ..., K\}$.

Lemma 15. *Suppose that $\mathbf{1} \cdot \underline{\theta} \leq \nu \leq \mathbf{1} \cdot \overline{\theta}$. The solutions of problems $M(a, \nu)$ and $M(\widetilde{a}, \nu)$ coincide if and only if a and $\widetilde{a}$ have the same order with respect to $\widetilde{k}$.*

Proof. Suppose that the order property holds. By the assumption and Definition 5 of $\widetilde{k}$, let $\widetilde{\theta}_{\widetilde{k}} \in [\underline{\theta}_{\widetilde{k}}, \overline{\theta}_{\widetilde{k}}]$ satisfy $\mathbf{1} \cdot (\overline{\theta}_1, ..., \widetilde{\theta}_{\widetilde{k}}, \underline{\theta}_{\widetilde{k}+1}, ..., \underline{\theta}_K) = \nu$. It is straightforward to see that $\theta^* = (\overline{\theta}_1, ..., \widetilde{\theta}_{\widetilde{k}}, \underline{\theta}_{\widetilde{k}+1}, ..., \underline{\theta}_K)$ is the optimal solution of $M(a, \nu)$. Moreover, for every $\widetilde{a}$ that preserves the order of a w.r.t. $\widetilde{k}$, the problem $M(\widetilde{a}, \nu)$ has exactly the same solution θ^*. Reciprocally, if a and $\widetilde{a}$ do not preserve the same order with respect to $\widetilde{k}$, then there exists a permutation φ of $\{1, ..., K\}$ such that $a_i \leq a_j$ if and only if $\widetilde{a}_{\varphi(i)} \leq \widetilde{a}_{\varphi(j)}$ and $\varphi(i) > \widetilde{k}$ for some $i \leq \widetilde{k}$. By the solution characterization above, we have that the solutions for problems $M(a, \nu)$ and $M(\widetilde{a}, \nu)$ must be different since φ permutes a type below $\widetilde{k}$ with a type above it. □

We now turn to the proof of the proposition. From the definition of strong adverse selection, we need to show that the following minimization problems have the same solutions:

$$\begin{array}{l} \min_{\widetilde{\theta} \in \Theta} \sum \widetilde{\theta}_i \phi_i''(x) \\ \text{s.t. } \sum \widetilde{\theta}_i \phi_i'(x) = \nu \end{array}$$

and

$$\begin{array}{l} \min_{\widetilde{\theta} \in \Theta} \sum \widetilde{\theta}_i \kappa_i(x) \\ \text{s.t. } \sum \widetilde{\theta}_i \phi_i'(x) = \nu \end{array}.$$

Applying a change of variables, we can write these problems as:

$$\begin{array}{ll}\min_{\widetilde{\theta}\in\Theta(x)} \sum \widetilde{\theta}_i a_i(x) \\ \text{s.t. } \sum \widetilde{\theta}_i = \nu\end{array} \quad \text{and} \quad \begin{array}{ll}\min_{\widetilde{\theta}\in\Theta(x)} \sum \widetilde{\theta}_i \widetilde{a}_i(x) \\ \text{s.t. } \sum \widetilde{\theta}_i = \nu\end{array}$$

where $a_i(x) := \phi_i''(x)/\phi_i'(x)$, $\widetilde{a}_i(x) := \kappa_i(x)/\phi_i'(x)$ and $\Theta(x)$ is an adjusted hyperrectangle.[26]

Fix x and, without loss of generality, suppose that $a_1(x) \leq a_2(x) \leq ... \leq a_K(x)$. By Lemma 15, the problems above have the same solutions if

$$a_i(x) \leq a_j(x) \iff \widetilde{a}_i(x) \leq \widetilde{a}_j(x).$$

Using the definition of these functions, we find that the problems above have the same solution

$$\frac{\phi_j''(x)}{\phi_j'(x)} - \frac{\phi_i''(x)}{\phi_i'(x)} \geq 0 \iff \frac{\kappa_j(x)}{\phi_j'(x)} - \frac{\kappa_i(x)}{\phi_i'(x)} \geq 0,$$

for all $i, j \in \{1, ..., K\}$ and all $x \in [0, 1]$.

A Numerical Simulations

A.1 Normal Distribution

We followed the algorithm from Azevedo and Gottlieb (2017) to calculate a competitive equilibrium. We used a perturbation with 26 evenly spaced contracts and a 1% mass of behavioral consumers that have zero cost. We then applied a fixed-point algorithm. In each iteration, consumers picked their favorite contracts taking prices as given. Then, we adjusted prices according to their profitability. Prices consistently converged to the same equilibrium for different initial values.[27]

Table 2 summarizes the parameters used in the calibration presented in the text, which are the same as in Azevedo and Gottlieb (2017). The monopolist and efficient allocations were computed numerically. As usual with arbitrary nonlinear maximization problems, it is impossible to guarantee that a local optimum is a global optimum. We calculated local optima starting from the equilibrium allocation, using both an ad-hoc procedure and the commercial optimization package KNITRO. We also calculated local optima from 300 random starting values in each simulation. The random starting values did not outperform the optimization starting at the equilibrium prices. Replication code is available at

[26] Notice that by changing variables $\widetilde{\theta}_i$ to $\widetilde{\eta}_i = \widetilde{\theta}_i \phi_i'(x)$, the hyperrectangle Θ becomes another hyperrectangle, which is now a function of x.

[27] Since the model may have multiple equilibria, it is important to consider different initial values to verify that the predictions do not depend on the initial value in the computations.

Table 2: Consumer Types (Normal Distribution)

	A	H	M	S
Mean	1×10^{-5}	$1,330$	$4,340$	$24,470$
Log covariance				
A	0.25	-0.01	-0.12	0
H		0.28	-0.03	0
M			0.20	0
S				0.25

Notes: The table presents the parameters used for the simulation described in the text. Consumer types are log normally distributed with the moments in the table.

https://github.com/diogowolff/ag-monopoly.

A.2 Truncated Normal Distribution

In example 1, we assumed that losses were normally distributed. While this assumption led to the transparent representation of preferences and costs in equations (3), it has the undesirable feature that it allows for negative losses. However, our results can still be applied for general distributions of losses.

Suppose losses are distributed according to a CDF F_κ, where κ is a vector of parameters that represent the consumer's private information about the distribution of losses. Preferences and costs can still be described with quasilinear preferences as in (1) and (2) with

$$\begin{aligned} u(\theta, x) &= \frac{\ln\left(\mathbb{E}\left[e^{Al}|l \sim F_\kappa\right]\right) - \ln\left(\mathbb{E}\left[e^{A(1-x)l}|l \sim F_\kappa\right]\right)}{A} + \frac{H}{2}x^2, \text{ and} \\ c(\theta, x) &= x\, \mathbb{E}\left[l|l \sim F_\kappa\right] + x^2 H, \end{aligned} \tag{36}$$

where $\theta = (\kappa, A, H)$ denotes the consumer's type.[28]

Table 3: Consumer Types (Truncated Normal Distribution)

	A	H	M	S
Mean	65×10^{-5}	$1,330$	$1,500$	$4,500$
Log covariance				
A	0.25	-0.01	-0.12	0
H		0.28	-0.03	0
M			0.20	0
S				0.25

Notes: The table presents the parameters used for the simulation described in the text. Consumer types are log normally distributed with the moments in the table.

[28] When using this model, one must ensure that the moments in (1) and (2) exist.

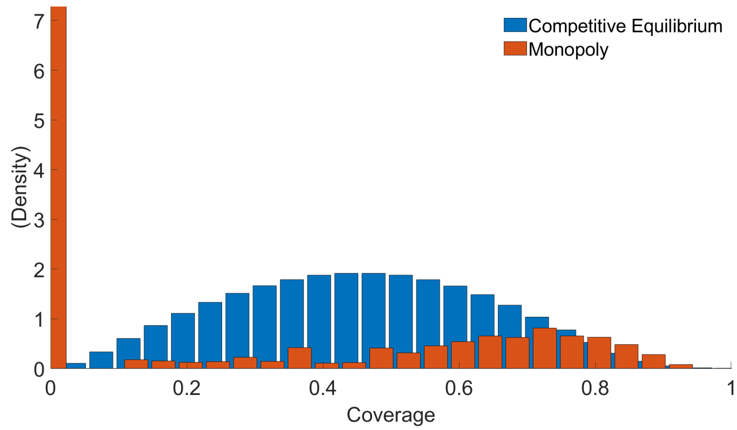

Figure 4: Coverage under Monopoly and Perfect Competition (Truncated Normal)

Notes: The figure depicts the distribution of coverage choices in example 1 with a truncated Normal distribution (losses are truncated at zero). The horizontal axis depicts the contracts chosen by consumers, with coverage ranging between 0% (uninsured) to 100% of expenses. The blue bars represent the distribution of coverage in the competitive equilibrium. The orange bars represent coverage under monopoly pricing. With monopoly pricing, approximately 78 percent of consumers remain uninsured. With perfect competition, all consumers purchase coverage. However, 7.9 percent of consumers purchase policies with coverage levels above 70 percent with monopoly, compared with only 5.8% with perfect competition.

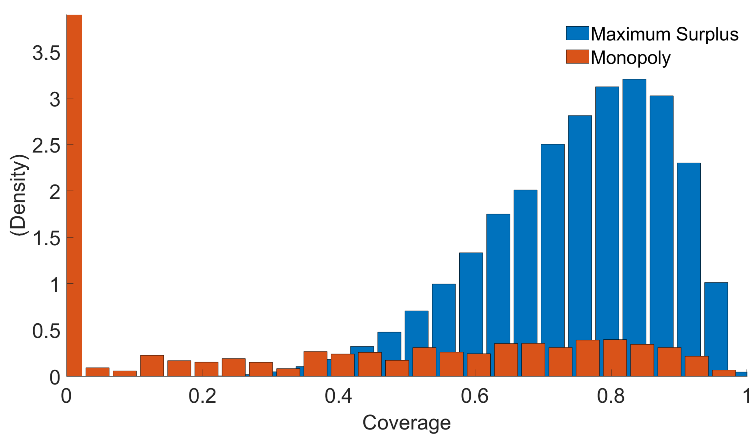

Figure 5: Coverage under Monopoly and Perfect Competition (Truncated Normal)

Notes: The figure depicts the distribution of coverage choices in example 1 with a truncated Normal distribution (losses are truncated at zero). The horizontal axis depicts the contracts chosen by consumers, with coverage ranging between 0% (uninsured) to 100% of expenses. The blue bars represent the distribution of coverage in the competitive equilibrium. The orange bars represent coverage under monopoly pricing. With monopoly pricing, over 70 percent of consumers remain uninsured. In the surplus-maximizing allocation, all consumers purchase coverage.

To ensure that losses are non-negative, we considered a truncated normal distribution, truncated at zero. We calculated the competitive equilibrium, the profit-maximizing allocation for the monopoly, and the surplus-maximizing allocation using the same algorithm described in the previous subsection. Table 3 describes the parameter values used.

Due to the truncation at zero, the parameter M no longer corresponds to the mean of the loss distribution. Keeping it at its original value while truncating losses at zero would lead to substantially higher average losses. We therefore chose to reduce the loss parameter M. Since the resulting distribution had considerably lower risk premia, we adjusted the coefficient of risk aversion A.

	No coverage	$\geq 70\%$ coverage	CS	Profit
Competition	0.0	5.8	5,453	0
Monopoly	78.3	7.9	2,406	1,393
Efficient	0.0	69.4	8,084	-2,133

Table 4: Coverage with Perfect Competition, Monopoly, and Maximum Surplus (Truncated Normal)

Notes: The table presents the coverage, consumer surplus (CS), and profits in the numerical example with a truncated normal distribution. The first column corresponds to the percentage of consumers excluded in each market structure. The second and third columns describe the proportion of consumer who purchase at least 70% and 80% coverage. The two last columns correspond to consumer surplus and profits.

The results are qualitatively similar to the ones with a normal distribution.

B Competitive Equilibrium in One-Dimensional Model

In this appendix, we characterize the competitive equilibrium under assumptions 4, 5, and 6, establishing the result stated in Lemma 2. We will first show that all competitive equilibria have degenerate allocations and pooling can only occur at zero. Then, we show that it features the least costly separating allocation.

B.1 Degenerate Allocations and No Pooling

We will show that all competitive equilibria have deterministic allocations and pooling can only occur at zero. Formally:

Proposition 5. *Let (p^*, α^*) be a competitive equilibrium. If (θ, x) and (θ, x') are in the support of α^*, then $x = x'$. If (θ, x) and (θ', x) are in the support of α^* and $x > 0$, then $\theta = \theta'$ (there is no pooling at positive contracts).*

The proof will follow a series of lemmas. Before presenting them, we note the following implication of single crossing (Assumption 4), which will be used throughout this appendix:

Remark 1. Let $x_1 > x_0$ and suppose that $u\left(\widetilde{\theta}, x_0\right) = u\left(\widetilde{\theta}, x_1\right)$ for some $\widetilde{\theta}$. Then, $u(\theta, x_1) < u(\theta, x_0)$ for all $\theta < \widetilde{\theta}$, and $u(\theta, x_1) > u(\theta, x_0)$ for all $\theta > \widetilde{\theta}$.

Note also that in any competitive equilibrium, prices are Lipschitz continuous (see Proposition 1 in Azevedo and Gottlieb, 2017). The first lemma establishes that consumer optimization implies in the monotonicity of the allocations:

Lemma 16. *Let (p^*, α^*) be an equilibrium. Suppose (θ, x_0) and (θ, x_1) are both in the support of α^*, where $x_0 < x_1$. Let $\theta_0 < \theta < \theta_1$. Then, (θ_0, x_1) and (θ_1, x_0) are not in the support of α^*.*

Proof. Suppose (θ, x_0) and (θ, x_1) are both in the support of α^*. Then, by consumer optimization, θ must be indifferent between them:

$$u(\theta, x_1) - u(\theta, x_0) = p(x_1) - p(x_0).$$

By single crossing, it follows that for any θ_0 and θ_1 with $\theta_1 > \theta > \theta_0$,

$$u(\theta_1, x_1) - u(\theta_1, x_0) > p(x_1) - p(x_0) > u(\theta_0, x_1) - u(\theta_0, x_0),$$

which implies that x_0 does not satisfy consumer optimization for θ_1 and vice-versa. □

The second lemma shows that any pooling allocation is an isolated point in the set of equilibrium allocations:

Lemma 17. *Let (p^*, α^*) be a competitive equilibrium and suppose (θ_0, x) and (θ_1, x) are in the support of α^*, where $\theta_0 \neq \theta_1$. There exists $\epsilon > 0$ such that $0 < |x' - x| < \epsilon$ implies that (θ, x') is not in the support of α^* for any θ.*

Proof. Let

$$\theta_* \equiv \inf \{\theta : (\theta, x) \text{ are in the support of } \alpha^*\},$$

and

$$\theta^* \equiv \sup \{\theta : (\theta, x) \text{ are in the support of } \alpha^*\}.$$

Since Θ is bounded and the set of types for which (θ, x) are in the support of α^* is non-empty (θ_0 and θ_1 belong to that set), θ_* and θ^* exist.

By zero profits and the fact that c is strictly increasing in types,

$$p^*(x) \in (c(\theta_*, x), c(\theta^*, x)). \tag{37}$$

Let $x' > x$ (the case where $x' < x$ is analogous, with θ^* substituted by θ_*) and suppose that (x', θ') is in the support of α^*. Consumer optimization gives

$$u(\theta', x') - p^*(x') \geq u(x, \theta') - p^*(x),$$

and

$$u(\theta^*, x) - p^*(x) \geq u(\theta^*, x') - p^*(x').$$

Combining these expressions, we obtain

$$u(\theta', x') - u(\theta', x) \geq p^*(x') - p^*(x) \geq u(\theta^*, x') - u(\theta^*, x),$$

which, by single crossing, yields $\theta' \geq \theta^*$. Since all types choosing x' are (weakly) greater than θ^*,

$$p^*(x') \geq c(\theta^*, x') \geq c(\theta^*, x),$$

where the first inequality follows from zero profits (and c increasing in types), while the second follows from $x' > x$ (and c non-decreasing in x).

Suppose that for any $n \in \mathbb{N}$, there exists (θ'_n, x'_n) is in the support of α^* such that $0 < |x'_n - x| < \frac{1}{n}$. Then, we can obtain a sequence $\{x'_n\}$ converging to x for which $p^*(x'_n) \geq c(\theta^*, x)$ for all n. By continuity of p^*, it follows that $p^*(x) \geq c(\theta^*, x)$, which contradicts $p^*(x) < c(\theta^*, x)$ (by 37). □

The next lemma establishes that the only contract that may exhibit pooling in equilibrium is the null contract:

Lemma 18. *Let (p^*, a^*) be a competitive equilibrium and suppose (θ_0, x) and (θ_1, x) are in the support of α^*, where $\theta_0 \neq \theta_1$. Then, $x = 0$.*

Proof. Suppose we have a competitive equilibrium in which multiple types choose the same allocation x. As in the proof of the previous lemma, let $\underline{\theta} \equiv \inf\{\theta : (\theta, x) \text{ are in the support of } \alpha^*\}$ and recall that, by zero profits, $p(x) > c(\underline{\theta}, x)$. Suppose, in order to obtain a contradiction, that $x > 0$. From the previous lemma, there exists $\epsilon > 0$ such that $0 < |x - x'| < \epsilon$ implies that (θ, x') is not on the support of a^* for any θ.

By continuity of u and the fact that $u(\theta, x) > 0$ for all $x > 0$, prices have to be strictly positive in a small neighborhood of $x > 0$. Thus, by condition 3 of Definition 1, for each contract x' in a neighborhood of x, there must be types who are indifferent between their equilibrium contracts and x'. By increasing difference, types who are indifferent between their equilibrium contracts and $x' < x$ satisfy $\theta \leq \underline{\theta}$ (otherwise, $\underline{\theta}$ would strictly benefit from picking x', thereby violating consumer optimality). Hence, again by condition 3 of Definition 1, we must have

$$p\left(x'\right) \leq c\left(\underline{\theta}, x'\right).$$

Taking the limit as $x' \nearrow x$ yields $p\left(x\right) \leq c\left(\underline{\theta}, x\right)$, contradicting $p(x) > c(\underline{\theta}, x)$. □

We now show that in any competitive equilibrium, there cannot be randomization involving two contracts with positive coverage.

Lemma 19. *Let (p^*, α^*) be a competitive equilibrium and suppose (θ, x) and (θ, x') are in the support of α^* for some θ, where $x < x'$ (i.e., the allocation of some type is non-degenerate). Then, $x = 0$.*

Proof. Let (p^*, α^*) be a competitive equilibrium and suppose (θ, x) and (θ, x') are in the support of α^*, for $0 < x < x'$. By the Lemma 18, θ is the only type who picks both x and x' in this competitive equilibrium (since they are both strictly positive and pooling can only occur at zero). Type θ has to be indifferent between these two allocations in order to mix:

$$u(\theta, x) - p\left(x\right) = u(\theta, x') - p\left(x'\right).$$

Moreover, fits must make zero profits:

$$p\left(x\right) = c\left(\theta, x\right), \; p\left(x'\right) = c\left(\theta, x'\right).$$

Combining both expressions, it follows that both allocations must yield the same surplus for type θ:

$$u(\theta, x) - c\left(\theta, x\right) = u(\theta, x') - c\left(\theta, x'\right). \tag{38}$$

Let $\widetilde{x} \in (x, x')$. By strictly concavity of the surplus, equation (38) implies that

$$u\left(\theta, \widetilde{x}\right) - c\left(\theta, \widetilde{x}\right) > u\left(\theta, x\right) - c\left(\theta, x\right). \tag{39}$$

By monotonicity (Lemma 16), either θ picks $\tilde{x}$ in this competitive equilibrium or no one does. In both cases, the equilibrium has to satisfy the following indifference and zero profit conditions:

$$u(\theta, x) - p\left(x\right) = u(\theta, \widetilde{x}) - p\left(\widetilde{x}\right), \quad \text{and} \quad p(\widetilde{x}) \leq c(\theta, \widetilde{x}).$$

To wit, if type θ picks $\widetilde{x}$ in equilibrium, then both must hold with equality. Otherwise, by condition 3 of Definition 1, some other type must be indifferent between $\tilde{x}$ and their equilibrium contract (otherwise, the price would be zero, which would contradict the optimality of picking $x < \tilde{x}$ while paying a non-negative price $p(x) = c(\theta, x) > 0$). By monotonicity, it must be type θ. Then, again by condition 3 of Definition 1, we must have $p(\widetilde{x}) \leq c(\theta, \widetilde{x})$.

Therefore, the two previous conditions also have to hold when $\tilde{x}$ is not chosen by any type in equilibrium. Combining them, we obtain:

$$u(\theta, x) - c(\theta, x) \geq u(\theta, \widetilde{x}) - c(\theta, \widetilde{x}),$$

which contradicts condition (39). Thus, we cannot have the same type picking two non-zero allocations with strictly positive probabilities.

Suppose two types $\theta < \theta'$ obtain non-degenerate allocations in a competitive equilibrium. Then, type θ must mix in $\{0, x\}$ and type θ' must mix in $\{0, x'\}$, where $x > 0$ and $x' > 0$. However, this contradicts Lemma 16, which states that if $(\theta, 0)$ and (θ, x) are in the support of α^*, then $(\theta', 0)$ cannot be in the support of α^*. Hence, there is at most one type that plays mixed strategies in any competitive equilibrium. □

Finally, we now show that the allocation in any competitive equilibrium must be deterministic:

Lemma 20. *Let (p^*, α^*) be a competitive equilibrium and suppose (θ, x) and (θ, x') are in the support of α^*. Then, $x = x'$.*

Proof. From the previous lemma, the only possible equilibrium with pooling has one single type randomizing between 0 and some $x > 0$. This type must be indifferent between these two allocations:

$$u(\theta, x) - p(x) = u(\theta, 0) - p(0).$$

By zero profits, we must have:

$$p(x) = c(\theta, x), \ \ p(0) = 0.$$

Using the fact that $u(\theta, 0) = 0$, we can write these conditions as

$$u(\theta, x) - c(\theta, x) = 0.$$

That is, the indifferent type must have zero surplus in both cases. Since the surplus $u(\theta, \cdot) - c(\theta, \cdot)$ is a strictly concave function of allocations,

$$u(\theta, x) - c(\theta, x) = u(\theta, 0) - c(\theta, 0) = 0$$

implies

$$u(\theta, \widetilde{x}) - c(\theta, \widetilde{x}) > 0 \tag{40}$$

for all $\widetilde{x} \in (0, x)$.

Condition 3 from Definition 1 implies that prices of contracts between 0 and x must be on type θ's indifference curve and cannot exceed type θ's cost. That is, for all $\widetilde{x} \in (0, x)$, we must have

$$u(\theta, \widetilde{x}) - p(\widetilde{x}) = 0, \tag{Indifference}$$

and

$$p(\widetilde{x}) \leq c(\theta, \widetilde{x}). \tag{Zero Profits}$$

Combining them, gives $u(\theta, \widetilde{x}) \leq c(\theta, \widetilde{x})$, which contradicts (40). □

B.2 Least Costly Separating Allocation

Since the allocation in any competitive equilibrium is deterministic, we can represent it by a pair $x : \Theta \to [0, 1]$ and $p : [0, 1] \to \mathbb{R}_+$.

Lemma 21. *Let (p, x) be a competitive equilibrium. The allocation function $x(\cdot)$ is continuous.*

Proof. By single crossing and the fact that there is no pooling outside of the null contract, $x(\cdot)$ is strictly monotone. Suppose there exists $\widetilde{\theta}$ such that

$$x_+ := \lim_{\theta \searrow \widetilde{\theta}} x(\theta) > \lim_{\theta \nearrow \widetilde{\theta}} x(\theta) =: x_-$$

(the argument for other types of discontinuities is analogous). Consumer optimization gives

$$u(\theta, x(\theta)) - p(x(\theta)) \geq u(\theta, x) - p(x) \;\; \forall x.$$

Take $x = x_-$ and consider the limit as $\theta \searrow \widetilde{\theta}$. Because both u and p are continuous, we have

$$u\left(\widetilde{\theta}, x_+\right) - p(x_+) \geq u\left(\widetilde{\theta}, x_-\right) - p(x_-).$$

Similarly, setting $x = x_+$ and taking the limit as $\theta \nearrow \widetilde{\theta}$, gives

$$u\left(\widetilde{\theta}, x_-\right) - p(x_-) \geq u\left(\widetilde{\theta}, x_+\right) - p(x_+).$$

Combining both conditions, gives

$$u\left(\widetilde{\theta}, x_+\right) - p(x_+) = u\left(\widetilde{\theta}, x_-\right) - p(x_-).$$

Zero profits implies that $p(x_+) = c\left(\widetilde{\theta}, x_+\right)$ and $p(x_-) = c\left(\widetilde{\theta}, x_-\right)$. Thus,

$$u\left(\widetilde{\theta}, x_+\right) - c\left(\widetilde{\theta}, x_+\right) = u\left(\widetilde{\theta}, x_-\right) - c\left(\widetilde{\theta}, x_-\right).$$

Since x is monotonic, it follows that $x \in (x_-, x_+)$ is not picked by any type. By single crossing, condition 3 from Definition 1 implies that they must be priced according to type $\widetilde{\theta}$'s indifference curve:

$$u\left(\widetilde{\theta}, x_+\right) - c\left(\widetilde{\theta}, x_+\right) = u\left(\widetilde{\theta}, x\right) - p(x), \tag{Indifference}$$

$$p(x) \leq c\left(\widetilde{\theta}, x\right). \tag{Zero Profits}$$

Combining the conditions, we obtain, for all $x_- < x < x_+$,

$$u\left(\widetilde{\theta}, x\right) - c\left(\widetilde{\theta}, x\right) \leq u\left(\widetilde{\theta}, x_+\right) - c\left(\widetilde{\theta}, x_+\right) = u\left(\widetilde{\theta}, x_-\right) - c\left(\widetilde{\theta}, x_-\right),$$

which violates the strict concavity of the surplus function. □

The next lemma shows that almost all types are distorted downwards:

Lemma 22. *Let (p, x) be a competitive equilibrium. Then, $x(\theta) \leq \arg\max_x \{u(\theta, x) - c(\theta, x)\}$, with strict inequality almost everywhere.*

Proof. Since $\Theta = [\underline{\theta}, \bar{\theta}]$ is an interval and x is a continuous and increasing function, it follows that $x(\Theta)$ is an interval. Consumer optimization on the equilibrium path can be written as:

$$x(\theta) \in \arg \max_{x(\underline{\theta}) \leq x \leq x(\bar{\theta})} u(\theta, x) - p(x).$$

Zero profits for allocations $x \in [x(\underline{\theta}), x(\bar{\theta})]$ gives

$$p(x(\theta)) = c(\theta, x(\theta)).$$

Note that x is a monotone function and, therefore, is differentiable almost everywhere. Consumer optimization requires that at all points of differentiability of x, we have

$$\frac{\partial u}{\partial x}(\theta, x(\theta)) = p'(x(\theta)) \tag{41}$$

for all $x(\theta) \in (x(\underline{\theta}), x(\bar{\theta}))$. In general, this local first-order condition is only necessary for consumer optimality. As usual, we ignore the sufficiency condition for now and verify it at the end.

From zero profits, at all points of differentiability of x, we have

$$p'(x(\theta))\dot{x}(\theta) = \frac{\partial c}{\partial \theta}(\theta, x(\theta)) + \frac{\partial c}{\partial x}(\theta, x(\theta))\dot{x}(\theta).$$

Combining both expressions, gives

$$\left[\frac{\partial u}{\partial x}(\theta, x(\theta)) - \frac{\partial c}{\partial x}(\theta, x(\theta))\right]\dot{x}(\theta) = \frac{\partial c}{\partial \theta}(\theta, x(\theta)).$$

Since $\frac{\partial c}{\partial \theta}(\theta, x(\theta)) > 0$, it follows that $\dot{x}(\theta) \neq 0$ and, since x is strictly increasing,

$$\frac{\partial u}{\partial x}(\theta, x(\theta)) > \frac{\partial c}{\partial x}(\theta, x(\theta))$$

at all points of differentiability of x. Since x is strictly increasing and satisfies the first-order condition (41), and the consumer's utility satisfies single crossing, it follows from standard arguments that global incentive compatibility holds.

Let $x^*(\theta) \equiv \arg\max_x u(\theta, x) - c(\theta, x)$ denote the surplus-maximizing allocation. Since the surplus is strictly concave, it follows from the Theorem of the Maximum that $x^*(\theta)$ is continuous. Then, the previous inequality implies that $x(\theta) < x^*(\theta)$ almost everywhere. Because x and x^* are both continuous, it follows that $x(\theta) \leq x^*(\theta)$ at all points. □

Lemma 23. *Let (p, x) be a competitive equilibrium. Then, $x(\bar{\theta}) = \arg\max_{x\in X} u(\bar{\theta}, x) - c(\bar{\theta}, x)$.*

Proof. From the previous lemma, $x(\bar{\theta}) \leq \arg\max_{x\in X} u(\bar{\theta}, x) - c(\bar{\theta}, x)$. Suppose, in order to obtain a contradiction, that the inequality is strict. By the monotonicity of $x(\cdot)$, it follows that no type picks $x > x(\bar{\theta})$. Then, because of single crossing, the competitive equilibrium requires all allocations $x > x(\bar{\theta})$ to be priced according to $\bar{\theta}$'s indifference curve:

$$p(x) = u(\bar{\theta}, x) - u(\bar{\theta}, x(\bar{\theta})) + p(x(\bar{\theta})),$$

and to satisfy the zero-profit condition:

$$p(x) \leq c\left(\bar{\theta}, x\right).$$

Combining both conditions and using the fact that $p(x(\bar{\theta})) = c\left(\bar{\theta}, x(\bar{\theta})\right)$, gives

$$u(\bar{\theta}, x) - c\left(\bar{\theta}, x\right) \leq u(\bar{\theta}, x(\bar{\theta})) - c(\bar{\theta}, x(\bar{\theta}))$$

for all $x > \bar{x}$, contradicting the assumption that $x(\bar{\theta}) \leq \arg\max_{x\in X} u(\bar{\theta}, x) - c(\bar{\theta}, x)$. □

These previous lemmas therefore establish that the equilibrium is determined by incentive compatibility with the boundary condition specifying that the highest type gets the efficient allocation (least costly separating allocation).

Next, we show that the slope of the equilibrium allocation becomes infinitely steep close to the top:

Lemma 24. *Let $x(\theta)$ be a competitive equilibrium. Then,* $\lim_{\theta \nearrow \overline{\theta}} \dot{x}(\theta) = +\infty$.

Proof. As shown in the previous lemma, the competitive equilibrium allocation solves:

$$\left[\frac{\partial u}{\partial x}(\theta, x(\theta)) - \frac{\partial c}{\partial x}(\theta, x(\theta))\right] \dot{x}(\theta) = \frac{\partial c}{\partial \theta}(\theta, x(\theta)),$$

with $x(\overline{\theta}) = \arg\max_x u(\overline{\theta}, x) - c(\overline{\theta}, x)$. As $\theta \nearrow \overline{\theta}$, we have

$$\underbrace{\left[\frac{\partial u}{\partial x}(\theta, x(\theta)) - \frac{\partial c}{\partial x}(\theta, x(\theta))\right]}_{\searrow 0} \dot{x}(\theta) = \underbrace{\frac{\partial c}{\partial \theta}(\theta, x(\theta))}_{\to \frac{\partial c}{\partial \theta}(\overline{\theta}, x(\overline{\theta}))>0},$$

so $\lim_{\theta \nearrow \overline{\theta}} \dot{x}(\theta) = +\infty$. □

C Monopolist Solution in One-Dimensional Model

In this appendix, we consider the monopolist's problem in the one-dimensional model. We first establish the result from Lemma 3, which considers deterministic mechanisms. Then, we generalize it to allow for stochastic mechanisms.

C.1 Deterministic Mechanisms and Lemma 3

Using standard arguments, we can write the monopolist's program as

$$\max_{x(\cdot)} \int_{\underline{\theta}}^{\overline{\theta}} \left[S(\theta, x(\theta)) - \frac{1 - F(\theta)}{f(\theta)} \frac{\partial u}{\partial \theta}(\theta, x(\theta))\right] f(\theta)\, d\theta$$

subject to $x(\cdot)$ non-decreasing. Let $x^m(\cdot)$ denote the solution to this program (i.e., the allocation that maximizes the monopolist's profits). It is helpful to separate into two cases:

- If the monotonicity constraint does not bind at the top (i.e., there exists $\epsilon > 0$ such that $x^m(\theta)$ is strictly increasing for $\theta \in (\overline{\theta} - \epsilon, \overline{\theta})$), the solution must satisfy $x^m(\overline{\theta}) = x^{FB}(\overline{\theta})$ and $\dot{x}^m(\overline{\theta}) < \infty$.

- If the monotonicity constraint binds at the top (i.e., there exists $\epsilon > 0$ such that $x^m(\theta) = x^*$ for all $\theta \in (\overline{\theta} - \epsilon, \overline{\theta})$), then $x^* \geq x^{FB}(\overline{\theta})$.

Case 1: Separation

In this case, the solution for types in $(\overline{\theta}-\epsilon,\overline{\theta})$ must satisfy the pointwise optimality condition:

$$\frac{\partial S}{\partial x}\left(\theta,x\left(\theta\right)\right)=\frac{1-F(\theta)}{f(\theta)}\frac{\partial^{2}u}{\partial\theta\partial x}\left(\theta,x\left(\theta\right)\right).$$

By Assumption 4, all types except for $\bar{\theta}$ get less coverage than the efficient amount. By the implicit function theorem, we have

$$\dot{x}(\overline{\theta})=-\frac{\frac{\partial^{2}S}{\partial\theta\partial x}\left(\overline{\theta},x(\overline{\theta})\right)-\frac{d}{d\theta}\left[\frac{1-F(\theta)}{f(\theta)}\frac{\partial^{2}u}{\partial\theta\partial x}\left(\theta,x\right)\right]\Big|_{\theta=\overline{\theta},x=x(\overline{\theta})}}{\frac{\partial^{2}u}{\partial x^{2}}\left(\overline{\theta},x(\overline{\theta})\right)-\frac{\partial^{2}c}{\partial x^{2}}\left(\overline{\theta},x(\overline{\theta})\right)}.$$

Differentiating and using the fact that $f'(\overline{\theta})<\infty$ (Assumption 7), we obtain

$$\frac{d}{d\theta}\left[\frac{1-F(\theta)}{f(\theta)}\frac{\partial^{2}u}{\partial\theta\partial x}\left(\theta,x\right)\right]\Big|_{\theta=\overline{\theta}}=-\frac{\partial^{2}u}{\partial\theta\partial x}\left(\overline{\theta},x\right).$$

Substituting back, we obtain:

$$\dot{x}(\overline{\theta})=-\frac{\frac{\partial^{2}S}{\partial\theta\partial x}\left(\overline{\theta},x(\overline{\theta})\right)+\frac{\partial^{2}u}{\partial\theta\partial x}\left(\overline{\theta},x(\overline{\theta})\right)}{\frac{\partial^{2}S}{\partial x^{2}}\left(\overline{\theta},x(\overline{\theta})\right)}<+\infty,$$

where the inequality follows from Assumption 6 ($\frac{\partial^{2}S}{\partial x^{2}}\left(\overline{\theta},x(\overline{\theta})\right)<0$) and the fact that S is twice continuously differentiable (so the numerator is bounded).

Case 2: Pooling

Introducing the auxiliary variable $z(\theta)=\dot{x}(\theta)$, we can write the monopolist's program as:

$$\max_{x(\cdot),z(\cdot)}\int_{\underline{\theta}}^{\overline{\theta}}\left[S\left(\theta,x\left(\theta\right)\right)-\frac{1-F(\theta)}{f(\theta)}\frac{\partial u}{\partial\theta}\left(\theta,x\left(\theta\right)\right)\right]f\left(\theta\right)d\theta$$

subject to

$$\dot{x}(\theta)=z(\theta)$$

$$z(\theta)\geq0$$

This is a standard optimal control problem, which has the following necessary optimality conditions:

$$z(\theta)\in\arg\max_{z\geq0}\left\{\left[S\left(\theta,x\left(\theta\right)\right)-\frac{1-F(\theta)}{f(\theta)}\frac{\partial u}{\partial\theta}\left(\theta,x\left(\theta\right)\right)\right]f\left(\theta\right)+\lambda(\theta)z\right\},$$

so that $\lambda(\theta) \leq 0$ for all θ with $z(\theta) = 0$ if $\lambda(\theta) < 0$,

$$\left[\frac{\partial S}{\partial x}\left(\theta, x\left(\theta\right)\right) - \frac{1-F(\theta)}{f(\theta)}\frac{\partial^2 u}{\partial\theta\partial x}\left(\theta, x\left(\theta\right)\right)\right] f\left(\theta\right) = -\dot{\lambda}(\theta), \tag{42}$$

and the transversality condition: $\lambda(\overline{\theta}) = 0$.

Since $\lambda(\overline{\theta}) = 0$ and $\lambda(\theta) \leq 0$ for all θ, we must have $\dot{\lambda}(\overline{\theta}) \geq 0$ (otherwise, we would have $\lambda(\theta) > 0$ for all $\theta < \overline{\theta}$ close enough to $\overline{\theta}$). Therefore, equation (42) gives:

$$\frac{\partial S}{\partial x}\left(\bar{\theta}, x\left(\bar{\theta}\right)\right) \leq 0,$$

which by the concavity of S implies $x(\bar{\theta}) > \bar{x}$ (that is, $x(\bar{\theta})$ is weakly above the first best).

C.2 Stochastic Mechanisms

We now generalize the analysis to allow for stochastic mechanisms. In addition to Assumptions 1 and 4-7 made in the deterministic case, we also assume that $\frac{\partial^2 u}{\partial \theta^2}(\theta, x) \geq 0$ for all (θ, x) and $\frac{\partial^3 u}{\partial\theta\partial x^2}(\theta, x) \geq 0$ for all x and all θ in a neighborhood of $\bar{\theta}$.

We consider direct mechanisms. Because the utility is quasi-linear, there is no loss of generality in restricting attention to deterministic prices. Therefore, a mechanism specifies a price $p(\theta)$ and distribution over coverages $\xi_\theta(x)$ for each type θ. Fix a mechanism and let

$$U(\theta) \equiv \int_x u\left(\theta, x\right) d\xi_\theta(x) - p(\theta)$$

denote the indirect utility of type θ.

Lemma 25. *Suppose* $\frac{\partial^2 u}{\partial \theta^2} \geq 0$. *A mechanism is incentive compatible if and only if* $\dot{U}(\theta)$ *is* non-decreasing and satisfies the envelope condition*:*

$$\dot{U}(\theta) = \int_x \frac{\partial u}{\partial \theta}\left(\theta, x\right) d\xi_\theta(x). \tag{43}$$

Proof. If the mechanism is incentive compatible, then

$$U(\theta) = \max_{\widehat{\theta}} \left\{ \int_x u(\theta, x) d\xi_{\widehat{\theta}}(x) - p(\widehat{\theta}) \right\}$$

where, given the assumption, the function $\theta \to \int u(\theta, x) d\xi_{\widehat{\theta}}(x) - p(\widehat{\theta})$ is convex in θ, for each $\widehat{\theta}$. Therefore, $U(\theta)$ is the upper envelope of convex functions and, therefore, convex, so that $\dot{U}(\theta)$ is non-decreasing. Moreover, condition (43) follows from the envelope theorem.

Conversely, suppose these two conditions hold. Then, for each $\theta > \widehat{\theta}$, we have that

$$U(\theta) - U(\widehat{\theta}) = \int_{\widehat{\theta}}^{\theta} \int_x \frac{\partial u}{\partial \theta}(\widetilde{\theta}, x) d\xi_{\widetilde{\theta}}(x) d\widetilde{\theta} \geq \int_{\widehat{\theta}}^{\theta} \int_x \frac{\partial u}{\partial \theta}(\widetilde{\theta}, x) d\xi_{\widehat{\theta}}(x) d\widetilde{\theta}$$

where the equality follows from the envelope condition (43) and the inequality follows from the convexity of $U(\theta)$. Therefore,

$$U(\theta) - U(\widehat{\theta}) \geq \int_x \left[u(\theta, x) - u(\widehat{\theta}, x) \right] d\xi_{\widehat{\theta}}(x),$$

which is equivalent to

$$U(\theta) \geq \int_x u(\theta, x) d\xi_{\widehat{\theta}}(x) - p(\widehat{\theta}).$$

□

We can therefore write the firm's program as:

$$\max_{\mu, U} \int_{\underline{\theta}}^{\bar{\theta}} \left\{ \int_x S(\theta, x)\, d\xi_\theta(x) - U(\theta) \right\} f(\theta)\, d\theta$$

subject to

$$\dot{U}(\theta) = \int_x \frac{\partial u}{\partial \theta}(\theta, x)\, d\xi_\theta(x)$$

$$\dot{U}(\theta) \text{ non-decreasing,}$$

and

$$U(\underline{\theta}) \geq 0.$$

The next lemma provides the key step for our main result. Let $\bar{x}_\theta = \arg\max_x S(\theta, x)$, which is unique by the strict concavity of S.

Lemma 26. *Let $\Theta_- \equiv \{\theta : \frac{\partial^3 u}{\partial \theta \partial x^2}(\theta, x) \geq 0\}$. Suppose a mechanism has a set of types with positive measure in Θ_- receiving non-degenerate allocations with $\mathbb{E}_\theta[x|\xi_\theta] < \bar{x}_\theta$. Then, the mechanism is not optimal.*

Proof. To build intuition, let x_θ^* be the deterministic payment that solves:

$$\int_x \frac{\partial u}{\partial \theta}(\theta, x)\, d\xi_\theta(x) = \frac{\partial u}{\partial \theta}(\theta, x_\theta^*).$$

Note that x_θ^* has the same informational rents as with the original allocation μ_θ, that is

$$\dot{U}(\theta) = \frac{\partial u}{\partial \theta}(\theta, x_\theta^*).$$

Therefore, replacing non-degenerate allocations ξ_θ by x^*_θ in any set of types in Θ_- leads to the same indirect utility and therefore preserves IC.

Since $\frac{\partial u}{\partial \theta}$ is convex, it follows that $x^*_\theta \geq \int x d\xi_\theta(x)$. Therefore, this replacement increases coverage. If $x^*_\theta \leq \bar{x}_\theta$, replacing the non-degenerate allocation ξ_θ by x^*_θ strictly increases surplus (pointwise) since S is strictly concave and strictly increasing in $x \leq \bar{x}_\theta$. If instead $x^*_\theta > \bar{x}_\theta$, the effect is ambiguous since S is no longer increasing in x.

For $x^*_\theta > \bar{x}_\theta > \mathbb{E}_\theta[x|\xi_\theta]$, we follow the same idea except we do not substitute ξ_θ by a deterministic contract x^*_θ. Instead, we construct a mean-preserving contraction that ensures that the effect is positive. First, note that for any mean-preserving contraction $\tilde{\xi}_\theta$ of ξ_θ, we have

$$\int_x \frac{\partial u}{\partial \theta}(\theta, x)\, d\xi_\theta(x) \geq \int_x \frac{\partial u}{\partial \theta}(\theta, x)\, d\tilde{\xi}_\theta(x)$$

by the weak convexity of $\frac{\partial u}{\partial \theta}$. Since $\frac{\partial^2 u}{\partial \theta \partial x} > 0$ and $\frac{\partial u}{\partial \theta}$ is continuous, there exists $\delta \geq 0$ such that

$$\int_x \frac{\partial u}{\partial \theta}(\theta, x)\, d\xi_\theta(x) = \int_x \frac{\partial u}{\partial \theta}(\theta, x + \delta)\, d\tilde{\xi}_\theta(x).$$

Intuitively, any mean-preserving contraction generates “additional coverage” $\delta \geq 0$ that can be used to increase surplus while keeping informational rents unchanged. The highest mean-preserving contraction replaces $\xi_\theta(x)$ by a degenerate distribution concentrated at its mean. But as seen above the additional coverage δ that it generates may be too large, leading to coverage beyond the surplus-maximizing point. However, by picking a small enough mean-preserving contraction (formally, any mean-preserving contraction with $\delta \leq \bar{x}_\theta - \mathbb{E}_\theta[x|\xi_\theta]$, which is positive by the assumption of the lemma), the previous argument gives a strict gain for the monopolist since it this keeps informational rents constant and increases surplus. Since $\bar{x}_\theta > \mathbb{E}_\theta[x|\xi_\theta]$ and ξ_θ is non-degenerate, such a small enough mean-preserving contraction always exists. □

By the previous lemma, non-degenerate contracts must offer expected coverage above the surplus-maximizing level, $\mathbb{E}_\theta[x|\xi_\theta] \geq \bar{x}_\theta$, for almost all types in Θ_-. Then, by the assumption that $\frac{\partial^3 u}{\partial \theta \partial x^2}(\theta, x) \leq 0$ for all x and all θ in a neighborhood of $\bar{\theta}$, it follows that (almost) any θ in that neighborhood that is offered a non-degenerate contract must have $\mathbb{E}_\theta[x|\xi_\theta] \geq \bar{x}_\theta > x^*(\theta)$, where $x^*(\theta)$ is the competitive equilibrium allocation (least-costly separating allocation). If types around $\bar{\theta}$ are offered deterministic contracts, it follows from the analysis from the last subsection (deterministic contracts) that the monopolist offers more coverage than the competitive equilibrium. The following proposition formalizes this argument:

Proposition 6. *Consider an incentive-compatible mechanism such that* $\mathbb{E}_\theta[x|\xi_\theta] < x^*(\theta)$ *for almost all types in an interval* $(\bar{\theta} - \epsilon, \bar{\theta}]$ *for some* $\epsilon > 0$. *Then, the mechanism does not maximize the monopolist's profits.*

Proof. If the mechanism offers a non-degenerate allocation for a positive measure of types in this interval, the previous lemma implies that this mechanism is not optimal. If the set of types obtaining non-degenerate allocations has measure zero, there is no loss in replacing their allocations by their certainty equivalents, which preserves incentive compatibility and does not affect the firm's profits (since this set has zero measure). But, from Lemma 3, the resulting allocation is also sub-optimal (since the optimal deterministic mechanism has $x^m(\theta) > x^*(\theta)$ for all $\theta \in (\bar{\theta} - \delta, \bar{\theta})$ for some $\delta > 0$). □

D Omitted Proof of Proposition 2

The proof incorporates mixed strategies using the same approach as the proof of Theorem 2. By quasi-linearity, there is no loss of generality in focusing on mechanisms with deterministic prices.

Let $U(\theta) \equiv \mathbb{E}_\theta \left[u(\theta, x))\right] - p(\theta)$. By the envelope condition, we must have

$$\dot{U}(\theta) \equiv \mathbb{E}_\theta \left[\frac{\partial u}{\partial \theta}(\theta, x))\right] > 0, \tag{44}$$

where the Dominated Convergence Theorem justifies differentiating under the expectation. Since $U(\cdot)$ is increasing, the exclusion region is an interval: $[\underline{\theta}, \theta^*]$.

If all types participate ($\theta^* = \underline{\theta}$), any allocation that maximizes the firm's profits must give zero utility to the lowest type. If there is exclusion ($\theta^* > \underline{\theta}$), all types who do not participate get zero utility. In either case we have $U(\theta^*) = 0$.

Integrate equation (44) to obtain:

$$U(\theta) - \underbrace{U(\theta^*)}_{0} = \int_{\theta^*}^{\theta} \mathbb{E}_{\tilde{\theta}} \left[\frac{\partial u}{\partial \theta}(\tilde{\theta}, x))\right] d\tilde{\theta}.$$

Substitute in the firm's expected profits and use integration by parts to rewrite the firm's profits as

$$\int_{\theta^*}^{\bar{\theta}} \left\{\mathbb{E}_\theta \left[S(\theta, x)\right] - \mathbb{E}_\theta \left[\frac{\partial u}{\partial \theta}(\theta, x))\right] \cdot \frac{1 - F(\theta)}{f(\theta)}\right\} f(\theta) d\theta.$$

For the moment, suppose $\underline{\theta} = 0$. To show that any optimal allocation excludes some types, we verify that the integrand of the expression above is negative for θ close enough to zero.

First, note that the integrand evaluated at $\theta = 0$ equals:

$$\mathbb{E}_{\theta=0} \left[S(\theta, x) - \frac{\partial u}{\partial \theta}(\theta, x)) \cdot \frac{1}{f(0)}\right].$$

If the mechanism assigns positive mass to $x > 0$ when $\theta = 0$, then this expression is strictly negative (since $S(0, x) \leq 0$ for all x and $\frac{\partial u}{\partial \theta}(0, x) \geq 0$ with $>$ for $x > 0$ (Assumption 1). Thus, setting $\theta^* > 0$ is optimal.

If the mechanism assigns full mass to $x = 0$ when $\theta = 0$, the expression above equals zero. Evaluating the derivative of the term inside expectation with respect to x, we obtain

$$\frac{\partial S}{\partial x}(0, 0) - \frac{\partial^2 u}{\partial x \partial \theta}(0, 0) \cdot \frac{1}{f(0)} < 0,$$

where the inequality follows from $\frac{\partial S}{\partial x}(0, 0) \leq 0$ and $\frac{\partial^2 u}{\partial x \partial \theta}(0, 0) > 0$ (Assumption 3). Since the mechanism assigns positive mass to $x > 0$ for all $\theta > \theta^* = 0$ (types above θ^* are not excluded), by the continuity of the expression on the LHS, it follows that the integrand is strictly negative for all θ in a neighborhood of $\theta^* = 0$. Since S is continuously differentiable and u is twice continuously differentiable, and the integrand remains negative if $\underline{\theta}$ is sufficiently small.