Perfect Competition in Markets with Adverse Selection*

Eduardo M. Azevedo and Daniel Gottlieb†

This version: November 28, 2014
First version: September 2, 2014

Abstract

This paper proposes a competitive model of adverse selection. The model allows for endogenous contract characteristics and multidimensional heterogeneity among consumers. Equilibria exist and yield sharp predictions under general conditions. Equilibria are inefficient in the Kaldor-Hicks sense, and government intervention can increase welfare.

We apply the model to the equilibrium effects of a mandate in a calibrated health insurance model. In the calibration the mandate increases welfare by $56 per consumer, but has large unintended consequences. The influx of low-cost types into the market exacerbates adverse selection in the intensive margin, above the minimum mandated coverage, driving many consumers to lower quality products in equilibrium. Policies that address selection in the intensive margin are closely related to risk-adjustment, and increase welfare by $291 per consumer.

*Preliminary, please do not cite. We would like to thank Alberto Bisin, Pierre-André Chiappori, Alex Cittana, Vitor Farinha Luz, Piero Gotardi, Lucas Maestri, Humberto Moreira, Roger Myerson, Kent Smetters, Phil Reny, Bernard Salanié, André Veiga, Glen Weyl, and seminar participants at the Chicago, Columbia, EUI, Oxford, UCL and Upenn for helpful discussions and suggestions.

†Wharton. eazevedo@wharton.upenn.edu (Azevedo) and dgott@wharton.upenn.edu (Gottlieb).

1 Introduction

Policy makers and market participants both consider adverse selection a first-order concern in many markets. These markets are often heavily regulated, if not subject to outright government provision, as in social programs like unemployment insurance and Medicare. Government interventions are typically complex, involving regulation of contract characteristics, personalized subsidies, community rating, risk-adjustment, and mandates. However, most models of competition take contract characteristics as given, considerably limiting the scope of normative and even positive analyses of these policies.

Standard models face three basic limitations. The first limitation arises in the Akerlof (1970) model, which, following Einav et al. (2010a), is used by most recent applied work.[1] The Akerlof model considers a market for a single contract with exogenous characteristics, so that it is impossible to consider the effect of policies that affect contract terms.[2] In contrast, the Spence (1973) and Rothschild and Stiglitz (1976) models do allow for endogenous contract characteristics. However, they restrict consumers to be heterogeneous along a single dimension,[3] despite evidence on the importance of multiple dimensions of heterogeneity.[4] Moreover, the Spence model suffers from the issue of rampant multiplicity of equilibria, while the Rothschild and Stiglitz model often has no equilibrium.[5]

This paper develops a competitive model of adverse selection. Motivated by the central role of contract characteristics in policy discussions and by recent empirical findings, the model incorporates three key features. First, the set of transacted contracts is endogenous,

[1] Recent papers using this framework include Handel et al. (2013); Hackmann et al. (2014); Mahoney and Weyl (2014); Scheuer and Smetters (2014).

[2] Despite the lack of existing theory, many authors have suggested that determination of contract characteristics is important. For example, Einav and Finkelstein (2011) say that "abstracting from this potential consequence of selection may miss a substantial component of its welfare implications [...]. Allowing the contract space to be determined endogenously in a selection market raises challenges on both the theoretical and empirical front. On the theoretical front, we currently lack clear characterizations of the equilibrium in a market in which firms compete over contract dimensions as well as price, and in which consumers may have multiple dimensions of private information." According to Einav et al. (2009), "analyzing price competition over a fixed set of coverage offerings [...] appears to be a relatively manageable problem, characterizing equilibria for a general model of competition in which consumers have multiple dimensions of private information is another matter. Here it is likely that empirical work would be aided by more theoretical progress."

[3] Chiappori et al. (2006) highlight this shortcoming: "Theoretical models of asymmetric information typically use oversimplified frameworks, which can hardly be directly transposed to real-life situations. Rothschild and Stiglitz's model assumes that accident probabilities are exogenous (which rules out moral hazard), that only one level of loss is possible, and more strikingly that agents have identical preferences which are moreover perfectly known to the insurer. The theoretical justification of these restrictions is straightforward: analyzing a model of "pure," one-dimensional adverse selection is an indispensable first step. But their empirical relevance is dubious, to say the least."

[4] See Finkelstein and McGarry (2006); Cohen and Einav (2007); Fang et al. (2008).

[5] See Chiappori et al. (2006) and Myerson (1995). According to Chiappori et al. (2006), "As is well known, the mere definition of a competitive equilibrium under asymmetric information is a difficult task, on which it is fair to say that no general agreement has been reached."

allowing us to study policies that affect contract characteristics. Second, consumers may be heterogeneous in several dimensions, engage in moral hazard, and exhibit deviations from rational behavior such as inertia and overconfidence.[6] Third, equilibria always exist and, at the same time, yield sharp predictions. We find that equilibria are often inefficient. These features arise due to three key modeling choices. First, we take the set of potential contracts as given and model trade as in standard competitive goods markets. Second, firms and consumers are price-takers, and prices are determined by zero profits conditions. Finally, which contracts are transacted is determined by whether entry into markets for each contract is profitable.

We first describe the model and discuss implications for market failure and policy below. We consider a set of consumers and a set of exclusive contracts. Consumers have preferences over contracts and the price paid, in terms of a numeraire. To model asymmetric information, we assume that the cost of supplying a contract depends on the identity of the buyer. Few restrictions are placed on cost functions, preferences, or contracts.

We consider a price-taking equilibrium notion, as in Akerlof (1970) and Spence (1973). A weak equilibrium is a set of prices and an allocation such that all consumers optimize taking prices as given, and prices equal the expected cost of supplying each contract. As in Spence's model, this leads to a large number of weak equilibria because the expected cost of non-transacted contracts is arbitrary. For example, there are always weak equilibria where no contracts are bought because prices are high, and prices are high because the expected cost of a non-transacted contract is arbitrary. These equilibria would be inconsistent with the idea of free entry, however, if some firms were able to enter this market, driving down prices and making positive profits.

To incorporate free entry, we consider one additional equilibrium requirement. Unlike most of the literature on signaling and competitive screening games, we do not attempt to refine the equilibrium with a game-theoretic model. Instead, we ask that equilibria are robust to a small perturbation of fundamentals. Namely, equilibria must survive in economies with a set of contracts that is similar to the original, but finite, and with a small mass of agents who demand all contracts and have low costs. In such a perturbation, all contracts are traded in equilibrium, sidestepping the pathologies associated with conditional expectation. The second part of the refinement is similar to that used by Dubey and Geanakoplos (2002) in a model of competitive pools. We find that equilibria always exist, a conclusion that does not depend on the pool structure. Moreover, our equilibrium notion makes sharp predictions in the wide range of applied models that are particular cases of our framework. The relationship between this equilibrium and existing concepts is discussed in detail below.

[6] See Spinnewijn (forthcoming) on overconfidence, Handel (2013) on inertia and Cutler and Zeckhauser (2004); Kunreuther and Pauly (2006) for a general discussion of behavioral biases in insurance markets.

After covering the model (Section 2) and the competitive equilibrium (Section 3), we turn to efficiency. Section 5 shows that competitive equilibria are often inefficient in the Kaldor-Hicks sense. Because every individual faces the same prices, price differences may not reflect social costs. This wedge can be divided into two components. First, because of multidimensional heterogeneity, agents purchasing the same contract can have different cost functions. Therefore, the social cost of providing more insurance to any particular agent does not equal the incremental price. Moreover, because of adverse or advantageous selection, even on average, incremental costs do not equate price differences.

This point implies that there is considerable scope for government intervention in markets with adverse selection, even if the government is not better informed than firms. We consider, in a stripped-down setting, government intervention aiming to maximize Kaldor-Hicks efficiency. We find that optimal regulation involves a modified risk-adjustment formula, similar to policies that are used in many insurance markets. A limitation of this basic optimal regulation formula is that it ignores important practical issues such as redistribution. Taking these issues into account would change the optimal regulation, requiring more than just risk adjustment. In the interest of space, we relegate a detailed discussion of optimal regulation and policy implications to Azevedo and Gottlieb (in preparation).

2 Model

2.1 The Model

We consider competitive markets, with a large number of consumers and free entry of identical firms operating at an efficient scale that is small relative to the market. To model the gamut of behavior relevant to policy discussions in a simple way, we take as given a set of potential contracts, preferences, and costs of supplying contracts. This approach is similar to Veiga and Weyl (2014a,b) and to the value and cost functions of Einav et al. (2009, 2010a). Without loss of generality, we restrict our analysis to a group of consumers who are indistinguishable with respect to characteristics over which firms are able to price discriminate.

Formally, the model is as follows. Firms offer **contracts** (or products) x in X. Consumer **types** are denoted θ in Θ. Consumer type θ derives utility $U(x, p, \theta)$ from buying contract x at a price p, and it costs a firm $c(x, \theta) \geq 0$ measured in units of the numeraire to supply it. Utility is strictly decreasing in price. There is a positive mass of consumers, and the **distribution of types** is a measure μ.[7] An **economy** is defined as $E = [\Theta, X, \mu]$.

[7] The relevant σ-algebra is defined below, along with the necessary assumptions.

2.2 Discussion and Clarifying Examples

The following examples clarify the definitions, limitations of the model, and the goal of deriving robust predictions in a wide range of selection markets. Note that the parametric assumptions in the examples are of little consequence to the general analysis. Some readers may therefore prefer to skim over details. We begin with the classic Akerlof (1970) model, which is the dominant framework in applied work. It is simple enough that the literature mostly agrees on equilibrium predictions.

Example 1. (Akerlof) Consumers choose whether to buy a single insurance product, so that $X = \{0, 1\}$. Utility is quasilinear,

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

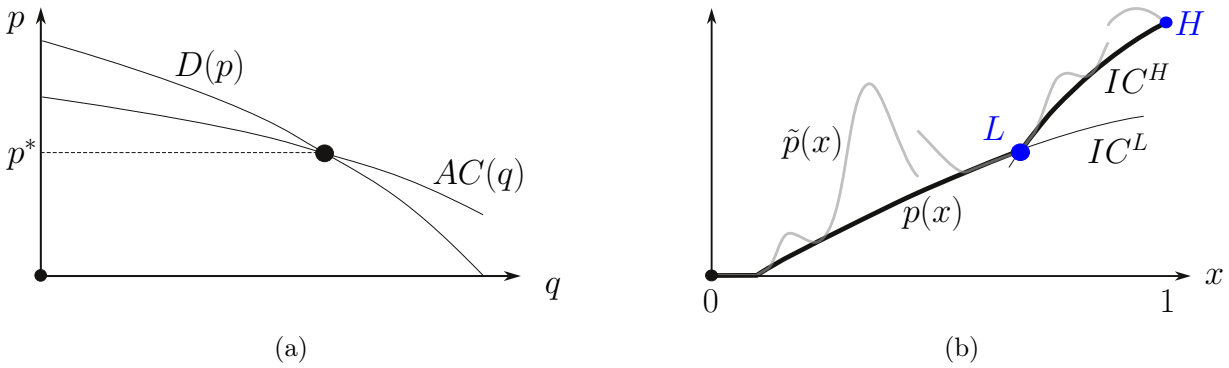
and the contract $x = 0$ generates no cost or utility, $u(0, \theta) \equiv c(0, \theta) \equiv 0$. Thus, in equilibrium it has a price of 0. Under these assumptions, all that matters is the joint distribution of willingness to pay $u(1, \theta)$ and costs $c(1, \theta)$, induced by the measure μ.

A competitive equilibrium in the Akerlof model has a compelling definition and is amenable to an insightful graphical analysis. Following Einav et al. (2010a), let the demand curve $D(p)$ be the mass of agents with willingness to pay higher than p, and let $AC(q)$ be the average cost of the q consumers with highest willingness to pay.[8] Then, an equilibrium in the Akerlof (1970) and Einav et al. (2010a) sense is determined by the intersection between the demand and average cost curves, depicted in Figure 1a. At this price and quantity, consumers behave optimally, and the price of insurance equals the expected cost of providing coverage. □

Because of its simplicity and tractability, the Akerlof model became the dominant framework in applied research. However, the model is restrictive in two important ways. First, contract terms are exogenous. This is important because market participants and regulators often see distortions in contract terms as being crucial. In fact, many of the interventions in markets with adverse selection regulate contract dimensions directly, aim to affect them indirectly, or try to shift demand from some type of contract to another. It is impossible to consider the effect of these policies in the Akerlof model. Second, the Akerlof model assumes a single non-null contract. This is also considerably restrictive. For example, Handel et al. (2013) approximate real-life health exchanges by assuming that they offer only two types of plans (corresponding to $x = 0$ and $x = 1$), and that consumers are forced to choose one of

[8] Formally, under appropriate assumptions, the definitions are

$$\begin{aligned} D(p) &= \mu(\{\theta : u(1, \theta) \geq p\}) \\ AC(q) &= \mathbb{E}[c(1, \theta) | \mu, u(1, \theta) \geq D^{-1}(q)]. \end{aligned}$$



(a) (b)

Figure 1: Weak equilibria in the (a) Akerlof and (b) Rothschild and Stiglitz models.

Notes: Panel (a) depicts demand $D(p)$ and average cost $AC(p)$ curves in the Akerlof model, with quantity in the horizontal axis, and prices in the vertical axis. The equilibrium price of contract $x = 1$ is denoted by p^*. Panel (b) depicts two weak equilibria of the Rothschild and Stiglitz model. The horizontal axis depicts contracts, and the vertical axis prices. IC^L and IC^H are indifference curves of type L and H consumers. L and H denote the contract-price pairs chosen by each type in these weak equilibria, which are the same as in Rothschild and Stiglitz (1976) when Bertrand equilibrium exists. The relative size of the circles L and H represents the mass of agents of each type. The bold curves $p(x)$ (black) and $\tilde{p}(x)$ (gray) depict two weak equilibrium price schedules. $p(x)$ is also a (refined) equilibrium price, but $\tilde{p}(x)$ is not.

them.[9] Likewise, Hackmann et al. (2014) and Scheuer and Smetters (2014) lump the choice of buying any health insurance as $x = 1$.

We now consider another classic example, the Rothschild and Stiglitz (1976) model. Here, the comprehensiveness of coverage is endogenous, although preferences are highly stylized. Still, this model already exhibits problems with existence of equilibrium, and there is no consensus about equilibrium predictions.

Example 2. (Rothschild and Stiglitz) Each consumer may buy an insurance contract in $X = [0, 1]$, which insures her for a fraction x of a possible loss of l. Consumers only differ in the probability θ of a loss. Their utility is

$$U(x, p, \theta) = \theta \cdot v(W - p - (1 - x)l) + (1 - \theta) \cdot v(W - p),$$

where $v(\cdot)$ is a Bernoulli utility function and W is wealth, both of which are constant in the population. The cost of insuring individual θ with policy x is $c(x, \theta) = \theta \cdot x \cdot l$. The set of types is $\Theta = \{L, H\}$, with $0 < L < H < 1$. The definition of an equilibrium in this model is a matter of considerable debate, which we address in the next section. □

We now illustrate how the model can fit more realistic multidimensional heterogeneity

[9] The Affordable Care Act allows insurers to offer four types of plan in exchanges, bronze, gold, silver and platinum, with approximate actuarial values ranging from 60% to 90%.

with an empirical model of preferences for health insurance used by Einav et al. (2013).

Example 3. (Einav et al.) Consumers are subject to a stochastic health shock l and, after the shock, decide the amount e they wish to spend on health services. Consumers are heterogeneous in their distribution of health shocks F_θ, risk aversion parameter A_θ, and moral hazard parameter H_θ.

For simplicity, we assume that insurance contracts specify the fraction $x \in X = [0, 1]$ of health expenditures that are reimbursed. Utility after the shock equals

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

where W is the consumer's initial wealth. The privately optimal health expenditure is $e = l + H_\theta \cdot x$, so that in equilibrium

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

Einav et al. (2013) assume constant absolute risk aversion (CARA) preferences before the health shock, so that ex-ante utility equals

$$U(x, p, \theta) = \mathbb{E}[-\exp\{-A_\theta \cdot CE^*(l; x, p, \theta)\}|l \sim F_\theta].$$

For our numerical examples below, we assume that losses are normally distributed with mean M_θ and variance S_θ^2, which leaves us with four dimensions of heterogeneity.[10] Calculations show that the model can be described with quasilinear preferences as in equation (1), with willingness to pay and cost functions

$$\begin{aligned} u(x, \theta) &= x \cdot M_\theta + \frac{x^2}{2} \cdot H_\theta + \frac{1}{2}x(2 - x) \cdot S_\theta^2 A_\theta, \text{ and} \\ c(x, \theta) &= x \cdot M_\theta + x^2 \cdot H_\theta. \end{aligned} \tag{2}$$

The formula decomposes willingness to pay into three terms. The first term xM_θ equals average covered expenses, which is also part of firm costs. The second term $x^2 H_\theta/2$ represents the utility from overconsumption of health services, which is caused by moral hazard. Note that the moral hazard component costs twice as much to firms as consumers are willing to pay for it. Finally, the last term represents the insurance value of policy x, which is greater for agents who are more risk averse and who have higher variance in health status and costs zero to risk-neutral firms. □

[10] In particular, losses and expenses are allowed to be negative in the numerical example. We report this parametrization because the closed form solutions for utility and cost functions make the model more transparent, and because qualitative features of equilibria are similar to more realistic parametric assumptions.

The previous example illustrates that the framework can fit multidimensional heterogeneity in a realistic empirical model. Moreover, it shows that it can incorporate ex-post moral hazard through the definitions of the utility and cost functions. It should be clear that the model can fit many other types of consumer behavior, such as ex-ante moral hazard, non-expected utility preferences, overconfidence, or inertia to abandon a default choice. It can also incorporate administrative or other per-unit costs on the supply side. Moreover, it is straightforward to consider more complex contract spaces, such as deductibles, copays, stop-losses, franchises, network quality and so on.

Note that in this last example, and in other models with complex contract spaces and rich heterogeneity, there is no agreement on a reasonable equilibrium prediction. Unlike the Rothschild and Stiglitz model, where there is controversy about what the correct prediction is, in this case the literature offers almost no possibilities.[11]

2.3 Assumptions

We now discuss the assumptions used in our analysis. The restrictions are mild enough to include all examples above, so applied readers may wish to skip this section. On a first read, it may be useful to keep in mind the particular case where X and Θ are compact subsets of Euclidean space, utility is quasilinear as in equation (1), and u and c are continuously differentiable. These assumptions are considerably stronger than what is necessary for our results, but they are weak enough to incorporate most models in the literature.

We begin with technical assumptions. Note that X and Θ can be infinite dimensional, and the distribution of types can admit a density with infinite support, may be a sum of point masses, or a combination of the two.

Assumption 1. *(Technical Assumptions) X and Θ are compact and separable metric spaces. Whenever referring to measurability we will consider the Borel σ-algebra over X and Θ, and the product σ-algebra over the product space. In particular, we take μ to be measurable with respect to the Borel σ-algebra.*

We now consider a more substantive assumption.

Assumption 2. *(Bounded Marginal Rates of Substitution) There exists a constant L with the following property. Take any $p \leq p'$ in the image of c, any x, x' in X, and any $\theta \in \Theta$. Assume that*

$$U(x, p, \theta) \leq U(x', p', \theta),$$

[11] The only competitive model we know of that is general enough to encompass this example is Veiga and Weyl's (2014b) oligopoly model that we discuss below.

that is, that a consumer prefers to pay more to purchase contract x' instead of x. Then, the price difference must be bounded by

$$p' - p \leq L \cdot d(x, x').$$

The bounded marginal rate of substitution assumption is simpler to understand in the quasilinear case, where it is equivalent to u being Lipschitz in x. In the particular case that u is differentiable, this is equivalent to the absolute value of the derivative of u being uniformly bounded. That is, there exists some maximum willingness to pay for an additional unit of any given dimension of insurance.

Assumption 3. *(Continuity) The functions U and c are continuous.*

Continuity of the utility function is not very restrictive because of Berge's Maximum Theorem. That is, even in a model with moral hazard, if contract dimensions vary continuously, under standard assumptions, the maximum utility an individual can obtain varies continuously with the contract and with one's characteristics. Continuity of the cost function is more restrictive. It means we can only consider models with moral hazard where the payoffs to the firm vary continuously with agent types and with the contract offered. This will typically fail in models where agents change their actions discontinuously with small changes in a contract. Nevertheless, it is still possible to include some models with moral hazard in our framework. See Kadan et al. (2014) Section 9 for a discussion of how to define a metric over a contract space, starting from a description of actions and states.

3 Competitive Equilibrium

3.1 Definition of a Weak Equilibrium

We now define a minimalistic equilibrium notion, a weak equilibrium, requiring only that firms make no profits and consumers optimize. We need the following definitions. A vector of **prices** is a measurable function $p : X \to \mathbb{R}$, with $p(x)$ denoting the price of contract x. An **allocation** is a measure α over $\Theta \times X$ such that the marginal distribution satisfies $\alpha|\Theta = \mu$. That is, $\alpha(\{\theta, x\})$ is the measure of θ types purchasing contract x.[12] Because we are often interested in the expected cost of supplying a contract x, we use the following shorthand notation for conditional moments:

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

[12]Note that this formalization is somewhat different than the traditional way of denoting an allocation as a map from types to contracts. We take this approach because we allow different agents of the same type to buy different contracts in equilibrium.

That is, $\mathbb{E}_x[c|\alpha]$ denotes the expectation of $c(\tilde{x}, \tilde{\theta})$ according to the measure α and conditional on $\tilde{x} = x$. Note that such expectations depend on the allocation α. When there is no risk of confusion we omit α, writing simply $\mathbb{E}_x[c]$. Similar notation is used for other moments. We can now formally define a weak equilibrium.

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

1. For each contract x firms make no profits. Formally,

$$p^*(x) = \mathbb{E}_x[c|\alpha^*]$$

   almost everywhere according to α^*.

2. Consumers select contracts optimally. Formally, for almost every (θ, x) with respect to α^*, we have

$$U(x, p^*(x), \theta) = \sup_{x' \in X} U(x', p^*(x'), \theta).$$

We clarify some points on the definition. Note first that this is a price-taking definition, not a game-theoretic definition. Consumers optimize taking prices as given, and so do firms, who also take the average costs of buyers as given. Note also that we did not include an individual rationality requirement. This can be modeled by assuming that there exists a null contract that costs nothing and provides zero utility as in the Akerlof Example.

A final observation is that we only ask that prices equal expected costs almost everywhere according to α.[13] In particular, this means that weak equilibria place no restrictions on the prices of contracts that are not purchased. As demonstrated in the examples below, this is a serious problem with this definition, and the reason why a stronger equilibrium notion is necessary to make sharp predictions.

3.2 Equilibrium Multiplicity and Free Entry

We now illustrate that weak equilibria are compatible with a wide variety of outcomes, most of which are unreasonable in a competitive marketplace.

Example 2′. (Rothschild and Stiglitz - Multiplicity of Weak Equilibria) We first revisit Rothschild and Stiglitz's (1976) original equilibrium prediction. They set up a Bertrand game with identical firms and showed that, when a Nash equilibrium exists, it has allocations given by the points L and H in Figure 1b. High-risk consumers buy full insurance $x_H = 1$ at

[13] The reason is that conditional expectation is only defined almost everywhere. Although it is possible to understand all of our substantive results without recourse to measure theory, we refer interested readers to Billingsley (2008) for a formal definition of conditional expectation.

actuarially fair rates $p_H = H \cdot l$. Low-risk types purchase partial insurance, with actuarially fair prices reflecting their lower risk. Note that the level of insurance x_L is just low enough so that high-risk agents do not wish to purchase contract x_L. That is, L and H are in the same indifference curve IC^H of high types.

Note that we can find weak equilibria (α^*, p^*) with the same allocation. One example of weak equilibrium prices is the curve $p(x)$ is Figure 1b. The zero profits condition is satisfied because the prices of the two contracts that are transacted, x_L and x_H, equal the average cost of providing them. The optimization condition is also satisfied because the price schedule $p(x)$ is above the willingness to pay given by the curves IC^L and IC^H. Therefore, no consumer wishes to purchase a different bundle.

By the same logic, however, many other weak equilibria exist. One example is given by the same allocation and the $\tilde{p}(x)$ curve in Figure 1b. Again, firms make no profits because the prices of x_H and x_L are actuarially fair, and consumers are optimizing because the price of other contracts is higher than their indifference curves. One may find that the prices of other contracts are unreasonable in some sense, but they do not violate the zero profits condition because the expected cost of contracts that are not transacted is arbitrary.

There are multiple weak equilibria with completely different allocations. For example, it is a weak equilibrium for all agents to purchase full insurance, and for all other contracts to be priced so high that no one wishes to buy them. This does not violate the zero profits condition because, again, the conditional expectation of costs of contracts that are not transacted is arbitrary. This weak equilibrium has full insurance, which is the first-best outcome in this model. It is also a weak equilibrium for no insurance to be sold, and for prices of all contracts with positive coverage to be prohibitively high. Therefore, weak equilibria provide very coarse predictions, with the Bertrand solution, full insurance, complete unraveling, and many other outcomes all being possible. □

In a market with free entry, however, some weak equilibria are more reasonable than others. Take the weak equilibrium associated with $\tilde{p}(x)$, which many would say is less reasonable than that associated with $p(x)$. Let $x_0 < x_L$ be a non-transacted contract with $\tilde{p}(x_0) > p(x)$. Suppose firms enter the market for x_0, driving down its price. Initially no consumers would purchase x_0, and firms would continue to break even. As prices decrease enough to reach $p(x_0)$, the L types become indifferent between purchasing x_0 or not. If they decrease any further, all the L types would purchase contract x_0. At this point, firms would lose money because the average cost would be higher than the price.[14] This logic suggests that the price of x_0 would be driven down to $p(x_0)$, at which point it is no longer advantageous for firms to enter. By a similar argument, the price of other contracts would be driven down, and we

[14] To see why, note that in this weak equilibrium the L types buy x_L at an actuarially fair price. Therefore, the only reason why they would purchase less insurance is if firms sell it at a loss.

would end up with the equilibrium associated with $p(x)$. Finally, note that a similar tâtonnement would eliminate the full insurance and the complete unravelling weak equilibria. We are then left with $p(x)$ as the unique equilibrium price schedule consistent with free entry.

3.3 Definition of an Equilibrium

We now define an equilibrium concept that formalizes the free entry argument, eliminating the unreasonable predictions. The definition works by considering small perturbations of a given economy. A perturbation has a large but finite set of contracts approximating X. The perturbation adds a small measure of behavioral agents, who always purchase each of the existing contracts and impose no costs on firms. The point of considering the perturbed economy is that all contracts are transacted, eliminating the paradoxes associated with defining the average cost of non-transacted contracts.

We introduce, in the perturbed economies, for each contract x, a behavioral consumer type who always demands contract x. We will use the notation x for such a behavioral agent, and extend the utility and cost functions so that $U(x, p, x) = \infty$, $U(x', p, x) = 0$ if $x' \neq x$, and $c(x, x) = 0$. For clarity, we refer to the non-behavioral types as the standard types.

Definition 2. Consider an economy $E = [\Theta, X, \mu]$. A perturbation of the economy E is an economy where the set of contracts is a finite set $\bar{X} \subseteq X$, and there is a small mass of behavioral agents demanding each contract in $\bar{X}$. Formally, a **perturbation** $(E, \bar{X}, \eta)$ is an economy $[\Theta \cup \bar{X}, \bar{X}, \mu + \eta]$, where $\bar{X} \subseteq X$ is a finite set, and η is a strictly positive measure over $\bar{X}$.

Note that both an economy and its perturbations have a set of types contained in $\Theta \cup X$, and contracts contained in X. To save on notation, we extend distributions of types to be defined over $\Theta \cup X$ and allocations to be defined over $(\Theta \cup X) \times X$. With this notation, measures pertaining to different perturbations are defined in the same space. We now define convergence of a sequence of perturbations to the original economy.

Definition 3. A sequence of perturbations $(E, \bar{X}^n, \eta^n)_{n\in\mathbb{N}}$ **converges to** E if

1. Every point in X is the limit of a sequence $(x^n)_{n\in\mathbb{N}}$ with each $x^n \in \bar{X}^n$.

2. The mass of behavioral types $\eta^n(\bar{X}^n)$ converges to 0.

The first condition says that the finite sets of contracts in the perturbations approach the actual set of contracts, in the sense that every contract in X is the limit of a sequence of contracts in the perturbations. The second condition says that the total mass of behavioral types converges to 0. We now define the convergence of weak equilibria of perturbations.

Definition 4. Take an economy E and a sequence of perturbations $(E, \bar{X}^n, \eta^n)_{n\in\mathbb{N}}$ converging to E, with weak equilibria (p^n, α^n). We say that the **sequence of perturbed equilibria** $(p^n, \alpha^n)_{n\in\mathbb{N}}$ **converges to a price allocation pair** (p^*, α^*) of E if

1. The allocations α^n converge weakly to α^*.

2. For every sequence $(x^n)_{n\in\mathbb{N}}$, with each $x^n \in \bar{X}^n$ and limit $x \in X$, $p^n(x^n)$ converges to $p^*(x)$.

Condition 1 asks that the allocations in the sequence of equilibria converge. Condition 2 demands that the price schedules converge. Note that, in a perturbation, prices are only defined for a finite subset $\bar{X}^n$ of contracts. The definition of convergence is strict in the sense that, for a given contract x, prices must converge to the price of x for any sequence of contracts x^n converging to x. We are now ready to define an equilibrium.

Definition 5. The pair (p^*, α^*) is a (refined) **equilibrium** of E if there exists a sequence of perturbations that converges to E, and an associated sequence of weak equilibria that converges to (p^*, α^*).

This is still a minimalistic definition. Besides the weak equilibrium requirements of zero profits and consumer optimization,[15] we only ask that an equilibrium is robust to perturbations of the original economy. This last requirement imposes discipline on the prices of contracts that are not transacted, because all contracts are transacted in the perturbations. More specifically, in any perturbation, behavioral agents pull down the prices of contracts that have small volume, ruling out unreasonable weak equilibria such as $\tilde{p}$ in Figure 1b.

This definition makes reasonable predictions in all of our examples. The only equilibrium of the Rothschild and Stiglitz model (Example 2) is the one displayed in Figure 1b. Likewise, the only equilibrium in the Akerlof model (Example 1) is the one in Figure 1a. More interestingly, this notion makes predictions in more complex multidimensional models. This is illustrated in our numerical calibration of the Einav et al. health insurance model (Example 3), which is described in Section 4 and depicted in Figure 2a.

3.4 Existence and Properties of Equilibria

We now show that equilibria always exist. Existence only depends on the assumptions of Section 2.3. Therefore, in a broad range of theoretical and empirical models, equilibria are always well defined.

Theorem 1. *Every economy has an equilibrium.*

[15] While the definition does not formally require than an equilibrium be a weak equilibrium, it follows from Proposition 1 that this is always the case.

The proof is based on three lemmas. Lemma A1 uses a standard fixed point argument to show that every perturbation has a weak equilibrium. The crucial observation is Lemma A2, which shows that weak equilibrium price schedules in any perturbation are uniformly Lipschitz. Informally, this means that prices of similar contracts are not too different. This lemma is a consequence of the bounded marginal rate of substitution (Assumption 2). The intuition is that, if prices increased too fast with x, no standard types would be willing to purchase more expensive contracts. This is impossible, however, because a contract cannot have a high equilibrium price if it is only purchased by the low-cost behavioral types. Lemma A3 then applies this result and the Arzelà–Ascoli Theorem to demonstrate existence of equilibria.

Our existence result relies on the price-taking assumption in the equilibrium definition. Most of the literature has studied game-theoretic models of Bertrand competition. In general, equilibria of Bertrand games may not correspond to competitive equilibria, even without adverse selection. In particular, it is not true that equilibria in our model correspond to Bertrand equilibria, nor that Bertrand equilibria always exist. This can be seen clearly in Example 2, where Bertrand equilibria may not exist.

We now turn to basic properties of equilibria.

Proposition 1. *Let (p^*, α^*) be an equilibrium of economy E. Then:*

1. *The pair (p^*, α^*) is a weak equilibrium of E.*

2. *For every contract $x' \in X$ with strictly positive price there exists (θ, x) in the support of α^* such that*

$$U(x, p^*(x), \theta) = U(x', p^*(x'), \theta),$$

   and moreover $c(x', \theta) \geq p(x')$. That is, every such contract that is not transacted in equilibrium has a low enough price so that some consumer is indifferent between buying it or not, and such a consumer has costs at least as high as the price.

3. *The price function is L-Lipschitz, and, in particular, continuous.*

4. *If X is a subset of Euclidean space, then p^* is Lebesgue almost everywhere differentiable.*

The proposition shows that equilibria have several regularity properties that were not required in the definition. For example, we did not assume that equilibrium price schedules are continuous. However, parts 3 and 4 show that in equilibrium price schedules do vary continuously with contract characteristics. Moreover, if the set of contracts is finite dimensional, then prices are differentiable almost everywhere. This is the case in the Rothschild and Stiglitz example, despite the distribution of types having point masses. Likewise, equilibrium prices are continuous in Example 3, as depicted in Figure 2a. In turn, part 1 shows that every

equilibrium is a weak equilibrium. Finally, part 2 shows that prices of any out-of-equilibrium contract must be either 0 or low enough for some type to be indifferent between buying it or not. Moreover, this consumer type must have costs that are at least as high as the price. If the price of a contract were so high that all agents strictly preferred not to purchase it, then no agents would buy that contract in a perturbation, which would drive prices down to 0. Intuitively, if prices exceeded the highest willingness to pay, firms would be able to enter, driving prices down without any of them making losses. The following result notes that the proposition offers a simple way to find equilibria.[16]

Corollary 1. *If (p^*, α^*) is the unique weak equilibrium satisfying any subset of conditions 2-4 of Proposition 1, then it is the unique equilibrium.*

Even though the corollary is a direct consequence of Theorem 1 and Proposition 1, it is a useful result from an applied perspective. The reason is that one may use it to find equilibria without having to consider a sequence of perturbations. For instance, applying the Corollary to the Rothschild and Stiglitz (1976) Example 2 is a simple way to prove that the unique equilibrium is that associated with $p(x)$ in Figure 1b.

3.5 Relationship with the Literature

We now discuss the relationship between our model and previous work on markets with adverse selection. Our price-taking approach is reminiscent of the early work by Akerlof (1970) and Spence (1973). In fact, the multiplicity of competitive equilibria is well-known since Spence's (1973) analysis of labor market signaling. Much like the discussion in Section 3.2, Spence (1973) noted that a large variety of outcomes are possible in unrefined price-taking equilibria (weak equilibria).

The literature addressed this multiplicity issue in three different ways. One strand of the literature employed game theoretic equilibrium notions and restrictions on consumer heterogeneity, typically in the form of ordered one-dimensional sets of types. This is the case in the competitive screening literature, initiated with Rothschild and Stiglitz's (1976) Bertrand game, which led to the issue of non-existence of equilibria. Subsequently, Riley (1979) showed that Bertrand equilibria do not exist for a broad (within the one-dimensional setting) class of preferences, including the standard Rothschild and Stiglitz model with a continuum of types. Among others, Wilson (1977), Miyazaki (1977), Riley (1979), and Netzer

[16] We note that Proposition 1 is not an alternative definition of equilibrium because not every weak equilibrium satisfying conditions 2-4 of Proposition 1 is an equilibrium. To see this, consider the case where X contains a null contract and a single non-trivial contract. At any point where the willingness to pay curve is tangent to the average cost curve, touching from the top, we have a weak equilibrium that is not an equilibrium.

and Scheuer (2014) proposed modifications of Bertrand equilibrium so that an equilibrium always exists.[17]

The literature on refinements in signaling games shares the features of game-theoretic equilibrium notions and restrictive type spaces. In order to deal with the multiplicity of price-taking equilibria described by Spence, this literature modeled signaling as a dynamic game. However, since signaling games typically have too many sequential equilibria, Banks and Sobel (1987), Cho and Kreps (1987), and several subsequent papers proposed equilibrium refinements that eliminate multiplicity.

Another strand of the literature considers price-taking equilibrium notions, like our work, but imposes additional structure on preferences, such as Bisin and Gottardi (1999, 2006). Notably, Dubey and Geanakoplos' (2002) model of competitive pools introduced a refinement of weak equilibria using behavioral agents. This is similar to our construction in the case of a finite set of contracts. The main difference between our work is that they consider the additional pool structure, a finite set of pools, and specification of states of the world and endowments for each household. Crucially, our framework imposes no such structure, and can thus accommodate standard applied models such as Example 3.

Our results are related to this previous work as follows. In standard one-dimensional models with single crossing, the unique equilibrium corresponds to what is usually called the "least costly separating equilibrium" in the signaling literature. Thus, our equilibrium prediction is the same as in models without cross-subsidies, such as Riley (1979) and Bisin and Gottardi (2006), and the same as Rothschild and Stiglitz (1976) when their equilibrium exists. It also coincides with Banks and Sobel (1987) and Cho and Kreps (1987) in the settings they consider. It differs from equilibria that involve cross-subsidization, such as Wilson (1977), Miyazaki (1977), Hellwig (1987), and Netzer and Scheuer (2014). In the case of a pool structure and finite set of contracts, our equilibria are the same as in Dubey and Geanakoplos (2002).

Another strand of the literature considers preferences with less structure. Closely related to the present paper, Chiappori et al. (2006) consider a very general model of preferences within an insurance setting. This paper differs from our work in that they consider general testable predictions without specifying an equilibrium concept, while we derive sharp predictions within an equilibrium framework. Rochet and Stole (2002) consider a competitive screening model with firms differentiated as in Hotelling (1929), where there is no adverse selection. Their Bertrand equilibrium converges to competitive pricing as differentiation vanishes, which is the outcome of our model. However, Riley's (1979) results imply that no Bertrand equilibrium would exist if one generalizes their model to include adverse selection.

[17] Luz (2013) characterizes the symmetric mixed strategy equilibrium in the Rothschild and Stiglitz model. There has also been work on this type of game with nonexclusive competition. Attar et al. (2011) show that nonexclusive competition leads to outcomes similar to the Akerlof model.

Like our paper, the model proposed by Veiga and Weyl (2014b) makes sharp predictions in a general multidimensional framework. Crucially for policy applications, product characteristics are determined endogenously. They consider an oligopoly model of competitive screening in the spirit of Rochet and Stole (2002), but where each firm can offer a single contract, and where contract characteristics are finite dimensional. This means that contract characteristics are determined by a simple first-order condition, as in the Spence (1975) model, yielding a tractable multidimensional model of competitive screening. Moreover, their model can incorporate imperfect competition, which is beyond the scope of the present work.

Interestingly, Veiga and Weyl's (2014b) predictions are strikingly different from ours. In their model, when perfectly competitive equilibria exist,[18] all firms offer the same contract.[19] In contrast, a rich set of contracts is offered in our equilibrium. For example, in the case of no adverse selection (when costs are independent of types), our equilibrium is for firms to offer all products priced at cost, which corresponds to the standard notion of a perfectly competitive outcome. A colorful illustration of this difference is a market like tomato sauce. While our prediction is for a broad offering of different types of tomato sauce being sold at cost, the Veiga and Weyl (2014b) model predicts a single type of tomato sauce being offered cheaply, with characteristics determined by the preferences of marginal consumers. The reason for this difference is that they model perfect competition as a market with few firms, and prohibitively high costs of introducing an additional product variety, in the limit where firms become undifferentiated. In contrast, we consider a perfectly competitive model with many buyers and sellers who are price-takers, with potential entrants for any product variety. Thus, we view the two models as complementary, describing different types of competitive limit, which are better suited to different markets.

4 Application: Equilibrium Effects of a Mandate

4.1 The Calibrated Model

To illustrate the equilibrium concept, we calibrated the multidimensional health insurance model in Example 3. In that example, consumer types vary along four dimensions: expected health loss, standard deviation of health losses, moral hazard, and risk aversion. Except for risk aversion, all parameters are measured in dollars and are interpretable from equation (2). We assumed that the distribution of parameters in the population is lognormal.[20] Moments

[18] Perfectly competitive equilibria do not always exist in their model and, in fact, do not exist in a model calibrated from the results of Handel et al. (2013).

[19] This is so in the more tractable case of symmetrically differentiated firms. In general, the number of contracts offered is no greater than the number of firms.

[20] Note that this means that the set of types is not compact in our numerical simulations. Restricting the set of types to a large compact set does not meaningfully impact the numerical results.

Table 1: Calibrated distribution of consumer types

	A	H	M	S
Mean	1.5E-5	1,330	4,340	24,474
Log covariance				
A	0.25	-0.01	-0.12	0
H		$\sigma^2_{\log H}$	-0.03	0
M			0.20	0
S				0.25

Notes: In all simulations of the health insurance example consumer type parameters are distributed lognormally, with the moments in the table.

of the type distribution were calibrated to match the central estimates of Einav et al. (2013), as in Table 1. Their data comes from health insurance choices by the US employees of a large corporation (Alcoa Inc.). The goal of our calibration is to qualitatively match the demand for insurance from this corporation, which in principle is representative of large employers. There were two moments that we did not simply take from their estimates. First, because they estimated their structural model on nonlinear contracts, whereas we use simple linear contracts, we reduced risk-aversion to better match substitution patterns in our setting. Second, we vary the the log variance of moral hazard in our simulations. In our baseline we set σ^2_H to 0.28, and later discuss how equilibrium changes with their higher central estimate of $\sigma^2_H = 0.98$. See Appendix B for details on the calibration.

4.2 Equilibrium without Government Intervention

To calculate an equilibrium, we used the following perturbation. We restricted ourselves to 26 evenly spaced contracts and added a mass equal to 1% of the population as behavioral agents. We then used a standard fixed point algorithm. Prices were set to an initial value. In each iteration, consumers chose optimal contracts taking prices as given. Prices are then adjusted up for unprofitable contracts and down for profitable contracts. Prices quickly and consistently converge to the same equilibrium for different initial values.

The equilibrium is depicted in Figure 2a. It features adverse selection with respect to average loss, in the sense that, on average, consumers who purchase more coverage have higher losses. Moreover, agents sort in complicated ways across contracts. This is illustrated in Figure 5a, which displays the contracts purchased by agents with different expected loss and risk aversion parameters. The figure corroborates the existence of adverse selection on average loss, as agents with higher expected losses tend to choose higher contracts. However, even for the same levels of risk aversion and expected loss, different consumers choose different contracts due to other dimensions of heterogeneity.

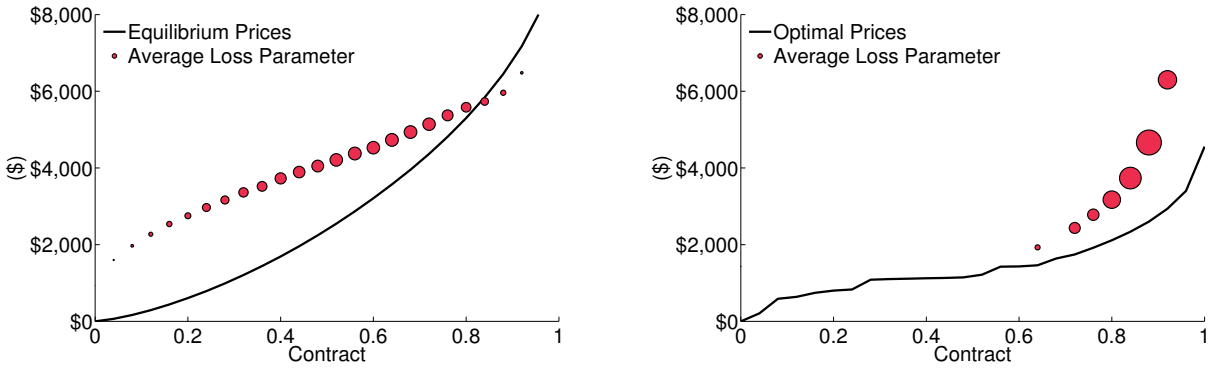


Figure 2: Equilibrium (a) and optimal (b) price schedules in the multidimensional health insurance model, Example 3.

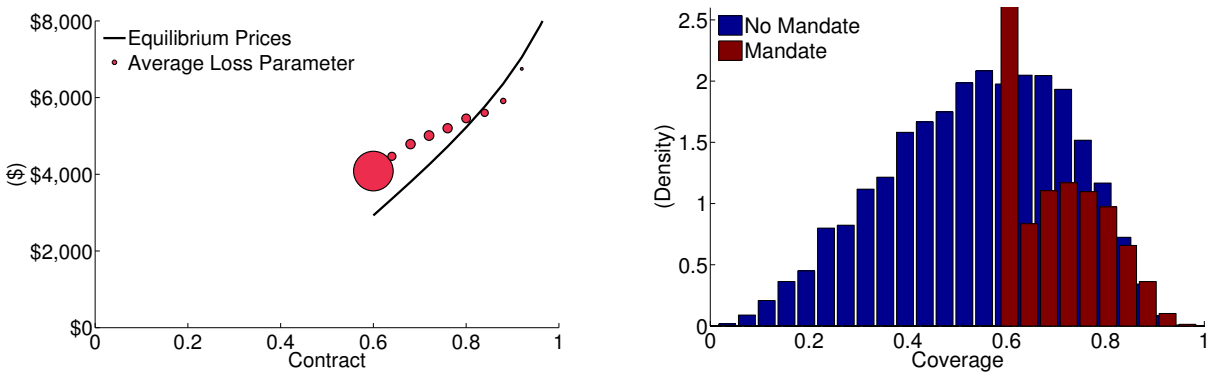
Notes: Panel (a) illustrates equilibrium, while panel (b) illustrates an optimal menu in our Example 3 based on the health insurance model of Einav et al. (2013). In both graphs the solid curve denotes prices. The size of the circles represent the mass of consumers purchasing each contract, and it's height represents the average loss parameter of such consumers, that is $\mathbb{E}_x[M]$. Calculations were based on a perturbation with a mass of behavioral agents equal to 5% of the population, and 26 evenly spaced contracts.

Although there is adverse selection, equilibria do not feature a complete "death spiral," in that all consumers purchase some contract. This is not true in general, and it is possible that no nontrivial contracts are purchased in equilibrium.[21] It is also possible that the support of traded contracts is a strict subset of all contracts. Whenever this is the case, the buyers with the highest willingness to pay for each contract that is not traded value it below their own average cost (Proposition 1). That is, markets for a the non-transacted contract must be shut down by an Akerlof-type death spiral within that market.

4.3 Equilibrium with a Mandate

We now apply the model to gauge the effects of a standard form of intervention, an insurance mandate. We consider the effects of mandating coverage of at least 60%. We chose 60% because it is the actuarial value of the lowest level of insurance offered in the Affordable Care Act (ACA) exchanges, a "bronze" plan. We highlight that we are taking the Einav et al. (2013) model too far from the data to reliably evaluate quantitative aspects of the policy – the point of calibrating their model is to evaluate qualitative effects in a demand model that still matches general features of demand in a typical US employer. Nevertheless, an advantage of choosing 60% is that our simulation is comparable to Handel et al. (2013)'s empirical model of health exchanges. They considered Riley equilibria in a market with a

[21] Recall that the Akerlof model, in which a death spiral sometimes occurs, is a particular case of our framework.



(a) (b)

Figure 3: Equilibrium with a mandate (a) and comparison of quantities with and without a mandate (b).

Notes: Panel (a) depicts prices in the black line, with the size of the circles proportional to quantity, and their height proportional to the average loss parameter of consumers purchasing each contract. Panel (b) compares the distribution of quantities under the mandate and the laissez-faire equilibrium, with the density at exactly 60% coverage censored for readability.

60% and a 90% policy, aiming to simulate the effects of the ACA. As we discuss below, despite several differences in our approaches, our results are remarkably consistent. As such we will focus the discussion on comparing the laissez-faire equilibrium with the mandate, an exercise that is only possible in our model with a rich set of products.

We computed equilibria in the same perturbation as in the laissez-faire case, but allowing consumers to purchase only contracts offering at least 60% coverage. The results are plotted in Figure 3. The first result is that about 75% of consumers purchase the minimum coverage of 60%. This is consistent with Handel et al. (2013)'s result that there is a high level of unravelling in their simulated health insurance exchanges. Upon some reflection, this result is not so surprising. Many consumers were purchasing less than 60% coverage in the laissez-faire equilibrium, and it is reasonable that a large number of these consumers will bunch at the minimum level of coverage.

A second finding is that many agents that originally purchased more than 60% purchase lower coverage after the mandate. This can be seen clearly from Figure 3b. For coverage just above 60%, demand falls to less than half of the laissez-faire quantities. The reason for this decrease in demand is that mandates can increase adverse selection in the intensive margin. The mandate forces many consumers with low costs to purchase 60% coverage, which reduces the price of that contract. In turn, this reduction in price induces agents who were purchasing higher coverage to switch to 60%. Adverse selection can increase even more if the consumers who move down to 60% tend to be cheaper.

Despite these unintended consequences, the mandate had a positive, albeit modest, impact

on welfare. Kaldor-Hicks welfare (defined formally in the next section) rises by \$56 per consumer. The clear winners of the policy were consumers who were already purchasing insurance close to 60%. The price of that contract was reduced with the mandate, benefitting those consumers. In fact, for our parameter values prices of all contracts above 60% coverage weakly decreased, so that all consumers who were previously purchasing 60% coverage or more gained from the influx of low-cost agents. We note that many agents are worse off under the mandate. This is often the case for agents who were choosing to purchase a low level of insurance. Many of these consumers would have preferred the lower coverage they initially had instead of switching to contracts above 60% coverage.

Summarizing, the calibration leads to three conclusions. First, mandates can induce a large number of agents to purchase the minimum level of insurance, especially when an unregulated equilibrium had many agents purchasing lower coverage (a similar finding to that of Handel et al. (2013)). Second, mandates can have large unintended consequences, considerably increasing adverse selection and reducing demand for coverage above the minimum level. Third, even with these unintended consequences, mandates can increase welfare (albeit this is not always the case, as demonstrated by Einav et al. (2010b)).

These three points lead to the question of whether more sophisticated policies can improve on mandates. The next section explains the reason for market failure in our model. We will then return to the calibration and show that policies that address adverse selection on the intensive margin can considerably increase welfare over mandates.

5 Market Failure

5.1 Inefficiency

We now show that information asymmetries about cost robustly generate market failures. In this sense, our model echoes the standard view of economists, regulators, and market participants that adverse and advantageous selection causes market outcomes to be inefficient.[22]

For clarity and to save on definitions, we proceed informally in this section and refer readers to Azevedo and Gottlieb (in preparation) for a formal mathematical treatment of efficiency and government intervention. We consider a simple set of contracts: $X = [0, 1]$. Moreover, we restrict attention to quasilinear preferences as in equation (1), ignoring income effects. We abstract from redistributive issues and investigate whether equilibria are efficient

[22]The reason for inefficiency is that, whenever a consumer purchases a contract and has a different cost than average buyers of the contract, she exerts an externality by affecting the average cost of the contract, which in turn affects firm profits. This externality is potentially present whenever costs are heterogeneous across types.

in the Kaldor-Hicks sense. In each given equilibrium (p, α), total welfare is

$$\int u - c \, d\alpha. \tag{3}$$

Assume that cost, utility, and the endogenously determined price functions are continuously differentiable.

To demonstrate the divergence between privately and socially optimal choices, fix a consumer θ. Denote marginal utility and marginal cost functions as $mu(x, \theta) = \partial_x u(x, \theta)$ and $mc(x, \theta) = \partial_x c(x, \theta)$. The consumer trades off the price paid against her preferences for contracts, as in Figure 4. In an interior choice such as x^{eq}, the consumer equates marginal utility and the price $p'(x)$ of an additional unit of coverage. However, the socially optimal choice x^{eff} equates marginal utility and marginal cost. Therefore, whenever costs vary among consumers, there is no reason to expect market outcomes to be efficient.

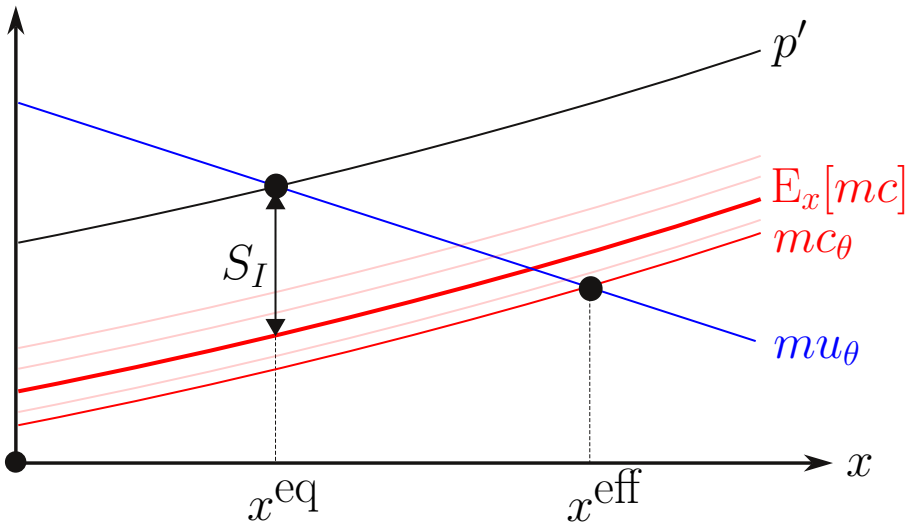


Figure 4: Inefficiency of private choices.

Notes: The figure illustrates the decision of an individual consumer θ. The privately optimal decision x^{eq} is to equate marginal utility mu with the price of an additional unit of coverage p'. The socially optimal choice x^{eff} equates marginal utility with marginal cost mc. The figure also depicts the average marginal cost curve of consumers purchasing each contract x, in bold and denoted $\mathbb{E}_x[mc]$, and the marginal cost curves of other consumers purchasing x^{eq} as faded curves.

To understand the sources of inefficiency, define the intensive margin selection coefficient as[23]

$$S_I(x) = \partial_x \mathbb{E}_x[c] - \mathbb{E}_x[mc].$$

[23]This is not a mathematically rigorous definition because we treat $\mathbb{E}_x[c]$ and $\mathbb{E}_x[mc]$ as differentiable functions even though conditional expectations are only defined almost everywhere according to α.

That is, $S_I(x)$ is the cost increase per additional unit of insurance, minus the average marginal cost of a unit of insurance. In other words, $S_I(x)$ is how much costs increase with an increase in coverage due to selection. This coefficient is positive if, locally around a contract x, agents who purchase more coverage are more costly, and it is negative if agents who purchase more coverage have lower costs. Thus, S_I is closely related to the positive correlation test of Chiappori and Salanié (2000). Moreover, it is natural to say that there is adverse selection around x if $S_I(x)$ is positive, and advantageous selection if $S_I(x)$ is negative. Notice that this is a local property, as it is possible to have adverse selection in some regions of the contract space and advantageous selection in other regions.

With this notation, the wedge can be divided in two components. First, in a multidimensional model many different types purchase contract x. This fact is illustrated by the equilibrium demand profile in the calibrated example, depicted in Figure 5a. Therefore, there is no reason for the marginal cost of any single buyer of x to equal $p'(x)$. Second, even the average marginal cost $\mathbb{E}_x[mc]$ of agents buying x may differ from $p'(x)$. The difference between $p'(x)$ and $\mathbb{E}_x[mc]$ is equal to $S_I(x)$, so that these two quantities differ whenever there is advantageous or adverse selection. These two components are depicted in Figure 4, where $p'(x)$ is higher than the average marginal cost, implying that there is adverse selection.

5.2 Optimal Regulation and Risk Adjustment

We now consider optimal government intervention. We focus on the simplified assumptions from the previous section. In addition, we assume that firms are unable to condition prices on types only because of private information. A government with no excess burden of public funds wishes to regulate the market to maximize welfare in the Kaldor-Hicks sense, as defined by equation (3). The government is capable of regulating the product space, prices, and to employ taxes and subsidies. However, we assume that the government faces the same informational constraints as firms do, and, therefore, cannot treat consumers differently conditional on their type.

Without loss of generality,[24] the government's problem is to choose an optimal price function $p(x)$, let consumers optimize reaching an allocation α, and make transfers $t(x)$ per unit of contract x to firms so that

$$p(x) + t(x) = \mathbb{E}_x[c].$$

That is, setting prices subsumes all other policy instruments. For example, forbidding a contract is tantamount to setting a very high price. Likewise, setting a system of subsidies

[24] We refer readers to Azevedo and Gottlieb (in preparation) for a formal statement of the planner's problem and its solution.

so that prices reach some equilibrium configuration is equivalent to mandating those prices and offering subsidies so that firms break even.

We approach the problem of finding optimal prices from a sufficient statistics approach similar to Wilson (1993), Piketty (1997), Roberts (2000), and Saez (2001). Namely, we consider a price schedule $p(x)$, and perform a small perturbation $\tilde{p}(x)$. If $p(x)$ was originally optimal, then the perturbation $\tilde{p}(x)$ should have a higher-order effect on welfare. Our perturbation works as follows. Fix a contract x_0. We increase the derivative $p'(x)$ by dp' in a small interval $[x_0, x_0 + dx]$, and leave it unchanged otherwise. For simplicity, we make the additional assumption that only adjustments in the intensive margin matter. That is, we assume that most of the welfare effects of this perturbation come from agents who adjust their consumption in the intensive margin, in face of the increase in p'.

Under these assumptions, the welfare effect of the perturbation can be calculated as follows. There are approximately $f(x_0) \cdot dx$ agents who are affected by the change, where $f(x)$ is the density of agents purchasing contract x. Each one of these agents reduces her coverage by

$$\epsilon(x_0, \theta) \cdot \frac{dp'}{p'(x_0)} \cdot x_0,$$

where $\epsilon(x, \theta)$ is the intensive margin elasticity of coverage with respect to p'.[25] Note that, although the perturbation extracts revenues from agents purchasing contracts higher than x_0, this is moot because we are ignoring redistribution and cost of public funds. Therefore, the total welfare effect equals

$$\mathbb{E}_{x_0}[\epsilon \cdot (mu - mc)] \cdot \frac{f(x_0) \cdot x_0}{p'(x_0)} \cdot dx dp'$$

plus higher-order terms, which become insignificant as dx, dp', and dp'/dx converge to 0. Consequently, optimal subsidies satisfy, for all x,

$$\mathbb{E}_x[\epsilon \cdot (p' - mc)] = 0.$$

This expectation can be decomposed as

$$\begin{aligned} 0 &= \mathbb{E}_x[\epsilon] \cdot \mathbb{E}_x[p' - mc] - \mathrm{Cov}_x[\epsilon, mc] \\ &= \mathbb{E}_x[\epsilon] \cdot (S_I(x) - t'(x)) - \mathrm{Cov}_x[\epsilon, mc]. \end{aligned}$$

[25] Formally, $\epsilon(x, \theta)$ is defined as

$$\epsilon(x, \theta) = \frac{1}{x} \cdot \frac{mu(x, \theta)}{\partial_{xx} u(x, \theta) - p''(x)}.$$

This can be calculated as the percentage change in the choice of x for a one percent of $p'(x)$ shift in the $p'(\cdot)$ curve.

Rearranging this equation we find that optimal subsidies are given by the simple formula[26]

$$\underbrace{t'(x)}_{\text{incremental subsidy}} = \underbrace{S_I(x)}_{\text{risk adjustment}} - \underbrace{\frac{\text{Cov}_x[\epsilon, mc]}{\mathbb{E}_x[\epsilon]}}_{\text{sorting}}. \tag{4}$$

The intuition for the optimal regulation formula (4) is as follows. The term $S_I(x)$ is a risk-adjustment subsidy. Namely, firms selling contracts to more costly consumers should be completely compensated for the effects of selection. Likewise, firms serving cheaper consumers should be taxed (or less subsidized), with the adjustment term $S_I(x)$ fully offsetting the effects of selection. Risk-adjustment is an extremely common policy intervention in insurance markets. However, practical versions of policies described under the umbrella of risk adjustment are quite different from each other, and sometimes quite different from the prescription of the $S_I(x)$ term. Nevertheless, policymakers and market participants seem to believe that it is important for cross subsidies to lean against the wind of advantageous or, more commonly,[27] adverse selection.

To understand the intuition for the covariance term, recall the two sources of inefficiency identified in Section 5.1: marginal costs of any given agent are different than the price of additional coverage, and marginal costs and prices differ even on average. The risk-adjustment term $S_I(x)$ addresses this second source of inefficiency, by eliminating adverse or advantageous selection from the point of view of the firms. The covariance term deals with the first source of inefficiency, which is due to the variation in marginal costs among different buyers of contract x. Hence, this term sorts consumers more efficiently. If buyers of x with higher marginal costs also have more elastic demands, then the formula calls for a higher slope of prices $p'(x)$ than under pure risk adjustment, which induces the high cost consumers to purchase lower levels of coverage.[28] If this correlation is negative, then incremental prices should be lower, inducing consumers with lower costs to purchase more insurance.

To grasp the intuition for government intervention, we return to the calibration of Exam-

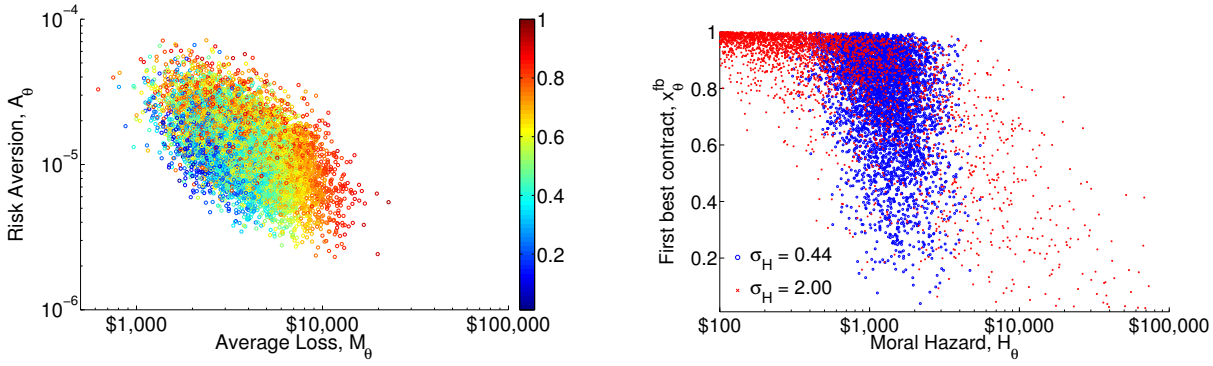
[26] We caution readers that this formula is not mathematically rigorous because conditional expectations are only defined almost everywhere according to α. Interested readers are again referred to Azevedo and Gottlieb (in preparation) for a formal mathematical analysis.

[27] The fact that adverse selection is considered a problem more often than advantageous selection is consistent with theoretical results of Chiappori et al. (2006). This suggests (although by no means conclusively), that optimal levels of t' are typically positive.

[28] This is the more common case in our simulations. The reason is that, in Example 3, the elasticity is given by

$$\epsilon = -\frac{1}{x}\frac{p'}{p'' + S^2 A - H}.$$

Therefore, consumers with higher moral hazard and lower insurance value ($S^2 A$) have more elastic demand. This creates a positive correlation because moral hazard increases marginal cost. Moreover, conditional on purchasing x, consumers with higher lower insurance value typically have higher mean loss, which also increases marginal cost.



(a) (b)

Figure 5: Equilibrium demand profile (a), and first-best allocation as a function of moral hazard (b).

Notes: Panel (a) depicts the equilibrium demand profile in the health insurance model calibrated in Section 4. Each of the 10,000 points represents a randomly drawn type from the population. The horizontal axis represents expected health shock M_θ, and the vertical axis represents the absolute risk aversion coefficient A_θ. The colors represent the level of coverage purchased in equilibrium. Panel (b) illustrates the first-best contracts versus moral hazard parameter for randomly drawn types. Blue circles were drawn from a distribution with low heterogeneity in moral hazard, and red xs were drawn from a distribution with high heterogeneity in moral hazard. All other moments of the parameter distribution are distributed according to Table 1.

ple 3. As previously shown, there is considerable adverse selection in equilibrium, suggesting under-provision of insurance. To understand optimal regulation, we numerically calculated an optimal price schedule, and the associated allocation. The results are displayed in Figure 2b. The optimal regulation involves substantial subsidies, which make the optimal price schedule much flatter than the competitive price. This shifts a considerable number of agents to purchase more insurance, increasing levels of coverage.

To understand the quantitative importance of offering an optimal menu of contracts, we contrast the welfare gains of the optimal policy with the 60% mandate. We find that the optimal policy increases welfare by $291 per consumer, compared with $56 for the mandate. Hence, for the parameter values in our calibration, an optimal menu achieves considerably higher welfare than a mandate.[29]

[29]In the next section we will show that in some extreme cases it is optimal to offer a single contract. More generally, sometimes it is optimal to offer a very narrow range of products in equilibrium. In those cases it is possible to implement the optimal regulation with a mandate plus the regulation of an uniform level of coverage. We highlight that, even in those cases, a simple mandate without a cap can do considerably worse than the optimal regulation.

5.3 Caveats

Our informal analysis glossed over some important issues related to the regulation of markets with adverse selection. First, there are cases where private and public incentives diverge to such a degree that the regulator may choose not to screen agents, in a degenerate case of equation (4). To see this, return to the numerical example. The first-best level of coverage for each agent can be calculated by equating marginal utility and marginal cost in equation (2). It equals

$$x_\theta^{\text{fb}} = \frac{A_\theta S_\theta^2}{A_\theta S_\theta^2 + H_\theta}.$$

Consider now the particular case where all the heterogeneity is on the moral hazard parameter H_θ. In this case, agents with higher H_θ have both a higher marginal utility, mu_θ, and a lower level of efficient coverage, x_θ^{fb}. In fact, the surplus from additional coverage $mu_\theta - mc_\theta$ is decreasing in H_θ. This situation is what Guesnerie and Laffont (1984) call non-responsiveness. They show that the optimal regulation in such one-dimensional models is to enforce a uniform mandate, where all consumers are forced to purchase the same contract. Because the regulator would like to offer coverage that decreases on types, whereas only increasing coverage is incentive compatible, the optimal coverage is constant.

A similar phenomenon can happen in multidimensional models. We computed optimal prices with a higher heterogeneity in moral hazard by raising σ_H^2 from 0.20 to 0.98. This later value is the central estimate of Einav et al. (2009) for the log variance of their moral hazard parameter. We then calculated optimal prices in a perturbation with 26 contracts. We found that the optimum had almost all consumers purchasing the same contract, similarly to the optimum in a one-dimensional model with non-responsiveness. That is, selection on moral hazard can, in extreme cases, lead to mandates being optimal. To gain some intuition, Figure 5b plots the first-best contracts versus moral hazard parameters for consumers drawn according to a distribution with a high and a low value of σ_H^2 and, otherwise, our baseline parameters. As can be seen, the case with higher dispersion in moral hazard is similar to a one-dimensional model, where agents in high percentiles of the moral hazard distribution are extremely likely to have a very low level of first-best coverage. In contrast, when dispersion is low, even agents in high percentiles of the moral hazard distribution have a broad range of first-best levels of coverage.

Besides this issue, our analysis ignored the possibility of richer substitution patterns, such as extensive margin responses.[30] We ignored redistribution, including the evaluation of ex-ante welfare, leaving aside issues such as reclassification risk. In practice, these distributional concerns are important and, perhaps, a more common impetus for regulations such as mandates and subsidies than non-responsiveness and risk adjustment. Finally, regulators

[30] See Araujo and Moreira (2010).

can often distinguish between some groups of consumers and decide whether to impose some degree of community rating.

The analysis of these issues is beyond the scope of the present paper. Therefore, we emphasize the basic points that competitive markets with adverse selection are inefficient, and that risk adjustment should play an important role in regulation aimed at restoring efficiency. We leave a more detailed analysis of optimal regulation to future work.

6 Conclusion

This paper builds a competitive model of adverse selection that allows us to study key features of standard market interventions. Importantly, the model has endogenously determined contract characteristics and permits rich consumer heterogeneity. Our basic finding is that the assumptions of price taking and free entry robustly make sharp predictions in a number of settings. More specifically, Theorem 1 guarantees that at least one equilibrium exists, and the examples illustrate that equilibria often make sharp predictions. Moreover, the model supports the view that adverse and advantageous selection cause market failures.

We calibrated a health insurance model based on demand for insurance in a large US employer studied by Einav et al. (2013), and used our equilibrium model to simulate the effects of a mandate. We found that mandates can induce a substantial number of consumers to purchase minimum coverage, that mandates can have important unintended consequences to adverse selection in the intensive margin, and that they can increase welfare despite these drawbacks. Finally, we showed that policies that address adverse selection in the intensive margin can substantially increase welfare over mandates, and that optimal policies are closely related to risk-adjustment schemes used in practice.

An natural direction of future research is that our framework leads naturally to questions on optimal government intervention, positive analyses of government policies, and optimal pricing by a monopolist. We leave these questions to future work (see Azevedo and Gottlieb, in preparation). Another important extension is incorporating imperfect competition, which is often an important feature of markets with adverse selection.

Finally, it would be interesting to test the explanatory power of the present model. That is, future empirical and experimental evidence can be used to test whether the predictions of price-taking and free entry match behavior in markets with adverse selection. Hopefully, even if the evidence does not match the predictions of the model, it can be used to develop more accurate and applicable models.

References

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

Araujo, Aloisio and Humberto Moreira, "Adverse Selection Problems without the Spence–Mirrlees Condition," *Journal of Economic Theory*, 2010, *145* (3), 1113–1141.

Attar, Andrea, Thomas Mariotti, and François Salanié, "Nonexclusive Competition in the Market for Lemons," *Econometrica*, 2011, *79* (6), 1869–1918.

Azevedo, Eduardo M and Daniel Gottlieb, "Optimal Regulation of Markets with Adverse Selection," *Mimeo, University of Pennsylvania*, in preparation.

Banks, Jeffrey S and Joel Sobel, "Equilibrium Selection in Signaling Games," *Econometrica*, 1987, pp. 647–661.

Billingsley, Patrick, *Probability and Measure*, John Wiley & Sons, 2008.

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

_ and _ , "Efficient Competitive Equilibria with Adverse Selection," *Journal of Political Economy*, 2006, *114* (3), 485–516.

Chiappori, Pierre-André and Bernard Salanié, "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 2000, *108* (1), 56–78.

_ , Bruno Jullien, Bernard Salanié, and Francois Salanié, "Asymmetric Information in Insurance: General Testable Implications," *Rand Journal of Economics*, 2006, *37* (4), 783–798.

Cho, In-Koo and David M Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 1987, pp. 179–221.

Cohen, Alma and Liran Einav, "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 2007, *97* (3), 745–788.

Cutler, David M and Richard Zeckhauser, "Extending the Theory to Meet the Practice of Insurance," *Brookings-Wharton Papers on Financial Services*, 2004, *2004* (1), 1–53.

Dubey, Pradeep and John Geanakoplos, "Competitive Pooling: Rothschild-Stiglitz Reconsidered," *Quarterly Journal of Economics*, 2002, pp. 1529–1570.

Einav, L, A Finkelstein, and J Levin, "Beyond Testing: Empirical Models of Insurance Markets.," *Annual Review of Economics*, 2009, *2*, 311–336.

Einav, Liran, Amy Finkelstein, and Mark R Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 2010, *125* (3), 877–921.

_ , _ , and Paul Schrimpf, "Optimal mandates and the welfare cost of asymmetric information: Evidence from the uk annuity market," *Econometrica*, 2010, *78* (3), 1031–1092.

_ , _ , Stephen P Ryan, Paul Schrimpf, and Mark R Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 2013, *103* (1), 178–219.

_ and _ , "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 2011, *25* (1), 115–138.

Fang, Hanming, Michael P Keane, and Dan Silverman, "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market," *Journal of Political Economy*, 2008, *116* (2), 303–350.

Finkelstein, Amy and Kathleen McGarry, "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, 2006, *96* (4), 938–958.

Guesnerie, Roger and Jean-Jacques Laffont, "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm," *Journal of Public Economics*, 1984, *25* (3), 329–369.

Hackmann, Martin B, Jonathan T Kolstad, and Amanda E Kowalski, "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *Mimeo, Yale University*, 2014.

Handel, Benjamin, Igal Hendel, and Michael D Whinston, "Equilibria in Health Exchanges: Adverse Selection Vs. Reclassification Risk," *NBER Working Paper*, 2013, (w19399).

Handel, Benjamin R, "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

Heinonen, Juha, *Lectures on Analysis on Metric Spaces*, Springer, 2001.

Hellwig, Martin, "Some Recent Developments in the Theory of Competition in Markets with Adverse Selection*," *European Economic Review*, 1987, *31* (1), 319–325.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39* (153), 41–57.

Kadan, Ohad, Philip Reny, and Jeroen Swinkels, "Existence of Optimal Mechanisms in Principal-Agent Problems," *Mimeo, University of Chicago*, 2014.

Kunreuther, Howard and Mark Pauly, *Insurance Decision-Making and Market Behavior*, Now Publishers Inc, 2006.

Luz, Vitor F., "Characterization and Uniqueness of Equilibrium in Competitive Insurance," *Mimeo, Yale University*, 2013.

Mahoney, Neale and E Glen Weyl, "Imperfect Competition in Selection Markets," *Mimeo, University of Chicago*, 2014.

Miyazaki, Hajime, "The Rat Race and Internal Labor Markets," *Bell Journal of Economics*, 1977, pp. 394–418.

Myerson, Roger B, "Sustainable Matching Plans with Adverse Selection," *Games and Economic Behavior*, 1995, *9* (1), 35–65.

Netzer, Nick and Florian Scheuer, "A Game Theoretic Foundation of Competitive Equilibria with Adverse Selection," *International Economic Review*, 2014, *55* (2), 399–422.

Piketty, Thomas, "La Redistribution Fiscale Face Au Chômage," *Revue Française D'Économie*, 1997, *12* (1), 157–201.

Riley, John G, "Informational Equilibrium," *Econometrica*, 1979, pp. 331–359.

Roberts, Kevin, "A Reconsideration of the Optimal Income Tax," in Peter Hammond and Gareth Myles, eds., *Incentives, Organization, and Public Economics: Papers in Honour of Sir James Mirrlees: Papers in Honour of Sir James Mirrlees*, Oxford University Press, 2000.

Rochet, Jean-Charles and Lars A Stole, "Nonlinear Pricing with Random Participation," *Review of Economic Studies*, 2002, *69* (1), 277–311.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

Saez, Emmanuel, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, *68* (1), 205–229.

Scheuer, Florian and Kent A Smetters, "Could a Website Really Have Doomed the Health Exchanges? Multiple Equilibria, Initial Conditions and the Construction of the Fine," *NBER Working Paper*, 2014, (w19835).

Spence, A Michael, "Monopoly, Quality, and Regulation," *Bell Journal of Economics*, 1975, pp. 417–429.

Spence, Michael, "Job Market Signaling," *Quarterly Journal of Economics*, 1973, pp. 355–374.

Spinnewijn, Johannes, "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, forthcoming.

Veiga, André and E Glen Weyl, "The Leibniz Rule for Multidimensional Heterogeneity," *Mimeo, University of Chicago*, 2014.

_ **and** _ , "Product Design in Selection Markets," *Mimeo, University of Chicago*, 2014.

Wilson, Charles, "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 1977, *16* (2), 167–207.

Wilson, Robert B, *Nonlinear Pricing*, Oxford University Press, 1993.

Appendix

A Proofs

A.1 Existence of Equilibrium

The proof of Theorem 1 follows from three lemmas. The first lemma uses a standard fixed point argument to show that every perturbation has a weak equilibrium. The second lemma establishes, moreover, that price vectors in any perturbation are uniformly Lipschitz. The third lemma, using this fact, shows that every sequence of weak equilibria of perturbations has a converging subsequence. In what follows we fix an economy $E = [\Theta, X, \mu]$.

Lemma A1. *Every perturbed economy has an equilibrium.*

Proof. Preliminaries.

Consider the perturbed economy $[\Theta \cup \bar{X}, \bar{X}, \mu + \eta]$. Denote by $\bar{\alpha} \in \Delta((\Theta \cup \bar{X}) \times \bar{X})$ the measure describing the allocation of the behavioral types. That is, for each $x \in \bar{X}$,

$$\bar{\alpha}(x, x) = \eta(x),$$

and $\bar{\alpha}$ has no mass in the complement of these points. We will write allocations as $\alpha + \bar{\alpha}$, where α denotes the allocation for the standard types, has support contained in $\Theta \times \bar{X}$, and $\alpha|\Theta = \mu$.

We can now define a tâtonnement correspondence

$$T : P \times A \rightrightarrows P \times A.$$

Here we let A be the set of all allocations for the standard types, with the topology of weak convergence of measures. That is,

$$A = \{\alpha \in \Delta((\Theta \cup \bar{X}) \times \bar{X}) : \text{support}(\alpha) \subseteq \Theta \times \bar{X}, \alpha|\Theta = \mu\}.$$

Let P be the set of all prices vectors in (convex closure of the image of $c)^{\bar{X}}$, with the standard topology of Euclidean space. The tâtonnement is defined in terms of two maps,

$$T(p, \alpha) = \Phi(\alpha) \times \Psi(p).$$

Let

$$\begin{aligned}\Phi(a) &= \{p \in P : p(x) = E_x[c|\alpha + \bar{\alpha}] \ \forall x \in \bar{X}\}, \text{ and} \\ \Psi(p) &= \arg\max_{\alpha \in A} \int U(x, p(x), \theta)\, d\alpha.\end{aligned}$$

That is, given an allocation α, $\Phi(\alpha)(x)$ is the expected cost of supplying contract x. Given p, $\Psi(p)$ is the set of allocations for the standard types where they choose optimally given p.

Note that the fixed points of T correspond to the equilibria of the perturbed economy. To see this, note that $p \in \Phi(\alpha)$ is equivalent to firms making 0 profits, and $\alpha \in \Psi(p)$ is equivalent to the standard types optimizing. Therefore, (p^*, α^*) is a fixed point of T if and only if $(p^*, \alpha^* + \bar{\alpha})$ is an equilibrium. We will now prove the existence of a fixed point. The proof has three steps. Steps 1 and 2 establish basic properties of Φ and Ψ, and step 3 applies a fixed point theorem.

Step 1: Φ is nonempty, convex valued, and has a closed graph.

To establish that it has a closed graph, consider a sequence $(\alpha^n, p^n)_{n\in\mathbb{N}}$ in the graph of Φ, with limit (α, p). We will show that $p(x)$ is the conditional expectation of cost given $\alpha + \bar{\alpha}$. To see this, take an arbitrary set $S \subseteq \bar{X}$. Let $\tilde{S} = (\Theta \cup \bar{X}) \times \bar{X}$. We have

$$\begin{aligned}\int_{\tilde{S}} p(x)\, d(\alpha + \bar{\alpha}) &= \sum_{x \in S} p(x) \cdot [\alpha(\Theta \times x) + \eta(x)] \\ &= \lim_{n\to\infty} \sum_{x \in S} p^n(x) \cdot [\alpha^n(\Theta \times x) + \eta(x)] \\ &= \lim_{n\to\infty} \int_{\tilde{S}} p^n(x)\, d(\alpha^n + \bar{\alpha}) \\ &= \lim_{n\to\infty} \int_{\tilde{S}} c(x, \theta)\, d\alpha^n \\ &= \int_{\tilde{S}} c(x, \theta)\, d\alpha.\end{aligned}$$

The first and third equations follow from decomposing the integral as a sum. The second equation follows from the convergence of (p^n, α^n). The fourth equation from the definition of conditional expectation, and the fifth equation from the fact that c is continuous and α^n converges weakly to α. Convex-valuedness and non-emptiness follow directly from the definition of Φ.

Step 2: Ψ is nonempty, convex valued, and has a closed graph.

To see that it has a closed graph, consider a sequence $(p^n, \alpha^n)_{n\in\mathbb{N}}$ in the graph of Ψ with

limit (p, α). For any $\alpha' \in A$, we have

$$\int U(x, p^n(x), \theta)\, d\alpha'(\theta, x) \leq \int U(x, p^n(x), \theta)\, d\alpha^n(\theta, x).$$

Taking the limit, we have

$$\int U(x, p(x), \theta)\, d\alpha'(\theta, x) \leq \int U(x, p(x), \theta)\, d\alpha(\theta, x).$$

The LHS limit follows from the Dominated Convergence Theorem. To see the convergence of the RHS term, it is helpful to decompose it as

$$\begin{aligned} & \int U(x, p^n(x), \theta) - U(x, p(x), \theta)\, d\alpha^n(\theta, x) \\ + & \int U(x, p(x), \theta)\, d\alpha^n(\theta, x). \end{aligned}$$

The first integrand converges to 0 uniformly in x and θ because $\bar{X}$ is finite, and hence p^n converges uniformly to p, and because the continuous function U is uniformly continuous in the compact set where prices belong to the image of c. Therefore, the first integral converges to 0. The second integral converges to 0 by the continuity of U and weak convergence of α^n to α.

Ψ is nonempty because X is finite, and therefore the expression $U(x, p(x), \theta)$ attains a maximum for every θ. Convexity follows from the definition of Ψ.

S*tep 3: Existence of a fixed point.*

We can now complete the proof. The claims about Φ and Ψ imply that T is convex valued, nonempty, and has a closed graph. We have that the set $P \times A$ is compact, convex and a subset of a locally convex topological vector space. Therefore, by the Kakutani-Glicksberg-Fan Theorem, T has a fixed point. □

The next result shows that, in a weak equilibrium of a perturbed economy, prices are a Lipschitz function with constant L. The intuition is that, if prices of similar contracts differed too much, no agent would be willing to purchase the most expensive contract.

Lemma A2. *Let (p^*, α^*) be a weak equilibrium of a perturbed economy. Then p^* is a L-Lipschitz function.*

Proof. Consider two contracts x, x'. Assume, without loss of generality, that $p^*(x) > p^*(x')$. In particular, $p^*(x) > 0$, and therefore there exists a standard type θ who prefers x to x'. That is, there exists $\theta \in \Theta$ such that

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

The assumption that marginal rates of substitution are bounded then implies

$$|p^*(x) - p^*(x')| \leq d(x, x') \cdot L.$$

□

The next lemma uses this observation to show that every sequence of perturbations of an economy E has a subsequence of equilibria that converges to an equilibrium of E.

Lemma A3. *Consider a sequence of perturbations $E(\bar{X}^n, \eta^n)$ converging to E with weak equilibria (p^n, α^n). Then $(p^n, \alpha^n)_{n\in\mathbb{N}}$ has a subsequence that converges to an equilibrium (p^*, α^*) of E. Moreover, p^* is L-Lipschitz.*

Proof. We begin by defining α^* and p^*. First note that the set of allocations is compact. Therefore, without loss of generality, passing to a subsequence, we can take $(\alpha^n)_{n\in\mathbb{N}}$ to converge to a measure $\alpha^* \in \Delta((\Theta \cup X) \times X)$. Moreover, α^* has support contained in $\Theta \times X$, and $\alpha^*|\Theta = \mu$.

As for p^*, take for each n a function $\tilde{p}^n$ with domain X, which coincides with p^n in $\bar{X}$ and is L-Lipschitz. Lemma A2 and Theorem 6.2 of Heinonen (2001) p. 43 guarantee the existence of these functions. Without loss of generality, passing to a subsequence, we may take the sequence $(\tilde{p}^n)_{n\in\mathbb{N}}$ to converge pointwise to a limit p^*. Note that because the sequence $(\tilde{p}^n)_{n\in\mathbb{N}}$ is uniformly L-Lipschitz, it is, in particular, equicontinuous. By the Arzelà–Ascoli Theorem, the sequence converges uniformly to p^*. This implies convergence in the sense of definition 4 and the Lipschitz property. □

Note that the previous lemma directly implies Theorem 1.

Proof of Theorem 1. Take any sequence of perturbations of economy E. By Lemma A3, there exists a subsequence with a converging sequence of equilibria. Hence, the limit of this sequence is a refined equilibrium of E.

□

A.2 Properties of Equilibria

We begin by establishing two of the properties in Proposition 1 as lemmas.

Lemma A4. *Every equilibrium is a weak equilibrium.*

Proof. Consider an economy $E = [\Theta, X, \mu]$, equilibrium (p^*, α^*), and a sequence of perturbations $E(\bar{X}^n, \eta^n)$ converging to E with weak equilibria (p^n, α^n) converging to the equilibrium (p^*, α^*).

To verify that prices are the conditional expectation of the cost, take a measurable set of contracts $S \subseteq X$. Let $\tilde{S} = (\Theta \cup X) \times S$. Let $\tilde{p}^n$ be an L-Lipschitz function extending p^n to X, which exists by the argument in the proof of Lemma A2. We have

$$\begin{aligned}\int_{\tilde{S}} p^*(x)\, d\alpha^* &= \lim_{n\to\infty} \int_{\tilde{S}} \tilde{p}^n(x)\, d\alpha^n \\ &= \lim_{n\to\infty} \int_{\tilde{S}} c(x,\theta)\, d\alpha^n \\ &= \int_{\tilde{S}} c(x,\theta)\, d\alpha^*.\end{aligned}$$

The first equality follows because $\{\alpha^n\}_{n\in\mathbb{N}}$ converges weakly to α^*, and $\{\tilde{p}^n\}_{n\in\mathbb{N}}$ converges uniformly to p^*. The second equality follows because p^n is the conditional expectation of c given α^n. The third equation follows because c is continuous and α^n converges weakly to α^*. From this argument, we have that p^* is the conditional expectation of c under the measure α^*.

To see that agents are optimizing, take an allocation α'. Since the (p^n, α^n) are weak equilibria, for all n we have

$$\int_{\Theta\times X} U(x, \tilde{p}^n(x), \theta)\, d\alpha^n \geq \int_{\Theta\times X} U(x, \tilde{p}^n(x), \theta)\, d\alpha'.$$

Because the $\tilde{p}^n$ converge uniformly to p^* and U is uniformly continuous in the relevant set, we can take limits on both sides, obtaining

$$\int_{\Theta\times X} U(x, p^*(x), \theta)\, d\alpha^* \geq \int_{\Theta\times X} U(x, p^*(x), \theta)\, d\alpha'.$$

Because this inequality holds for any α', we have that, for α^* almost every (θ, x),

$$U(x, p^*(x), \theta) = \sup_{x'\in X} U(x', p^*(x'), \theta),$$

as desired. □

Lemma A5. *Consider a refined equilibrium (p^*, α^*) of an economy E. Let x' be a contract with $p^*(x') > 0$. Then there exists (θ, x) in the support of α such that*

$$U(x, p^*(x), \theta) = U(x', p^*(x'), \theta). \tag{5}$$

Proof. Take a sequence of perturbations $(E(\bar{X}^n, \eta^n))_{n\in\mathbb{N}}$ converging to E with equilibria (p^n, α^n) converging to (p^*, α^*). Take $x'^n \in \bar{X}^n$ converging to x'. Since $p^n(x'^n)$ converges to $p^*(x') > 0$, we must have $p^n(x'^n) > 0$ for sufficiently large n. This implies that there exists a

standard type θ^n who buys contract x'^n. We can take a subsequence such that θ^n converges to a type θ because the set of types is compact. Let x be a contract that is optimal for θ at prices p^*, and take a sequence $(x^n)_{n\in\mathbb{N}}$ with each $x^n \in \bar{X}^n$ converging to x. Since x'^n is optimal for θ^n in the perturbed economy, for all sufficiently large n we have

$$U(x'^n, p^n(x'^n), \theta^n) \geq U(x^n, p^n(x^n), \theta^n).$$

Taking the limit, we have

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

Moreover, since x is optimal for θ at prices p^*, this implies equation (5). □

We can now establish Proposition 1.

Proof of Proposition 1. Parts 1 and 2 follow from lemmas A4 and A5. Part 3 follows from Lemma A3, and part 4 follows from part 3 and Rademacher's theorem.

□

B Details on the Calibration of Example 3

The model in Einav et al. (2013) differs from ours in three key ways, which keep us from simply using their estimates. First, they considered contracts with more complex characteristics such as out of pocket maximums. Second, they assumed that losses were distributed according to a shifted lognormal distribution, whereas we assume that losses are normally distributed. The reason we modified these two assumptions is transparency: our simpler model admits a simple closed-form expression for willingness to pay and costs – equation (2).

The third difference is that they estimate an empirical model for the distribution of types, letting it depend on characteristics of the population in their data. In contrast, we assume that the distribution of types is lognormally distributed. This also makes the calibration less direct.

We calibrated the distribution of types as follows. For the means of H and M, we used the numbers they report in Table 7B. For the mean of S, we used a standard deviation of total losses of 25,000 and a standard deviation of expected losses of 5,100, as they report on p. 204 paragraph 2. This implies a mean of 24,474 for S, were variance to be constant in the population. We simply take this value as the mean.

As for A, we found that using their central estimate of 1.9E-3 created implausible substitution patterns in our setting. As an illustration, using this value and the other means, we found that an average consumer would be willing to pay \$569,027 (equal to $AS^2/2$) for full insurance, even without taking moral hazard into account. Simulations assuming the

mean of A to equal their central estimate unsurprisingly lead to both equilibria and optimal allocations involving full coverage.

Based on this, we decided to use lower values of mean risk aversion in our calibrations. We believe this is reasonable because models with constant absolute risk aversion are of limited external validity outside the range of losses where they are estimated. This is the case in our setting because we have different loss distributions, contracts with different structures, and, in particular, no stop losses. We chose the mean of A to ensure that risk aversion was within the range where substitution patterns are intuitively plausible in equation 2. Although this still left a wide range of possible choices, we ran estimates with different values, and found that the qualitative features of equilibria were similar. For our calibration we set mean risk aversion to 1.5E-5, which makes the value of full insurance equal to approximately \$4,500. In particular, the value of insurance is of the same order of magnitude as the moral hazard costs, which, for full insurance, are equal to $H/2$. Thus, the average inefficiency from moral hazard is worth \$665, which is still lower but comparable to the gains from insurance. Estimates in this range place the surplus generated by a full insurance contract at about \$3,800.

The covariance matrix we used for A, H, and M was based on the log covariance matrix from Table 7A for the most closely associated parameters. Note that there is a certain extrapolation here, especially for mean losses, because they consider a lognormal shifted distribution of losses. Moreover, this made it unclear what assumptions we should make about the correlation with respect to S. Due to the fact that S is conflated with A for all practical purposes in our model, we decided to assume 0 correlation and leave the correlation between willingness to pay for insurance and other parameters to be determined more transparently by the correlations with A. As for the variability in S, we assumed the same log standard deviation as risk aversion.