Online Appendix to "Perfect Competition in Markets with Adverse Selection"

Eduardo M. Azevedo and Daniel Gottlieb[*]

This version: August 31, 2016
First version: September 2, 2014

A   Strategic Foundations when Firms Offer a Single Contract

This appendix shows that our price-taking equilibrium concept corresponds to the limit of Bertrand competition between firms selling differentiated varieties of each contract. Here, we consider the case where each firm offers one contract. In the supplementary material, we consider the case where each firm offers a menu of contracts.

A.1   Definition of the differentiated products Bertrand Game

Throughout this section, fix an economy $E = [\Theta, X, \mu]$ and a perturbation $(E, \bar{X}, \eta)$. To simplify the notation, take the total mass of consumers $\mu(\Theta) + \eta(\bar{X})$ to equal 1. Assume that preferences are quasilinear, so that $U(x, p, \theta) = u(x, \theta) - p$.

Consider Bertrand competition between differentiated firms selling varieties of each contract x. For each contract x, firm (x, i) with $i \in \{1, 2, \cdots, n\}$ offers a differentiated variety. Consumers have logit demand with semi-elasticity parameter σ.[1] Consider the case where all firms but one set the same price for their variety of the product. It is sufficient to define demand in these situations because we consider symmetric equilibria. Assume that firm (x, i) sets a price of P, while all other firms set prices that only depend on the product, according

---

[*]Azevedo: Wharton, eazevedo@wharton.upenn.edu. Gottlieb: Washington University in St. Louis, dgottlieb@wustl.edu.
[1]σ is approximately equal to the semi-elasticity of demand of a firm with small market share.

to a vector p. Then, the share of standard types θ purchasing variety (x, i) equals

$$S(P, p, x, \theta) = \frac{e^{\sigma \cdot (u(x, \theta) - P)}}{\sum_{x' \neq x} n \cdot e^{\sigma \cdot (u(x', \theta) - p(x'))} + (n - 1) \cdot e^{\sigma \cdot (u(x, \theta) - p(x))} + e^{\sigma \cdot (u(x, \theta) - P)}}.$$

For a behavioral type $\theta = x$, this share equals

$$S(P, p, x, x) = \frac{e^{-\sigma \cdot P}}{(n - 1) \cdot e^{-\sigma \cdot p(x)} + e^{-\sigma \cdot P}},$$

and 0 for all other products. The quantity supplied by the firm is

$$Q(P, p, x) = \int_\theta S(P, p, x, \theta) \, d(\mu + \eta).$$

Each firm has constant returns to scale up to a capacity limit of k consumers and infinite costs of supplying more than k consumers. We consider this production function because it is a simple way to model a small efficient scale. Profits of a firm that produces a variety of x at price P while all other firms price according to p equal

$$\Pi(P, p, x) = \int_\theta (P - c(x, \theta)) \cdot S(P, p, x, \theta) \, d(\mu + \eta), \tag{A1}$$

if $Q(P, p, x) \leq k$ and $-\infty$ otherwise.

A **symmetric Bertrand equilibrium** under parameters (n, k, σ) is a vector p^* such that, for all $x \in \bar{X}$,

$$p^*(x) = \arg \max_{P \geq 0} \Pi(P, p^*, x).$$

A.2   Existence of a Bertrand Equilibrium and Bound on Profits

The key result in the analysis is the following proposition, which gives conditions for Bertrand equilibria to exist and for firms to earn low profits in equilibrium. Define the constants

$$\bar{c} = \max_{(x, \theta)} c(x, \theta)$$

and

$$\bar{\eta} = \min_{x \in \bar{X}} \eta(\{x\}).$$

Proposition A1. *Assume that there are enough firms selling each variety to serve the entire market,*

$$n \cdot k > 1,$$

and that the capacity of individual firms is sufficiently small,

$$k < \frac{\bar{\eta}}{2} \cdot \frac{1}{1 + \bar{c} \cdot \sigma}. \tag{A2}$$

Then, a symmetric Bertrand equilibrium p^* exists. Moreover, profits per sale are bounded above by

$$\frac{\Pi(p^*(x), p^*, x)}{Q(p^*(x), p^*, x)} < \frac{2}{\sigma} + \frac{\bar{c}}{\bar{\eta}} \cdot k.$$

The key restriction on parameters is that firm scale is small. Moreover, the maximum scale that guarantees existence of equilibrium is decreasing in the semi-elasticity parameter and of the order of the inverse of semi-elasticity.

Note that Proposition A1 does not immediately imply that Bertrand equilibrium allocations converge to competitive equilibria. If we take a sequence of parameters with semi-elasticities converging to infinity, so that we approach perfect competition, to guarantee existence we need the number of firms to converge to infinity. In principle these Bertrand equilibria might not converge to perfectly competitive equilibria because, when the number of firms and the semi-elasticity both grow, logit demand does not necessarily become similar to demand for undifferentiated products. Proposition A2 below shows that this potential issue does not arise here, and that there are sequences of parameters such that Bertrand equilibria converge to perfectly competitive equilibria.[2]

Small firm scale is crucial to guarantee existence of an equilibrium. For example, an equilibrium may not exist in the original Rothschild and Stiglitz (1976) model because firms can offer a low-priced contract that attracts the whole market and gives positive profits. In our Bertrand game with preferences as in Rothschild and Stiglitz (1976), if firms have no scale restrictions, then, with high enough elasticities, the deviation that they point out would still be profitable.

A.3 Discussion of the Assumptions

The Bertrand game involves a number of modeling choices. Before proving Proposition A1, we discuss these choices. Decreasing returns to scale (which we model as a capacity constraint) is important to guarantee existence, as discussed above. We also assume that firms must serve all consumers who come to the door. Firms cannot set a low price but ration some consumers. Our assumption is reasonable, for example, in markets with community rating regulations that forbid turning away consumers.

---

[2]The proposition does not imply that, for a given competitive equilibrium, there is a sequence of Bertrand equilibria the converge to it.

Our strategic model is closely related to those used in the industrial organization literature. For example, Starc (2014), Decarolis et al. (2012), and Tebaldi (2015) consider strategic models where firms offer differentiated contracts and play a Nash equilibrium in prices. Mahoney and Weyl (forthcoming) consider a model with differentiated products, but more general assumptions on conduct. Our results give conditions for these models to converge to the competitive model. In particular, our results relate these imperfectly competitive models to perfectly competitive models that have been used in the literature. For example, our results imply that the competitive models of Riley (1979) and Handel et al. (2015) are limiting cases of standard imperfectly competitive models.

An alternative assumption would be that consumers and firms find each other through a search and matching process with frictions, so that some consumers could be rationed. This has been pursued by Inderst and Wambach (2001) and Guerrieri et al. (2010). These rationing assumptions lead to competitive equilibria in some, but not all cases. For example, Guerrieri et al. (2010) consider a directed search model. Firms post contract-price pairs, and consumers choose what contract-price pairs to search for. The probability that a consumer finds a match depends on the market tightness for the contract-price pair. Guerrieri et al. (2010) show that, in a version of the Rothschild and Stiglitz (1976) model, as search frictions become small, equilibria converge to our competitive equilibria. However, there are examples where the equilibria of this search model do not converge to competitive equilibria. One simple example is the one-contract case, as in example 1. For some parameter values, the (single) contract is offered at many different prices, with longer queues for lower prices, and consumers self-select depending on the queue length. The advantage of our approach is that we obtain convergence to competitive equilibrium in the broad set of models considered in our paper.

Instead of our small efficient scale and rationing assumptions, one could alternatively require firms to set locally optimal prices. In the equilibrium that we construct below, prices are locally optimal. Hence, an "equilibrium in locally optimal strategies" exists, regardless of the small efficient scale. Moreover, in any such symmetric equilibrium, profits are low, as can be shown by the first-order conditions. Therefore, we could establish an alternative result of convergence to competitive equilibrium. A shortcoming of this approach is that there is no simple formula for how large price deviations are allowed to be. This is in contrast to the capacity constraints approach, where we can show that capacity constraints of the order of the inverse of the semi-elasticity parameter are sufficient to ensure existence. Finally, note that local optimality is not a viable alternative to our perfectly competitive equilibrium notion. Riley (1979) gives an example where firms can always profit from local deviations, implying that this kind of equilibrium does not always exist.[3]

---

[3]In fact, Rothschild and Stiglitz (1976) themselves suggested that focusing on "equilibria in locally optimal

We assumed that demand for varieties has a logit functional form. The logit functional form considerably simplifies the statement of Proposition A1 but does not play a key role in its proof. The proof is based on bounds on the first and second derivatives of the market share function, $\partial_P S$ and $\partial_{PP} S$. The logit demand has simple functional forms for these expressions, which depend on the semi-elasticity parameter σ. Thus, one could replicate the argument given below for any market share function that satisfies the same bounds. This non-parametric result would, however, depend on assumptions about $\partial_P S$ and $\partial_{PP} S$.

An important modeling choice, which is relaxed in the supplementary material section C, is that each firm sells a single variety.

A.4   Proof Proposition A1

The proof uses formulas for the derivatives of the logit market shares. We take note of the formulas for the first two derivatives:

$$\partial_P S = -\sigma \cdot S \cdot (1 - S)$$

and

$$\partial_{PP} S = \sigma^2 \cdot S \cdot (1 - S) \cdot (1 - 2S).$$

The proof is organized in a series of claims. Throughout the proof, we let

$$\delta = \frac{k}{\bar{\eta}}.$$

The first claim shows that, if a firm does not produce above capacity, then the share of any type that the firm attracts is bounded above by δ.

Claim 1. If $\Pi(P, p, x) > -\infty$, then

$$S(P, p, x, \theta) \leq \delta$$

for all $\theta \in \Theta \cup \bar{X}$.

Proof. It suffices to prove the result for behavioral types $\theta = x$ because their share of variety (x, i) is weakly higher than any other type. The total demand for variety (x, i) from behavioral types equals

$$S(P, p, x, x) \cdot \eta(\{x\}),$$

---

strategies" may solve the non-existence problem, a conjecture that was proved wrong by Riley (1979) in the context of their (undifferentiated-contracts) setting.

which is no greater than k because firm (x, i) is not producing over capacity. This gives us the desired inequality. $\square$

Claim 2. Let

$$\bar{P} = \frac{1}{\sigma(1 - \delta)} + \bar{c}.$$

If $P > \bar{P}$, then either $\Pi(P, p, x) = -\infty$ or $\partial_P \Pi(P, p, x) < 0$.

Proof. If $\Pi(P, p, x) \neq -\infty$, then the profits of firm (x, i) are differentiable with respect to own price at P, and the derivative equals

$$
\begin{aligned}
\partial_P \Pi(P, p, x) &= \int S + (P - c) \cdot \partial_P S \, d(\mu + \eta) \\
&= \int S + (P - c) \cdot (-\sigma \cdot S \cdot (1 - S)) \, d(\mu + \eta) \\
&= \int S \cdot [1 - (P - c) \cdot \sigma \cdot (1 - S)] \, d(\mu + \eta).
\end{aligned}
$$

The equation in the first line follows from the definition of profits. The second line follows from the formulas for the derivative of the logit shares. The third line multiplies S out of the formula. Note that the expression inside the integral is strictly negative for all θ because

$$P - c > \bar{P} - c \geq \frac{1}{\sigma(1 - S)}.$$

Therefore, $\partial_P \Pi(P, p, x) < 0$. $\square$

The next claim is the key step to prove existence. It shows that the profit functions do not have interior local minima as long as the maximum scale of each firm is below a constant. The claim is based on Aksoy-Pierson et al. (2013), who establish existence of Bertrand equilibrium in a setting related to Caplin and Nalebuff (1991), but where firms have small market shares.

Claim 3. If $\Pi(P, p, x) > -\infty$ and $\partial_P \Pi(P, p, x) = 0$ then $\partial_{PP} \Pi(P, p, x) < 0$.

Proof. The first derivative of the profit function equals

$$
\begin{aligned}
\partial_P \Pi(P, p, x) &= \int S + (P - c) \cdot \partial_P S \, d(\mu + \eta) \\
&= \int S + (P - c) \cdot (-\sigma \cdot S \cdot (1 - S)) \, d(\mu + \eta) \\
&= Q(P, p, x) - \sigma \int (P - c) \cdot S \cdot (1 - S) \, d(\mu + \eta).
\end{aligned}
$$

The first equation follows from differentiating the profit function (A1). The second equation is derived by substituting the formula for the derivative of the logit share function. The third

equation follows from evaluating part of the integral. The first-order condition $\partial_P \Pi(P, p, x) = 0$ implies that

$$Q(P, p, x) = \sigma \int (P - c) \cdot S \cdot (1 - S) \, d(\mu + \eta). \tag{A3}$$

The second derivative of profits equals

$$
\begin{aligned}
\partial_{PP}\Pi(P, p, x) &= \int 2\partial_P S + (P - c) \cdot \partial_{PP} S \, d(\mu + \eta) \\
&= \int -2\sigma \cdot S \cdot (1 - S) + \sigma^2 \cdot (P - c) \cdot S \cdot (1 - S) \cdot (1 - 2S) \, d(\mu + \eta) \\
&= -2\sigma Q + 2\sigma \int S^2 \, d(\mu + \eta) \\
&\quad + \sigma^2 \int (P - c) S (1 - S) \, d(\mu + \eta) \\
&\quad - 2\sigma^2 \int (P - c) S^2 (1 - S) \, d(\mu + \eta).
\end{aligned}
$$

The first equation follows from differentiation of $\partial_P \Pi$. The second equation follows from substituting the formulae for derivatives of the logit share. The third equation follows from rearranging the expression and computing the integral for quantity. If we substitute expression (A3) we obtain

$$
\begin{aligned}
\partial_{PP}\Pi(P, p, x) &= -\sigma \cdot Q \\
&\quad + 2\sigma \int S^2 \, d(\mu + \eta) \\
&\quad - 2\sigma^2 \int (P - c) S^2 (1 - S) \, d(\mu + \eta). \tag{A4}
\end{aligned}
$$

We need two bounds to show that this expression is negative. Using Claim 1, we have

$$2\sigma \int S^2 \, d(\mu + \eta) < 2\sigma \int \delta \cdot S \, d(\mu + \eta) = 2\delta\sigma Q.$$

Moreover, we can bound the last term

$$
\begin{aligned}
-2\sigma^2 \int (P - c) S^2 (1 - S) \, d(\mu + \eta) &= 2\sigma^2 \int S \cdot (c - P) S (1 - S) \, d(\mu + \eta) \\
&< 2\sigma^2 \int S \cdot \bar{c} \cdot \delta \, d(\mu + \eta) \\
&= 2\delta\bar{c}\sigma^2 Q.
\end{aligned}
$$

The first line follows from rearranging the terms in the integral. The inequality in the second line follows from $c - P \le \bar{c}$ and $S(1 - S) < \delta$. The third line follows from evaluating the

integral.

Substituting these two bounds on equation (A4), we obtain

$$\partial_{PP}\Pi(P, p, x) < (-1 + 2\delta + 2\delta\bar{c}\sigma) \cdot \sigma \cdot Q(P, p, x).$$

The second derivative is negative if

$$-1 + 2\delta(1 + \bar{c}\sigma) \leq 0,$$

or

$$\delta(1 + \bar{c}\sigma) \leq \frac{1}{2},$$

which follows from inequality (A2). $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

The next claim shows that profits per unit sold are bounded by a multiple of the inverse of the semi-elasticity parameter σ plus a multiple of the firm capacity. This claim and its proof are based on Vives (1985).[4]

Claim 4. Assume that $\partial_P\Pi(P, p, x) = 0$. Then

$$\frac{\Pi(P, p, x)}{Q(P, p, x)} < \frac{2}{\sigma} + \frac{\bar{c}}{\bar{\eta}} \cdot k.$$

Proof. Equation (A3), which follows from the firm's first-order condition, can be written as

$$Q = \sigma \cdot \Pi - \sigma \cdot \int (P - c) \cdot S^2 \, d(\mu + \eta),$$

or equivalently as

$$\frac{\Pi}{Q} = \frac{1}{\sigma} + \frac{\int (P - c) \cdot S^2 \, d(\mu + \eta)}{Q}. \qquad\qquad\qquad (A5)$$

---

[4]This footnote clarifies how Vives's bound for profits in terms of the number of firms is related to our bound. In the particular case where there is no adverse selection (that is, costs do not depend on θ), our argument bounds profits above by $2/\sigma$ because $P - c$ is bounded above by $\delta/(1-\delta)\sigma$. Vives's Proposition 3 shows that prices in a differentiated-goods Bertrand model converge to marginal costs at rate $1/n$. His key assumption (A.5) is that inverse demand curves have bounded slopes. This is closely related to demand curves having slopes that are bounded away from 0. In our model the slope of demand is approximately equal to σq, where q is the quantity produced. Moreover, q is of order $1/n$, and therefore assumption (A.5) is similar to asking that σ/n is bounded away from 0 in the logit model. When this is the case, our bound of $2/\sigma$ implies the bound of order $1/n$ that Vives finds.

Moreover, note that

$$\int (P - c) \cdot S \cdot S \, d(\mu + \eta) \ \leq \ \int \bar{P} \cdot \delta \cdot S \, d(\mu + \eta)$$

$$\leq \ (\frac{1}{(1 - \delta) \cdot \sigma} + \bar{c}) \cdot \delta \cdot Q(P, p, \delta).$$

The first inequality uses the fact that $P \leq \bar{P}$, by Claim 2, and that $S \leq \delta$ by Claim 1. The second inequality follows from evaluating the integral.

Substituting this bound on equation (A5), we have

$$\frac{\Pi}{Q} \leq (1 + \frac{\delta}{1 - \delta}) \cdot \frac{1}{\sigma} + \delta \cdot \bar{c}.$$

Substituting $\delta < 1/2$, we have

$$\frac{\Pi}{Q} < \frac{2}{\sigma} + \delta \cdot \bar{c}.$$

Substituting $\delta = k/\bar{\eta}$ then gives the desired bound on profits. $\qquad \square$

We now employ these results to establish the proposition.

Proof. (Proposition A1) We begin by proving that an equilibrium exists. Define the modified profit function as

$$\tilde{\Pi}(P, p, x) = \Pi(P, p, x)$$

if $\Pi(P, p, x) > -\infty$ and

$$\tilde{\Pi}(P, p, x) = \Pi(\underline{P}(p, x), p, x) - (\underline{P}(p, x) - P)$$

if $\Pi(P, p, x) = -\infty$, where

$$\underline{P}(p, x) = \min_{\tilde{P}} \Pi(\tilde{P}, p, x) > -\infty.$$

The function $\tilde{\Pi}$ is continuous in P and p. Moreover, Claims 2 and 3 imply that $\tilde{\Pi}$ is quasiconcave in P. Therefore, the game defined by $\tilde{\Pi}$ has an equilibrium p^* if strategies are restricted to the compact interval $[0, \bar{P}]$ (Fudenberg and Tirole, 1991 p. 34 Theorem 1.2). Claim 2 implies that p^* is a best response to itself even if players can set any price because there are no best responses greater than $\bar{P}$. Finally, we defined $\tilde{\Pi}$ so that best responses are also best responses according to Π, so p^* is an equilibrium of the original game.

We now establish the bound on profits. In equilibrium, each firm serves at most $1/n$ consumers and, consequently, is below capacity. Therefore, the first-order condition must hold. Claim 4 implies the bound on profits. $\qquad \square$

A.5 Convergence of Bertrand Equilibria to Competitive Equilibria

Consider parameters (n, k, σ) of a Bertrand game with an equilibrium p^*. Define the allocation α^* **associated** with (n, k, σ) and p^* as, for any $x \in \bar{X}$ and $\tilde{\Theta} \subseteq \bar{\Theta}$,

$$\alpha^*(\{x\} \times \tilde{\Theta}) = \int_{\tilde{\Theta}} n \cdot S(p^*_x, p^*, x, \theta) \, d(\mu + \eta).$$

Proposition A2. *Consider a sequence of parameters $(n_l, k_l, \sigma_l)_{l \in \mathbb{N}}$ such that*

$$n_l \cdot k_l > 1,$$

$$k_l < \frac{\bar{\eta}}{2} \cdot \frac{1}{1 + \bar{c} \cdot \sigma_l},$$

and the semi-elasticities σ_l converge to infinity. Then:

1. *There exists a sequence $(p^*_l)_{l \in \mathbb{N}}$ where p^*_l is a symmetric Bertrand equilibrium associated with the parameters (n_l, k_l, σ_l).*

2. *Let α^*_l be the allocation associated with (n_l, k_l, σ_l) and p^*_l. Then the sequence $(p^*_l, \alpha^*_l)_{l \in \mathbb{N}}$ has a convergent subsequence.*

3. *If (p^*, α^*) is an accumulation point of $(p^*_l, \alpha^*_l)_{l \in \mathbb{N}}$, then (p^*, α^*) is a weak equilibrium of the perturbation $(E, \bar{X}, \eta)$.*

Proof. Part 1 follows from Proposition A1. Moreover, we can take the sequence $(p^*_l)_{l \in \mathbb{N}}$ to be bounded. Part 2 follows from compactness of the set of allocations and the fact that the sequence of prices is bounded.

We now prove part 3. We first show that consumers optimize under (p^*, α^*). Consider (θ, x) such that there exists $x' \in \bar{X}$ with

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

Then there exists $\epsilon > 0$ such that, for all θ' in a neighborhood of θ, and all sufficiently large l,

$$u(x, \theta') - p^*_l(x) < u(x', \theta) - p^*_l(x') - \epsilon.$$

This implies that

$$
\begin{aligned}
n \cdot S(p^*_l(x), p^*_l, x, \theta') \quad &\leq \quad \frac{\exp\{\sigma_l \cdot (u(x, \theta') - p^*_l(x))\}}{\exp\{\sigma_l \cdot (u(x, \theta') - p^*_l(x))\} + \exp\{\sigma_l \cdot (u(x', \theta') - p^*_l(x))\}} \\
&< \quad \frac{1}{1 + \exp\{\sigma_l \cdot \epsilon\}}
\end{aligned}
$$

converges to 0 uniformly in this neighborhood. Therefore, (x, θ) is not in the support of α^*.

It only remains to show that the price of each contract equals average cost. Note that

$$p_l^*(x) - \mathbb{E}_x[c|\alpha_l^*] = \frac{\int_{\{x\} \times (\Theta \cup \bar{X})} p - c \, d\alpha_l^*}{\int_{\{x\} \times (\Theta \cup \bar{X})} 1 \, d\alpha_l^*} = \frac{\Pi(p_l^*(x), p_l^*, x)}{Q_l(p_l^*(x), p_l^*, x)} < \frac{2}{\sigma_l} + \frac{\bar{c}}{\bar{\eta}} k_l,$$

where the inequality follows from Proposition A1. Taking limits (and noting that equilibrium profits are weakly positive) we have that

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

completing the proof. □

B   Details on the Calibration of Section 5

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. We modified these assumptions to make the results more transparent. Our model admits a 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 log normally distributed.

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 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 between linear contracts. As an illustration, using this value, 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 lead to both equilibria and optimal allocations involving essentially all consumers purchasing full coverage.

We used lower values of mean risk aversion in our calibrations. 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 linear contracts

have no stop losses. We chose the mean of A to ensure that risk aversion was within the range where substitution patterns are 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 1.0E-5, which makes the value of full insurance equal to approximately \$3,125. 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 lower but comparable to the gains from insurance. Estimates in this range place the surplus generated by a full insurance contract at about \$2,500.

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. Table B1 summarizes the parameters used in the calibration.

For the computations, we calculated equilibrium based on a simulated population with 1% behavioral types and 26 evenly spaced contracts. One important issue is that our model may have multiple equilibria. Thus, we want to guarantee that our predictions do not depend on the initial value in the computations. This is a particularly important issue because, as discussed in Scheuer and Smetters (2014), in the case of multiple equilibria, market outcomes can be sensitive to initial conditions. To assuage these concerns, we calculated equilibria from different starting values. We robustly found that our numerical method converged to the same equilibrium.

We calculated the optimal allocation numerically. The standard caveat applies that, in a complex nonlinear maximization problem, it is impossible to guarantee that a local optimum is a global optimum. We calculated local optima starting from the equilibrium allocation, both using an ad-hoc procedure and the commercial optimization package KNITRO. We also calculated local optima from 300 random starting values in each simulation. The random starting values tend not to do better than the optimization starting at the equilibrium prices. We used a similar procedure for the calibration with nonlinear contracts. Replication code is available at https://github.com/rafaelmourao/ag-competition.

References

Aksoy-Pierson, Margaret, Gad Allon, and Awi Federgruen, "Price Competition under Mixed Multinomial Logit Demand Functions," *Management Science*, 2013, *59* (8), 1817–1835.

Caplin, Andrew and Barry Nalebuff, "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 1991, pp. 25–59.

Decarolis, Francesco, Maria Polyakova, and Stephen Ryan, "The Welfare Effects of Supply Side Regulations in Medicare Part D," *Mimeo, Stanford University*, 2012.

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

Fudenberg, Drew and Jean Tirole, *Game Theory*, MIT Press, 1991.

Guerrieri, Veronica, Robert Shimer, and Randall Wright, "Adverse Selection in Competitive Search Equilibrium," *Econometrica*, 2010, *78* (6), 1823–1862.

Handel, Benjamin, Igal Hendel, and Michael D Whinston, "Equilibria in Health Exchanges: Adverse Selection Vs. Reclassification Risk," *Econometrica*, 2015, *83* (4), 1261–1313.

Inderst, Roman and Achim Wambach, "Competitive Insurance Markets under Adverse Selection and Capacity Constraints," *European Economic Review*, 2001, *45* (10), 1981–1992.

Mahoney, Neale and E Glen Weyl, "Imperfect Competition in Selection Markets," *Review of Economics and Statistics*, forthcoming.

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

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.

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 19835*, 2014.

Starc, Amanda, "Insurer Pricing and Consumer Welfare: Evidence from Medigap," *The Rand Journal of Economics*, 2014, *45* (1), 198–220.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *Mimeo, Stanford University*, 2015.

Vives, Xavier, "On the Efficiency of Bertrand and Cournot Equilibria with Product Differentation," *Journal of Economic Theory*, 1985, *36* (1), 166–175.

Table B1: Calibrated distribution of consumer types (linear contracts)

	A	H	M	S
Mean	1.0E-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: Consumer types are log normally distributed with the moments in the table. The log variance of moral hazard $\sigma^2_{\log H}$ is set equal to 0.28 and 0.98.