Competitive Nonlinear Pricing under Adverse Selection*

Andrea Attar† Thomas Mariotti‡ François Salanié§

January 2, 2023

Abstract

This article surveys recent attempts at characterizing competitive allocations under adverse selection when each informed agent can privately trade with several uninformed parties: that is, trade is nonexclusive. We first show that requiring market outcomes to be robust to entry selects a unique candidate allocation, which involves cross-subsidies. We then study how to implement this allocation as the equilibrium outcome of a game in which the uninformed parties, acting as principals, compete by making offers to the informed agents. We show that equilibria typically fail to exist in competitive-screening games, in which these offers are simultaneous. We finally explore alternative extensive forms, and show that the candidate allocation can be implemented through a discriminatory ascending auction. These results yield sharp predictions for competitive nonexclusive markets.

Keywords: Adverse Selection, Entry-Proofness, Discriminatory Pricing, Nonexclusive Markets, Ascending Auctions.
JEL Classification: D43, D82, D86.

---

*The material in this article has been presented at the semi-plenary lecture "Constests and Contracts" of the 2020 World Congress of the Econometric Society (Università Bocconi, Milano, Italy). This research has benefited from financial support of the ANR (Programme d'Investissement d'Avenir ANR-17-EURE-0010), the MIUR (PRIN 2015), and the research foundation TSE-Partnership.

†Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, and Università degli Studi di Roma "Tor Vergata," Roma, Italy. Email: `andrea.attar@tse-fr.eu`.

‡Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, CEPR, and CESifo. Email: `thomas.mariotti@tse-fr.eu`.

§Toulouse School of Economics, INRAE, University of Toulouse Capitole, Toulouse, France. Email: `francois.salanie@tse-fr.eu`.

1 Introduction

On the vast majority of markets, nonexclusivity is the rule: agents can privately trade with different parties, without having to inform each of them of these multiple relationships. Fortunately, detailed knowledge of individual trades is generally useless for the involved parties. Nor is it usually needed, moreover, to predict market outcomes: instead, by solely relying on the aggregate equality of supply and demand, general-equilibrium theory elegantly sidesteps the lack of such information to focus on the determination of equilibrium prices. These prices in turn convey all the information agents need to know in order to formulate their supplies and demands.

However, there are important cases in which information about the characteristics of the goods for trade is not symmetrically distributed, though it directly matters to the parties involved. The chief example is that of goods whose quality is privately known to a party but unknown to his trading partners, including, for instance, the sale of shares in a firm with different returns, the supply of labor services by workers with different productivities, or the design of insurance contracts for consumers with different riskiness.

These common-value situations raise important difficulties for the performance of markets. They also question the relevance of general-equilibrium theory, because, for instance, a buyer may now try to infer the missing information from the seller's behavior and not only from prices. This gives rise to new phenomena such as adverse selection, reflecting that agents who are more eager to trade are often endowed with goods of lower quality; moreover, these inferences are made even more difficult when trade is nonexclusive. Game-theoretical tools then become useful to precisely describe who trades with whom and how the distribution of information impacts behaviors and outcomes.

This article surveys recent developments in the theory of competitive markets under adverse selection and nonexclusivity. We are particularly interested in how price competition and the threat of entry lead to sharp predictions for equilibrium outcomes, in spite of the complexities associated with nonexclusive trading. We will show how different approaches often lead to the same prediction, though we also emphasize difficulties for the existence of an equilibrium, depending on the game studied and the solution concept adopted. These results may be of interest for many financial and insurance markets, including over-the-counter, life-insurance, and annuity markets; for labor markets such as the markets for professionals or freelance workers; and, more generally, for many markets in goods or services whose quality is the private information of agents on one side of the market.

Because this combination of adverse selection and nonexclusivity is the hallmark of the literature reviewed in this survey, it may be helpful to briefly recall how these two topics have been addressed in now classical works.

On the one hand, adverse selection has so far been typically studied under exclusive competition. This is by design in Akerlof's (1970) example of a market for an indivisible good. Subsequent works focusing on markets for perfectly divisible goods assume that exclusive contracts are enforceable. In Rothschild and Stiglitz (1976), this allows an insurance company to screen its customers by making low-risk consumers self-select into contracts with higher deductibles. In Leland and Pyle (1977), this allows an entrepreneur to signal the profitability of her project by retaining a greater or lesser equity share. In both cases, the observability of agents' aggregate trades is required to support equilibrium outcomes. Extensions by Prescott and Townsend (1984), Kehoe, Levine, and Prescott (2002), Bisin and Gottardi (2006), and Rustichini and Siconolfi (2008) of the standard existence and welfare theorems of general-equilibrium theory to private-information economies similarly restrict feasible allocations to those satisfying incentive-compatibility of individual trades on each market, which again requires strong observability assumptions.

On the other hand, nonexclusive competition has so far been mostly studied in private-values environments. As noticed above, in such cases, the functioning of Walrasian markets is unaffected by private information. The side-trading literature, from the early works of Hammond (1979, 1987), Allen (1985), and Jacklin (1987) to the more recent contributions of Cole and Kocherlakota (2001), Golosov and Tsyvinski (2007), and Farhi, Golosov, and Tsyvinski (2009), has accordingly investigated the limits that side trading on Walrasian markets, outside of the central planner's control, imposes on the set of allocations that can be achieved in standard Mirrlees (1971) or Diamond and Dybvig (1983) economies, which feature private information but not adverse selection.

By contrast, in common-values environments, assuming a priori that a Walrasian market exists and perfectly balances supply and demand is too much to ask for; in fact, we shall see in Section 2.2 that sellers on such a market, anticipating adverse selection, would like to reduce their competitive supply and thus choose to ration demand. In addition, nonexclusivity means that trades cannot be monitored, so that even a centralized market authority would not be able to ensure that all agents trade at the same price. Instead, the recent literature on nonexclusive markets under adverse selection assumes that contracting is bilateral, and allows for general contracts.

We now present the main findings of this literature. In line with Riley (2001), we adopt in this survey a strategic approach to the determination of market outcomes; in contrast with him, however, we exclusively focus on screening models.[1] A market is, therefore, described by a set of uninformed sellers competing through menus of contracts, or nonlinear tariffs, to serve the demand emanating from privately informed buyers. Nonexclusivity is captured by the assumption that, while each seller can monitor the trades each buyer conducts with him—which is what makes nonlinear pricing possible—he can monitor none of the trades this buyer makes with his competitors.

A useful entry point into the literature is to first abstract from the determination of individual tariffs and, in a reduced-form way, to directly impose properties on the market tariff that is obtained from them by aggregation. In line with Rothschild and Stiglitz (1976), who characterize the set of exclusive contracts preventing an entrant from making a profit, a desirable property of the market tariff is that it be entry-proof. Under nonexclusivity, this property means that no entrant can make a profit by offering a menu of contracts, given that each buyer is free to combine a contract offered by the entrant with a trade along the market tariff. Restricting attention to entry-proof market tariffs allows us to identify robust predictions for nonexclusive markets under adverse selection, which do not depend on the details of a specific extensive-form game. Section 3 of this article is devoted to the characterization of such tariffs.

The following insights emerge from the analysis (AMS (2020, 2021)).[2] First, neither the Rothschild and Stiglitz (1976) allocation, nor any of the second-best allocations characterized by Prescott and Townsend (1984) and Crocker and Snow (1985) can be implemented by an entry-proof market tariff. Second, entry-proofness, along with budget-feasibility, singles out a unique market tariff, which generally turns out to be nonlinear. The defining features of this tariff are that each marginal quantity is priced at the expected cost of serving the buyer types who optimally choose to trade it, and that gains from trade are exhausted subject to this constraint. For instance, in the two-type case, the low-cost type purchases her demand at a price equal to the expected cost, while the high-cost type in addition trades the quantity she prefers at a price that equals her own cost. The corresponding allocation, first described by Jaynes (1978), Hellwig (1988), and Glosten (1994)—and henceforth referred to as the JHG allocation—can thus be interpreted as a marginal version of the Akerlof (1970) competitive-equilibrium allocation, and stands out as a focal prediction for nonexclusive competitive markets under adverse selection. The entry-proofness criterion can also be seen as an external constraint imposed on the decisions of a central planner. AMS (2020)

[1] Extending the theory of signaling or, more generally, the informed-principal paradigm of Myerson (1983) to nonexclusive markets is a fascinating task for future research.

[2] To avoid repetitions, AMS hereafter stands for Attar, Mariotti, and F. Salanié.

emphasize that this constraint is so strong so as to limit the planner's feasible policies to a unique policy, which is not second-best. The social costs of side trading thus appear to be particularly severe under common values.

In line with the program outlined by Wilson (1989), a natural question is whether the JHG allocation can be implemented as an—ideally, unique—equilibrium outcome of a decentralized trading protocol. The contributions reviewed in Section 4 focus on competitive-screening games in which offers are simultaneously made by the uninformed sellers; they thus bring to bear insights and methods from the common-agency literature, in which several uninformed principals compete by posting menus of contracts to deal with a privately informed agent. While, as surveyed by Martimort (2006), common-agency games have mainly been used to tackle issues in industrial organization, public economics, and in the theory of organizations—from vertical contracting and the internal structure of the firm to lobbying and the relationships between governments and regulatory agencies—they have also provided, following the seminal contribution of Biais, Martimort, and Rochet (2000), a powerful tool for modeling competition under common values.

The results of this approach, however, are mixed. The only case in which the JHG allocation can be robustly implemented through a competitive-screening game is when the buyer's preferences are linear, subject to a capacity constraint. The JHG allocation then coincides with the Akerlof (1970) competitive-equilibrium allocation that maximizes the gains from trade, and it is the unique candidate-equilibrium allocation; moreover, there always exists an equilibrium in which sellers posts linear tariffs (AMS (2011)). These positive results thus extend the conclusions of Akerlof (1970) to the case of a divisible good. However, when the buyer has strictly convex preferences, the results crucially depend on fine modeling details such as the cardinality of the set of types. When the distribution of types is continuous, Biais, Martimort, and Rochet (2000) construct an equilibrium in which sellers posts strictly convex tariffs, and such that the resulting aggregate equilibrium allocation converges to the JHG allocation as the number of sellers grows large; however, as we show in Section 4.4, their existence result requires that some types be excluded from trade. When the distribution of types is discrete, exclusion is even more extreme, as an equilibrium exists only in the degenerate case where all types except possibly the highest-cost one do not trade in the JHG allocation (AMS (2014, 2019a)).

Section 5 provides more positive results, by exploring alternative extensive forms whereby uninformed sellers sequentially receive information about previously signed contracts or previously made offers. The bottom line is that transparency makes it easier for sellers to directly punish deviators, in contrast with competitive-screening games in which the burden of punishments entirely falls on the buyer. In this spirit, Beaudry and Poitevin (1995) and AMS (2021) implement the JHG allocation in a repeated game of signalling and in an ascending discriminatory auction with frequent offers, respectively. We also survey contributions by Jaynes (1978, 2011), Hellwig (1988), and Stiglitz, Yun, and Kosenko (2020) that allow for endogenous information disclosure.

Section 6 concludes on the empirical perspectives, in particular about tests for the presence of private information.

2 The Economy

We study a simple economy in which a single buyer (she) trades a divisible good with multiple, identical sellers (he). The buyer is endowed with private information about, for instance, the quality of the traded good. As usual, the buyer/sellers convention can be inverted thanks to a change of variables, so as to encompass a broad variety of situations:

- Insurance companies sell coverage to a consumer: Rothschild and Stiglitz (1976), Prescott and Townsend (1984), Crocker and Snow (1985), Hendren (2013).
- Market makers provide liquidity to an insider: Glosten (1989, 1994), Biais, Martimort, and Rochet (2000), Back and Baruch (2013).
- Investors purchase securities issued by a firm: Leland and Pyle (1977), Myers and Majluf (1984), DeMarzo and Duffie (1999), Biais and Mariotti (2005).
- Firms hire the services of a worker: Spence (1973), Miyazaki (1977).

In these situations, the private information of the buyer is directly relevant to the sellers, because it determines their production costs or their opportunity costs of selling. It is this common-value component that may generate adverse selection, as we now discuss in more formal terms.

2.1 The Model

Unless stated otherwise, the following assumptions are maintained throughout this article.

The Buyer The buyer's private information is represented by a type $i = 1, \ldots, I$ that takes a finite number of values with strictly positive probabilities m_i. Type i's preferences are represented by a utility function $u_i(Q, T)$ that is continuous and weakly quasiconcave in

(Q,T) and strictly decreasing in T, with the interpretation that Q is the nonnegative quantity of the good she purchases and T is the payment she makes in return. To define marginal rates of substitution without assuming differentiability, we let $\tau_i(Q,T)$ be the supremum of the set of prices p such that

$$u_i(Q,T) < \max\{u_i(Q+q, T+pq) : q \geq 0\}.$$

Thus $\tau_i(Q,T)$ is the slope of type i's indifference curve at the right of (Q,T). Quasiconcavity ensures that $\tau_i(Q,T)$ is finite, except possibly at $Q=0$, and that it is nonincreasing along an indifference curve of type i. For all i and $p>0$, we also define the demand $D_i(p)$ of type i as the set of quantities Q that maximize $u_i(Q,pQ)$. These demands are well-defined under the following Inada condition:

$$\text{For all } i,\ (Q,T), \text{ and } p>0,\ \arg\max\{u_i(Q+q,T+pq) : q \geq 0\} < \infty, \tag{1}$$

or if the domain of admissible quantities is compact. Types are ordered according to the weak single-crossing condition (Milgrom and Shannon (1994)), which states that higher types are at least as willing to increase their purchases as lower types are:

$$\text{For all } i<j,\ Q<Q',\ T, \text{ and } T',\ u_i(Q,T) \leq (<)\, u_i(Q',T') \text{ implies } u_j(Q,T) \leq (<)\, u_j(Q',T').$$

Weak single-crossing implies that $\tau_i(Q,T)$ and $D_i(p)$ are weakly increasing in i. For future reference, we also state the slightly stronger, strict single-crossing condition:

$$\text{For all } i<j,\ Q<Q',\ T, \text{ and } T',\ u_i(Q,T) \leq u_i(Q',T') \text{ implies } u_j(Q,T) < u_j(Q',T').$$

We shall occasionally make additional assumptions. Theorems 1 and 3, for instance, require that higher endowments of the good reduce the buyer's marginal rate of substitution:

Assumption 1 *For all i and T, $\tau_i(Q,T)$ is nonincreasing in Q.*

Our assumptions on the buyer's preferences hold, for instance, in a Rothschild and Stiglitz (1976) insurance economy in which the loss L is the same for all types: then i indexes the buyer's riskiness, Q is the amount of coverage she purchases, and T is the premium she pays in return. AMS (2021, Online Appendix C) show that these assumptions also hold in more general insurance economies, allowing for multiple loss levels or various forms of nonexpected utility. But our framework is relevant beyond insurance; in particular, first-best quantities may differ across types.

The Sellers On the supply side, sellers are identical, risk-neutral, and use the same linear technology. We denote by $c_i > 0$ the unit cost of serving type i, and by $\bar{c}_i$ the corresponding upper-tail conditional expectation of unit costs,

$$\bar{c}_i \equiv \mathbf{E}[c_j \,|\, j \geq i] = \frac{\sum_{j \geq i} m_j c_j}{\sum_{j \geq i} m_j}.$$

Adverse selection occurs if the unit cost c_i is nondecreasing in i. This case is the least conducive to trade, as types who are more willing to trade are also more costly to serve. In this article, we only rely on a slightly weaker assumption, namely, that $\bar{c}_i$ be nondecreasing in i. This weak adverse-selection condition is exactly equivalent to

$$\text{For all } j \leq i, \ c_j \leq \bar{c}_i. \tag{2}$$

Contracts A *contract* (q, t) between a seller and the buyer specifies a nonnegative quantity to be delivered by the seller and a transfer to be made in return by the buyer. Under nonexclusivity, two contracts (q, t) and (q', t') offered by different sellers can be added to form a trade $(q + q', t + t')$. In this article, we consider two different settings in turn.

- We first study when a nonexclusive market is entry-proof. We then only rely on the existence of a *market tariff* T, where $T(Q)$ is defined as the minimum transfer that allows the buyer to obtain a quantity Q. If an entrant proposes additional contracts, then the buyer can pick one of them, say (q, t), together with a trade $(Q, T(Q))$ along the market tariff, ending up with utility $u_i(Q + q, T(Q) + t)$.

- We next study competition in menus of nonexclusive contracts among $K \geq 2$ sellers. We then have to precisely specify the menus that are offered by the sellers $k = 1, \ldots, K$. We impose that each seller must always propose at least the null contract $(0, 0)$, so that the buyer may be seen as trading one contract (q^k, t^k) with each seller k, ending up with aggregate trade $(Q, T) \equiv (\sum_k q^k, \sum_k t^k)$ and utility $u_i(Q, T)$.

Let us emphasize that we focus on situations in which trade is not anonymous. Our view of nonexclusivity is thus that a seller can monitor the trades the buyer makes with him, though he cannot monitor the trades the buyer makes with the other sellers. This is a natural assumption to make in insurance markets, because an insurance contract must name a beneficiary. Preventing the buyer from making concealed repeat purchases from the same seller enables each seller to price the quantities he sells in a nonlinear way, charging different prices for different marginal units.

Our analysis and results extend to the case of multiple buyers, provided contracting is bilateral and the buyers' types are independent and identically distributed. Contracting is bilateral if trade between a seller and a buyer is only contingent on the information reported by the buyer to the seller, and not on the information this seller may obtain from other buyers.[3] Together with the linearity of costs, the independence of types across buyers then implies that the interactions between a seller and each of his potential customers can be studied separately. Finally, if the buyers' types are identically distributed, we can assume, using a symmetry argument, that each seller offers the same contracts to each buyer and that each type of each buyer facing the same choices behaves in the same way. In this way, the analysis of the multiple-buyer case can be reduced to that of the single-buyer case.

2.2 Benchmarks

The general question addressed in this article is how to define a notion of competitive allocation for the above-described economy. This section discusses a few benchmarks, with the purpose of introducing the main effects and difficulties.

The complete-information benchmark assumes that the buyer's type is made public before sellers make their supply decisions, so that information is symmetric about this collective risk (Malinvaud (1972)). The interpretation is that the good for trade comes in I observable varieties, each representing a different quality. When this quality i is revealed, the sellers learn their cost c_i; competition then implies that type i purchases her demand $D_i(c_i)$ at price c_i and that sellers make zero profits. Therefore, equilibria exist and are efficient.

When information becomes asymmetric, the now privately informed buyer optimally channels her demand to the market with the lowest price $p \equiv \min_i c_i$. But then aggregate profits $\mathbf{E}[(p - c_i)D_i(p)]$ are typically negative, except when unit costs are independent of the buyer's type, that is, $c_i \equiv c$ for all i. In this private-value case, sellers know their costs, while the tastes of the buyer are her private information. Nevertheless, a competitive market still plays its allocative role: in equilibrium, every type i purchases her demand $D_i(c)$ at price c, and sellers make zero profits. Once more, equilibria exist and are efficient.

We allow for private values as a limiting case, but our main focus is on the common-value case where the sellers' unit costs depend on the buyer's type. Proceeding as in Akerlof (1970), Pauly (1974) studies the case where linear pricing is imposed, independently of the quantity traded; that is, sellers stand ready to supply any quantity at the going price. Then an

[3]Under multilateral contracting, we enter the realm of multi-principal multi-agents models, in which unrestricted communication can be exploited to support many equilibrium allocations; see, for instance, Yamashita (2010).

equilibrium exists as soon as demand functions are continuous. Because equilibrium profits must be zero under constant returns to scale, the equilibrium price satisfies the equality

$$p = \mathbf{E}\left[c_i \, \frac{D_i(p)}{\mathbf{E}[D_j(p)]}\right]. \tag{3}$$

This formula is widely used in the annuity literature (Sheshinski (2008), Hosseini (2015), Rothschild (2015)). When the good is indivisible, each demand term in (3) is either zero or one, and we are back to the classical Akerlof (1970) formula that states that the price is equal to the expected unit cost of active types. With a divisible good, the formula in addition weighs the unit cost of serving each type by her demand. Because higher types have higher demands, it generally follows that the equilibrium price must lie above the expected unit cost of active types. Accordingly, the active types with the lowest costs subsidize the higher types, who are more costly to serve.

However, there exists a simple way to reduce the risk of having to sell too much to a high type at the going price: to this end, a seller need only post a *limit order* $(p, \overline{q})$ specifying the maximum quantity $\overline{q}$ he is ready to sell at price p. Such limit orders are commonly used on financial markets, and this may indeed be because they allow sellers to hedge against the risk of a high demand.[4] A well-chosen limit order, with a price just below p, is profitable because it reduces the loss-making sales to high types while preserving the profits from selling to low-cost types.[5] This, incidentally, shows that the Pauly (1974) outcome is not a competitive equilibrium: anticipating adverse selection, sellers would like to reduce their competitive supply, thereby collectively rationing demand.

We conclude that the linear-pricing construction is rather fragile under adverse selection. A natural step forward is to consider a competitive game in which sellers are allowed to post limit orders.[6] Notice that a collection of limit orders gives rise to a convex market tariff and, conversely, that any convex market tariff can be decomposed into a (possibly infinite) collection of (possibly infinitesimal) limit orders. In what follows, we sometimes impose that tariffs be convex, but we also explore cases where tariffs can be arbitrary.

[4] We follow the literature in using the term "limit order," although the maximum quantity $\overline{q}$ is here understood to apply to a single buyer, whereas a limit order on a financial market specifies the maximum aggregate quantity one is ready to sell to a set of traders.

[5] AMS (2019a, Lemma 4) formalize this intuition and show that, under adverse selection, such a limit order allows a seller to approximate the maximum profit he can earn when competing with a linear tariff.

[6] The game in which sellers compete by posting a single limit order implements the efficient competitive equilibria when information is symmetric and cost functions are weakly convex. In the present model with asymmetric information and linear cost functions, AMS (2018, pp. 1013–1014) show that this game has a pure-strategy equilibrium only in limiting cases where costs or demands do not depend on the buyer's type.

3 Entry-Proofness

The idea of using entry-proofness as a solution concept first originates in an attempt at simplicity, because this avoids the need to precisely describe the supply side of the economy or to fully specify the details of an extensive form. Moreover, entry-proof allocations, when they exist, are widely considered as capturing the idea of perfect competition.[7] We begin by studying when inactive markets are entry-proof. We then turn to active markets, for which the distinction between exclusive and nonexclusive competition becomes relevant, and we formulate a definition of entry-proofness consistent with nonexclusivity. We show in particular that entry-proofness selects a unique budget-balanced allocation, which exists under very general conditions. This requirement is thus more fruitful under nonexclusivity than under exclusivity.

3.1 Entry-Proofness in Inactive Markets

In this section, we describe the circumstances under which private information impedes trade altogether. We say that a market is *inactive* if the market tariff reduces to a single point, given by $T(0) = 0$, or, equivalently, if only the null contract $(0, 0)$ is available. Our goal is to find conditions ensuring that no entrant can make an offer leading to profitable trades. Accordingly, we say that an inactive market is *entry-proof* if, *for any menu of contracts offered by an entrant, the buyer has a best response such that the entrant earns at most zero expected profit.*

To characterize the inactive markets that are entry-proof, we first study the simplest case where the entrant offers a single contract. The key argument here is that, if this contract strictly attracts a type i, then it must also attract all types $j > i$: this is a simple consequence of weak single-crossing. Hence, from the entrant's viewpoint, the relevant unit cost is not the individual cost c_i of serving type i, but, rather, the expected cost $\bar{c}_i$ of serving types $j \geq i$. Notice that some other types $j < i$ may also be attracted by the entrant's offer, but the weak adverse-selection condition (2) ensures that this can only reduce the entrant's expected unit cost. This shows that the following condition is necessary for entry to be unprofitable.

Condition EP *For each i, $\tau_i(0, 0) \leq \bar{c}_i$.*

Notice that Condition EP does not rule out gains from trade, in the usual first-best sense of the term; that is, it may well be that $\tau_i(0, 0) > c_i$ for some i. AMS (2021, Theorem

[7] As argued by Rothschild and Stiglitz (1976, p. 642): "The basic idea underlying competitive markets involves free entry and noncollusive behavior among the participants in the market."

1) show that this necessary condition for entry-proofness is sufficient even when menus of contracts are allowed:

Theorem 1 *Under Assumption 1, an inactive market is entry-proof if and only if Condition EP is satisfied.*

The intuition is as follows. If the entrant offers an arbitrary menu of contracts, then, by weak single-crossing, the buyer has a best response with nondecreasing quantities, which we denote by $(q_i, t_i)_{i=1}^I$. Suppose that $q_i > q_{i-1}$ for some type i, and let us first locate the contract (q_{i-1}, t_{i-1}). We can safely assume that t_{i-1} is positive, as the entrant's profit can only be reduced by giving away costly production. We also know that type $i-1$ weakly prefers this contract to the null contract; by weak single-crossing again, so does type i. Therefore, the point (q_{i-1}, t_{i-1}) must lie in the north-east quadrant in Figure 1, at the right of the indifference curve of type i that goes through the origin.

Now, to be willing to trade the contract (q_i, t_i), it must be that type i, having already traded the contract (q_{i-1}, t_{i-1}), is willing to trade the additional *layer* $(q_i - q_{i-1}, t_i - t_{i-1})$. To evaluate her marginal rate of substitution at (q_{i-1}, t_{i-1}), we can use, in turn, the concavity of the indifference curve of type i, then Assumption 1, and finally Condition EP to obtain the following inequalities:

$$\tau_i(q_{i-1}, t_{i-1}) \leq \tau_i(\underline{q}_i, 0) \leq \tau_i(0, 0) \leq \overline{c}_i. \tag{4}$$

This implies that type i is not ready to pay more than $\overline{c}_i(q_i - q_{i-1})$ for the additional quantity $q_i - q_{i-1}$. Therefore,

$$t_i - t_{i-1} \leq \overline{c}_i(q_i - q_{i-1}).$$

Summing these inequalities over i with appropriate weights yields

$$\sum_i \left(\sum_{j \geq i} m_j \right) [t_i - t_{i-1} - \overline{c}_i(q_i - q_{i-1})] \leq 0.$$

Finally, rearranging terms in the spirit of Wilson (1993), we obtain

$$\sum_i m_i(t_i - c_i q_i) \leq 0,$$

which shows that entry cannot be profitable.

A noticeable feature of this proof is that it does not consider each contract (q_i, t_i) in isolation. Instead, the key role is played by layers of the form $(q_i - q_{i-1}, t_i - t_{i-1})$. Under

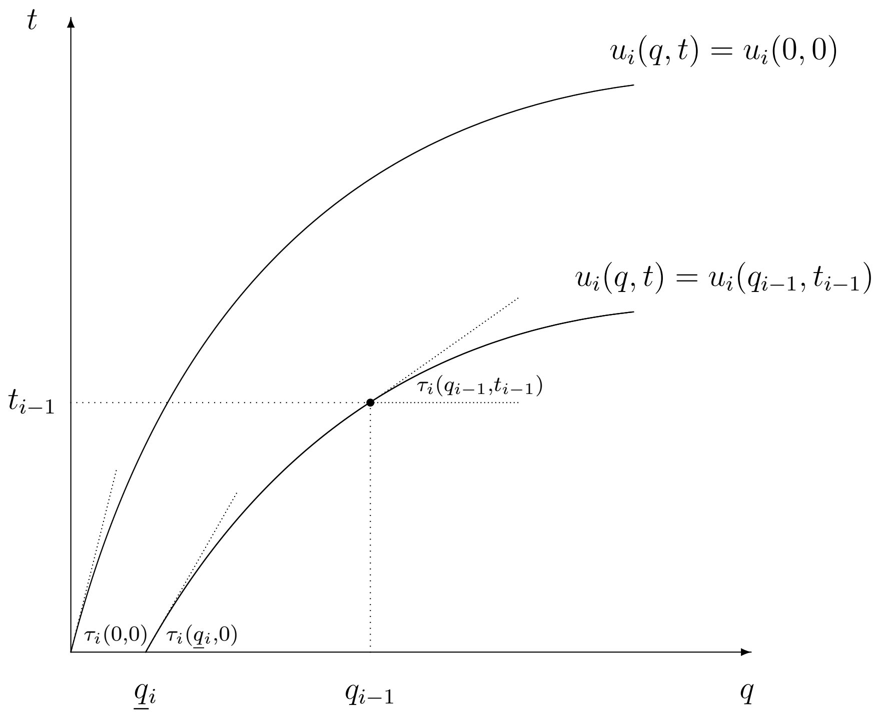

Figure 1: A graphical illustration of (4).

weak single-crossing, optimal quantities can be assumed to be nondecreasing in the buyer's type, so that the i^{th} layer can be thought of as traded by all types $j \geq i$, and thus has expected unit cost $\overline{c}_i$. Condition EP implies that, at this price, type i is not strictly willing to trade, so that each layer must yield a nonpositive expected profit. By contrast, some of the *contracts* proposed in a menu may yield positive profits. For instance, although the condition $t_1 \leq \overline{c}_1 q_1$ ensures that the expected profit on the first layer (q_1, t_1) is nonpositive, it may well be that $t_1 > c_1 q_1$.

AMS (2021) show that the assumptions of Theorem 1 can be weakened in several ways; however, the weak single-crossing condition and the seemingly innocuous Assumption 1 are tight. They also provide a result characterizing *market breakdown*, defined as a situation in which *any menu of contracts that strictly attracts at least some type yields a strictly negative expected profit, even if the buyer's best response is most favorable to the entrant.* Condition EP clearly remains necessary for this stronger concept, and it also remains sufficient under slightly stronger conditions on preferences. Earlier results were obtained by Mailath and Nöldeke (2008) for an economy in which the buyer has quadratic quasilinear preferences, and by Hendren (2013) for a Rothschild and Stiglitz (1976) insurance economy.

3.2 Entry-Proofness in Active Markets: The Two-Type Case

We now turn to *active* markets, on which nonnull contracts are available. In line with Rothschild and Stiglitz (1976), our goal is to characterize when entry on such a market is unprofitable, given the contracts available; in contrast with them, we suppose that the buyer can trade with several sellers. To this end, the proper object of study is the *market tariff*, which describes the frontier of the set of aggregate trades that can be achieved by trading on the market.

A market tariff specifies the minimum aggregate transfer $T(Q)$ required to purchase an aggregate quantity Q, with $T(Q) \equiv \infty$ if this is impossible; notice that we obviously have $T(0) = 0$. By assuming that T is lower semicontinuous, with a compact domain, we ensure that, for every type i, the problem of maximizing $u_i(Q, T(Q))$ admits a solution Q_i. We then say that the allocation $(Q_i, T(Q_i))_{i=1}^I$ is *implemented* by the tariff T. We assume that types are ordered according to the strict single-crossing condition, so that the optimal quantities Q_i are nondecreasing in i. Moreover, this allocation is *budget-feasible* if

$$\sum_i m_i[T(Q_i) - c_i Q_i] \geq 0. \tag{5}$$

Now, suppose an entrant can propose additional trades to the buyer, in the form of a menu of contracts that complement the market tariff. We say that the tariff T is *entry-proof* if, *for any menu of contracts offered by an entrant, the buyer has a best response such that the entrant earns at most zero expected profit, given that the buyer is free to combine any contract offered by the entrant with a trade along the tariff* T. The last clause of this definition is crucial, and captures the nonexclusivity of trade.

Our goal is to characterize the set of budget-feasible allocations that are implemented by entry-proof market tariffs. In this section, we focus on the two-type case $I = 2$, which is simple enough to allow for a precise discussion of the proof to the main result; the weak adverse-selection condition (2) then amounts to $c_1 \leq c_2$.

Thus consider an allocation $(Q_i, T_i)_{i=1}^2$ that is implemented by some market tariff T. Because this allocation is incentive-compatible, it satisfies $Q_2 \geq Q_1$ by strict single-crossing. Moreover, if T is entry-proof, then we must have

$$u_1(Q_1, T_1) \geq \max\{u_1(q, \bar{c}_1 q) : q \geq 0\}. \tag{6}$$

Otherwise, an entrant can offer a contract with unit price slightly above $\bar{c}_1$ that profitably attracts type 1, and remains profitable even if type 2 is attracted—recall that, by definition, $\bar{c}_1 = m_1 c_1 + m_2 c_2$. Similarly, we must have

$$u_2(Q_2, T_2) \geq \max\{u_2(Q_1 + q, T_1 + c_2 q) : q \geq 0\}. \tag{7}$$

Otherwise, an entrant can offer a contract with unit price slightly above c_2 that profitably attracts type 2 along with the contract (Q_1, T_1), and is even more profitable if type 1 is also attracted; notice that this second type of entry is specific to the nonexclusive case.[8] It follows from (6) that

$$T_1 \leq \bar{c}_1 Q_1. \tag{8}$$

Similarly, it follows from (7) that

$$T_2 \leq T_1 + c_2(Q_2 - Q_1). \tag{9}$$

However, rearranging terms in the spirit of Wilson (1993), the budget-feasibility constraint (5) can be rewritten as

$$T_1 - \bar{c}_1 Q_1 + m_2[T_2 - T_1 - c_2(Q_2 - Q_1)] \geq 0. \tag{10}$$

Thus, in light of (10), the equalities (8)–(9) are in fact equalities. That is, profits are zero on the first layer (Q_1, T_1), which is traded by both types; similarly, profits are zero on the second layer $(Q_2 - Q_1, T_2 - T_1)$, which is traded by type 2 only. This, in turn, implies that the inequalities (6)–(7) are also equalities. Overall, the four resulting equalities pin down the set of candidates for a budget-feasible allocation that is implemented by an entry-proof tariff. AMS (2020, Theorem 2) show that these necessary conditions for entry-proofness are sufficient even when menus of contracts are allowed:

Theorem 2 *Any budget-feasible allocation* $(Q_i^*, T_i^*)_{i=1}^2$ *that is implemented by an entry-proof market tariff satisfies*

$$Q_1^* \in \arg\max\{u_1(Q, \bar{c}_1 Q) : Q \geq 0\}, \tag{11}$$

$$T_1^* = \bar{c}_1 Q_1^*, \tag{12}$$

$$Q_2^* - Q_1^* \in \arg\max\{u_2(Q_1^* + q, T_1^* + c_2 q) : q \geq 0\}, \tag{13}$$

$$T_2^* - T_1^* = c_2(Q_2^* - Q_1^*). \tag{14}$$

Conversely, any allocation that satisfies (11)–(14) *can be implemented by the piecewise-linear convex market tariff*

$$T^*(Q) \equiv 1_{\{Q \leq Q_1^*\}} \bar{c}_1 Q + 1_{\{Q_1^* < Q \leq Q_2^*\}}[\bar{c}_1 Q_1^* + c_2(Q - Q_1^*)], \tag{15}$$

and this tariff is entry-proof.

[8] In a Rothschild and Stiglitz (1976) insurance economy, Stiglitz, Yun, and Kosenko (2020, Definition 1) base their analysis on the assumption that additional coverage is available without limits at price c_2, which implies an inequality similar to (7). They do not, however, state the first inequality (6).

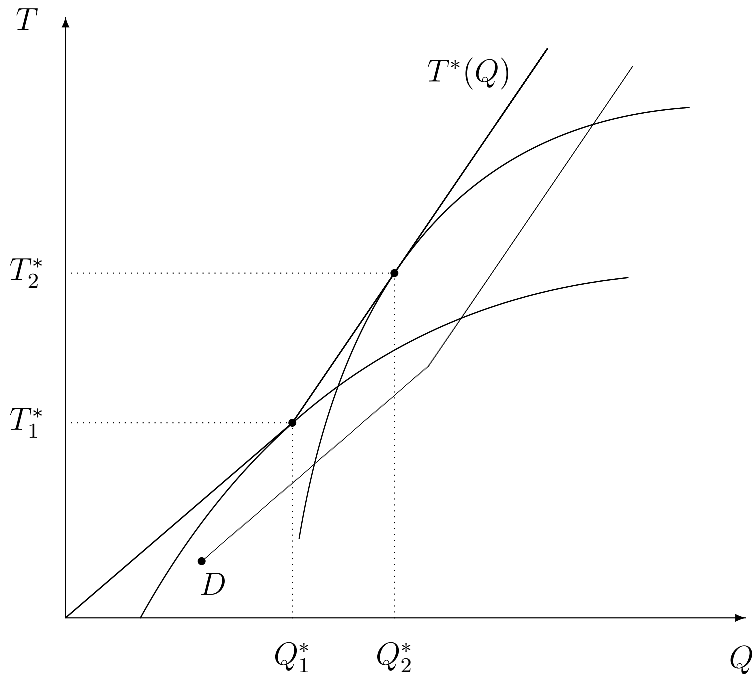

Figure 2: Blocking cream-skimming deviations.

When u_1 and u_2 are strictly quasiconcave—and also, generically, when they are only weakly quasiconcave—conditions (11)–(14) characterize a unique allocation. Notice that, because the low-cost type 1 obtains her demand at a price equal to the average cost $\bar{c}_1$, she subsidizes the high-cost type 2, though to a lesser degree than in the linear-pricing candidate with price (3) discussed in Section 2.2.

Concerning the second part of Theorem 2, it should be noted that the natural two-point tariff obtained by restricting the market tariff (15) to the quantities Q_1^* and Q_2^* does not generally resist entry, as an entrant may cream-skim type 1 and make a profit. To deter entry, we have to ensure that any such offer would also attract type 2. This is exactly what the convex tariff (15) achieves, by enabling type 2 to purchase any fraction of the first layer at price $\bar{c}_1$ and any additional quantity at price c_2. This ensures that any entrant's contract that would attract type 1, such as D in Figure 2, would also attract type 2, because type 2 can complement this contract by latent contracts made available by the market tariff and thereby reach aggregate trades that she strictly prefers to (Q_2^*, T_2^*). The need for latent contracts to block attempts at cream-skimming contrasts with the exclusive-competition case, where the revelation principle ensures that there is no need to distinguish a market tariff from the allocation it implements.

In summary, in the two-type case, entry-proofness singles out a generically unique budget-balanced allocation; moreover, the existence problem emphasized by Rothschild and Stiglitz (1976) under exclusive competition no longer arises, whatever the distribution of types.

The generality of this conclusion is striking: single-crossing is only used to ensure that the inequality $Q_2 \geq Q_1$ holds, Assumption 1 is not needed, and preferences and candidate tariffs can be arbitrary as long as the maximization problems in (6)–(7) admit a solution.

3.3 Entry-Proofness in Active Markets: The Convex-Tariff Case

We now extend these results to the case of an arbitrary number of types. We will see that this raises a subtle new difficulty; to deal with it, the key restriction we impose in this section is that the market tariff be *convex*. A case in point is when each seller k posts a convex tariff t^k such that $t^k(0) = 0$. An intuitive rationale is that this allows sellers to hedge against the risk of attracting high-cost types buying large quantities; for instance, in the market-microstructure literature, convex tariffs are often used to model collections of limit orders placed by strategic market makers and executed in order of price priority by an informed insider.[9] Then the market tariff $T(Q) \equiv \min\{\sum_k t^k(q^k) : \sum_k q^k = Q\}$, which incorporates the possibility of trading with several sellers on the market, is indeed a convex function of the aggregate quantity Q.[10]

For a seller contemplating entering on a market where existing trading opportunities are summarized by the market tariff T, everything is as if the market were inactive and every type i's preferences were represented by the indirect utility function

$$u_i^T(q,t) \equiv \max\{u_i(Q+q, T(Q)+t) : Q \geq 0\}. \tag{16}$$

Convexity of the market tariff ensures that, if the primitive utility functions $(u_i)_{i=1}^I$ satisfy the strict single-crossing property, then the indirect utility functions $(u_i^T)_{i=1}^I$ satisfy the weak single-crossing property. This allows AMS (2021, Theorem 2) to rely on Theorem 1, which deals with inactive markets, to tackle the case of an active market.

Theorem 3 *Under Assumption 1, an allocation* $(Q_i^*, T(Q_i^*))_{i=1}^I$ *is budget-feasible and is implemented by an entry-proof convex market tariff* T^* *with domain* $[0, Q_I^*]$ *if and only if they jointly satisfy the following recursive system:*

(i) $(Q_0^*, T^*(Q_0^*)) \equiv (0,0)$;

(ii) for each i, $Q_i^* - Q_{i-1}^* \in \arg\max\{u_i(Q_{i-1}^* + q, T^*(Q_{i-1}^*) + \bar{c}_i q) : q \geq 0\}$;

(iii) for each i, *if* $Q_{i-1}^* < Q_i^*$, *then* T^* *is affine with slope* $\bar{c}_i$ *over the interval* $[Q_{i-1}^*, Q_i^*]$.

[9] See, for instance, Biais, Martimort, and Rochet (2000, 2013), Back and Baruch (2013), AMS (2019a), and Baruch and Glosten (2019).

[10] This is because T is the *infimal convolution* of the convex tariffs t^k (Rockafellar (1970, Theorem 5.4)).

In particular, any such allocation is budget-balanced.

This result generalizes Theorem 2 to more than two types. While item (i) is merely a convention, (ii)–(iii) are substantial, and indicate how to recursively build a complete family of quantities, as well as the corresponding market tariff; by construction, this tariff is convex, because the upper-tail conditional expectation of unit costs is nondecreasing in the buyer's type. The proof parallels the argument provided in Section 3.2 for the two-type case: at each step, the entrant must be deterred from supplying a well-chosen quantity at a price slightly above $\bar{c}_i$. By single-crossing, if such an offer attracts type i, then it must also attract all types $j \geq i$, so that the offer is profitable as soon as type i is attracted. Therefore, entry-proofness implies the following inequalities:

$$\text{For each } i, \; u_i(Q_i^*, T^*(Q_i^*)) \geq \max\{u_i(Q_{i-1}^* + q, T(Q_{i-1}^*) + \bar{c}_i q) : q \geq 0\}. \tag{17}$$

It follows that no layer can be profitable,

$$\text{For each } i, \; T^*(Q_i^*) - T^*(Q_{i-1}^*) \leq \bar{c}_i(Q_i^* - Q_{i-1}^*). \tag{18}$$

Summing these inequalities as in Section 3.2, we obtain that the allocation $(Q_i^*, T(Q_i^*))_{i=1}^I$ is budget-balanced, so that the inequalities (18) are in fact equalities. Notice that these equalities can be interpreted as a marginal version of Akerlof (1970) pricing: each layer is priced at the expected cost of serving the types who trade it. As a result, the constraints (17) must all be binding, and the result follows.

Theorem 3 generalizes a similar but weaker entry-proofness result due to Glosten (1994, Proposition 7). His analysis of limit-order markets requires that the buyer's preferences be quasilinear, and that the entrant's tariff satisfy a property he dubs *single-crossing* and that captures a convexity requirement. By allowing for general preferences, Theorem 3 makes the result relevant for insurance markets, in which wealth effects may be significant.

Existence of an entry-proof convex market tariff obtains because each maximization problem in (ii) admits a solution under the Inada condition (1).[11] Hence budget-feasibility and entry-proofness are not conflicting requirements under nonexclusivity, in contrast with the pervasive nonexistence problems arising under exclusivity (Rothschild and Stiglitz (1976)). The difference is that, when competition is exclusive, the buyer's indirect utility functions no longer satisfy single-crossing: by offering a cream-skimming contract, the entrant can attract a type i without attracting types $j > i$, which allows him to target type i without worrying about adverse selection. The nonexistence of an entry-proof tariff is then arguably

[11] Or, when the buyer's preferences are linear, because of the imposition of a capacity constraint.

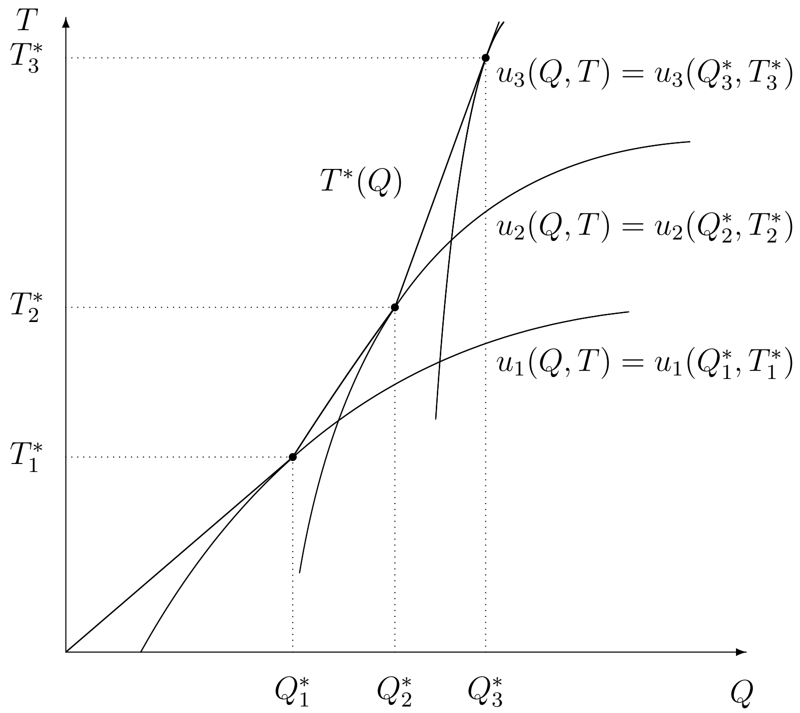

Figure 3: The JHG allocation and the JHG tariff for $I = 3$.

not due to private information or entry-proofness per se, but rather to this violation of single-crossing—or, to put it more provocatively, to the fact that the exclusive model does not capture the full extent of adverse selection.

Uniqueness of an entry-proof convex market tariff also follows if the solution to each maximization problem in (ii) is unique. This is the case if the buyer's preferences are strictly convex. If they are only weakly convex, multiple solutions may appear if the marginal rate of substitution of some type i is equal to $\overline{c}_i$ over a whole interval of quantities, but this is clearly a nongeneric phenomenon.

Theorem 3 thus characterizes an essentially unique allocation. Following AMS (2014, 2019a, 2021), we label this allocation, which was originally introduced in different contexts by Jaynes (1978), Hellwig (1988), and Glosten (1994), the *JHG allocation*.[12] Similarly, the *JHG tariff* T^* consists of a sequence of layers with unit prices $\overline{c}_i$, and features an upward kink at any quantity $Q_i^* \in (0, Q_I^*)$ such that $Q_{i+1}^* > Q_i^*$ and $\overline{c}_{i+1} > \overline{c}_i$. This sequence of layers can be interpreted as a family of limit orders with maximum quantities $Q_i^* - Q_{i-1}^*$ and unit prices $\overline{c}_i$. The JHG allocation and the JHG tariff are illustrated in Figure 3 in the case of three types with strictly convex preferences.

As an application, consider linear utility functions $u_i(Q,T) \equiv v_i Q - T$, subject to a capacity constraint $Q \in [0,1]$. Such linear preferences generalize those in Akerlof (1970) to the case of a divisible good; strict single-crossing requires that v_i be strictly increasing in

[12]The contributions of Jaynes (1978) and Hellwig (1988) are discussed in Section 5.3.

i. Each problem in (ii) admits a unique solution if $v_i \neq \overline{c}_i$ for all i, which we will assume for simplicity. To determine the JHG allocation and the JHG allocation, we apply (ii) in Theorem 3 recursively. By convention, $Q_0^* = 0$; then, beginning with type 1, Q_i^* remains zero as long as $v_i < \overline{c}_i$. If this inequality holds for all types, then the market is inactive; in that case, according to (iii), the essentially unique entry-proof convex market tariff is only defined at zero, with $T(0) = 0$. Otherwise, let i^* be the lowest type such that $v_i > \overline{c}_i$. Applying (ii) at i^* implies that type i^* trades up to capacity at unit price $\overline{c}_{i^*}$; moreover, types $i > i^*$ must also trade $Q_i^* = 1$, as the capacity constraint is binding in (ii). Finally, according to (iii), the unique entry-proof convex market tariff is linear, with $T(Q) = \overline{c}_{i^*}Q$ for all $Q \in [0, 1]$. The upshot from this discussion is that, when the buyer's preferences are linear, the JHG allocation generically features a single layer, and corresponds to the competitive-equilibrium allocation in Akerlof (1970) that maximizes the gains from trade.

The property that the indirect utility functions $(u_i^T)_{i=1}^I$ be ordered according to the weak single-crossing condition plays a key role in the above analysis. This property itself results from the two assumptions that the primitive utility functions $(u_i)_{i=1}^I$ be ordered according to the strict single-crossing condition, and that the market tariff be convex. Because this second assumption effectively constrains market outcomes, it is natural to ask whether it can be dispensed with. The answer is positive in the following three settings. In the two-type case, the proof of Theorem 2 follows from a direct argument that does not require that the market tariff be convex. When the buyer has linear preferences, as above, AMS (2011, p. 1888) also offers a direct proof. Finally, AMS (2021, Online Appendix F) show that the JHG allocation turns out to be the only budget-feasible allocation implemented by an entry-proof market tariff that is first convex and then concave. The general case raises however a difficult issue: in the absence of single-crossing, we do not know for sure whether a contract that attracts type i also attracts all types $j \geq i$, or only a subset of those with a more or less favorable expected cost; as a result, the entry-proofness constraint (17) need not hold. While entry-proofness per se selects a convex tariff in a large class of admissible tariffs allowing for quantity discounts, the general problem thus remains open.

3.4 Discussion

A noticeable feature of the JHG allocation is the relationship between demand and supply on each layer. On the first layer, the price is the expected cost of serving all types, and the quantity supplied is exactly the demand of the first type at this price. Indeed, supplying less would inefficiently ration demand, while supplying more would entail losses on the excess quantity. On the second layer, the first type is no longer active, and the same reasoning applies: the price is the expected cost of serving all types except the first, and the quantity supplied is exactly the residual demand of the second type at this price—and so on. Overall, the quantity supplied on each layer matches the residual demand of the marginal type, at a price equal to expected cost. On each layer but the last one, relatively low-cost types thus subsidize relatively high-cost ones, in contrast with the absence of cross-subsidies that characterizes candidate entry-proof allocations under exclusive competition.

It should also be noted that the JHG allocation typically allows for marginal rates of substitution to differ across types. For instance, in the two-type case under adverse selection, we have $\tau_1(Q_1^*, T_1^*) = \bar{c}_1 < c_2 = \tau_2(Q_2^*, T_2^*)$ when this allocation is interior and separating. This contrasts with private-value models where side trades take place on Walrasian markets, which calls for an equalization of marginal rates of substitution (Hammond (1979, 1987)). Yet, this difference does not create any opportunities for side trading, because the goods under consideration are not the same: for instance, in a Rothschild and Stiglitz (1976) insurance economy, insurance for a low-risk consumer is not the same good as insurance for a high-risk consumer. Indeed, supposing that consumers have access to the same constant-return-to-scale technology as firms, the opportunity cost for type 1 of selling additional coverage to type 2 is c_2, and at this price type 2 is not willing to buy. Similarly, the opportunity cost of selling coverage to type 1 is only c_1, but at this price all types would be attracted; hence the relevant unit cost is $\bar{c}_1$, and at this price type 1 is not willing to buy. The JHG allocation thus exhausts the incentive-compatible gains from trade. Together with entry-proofness, these features support the idea that the JHG allocation is a natural candidate for a competitive allocation.

Alternatively, AMS (2020) propose to reconsider this economy from the viewpoint of a social planner endowed with the same linear technology as the buyers and acting under asymmetric information. As in the classical setting of Harris and Townsend (1981), the planner is able to control all communication among the buyers, and in fact he optimally chooses to prohibit all forms of communication apart from a report each agent privately sends to him. Then the only constraints he faces are the incentive-compatibility constraints

$$\text{For all } i \text{ and } j,\ u_i(Q_i, T_i) \geq u_i(Q_j, T_j) \tag{19}$$

and the budget constraint (5). This leads to the classical definition of second-best allocations as Pareto-optima in the set of budget-feasible and incentive-compatible allocations. In general, such allocations form a non-degenerate continuum, according to the weight put on each type; moreover, under single-crossing, either the downward or the upward local incentive-compatibility constraints must be binding, apart from special cases.[13]

Let us now assume that the planner cannot monitor side trades between different buyers nor prevent the entry of a seller with the same technology. For simplicity, consider the two-type case, and suppose that preferences are strictly convex and satisfy strict single-crossing. Two consequences then follow for the set of allocations that the planner can implement. First, according to Theorem 2, this set collapses to a single allocation, namely, the JHG allocation. It is thus impossible for the planner to redistribute between types: both quantities and transfers are uniquely defined. Second, in general, this allocation is not second-best efficient, in the sense given above. Indeed, the incentive-compatibility constraints (19) are superseded by the entry-proofness constraints (17), which turn out to be necessary and sufficient to characterize the JHG allocation. When $Q_2^* > Q_1^*$, (17) implies that the local incentive-compatibility constraints do not bind, so that the JHG allocation is not second-best.

Overall, the uniqueness of the budget-balanced allocation robust to side trading contrasts with the multiplicity of second-best allocations, which form a nondegenerate frontier. The planner is thus severely constrained by his inability to monitor trades. As discussed in AMS (2020), this result has consequences for actual policies, because the possibility of side trading may undo their effects. For instance, the only possibility for public health insurance is to propose a single basic coverage, sold at average cost, and chosen so as to maximize the utility of low-risk consumers at that price. Private insurers can then compete to provide complementary coverage at price c_2. Another example is provided by bailout policies on financial markets. Under exclusivity, they aim at attracting only the least profitable borrowers, either through direct lending (Philippon and Skreta (2012)), or through the repurchasing of low-quality assets (Tirole (2012)). By contrast, when the borrower can complement a public program with private funds, the only possibility is to provide the same loan Q_1^* to all projects at expected cost, while the riskiest borrowers in addition turn to a competitive market for additional funding at price c_2.

4 Competitive Screening

Entry-proofness provides a parsimonious and tractable way of modeling perfect competition, which is relatively insensitive to the details of market interactions; for that reason, the

[13]See, for instance, Prescott and Townsend (1984) and Crocker and Snow (1985) for characterizations of second-best allocations in Rothschild and Stiglitz (1976) insurance economies.

JHG allocation characterized in Section 3 is arguably a natural and robust candidate for a competitive-equilibrium allocation of a nonexclusive market subject to adverse selection. Yet, by design, this approach does not shed light on how this allocation may be decentralized; a valuable complement to this approach would thus be to implement the JHG allocation as the unique equilibrium outcome of an extensive-form game in which strategic sellers compete to serve privately informed buyers.

To start with, and by way of comparison, we should observe that decentralization is easy to achieve in the standard case of exclusive competition. Indeed, in this context, the unique entry-proof allocation characterized in the insurance setting of Rothschild and Stiglitz (1976) or in the more general setting of Riley (1979) can be easily supported—as long as it exists—in a pure-strategy equilibrium of a competitive-screening game in which sellers first simultaneously post menus of contracts, from which the buyer then choose a single contract according to her type; specifically, there exists an equilibrium of this game in which two sellers offer a menu consisting of the trades comprised in this allocation. Therefore, the existence problem under exclusivity is not tied to decentralization per se, but to the fact that an entry-proof allocation may robustly fail to exist.[14]

By contrast, under nonexclusivity, we know that an entry-proof tariff exists, but its decentralization is much more delicate, because the buyer is now free to combine contracts issued by different sellers. As we shall now see, this generates novel strategic effects, which make it more difficult—indeed, in general, impossible—to implement the JHG allocation via competitive-screening games.

4.1 The Competitive-Screening Game

To clarify this issue, let us consider a general setting in which a finite number K of sellers simultaneously contract with a single buyer. We throughout assume that types are ordered according to the strict single-crossing condition and, when types are continuously distributed, that the mapping $(i, q, t) \mapsto u_i(q, t)$ is continuous. As discussed in Section 2.1, trade is nonanonymous and contracting is bilateral, in the sense that trade between a seller and a buyer can only be made contingent on the information reported by the buyer to this seller. In these situations, the menu theorems of Peters (2001), Martimort and Stole (2002), and Page and Monteiro (2003) allow us to restrict, with no loss of generality, to competition

[14] While the existence issue has been addressed by considering mixed strategy-equilibria (Rosenthal and Weiss (1984), Dasgupta and Maskin (1986), Farinha Luz (2017)) or by considering alternative extensive forms (Miyazaki (1977), Wilson (1977), Spence (1978), Riley (1979), Engers and Fernandez (1987), Netzer and Scheuer (2014), Mimra and Wambach (2019)), the corresponding equilibrium allocations typically do not coincide with the Rothschild and Stiglitz (1976) allocation; see Mimra and Wambach (2014) for a survey.

in menus or nonlinear tariffs. The corresponding extensive-form game, which we denote by G^{CS}, unfolds in two stages:

1. Each seller k offers a compact menu of contracts $C^k \subset \mathbb{R}_+ \times \mathbb{R}$ that contains at least the null trade $(0, 0)$.

2. After privately learning her type, the buyer selects a contract from each of the menus C^k offered by the sellers.

A pure strategy for type i is a function that maps every menu profile $(C^1, \ldots, C^K)$ into a contract profile $((q^1, t^1), \ldots, (q^K, t^K)) \in C^1 \times \ldots \times C^K$. The compactness of the sellers' menus ensures that every type i's utility-maximization problem

$$\max \left\{ u_i \left(\sum_k q^k, \sum_k t^k \right) : (q^k, t^k) \in C^k \text{ for each } k \right\}$$

always has a solution. The solution concept for G^{CS} is pure-strategy subgame-perfect Nash equilibrium. For future reference, we let

$$T(Q) \equiv \min \left\{ \sum_k t^k : (q^k, t^k) \in C^k \text{ for each } k \text{ and } \sum_k q^k = Q \right\}$$

be the market tariff associated to the equilibrium menus C^k, and we let, for each i,

$$U_i \equiv \max\{u_i(Q, T(Q)) : Q \geq 0\}$$

be the equilibrium utility of type i.

It should be noted that the set of strategies for the sellers in G^{CS} is the same as in a standard competitive-screening game under exclusivity. Yet the assumption that the buyer can simultaneously trade with several sellers has two implications for the set of potentially profitable contracts any seller may offer. On the one hand, it tends to expand this set, as this seller may choose to complement his competitors' offers by proposing additional trades to the buyer. On the other hand, it also gives his competitors more instruments to block his deviations, compared to when competition is exclusive; indeed, contracts that are not traded on the equilibrium path may become relevant in case a seller deviates, and in fact equilibria often require the presence of such latent contracts, as we shall now see.

4.2 Linear Preferences

Let us first assume that the buyer has linear preferences, subject to a capacity constraint. In this scenario, which extends Akerlof (1970) to the case of a divisible good, the JHG allocation features a single layer, and corresponds to the competitive-equilibrium allocation that maximizes the gains from trade. The following result, due to AMS (2011), shows that this allocation is uniquely supported in any equilibrium of G^{CS}:

Theorem 4 *Let $u_i(Q, T) \equiv v_i Q - T$ for $Q \in [0, 1]$ and suppose that $v_i \geq c_i$ for all i. Then, generically, any equilibrium of G^{CS} implements the JHG allocation, and there exists a linear-pricing equilibrium with price $\bar{c}_{i^*}$, where i^* is the first type i such that $v_i > \bar{c}_i$.*

In equilibrium, all the buyer types with valuations $v_i > \bar{c}_{i^*}$ trade up to capacity, while all the buyer types with valuations $v_i < \bar{c}_{i^*}$ do not trade at all. Sellers thus earn zero expected profits, and none of them is indispensable to serve any buyer type. Finally, all trades take place at the same price in equilibrium, despite the fact that sellers can propose arbitrary nonlinear tariffs. Thus Theorem 4 provides a game-theoretic foundation for Akerlof's (1970) predictions in a setting where the traded good is divisible and, besides nonexclusivity, few restrictions on feasible trades or instruments are imposed. In particular, low-valuation types such that $c_i < v_i < \bar{c}_{i^*}$ are excluded from trade in equilibrium, unlike what would happen under exclusive competition. The existence and uniqueness of the equilibrium allocation described in Theorem 4 paves the way for many applications—notably in finance, where the divisibility assumption is natural.[15]

The driving intuition for these results is that the unobservability of the buyer's aggregate purchases limits the sellers' ability to screen types and thereby the effectiveness of cream-skimming deviations. Suppose, for instance, that the equilibrium price is high, so that low-valuation, and hence on average low-cost types are not served. A cream-skimming deviation targeted at these types must involve trading a relatively small quantity q at a relatively low price. However, this contract becomes also attractive to high-valuation, and hence on average high-cost types if, along with it, they can trade the additional quantity $1 - q$ at the equilibrium price; this is exactly what the linear tariff allows for.

This reasoning illustrates the fact that deviations are blocked by latent contracts, that is, contracts that are not traded on the equilibrium path, but which the buyer may want to trade at the deviation stage.[16] In general, many such contracts are needed to support the equilibrium allocation. This is particularly striking when the distribution of types is discrete, because then only finitely many contracts are effectively traded, while infinitely many latent

[15] A case in point is the security-design model of DeMarzo and Duffie (1999), to which Theorem 4 directly applies; see Biais and Mariotti (2005) for an early result along these lines.

[16] The role of latent contracts has originally been stressed in moral-hazard environments; see, for instance, Hellwig (1983), Arnott and Stiglitz (1991), Bisin and Guaitoli (2004), and Attar and Chassagnon (2009).

contracts must be issued. In particular, no equilibrium can in this case be sustained through direct mechanisms, which provides a concrete example of a failure of the revelation principle in common-agency games (Peters (2001), Martimort and Stole (2002)).

4.3 Cournot-Convergence under Strictly Convex Preferences

Although Theorem 4 holds for general distributions of types,[17] it does not generally extend to the case of strictly convex preferences for the buyer. In a seminal article, Biais, Martimort, and Rochet (2000) consider a situation in which strategic market-makers (sellers) compete to serve a risk-averse insider (buyer) who has private but imperfect information about the value of an asset, and thus has both informational and hedging motives for trade. Assuming that the buyer has constant absolute risk-aversion α and faces residual Gaussian risk with variance σ^2, they show the following result:

Theorem 5 *Let* $u_i(Q,T) \equiv v_i Q - \frac{\alpha\sigma^2}{2} Q^2 - T$ *and let the buyer's type be continuously distributed. Then, under regularity conditions,* G^{CS} *admits a symmetric equilibrium in which sellers post the same strictly convex tariff and earn strictly positive expected profits. The equilibrium market tariff converges to the JHG tariff as the number* K *of sellers grows large.*

This equilibrium exhibits the Cournot-like feature that each seller is indispensable to serve any buyer type who trades a nonzero quantity in equilibrium. Specifically, the strict convexity and symmetry of the equilibrium tariffs implies that any such type has a unique best response that consists in evenly splitting her total purchases between the sellers. This contrasts with the equilibria that obtain in the linear case, in which no seller is indispensable and thus any buyer type who trades up to capacity has multiple best responses that involve trading with different sellers. Because sellers earn strictly positive expected profits in equilibrium, the aggregate equilibrium allocation does not coincide with the JHG allocation; yet, in analogy with classical Cournot-convergence theorems, it converges to the competitive JHG allocation as the number of sellers grows large.

The assumption in Theorem 5 that the buyer's type be continuously distributed is key to ensure that, despite being indispensable, a single seller cannot profitably raise his tariff. To illustrate this point, suppose that a seller deviates by replacing a portion of his strictly convex equilibrium tariff by the corresponding chord. This would increase his expected profit if the buyer's behavior remained the same. But such a change raises (lowers) the marginal price for relatively low-cost (high-cost) types who would choose trades in this portion of

[17] Specifically, it holds for any distribution for which $v_i \geq c_i$ at any atom.

the tariff. As a result, under adverse selection, trades change in an unfavorable way for the deviating market maker. This effect is reinforced by the fact that the buyer simultaneously trades with several sellers, as any increase in the quantity she purchases from a seller is compensated by a reduction in the quantity she purchases from his competitors.

4.4 Exclusion

We now argue that, in any game G^{CS}, exclusion is a robust feature of any equilibrium that shares two key properties of the Biais, Martimort, and Rochet (2000) equilibrium. To formulate these properties, we let, for each k,

$$T^{-k}(Q) \equiv \min \left\{ \sum_{l \neq k} t^l : (q^l, t^l) \in C^l \text{ for each } l \neq k \text{ and } \sum_{l \neq k} q^l = Q \right\}$$

be the submarket tariff associated to the equilibrium menus C^l, $l \neq k$, and we let, for all i and k,

$$z_i^{-k}(q,t) \equiv \max \left\{ u_i(q + Q^{-k}, t + T^{-k}(Q^{-k})) : Q^{-k} = \sum_{l \neq k} q^l \text{ for some } (q^l, t^l) \in C^l,\ l \neq k \right\}$$

be type i's indirect utility from trading (q,t) with seller k. The two properties we wish to emphasize can now be stated as follows.

P1 For each k, there exists i such that $U_i = z_i^{-k}(0,0)$.

An equilibrium satisfies P1 if, for each seller, there exists at least one type for whom trading with this seller is not indispensable for her to obtain her equilibrium utility; this reflects the relatively weak requirement that, in equilibrium, the buyer's individual-rationality constraint in her dealings with each seller must bind for at least one type.

P2 For all k and $Q > 0$, $T(Q) < T^{-k}(Q)$.

An equilibrium satisfies P2 if trading with each seller is indispensable for each type who purchases a nonzero aggregate quantity.

The symmetric equilibrium characterized by Biais, Martimort, and Rochet (2000) satisfies both P1 and P2 because all sellers offer the same strictly convex tariff. Indeed, this implies that each seller is indispensable to minimize the cost of purchasing any strictly positive aggregate quantity, whence P2. This also implies that the indirect utility functions z_i^{-k} satisfy the strict single-crossing condition for all k, so that any seller k for whom the individual-rationality constraint were not binding could raise his tariff without affecting the buyer's incentives, whence P1.

The following theorem, a formal proof of which is provided in the appendix, abstracts from the parametric assumptions of Biais, Martimort, and Rochet (2000) to show that exclusion must more generally take place in any equilibrium of any game G^{CS} that satisfies P1–P2:

Theorem 6 *Consider an equilibrium of a game G^{CS} that satisfies P1–P2. Then there exists some type i_1 such that, in equilibrium, every type $i \leq i_1$ trades $q_i^k = 0$ with every seller k and obtains utility $U_i = u_i(0, 0)$.*

In light of this result, it is worth noticing that Biais, Martimort, and Rochet (2000) assume that the continuous support of the buyer's type distribution includes an interior type i_0 such that $\tau_{i_0}(0, 0) = c_{i_0}$ and thus for which there are no gains from trade.[18] This in turns ensures that there exists an interval of types at the bottom of the type distribution who are excluded from trade in equilibrium, as requested by Theorem 6. However, this assumption is fairly restrictive: it does not hold, for instance, in standard Rothschild and Stiglitz (1976) insurance economies, because a risk-averse consumer is always willing to purchase full coverage at the fair price, equal to her riskiness. As we shall now see, the existence of equilibria of G^{CS} games then becomes problematic.

4.5 The Existence Conundrum

As a starting point, let us consider the two-type case with $c_2 > c_1$ and strictly convex preferences for the buyer, and let us examine a candidate equilibrium of G^{CS} in which both types 1 and 2 purchase strictly positive quantities. AMS (2014) show that the aggregate equilibrium allocation then has the same structure as the JHG allocation. First, sellers earn zero expected profits. Second, the quantity Q_1 purchased by type 1 is priced at the average cost $\bar{c}_1$, while the additional quantity $Q_2 - Q_1$ purchased by type 2 is priced at the marginal cost c_2. Third, the quantity Q_1 purchased by type 1 is equal to her demand $D_1(\bar{c}_1)$.

This sounds promising but, because type 1 is not excluded from trade, Theorem 6 implies that P1 or P2 has to give way. It is intuitive that the weak requirement P1 should be maintained, and thus that P2 should be dropped. Specifically, it can be shown that P1 holds for type 1 and that $T(Q_1) = T^{-k}(Q_1) = \bar{c}_1 Q_1$ for all k, so that no seller is indispensable to

[18] Specifically, the insider can trade on both sides of the market, so that types $i > i_0$ ($i < i_0$) are willing to buy (sell) at price c_i. As in Back and Baruch (2013) and Biais, Martimort, and Rochet (2013), we focus on the ask side of the market.

provide type 1 with her equilibrium aggregate trade. Because any seller k who trades with both types 1 and 2 on the equilibrium path makes a profit with type 1 and a loss with type 2, this opens the way to a *lemon-dropping* deviation that essentially consists for seller k in convincing type 2 to trade the layer $(Q_1, \bar{c}_1 Q_1)$ with his competitors and the complementary layer $(Q_2 - Q_1, c_2(Q_2 - Q_1))$ with him, thus neutralizing his losses with type 2.

Specifically, the deviation involves two contracts. When $Q_2 > Q_1 > 0$, the first one is essentially that traded by type 1 with seller k on the equilibrium path, while the second one makes the quantity $Q_2 - Q_1$ available at a unit price slightly lower than c_2. Thus type 2 can strictly increase her utility by trading the second contract on top of the layer $(Q_1, \bar{c}_1 Q_1)$ made available by the sellers other than k; besides, seller k can break ties to make sure that type 2 strictly prefers this contract to the first one. As a result, seller k can make his loss with type 2 arbitrarily small while securing, as $\bar{c}_1 > c_1$, a strictly positive profit with type 1; hence the deviation is profitable. A similar but slightly more involved argument shows that seller k has a profitable deviation also when $Q_2 = Q_1 > 0$.

As a consequence, trade can take place in equilibrium only if type 1 is excluded from trade: in short, the possibility of cross-subsidizing between contracts at the deviation stage makes it impossible to support cross-subsidies between types on the equilibrium path. Specifically, the following result holds (AMS (2014, Theorems 1–2)):

Theorem 7 *Suppose there are two buyer types with strictly convex preferences, and that* $c_2 > c_1$*. Then any equilibrium of* G^{CS} *implements the JHG allocation, but an equilibrium exists if and only if* $Q_1^* = 0$ *in that allocation, that is, if and only if* $\tau_1(0, 0) \leq \bar{c}_1$*. If an equilibrium exists, it can be sustained by each seller posting the JHG tariff, which consists of a single layer with unit price* c_2*.*

To allow for a finer comparison with the continuous-type model of Biais, Martimort, and Rochet (2000), let us now suppose as in Section 2.1 that there is an arbitrary but finite number I of types, each assumed to have strictly convex preferences. The same difficulty arises as for the characterization of entry-proof tariffs, however: when $I > 2$, the game G^{CS} with no restrictions on admissible menus is hardly tractable. As in Section 3.3, a convenient assumption is that sellers are restricted to post convex tariffs, which ensures that the indirect utility functions $(z_i^{-k})_{i=1}^I$ are quasiconcave and satisfy the weak single-crossing property for all k. The resulting convex-tariff game G^{CT} can be interpreted as a discriminatory auction in which sellers simultaneously bid quantities at each marginal price. In line with Back and Baruch (2013), this captures oligopolistic competition on a limit-order market where market-makers post collections of limit orders that are executed by an informed insider in order of price priority.

It should be noted that the set of deviations for the sellers is much smaller in G^{CT} than in G^{CS}; in particular, the deviation that led to Theorem 7 is no longer feasible. However, the following result, due to AMS (2019a, Theorems 2–3), shows that, in spite of this, the equilibrium-existence problem only becomes more acute when the number of types increases:

Theorem 8 *Suppose there are I buyer types with strictly convex quasilinear preferences, and that c_i is strictly increasing in i. Then any equilibrium of G^{CT} implements the JHG allocation, but an equilibrium exists if and only if $Q_i^* = 0$ for all $i < I$ in that allocation, that is, if $\tau_i(0,0) \leq \overline{c}_i$ for all $i < I$. If an equilibrium exists, it can be sustained by each seller posting the JHG tariff, which consists of a single layer with unit price c_I.*

The proof proceeds by showing that, in any candidate equilibrium, the market tariff T is piecewise linear and has a structure similar to that of the JHG tariff. That is, the quantity supplied on each layer but the last one matches the residual demand of the marginal type, at a price equal to the expected cost of serving the buyer types who trade along this layer. This implies that sellers earn zero expected profits, and also that the marginal type on any such layer exhausts the supply at the corresponding marginal price. As a result, each seller offering trades at this marginal price is indispensable for the marginal type and all higher types to reach their equilibrium utility. But one can hardly be indispensable and yet earn zero expected profit: hence any such seller could raise his tariff in a profitable way, a contradiction. This shows that the market tariff T must consists of a single layer and be such that no seller is indispensable to serve the buyer types who trade along it. However, each seller will then want to issue a limit order to hedge against the risk of large purchases emanating from the most costly types. This implies that all types except perhaps the last one must be excluded from trade.

The upshot from Theorems 5 and 7–8 is twofold. First, the structure of equilibria of discrete-type models, when they exist, is very different from that of the equilibria of continuous-type models that have been emphasized in the literature: namely, pricing is linear and only the last type can trade in equilibrium. Second, necessary and sufficient conditions for the existence of an equilibrium become increasingly stringent as the number of types increases: equilibria fail to exist when there are sufficiently many types with similar preferences, as when we approximate the continuous sets of types postulated by Biais, Martimort, and Rochet (2000) or Back and Baruch (2013). The pure-strategy-equilibrium correspondence thus fails to be lower hemicontinuous when we move from discrete-type models to continuous-type models. Overall, the predictions of competitive-screening models are very sensitive to fine modeling details, which makes them somewhat fragile.

4.6 Ways Out

A natural way to address the equilibrium-existence problem in discrete-type models is to weaken the equilibrium concept. This can be done in two ways.

First, we may consider mixed-strategy equilibria of G^{CS}, the existence of which follows from Carmona and Fajardo (2009). Preliminary investigations of the two-type case have led to a robust example of a mixed-strategy equilibrium that exists when the necessary and sufficient conditions for the existence of a pure-strategy equilibrium are not satisfied (Attar, Farinha Luz, Mariotti, and F. Salanié (2021)). The key point is that the strategic uncertainty faced by each seller regarding the tariffs offered by his competitors makes it now impossible for him to target specific types, unlike in the deviations used to derive Theorems 7–8. However, this equilibrium bears no obvious relationship with existing equilibrium candidates; the JHG allocation, in particular, does not emerge even when the number of sellers grows large. The systematic characterization of mixed-strategy equilibria nevertheless remains a fascinating—though hard—topic for future research.

Next, we may consider ε-equilibria of G^{CS}. AMS (2019a) show that, if every type i has quasilinear preferences $u_i(Q,T) \equiv v_i(Q) - T$, then, as the number K of sellers grows large, G^{CS} admits an ε-equilibrium, with ε of the order of $1/K^2$, that supports the JHG allocation. The intuition is that if $K-1$ sellers contribute to providing a fraction $1/K$ of the JHG tariff, the residual gains from trade for the remaining seller vanish when K grows large because the resulting market tariff is almost entry-proof. The reason why convergence takes place at rate $1/K^2$ is that these gains from trade are, for every type i, bounded above by

$$v_i(Q_i^*) - v_i\left(\frac{K-1}{K}\, Q_i^*\right) - \frac{1}{K}\, \bar{c}_i Q_i^*,$$

which is at most of the order of $1/K^2$ as $v_i'(Q_i^*) \leq \bar{c}_i$ at the JHG allocation. Thus we retrieve the convergence result of Theorem 5 for the competitive limit, albeit at the cost of relying on a notion of approximate equilibrium. Glosten (1994, Proposition 2) provides a similar result, assuming from the outset that there is an infinite number of sellers.

4.7 Regulation

An alternative route to decentralize the JHG allocation consists in explicitly introducing a market regulation. In this spirit, AMS (2019b) study a nonexclusive insurance market in which it is prohibited for sellers to cross-subsidize between contracts. The regulation thus bears on the total profit a seller earns on each contract, and is targeted at dumping practices; it can alternatively be interpreted as banning profits on basic-coverage contracts.[19] Specifically, let G^{CSR} be the regulated game that is obtained from G^{CS} by adding one final stage in which a seller's profit is confiscated whenever he makes a loss on any of the contracts he is trading. The following result then holds (AMS (2019b, Theorem 2)):

Theorem 9 *Suppose there are two buyer types with strictly convex preferences, and that* $c_2 > c_1$. *Then the JHG allocation is the unique candidate-equilibrium allocation of* G^{CSR}. *Moreover, under regularity conditions on the buyer's preferences,* G^{CSR} *has an equilibrium as long as there are sufficiently many sellers.*

There are two parts in this result. The necessity part states that the regulation has no anti-competitive implications. The intuition is that each seller aims at increasing his profit by complementing the aggregate coverage provided by its competitors, which gives rise to a form of Bertrand competition over each layer; as a result, the JHG allocation is the unique candidate-equilibrium outcome. The sufficiency part reflects that the regulation, by blocking the cross-subsidies between contracts that were at the root of Theorem 7, helps restore existence of an equilibrium even if both types trade in the JHG allocation.

A necessary feature of equilibrium is that sellers must issue latent contracts to discipline their competitors. As in AMS (2011), these contracts are not traded in equilibrium but are meant to block cream-skimming deviations. In the context of insurance, these contracts provide additional coverage that high-risk consumers are willing to combine with the coverage provided by any such deviation. This makes it impossible for a seller to profitably deviate by separating low-risk from high-risk consumers.

5 Alternative Extensive Forms

An important takeaway from the literature surveyed in Section 4 is that competitive-screening games generally fail to implement the JHG allocation. This failure can be traced back to a common source, namely, the paucity of instruments allowing to punish a deviating

[19]Several insurance markets are actually regulated along analogous lines. For instance, in health insurance, Germany and Switzerland rely on a central fund to redistribute costs among firms according to a risk-equalization scheme. These cost-sharing mechanisms, by pooling and redistributing costs among sellers of a standardized basic-coverage contract, prevent firms from earning abnormal profits on such coverage by dropping lemons on their competitors.

seller. Indeed, if sellers make their offers simultaneously, the only device available to block deviations consists in letting the buyer select latent contracts in the nondeviating sellers' menus. However, this device is effective only when the buyer has linear preferences, reflecting the very special property that, if latent contracts are issued at the equilibrium price, all the types who are willing to trade at this price have the same indirect utility function z_i^{-k}: they are willing to trade a contract issued by a deviating seller if and only if its unit price is less than the equilibrium price. This no longer holds when the buyer has strictly convex preferences and different types trade at different marginal prices.

The generic failure of latent contracts at sustaining equilibria in competitive-screening games—let alone at implementing the JHG allocation—suggests that we look for alternative extensive forms whereby the sellers sequentially receive information over the course of the game. Three kinds of extensive forms have been studied in the literature. The first one allows the buyer to signal her type by recontracting, which requires that all previously signed contracts be publicly observable. The second one lets the sellers bid through an ascending discriminatory auction, in which the offers made at previously quoted prices are publicly observable. The third one enables the sellers—and, possibly, the buyer as well—to voluntarily disclose information about the contracts selected by the buyer; which information eventually becomes available, and to whom, then depends on the agents' disclosure strategies.

5.1 Recontracting

Beaudry and Poitevin (1995) study a sequential game in which a risk-averse entrepreneur whose project can be of low or high riskiness can repeatedly solicit financing from successive cohorts of uninformed lenders before the realization of the project's return. An important feature of this game is that there is a potentially infinite number of recontracting rounds. Hence there is no last stage of the game in which the entrepreneur could commit to reject further offers: a lender can never be sure that she will not try to further diversify her risk by selling new claims on her project.[20] That is, nonexclusivity is distinctively linked to the absence of commitment and the sequential nature of contracting.

At each round of recontracting, the buyer can solicit further offers from a new cohort of lenders; if she does so, she has to provide a summary of all the contracts she has signed in previous rounds, though not of the contracts she has rejected. This observability assumption stands in contrast with the competitive-screening models surveyed in Section 4, in which, by design, no seller has information about existing contractual relationships. Another difference

[20]This feature is similar to Kahn and Mookherjee's (1998) moral-hazard model of nonexclusive contracting.

is that, at each round of recontracting, competition is exclusive, in the sense that the buyer can accept at most one offer from one lender. The assumption that lenders are short-lived and cannot observe previously rejected contracts is meant to limit the lenders' ability to sustain collusive outcomes.

In this setup, Beaudry and Poitevin (1995) show that, when the high-risk project has positive NPV and each project can be financed by riskless claims using the entrepreneur's collateralizable wealth, there exists a perfect Bayesian equilibrium that supports the JHG allocation: an entrepreneur with a low-risk project obtains the net claims corresponding to her preferred financial position among the contracts with nonnegative pooling profits, while an entrepreneur with a high-risk project manages to completely diversify her risk without pledging any of her wealth in the project.

This outcome can be supported without recontracting on the equilibrium path, with each seller posting in the first round a menu consisting of the trades comprised in the JHG allocation. If the entrepreneur solicits additional offers, then she is believed to have a high-risk project; for instance, she is offered $Q_2^* - Q_1^*$ at price c_2 if she initially accepted (Q_1^*, T_1^*), and she is offered $(0, 0)$ if she initially accepted (Q_2^*, T_2^*). No lender, therefore, can profit by deviating, for any offer that would be accepted by an entrepreneur with a low-risk project would also be accepted by an entrepreneur with a high-risk project in anticipation of future rounds of recontracting. Notice that it is essential for this reasoning that the buyer can only select a single contract at each round, that previously signed contracts be observable, and that there always be further opportunities of recontracting. It is fair to ask whether the first two assumptions are consistent with the intuitive notion of a nonexclusive market; in particular, a prediction of the model is that each entrepreneur trades with a single lender.

5.2 A Discriminatory Ascending Auction

An alternative approach consists in sticking more closely to competitive-screening games, while allowing punishments to be carried out by the sellers themselves. This requires, of course, that deviations be observable by the nondeviating sellers, in the spirit of the reactive-equilibrium literature cited in Footnote 14. In this respect, the recursive structure of the JHG allocation suggests that it be implemented sequentially, layer by layer, from the bottom up. To validate this intuition, AMS (2021) propose to model the strategic interactions between sellers as a discriminatory ascending auction.

In their model, the auctioneer quotes price sequentially, in increasing order, and according to a discrete price grid with a minimum tick size. Each time he quotes a new price, each seller publicly announces the maximum quantity he stands ready to trade with the buyer at this price; in other terms, he offers a limit order at the current price. Once this auctioning phase is completed, the buyer selects which quantities to purchase from which sellers at each price, according to her type. We denote by G^{DA} the corresponding extensive-form game. The solution concept for G^{DA} is pure-strategy subgame-perfect Nash equilibrium.

It should be noted that G^{DA} can be interpreted as a sequential version of the convex-tariff game G^{CT}. Indeed, as it is optimal for the buyer to take up the best price offers first, she in the end faces a collection of convex tariffs that aggregate into a convex market tariff T. From her perspective, the fact that T was built up sequentially is irrelevant.

For the sellers, by contrast, the fact that bids are made sequentially and publicly during the auctioning phase is crucial, as it allows them to react, at any price, to a deviation at a lower price; indeed, the key advantage of a sequential auction lies in its transparency, a point emphasized in other contexts by Milgrom (2000) and Ausubel (2004). Importantly, such reactions—which can take place almost immediately when the tick size is small—can only take the form of quantity increases or decreases at future prices, while the quantities supplied at lower prices cannot be withdrawn or augmented. This commitment assumption makes G^{DA} quite different from the Walrasian *tâtonnement* process.

From our implementation perspective, two questions immediately arise. First, does G^{DA} admit an equilibrium? Second, how do equilibrium allocations relate to the JHG allocation? The second question is especially pressing, because the dynamic nature of G^{DA} may perhaps allow to sustain equilibria with collusive outcomes.

The first result established by AMS (2021) is that G^{DA} admits a very simple Markov perfect equilibrium. The relevant states variables are the current price p and the aggregate quantity Q^- supplied at prices lower than p. Assuming for simplicity that the buyer has quasilinear preferences, with demand function D_i, the residual demand of type i in state (p, Q^-) is $\max\{D_i(p) - Q^-, 0\}$; observe that maximizing aggregate expected profits in any state (p, Q^-) asks for serving the residual demand of the type i such that $\overline{c}_i < p \leq \overline{c}_{i+1}$, which we shall call the *profitable residual demand* in state (p, Q^-). The following result then holds (AMS (2021, Theorem 3)):

Theorem 10 *Suppose there are I buyer types with strictly convex quasilinear preferences, and that $\overline{c}_i$ is strictly increasing in i. Then there exists a Markov perfect equilibrium of G^{DA} in which, in any state (p, Q^-),*

(i) if $p \leq \overline{c}_1$, each seller supplies a zero quantity;

(ii) if $\overline{c}_1 < p \leq \overline{c}_I$*, each seller supplies a share* $1/K$ *of the profitable residual demand*;

(iii) if $p > \overline{c}_I$*, each seller supplies an infinite quantity.*

Moreover, the resulting aggregate equilibrium allocation converges to the JHG allocation as the tick size goes to zero.

The mechanics of the equilibrium are very simple. First, at any price $p \leq \overline{c}_I$, no unilateral increase in supply is profitable if the profitable types at price p, that is, all the types j such that $p > \overline{c}_j$, rationally choose to ignore this deviation and carry on trading the same quantity with each seller; indeed, the deviation can then only lead to losses with unprofitable types at price p and reduce their residual demand at higher prices. Second, at any price $p \leq \overline{c}_I$, no unilateral decrease in supply is profitable, because the corresponding increase in the profitable residual demand at the next price will be shared with the other sellers in the continuation equilibrium. As a result, although each seller is indispensable to serve a strictly positive profitable residual demand, no seller has an incentive to wait for a higher price to be quoted. This stands in stark contrast with the convex-tariff game G^{CT}, in which a seller indispensable at price p can always secretly deviate by bidding a lower quantity at price p and a higher quantity at a slightly higher price.

Given a tick size Δ, the following equilibrium outcome obtains. As soon as the price reaches $\overline{c}_1 + \Delta$, the sellers serve the demand $D_1(\overline{c}_1 + \Delta)$ of type 1; this quantity will also be purchased by types $i > 1$. Then, as soon as the price reaches $\overline{c}_2 + \Delta$, the sellers serve the residual demand $\max\{D_2(\overline{c}_2 + \Delta) - D_1(\overline{c}_1 + \Delta), 0\}$ of type 2; this quantity will also be purchased by types $i > 2$—and so on, until the price reaches $\overline{c}_I + \Delta$, at which point the sellers flood the market. By construction, the resulting aggregate equilibrium allocation converges to the JHG allocation as Δ goes to zero.

The second result established by AMS (2021) is that every sequence of equilibria of G^{DA} satisfies this convergence property, modulo an intuitive refinement that can be described as follows. In any play of the game, every type i accepts all offers up to some price p_i. However, the sellers' aggregate supply at price p_i may well exceed type i's residual demand at this price; she can then break ties in many different ways, and her choice typically matters to the sellers. An equilibrium of G^{DA} is *robust to irrelevant offers* if every type i's trades at price p_i do not depend on offers made at prices $p > p_i$. Intuitively, the buyer never punishes a seller for deviating at a price at which she is not willing to trade.[21] The following result then holds (AMS (2021, Theorem 4)):

[21] The tie-breaking rules used in the proof of Theorem 10 are consistent with this refinement.

Theorem 11 *In any sequence of equilibria robust to irrelevant offers of G^{DA} associated to a sequence of tick sizes going to zero, the aggregate equilibrium allocations converge to the JHG allocation, and the equilibrium market tariffs converge to the JHG tariff.*

Leaving technical details aside, Theorem 11 results from a simple Bertrand undercutting argument. To see this, suppose, by way of contradiction, that, given the limit market tariff, strictly positive expected profits can be earned at some price p. Because the highest price at which trade takes place can be shown to be bounded along any sequence of equilibria when the tick size Δ goes to zero, we can focus on the highest such p. Continuation profits at higher prices must be zero: indeed, the robustness refinement ensures that, if they were strictly negative, then, for Δ small enough, some seller could profitably withdraw all his offers at such prices without affecting his expected profits up to price p. Now, the convergence of aggregate supply functions as Δ goes to zero implies that, for Δ small enough, aggregate supply in a left-neighborhood of p becomes negligible. Hence each seller can, almost without losing priority, undercut his competitors at a price arbitrarily close to p, and supply nothing afterwards; by doing so, he can appropriate almost all expected profits at price p, and the robustness refinement again ensures that his expected profits at lower prices remain the same. But then this deviation would be profitable for at least one seller, a contradiction. Overall, this argument shows that, given the limit tariff, strictly positive expected profits cannot be earned at any price p; given budget-balance, this exhaustion of gains from trade characterizes the JHG tariff, from which Theorem 11 follows.

Taken together, Theorems 10–11 provide an implementation of the JHG allocation as the essentially unique equilibrium outcome of competition when each seller can quickly react to his competitors' offers; an attractive feature of this implementation is that, in the spirit of Bertrand competition, it only requires that there be two competing sellers. From a market-design perspective, these positive results invite us to reconsider the role of continuous bidding for financial and insurance markets, and offer a useful complement to studies that advocate a transformation of continuous markets into batch auctions, so as to avoid possible inefficiencies linked to high-frequency trading (Budish, Cramton, and Shin (2015)).

5.3 Information Disclosure

A common feature of the recontracting game of Beaudry of Poitevin (1995) and of the ascending discriminatory auction of AMS (2021) is that the release of information to sellers about previously signed contracts or previously made offers is exogenous. By contrast, following the seminal contribution of Jaynes (1978), several authors have studied what happens when agents can voluntarily disclose information about the contracts they are engaged in, so that the information available to each seller is endogenously determined in equilibrium. Although the buyer can still in principle subscribe to multiple contracts issued by different sellers, each seller can then enforce exclusivity clauses contingent on the information disclosed by his competitors.[22]

Jaynes (1978) considers the following timing. First, each seller offers a menu of contracts, possibly specifying exclusivity clauses; besides, he commits to disclose to a subset of his competitors the contract selected by the buyer from his own menu. Next, the buyer selects a contract from each seller's menu, information is disclosed, and exclusivity clauses are enforced, which determines the contracts that are eventually executed. The JHG allocation turns out to be the only candidate-equilibrium allocation. In the two-type case, Jaynes' (1978) proposed equilibrium can be described as follows. First, two sellers offer a limit order at price $\overline{c}_1$ with maximum quantity Q_1^*; these sellers share their information and enforce exclusivity clauses, which ensures that they do not make losses by overselling to type 2. Second, two sellers offer a limit order at price c_2 with maximum quantity $Q_2^* - Q_1^*$; these sellers do not share their information.

As pointed out by Hellwig (1988), however, this candidate equilibrium is not robust to a cream-skimming deviation, whereby one of the sellers supposed to offer trades at price $\overline{c}_1$ secretly deviates by offering a contract at a price slightly less than $\overline{c}_1$ that, per se, attracts type 1, but would attract type 2 only in combination with contracts issued by the nondeviating sellers at price $\overline{c}_1$. Thanks to the information disclosed by the nondeviating sellers, the deviating seller is still able to enforce exclusivity on this contract; as a result, he is assured to only attract type 1, and the deviation is profitable. Intuitively, the idea is that sellers by themselves have no basis for treating deviating and nondeviating firms asymmetrically in their disclosure decisions.

In that respect, it should be noted that the above deviation is effective only if the nondeviating sellers are not aware of it. Otherwise, they could punish the deviating seller by concealing from him the contracts selected by the buyer from their menus. In that case, a seller attempting to cream-skim type 1 would no longer be able to enforce exclusivity; as a result, his contract would also become attractive for type 2, along with contracts issued by nondeviating sellers at price $\overline{c}_1$ and, possibly, c_2. Hence cream-skimming is impossible if the sellers' offers are public. This intuition is formalized by Hellwig (1988), who studies a multi-stage extensive-form game in which each seller can make his disclosure decisions

[22] In exclusive insurance markets, information sharing enables the construction of joint databases collecting information on each loss so as to ensure that the same loss is not indemnified twice.

contingent on his competitors' contract offers.

Jaynes (2011) and Stiglitz, Yun, and Kosenko (2020) have more recently argued that information sharing may allow to support the JHG allocation in equilibrium even if sellers cannot change their disclosure decisions in reaction to the offers of their competitors. The idea is that they can instead rely on information revealed by the buyer. In equilibrium, the firms' disclosure strategies induce the buyer to truthfully reveal her information to them, which in turn enables them to treat deviating and nondeviating firms asymmetrically.

6 Concluding Remarks: Empirical Perspectives

To conclude, we briefly examine the implications of the theoretical results surveyed in this article for empirical work. The discussion will focus on insurance markets, prominent examples of which—life-insurance, annuity, long-term-care, and, to some extent, health-insurance markets—are nonexclusive.

6.1 The Positive-Correlation Property

A standard way to test for the presence of adverse selection on insurance markets is to exploit the positive-correlation property, which states that, under adverse selection, the aggregate coverage purchased by a consumer and her riskiness should be positively correlated conditionally on observables (Chiappori and B. Salanié (2000)). This property is typically satisfied when a consumer's preferences over aggregate coverage-premia pairs (Q, T) only depend on her riskiness; indeed, it is then equivalent to the single-crossing condition, which precisely expresses the fact that riskier consumers are more eager to purchase more coverage. In our notation, this is the case if the riskiness c_i and the willingness-to-pay $\tau_i(Q, T)$ are both increasing in the consumer's type i, so that the demand Q_i for coverage is increasing in c_i.[23] Chiappori and B. Salanié (2000, 2003) have developed several econometric methods to test this prediction.

Under single-crossing, the positive-correlation property is a characteristic of consumer demand; as such, it is independent of whether competition on the market is exclusive or nonexclusive.[24] Yet the empiricist should care about the difference. Indeed, under

[23]This is the case, for instance, if consumer preferences have an expected-utility representation and face a binary loss (Rothschild and Stiglitz (1976)), as well as for many other specifications of consumer preferences and more than two loss levels (AMS (2021, Online Appendix C)).

[24]In a more general setting, Chiappori, Jullien, B. Salanié, and F. Salanié (2006) show that the positive-correlation property can alternatively be derived from a simple inequality on equilibrium profits, even if the single-crossing property is not satisfied.

exclusivity, it does not matter whether his data originate from firms or from consumers, as each consumer's aggregate coverage is provided by a single contract issued by a single firm. This explains the usual reliance on within-firm data, which are easier to obtain. However, under nonexclusivity, such data can be misleading, as low-coverage contracts may disproportionately attract high-risk consumers in combination with other contracts.[25] Thus the positive-correlation property may still hold at the consumer level, taking into account all sources of coverage; but the contracts sold by a firm may feature a negative correlation between the riskiness of its customers and the coverage it sells to them.

The validity of the positive-correlation property has been at the centre stage of empirical studies of nonexclusive insurance markets (Cawley and Philipson (1999), Finkelstein and Poterba (2004), Finkelstein and McGarry (2006)). However, because these studies typically take as a benchmark the exclusive-competition model, the above distinction between demand- and supply-side approaches is often overlooked. Rejecting adverse selection on these markets on the basis of the failure of the positive-correlation property is a decision that should, therefore, be taken with some care: in principle, we would need to collect comprehensive data at the consumer level about all sources of coverage. As pointed out by B. Salanié (2017), this is likely to be a demanding, though worthwhile task.

6.2 Exploiting Price and Cost Data

The analysis of entry-proof tariffs in Section 4 leads to a very sharp prediction for the competitive outcomes of nonexclusive insurance markets: each marginal unit of coverage available along the market tariff should be priced at the expected cost of serving the consumer types who choose to purchase it. This suggests an alternative empirical strategy exploiting price and cost data to compare the price of each layer of insurance to its average cost, as measured by the empirical loss frequency of the consumers who trade this layer.

To illustrate this approach, suppose that the loss is binary and that we have data on observationally equivalent consumers $n = 1, \ldots, N$, providing information about individual aggregate coverage-premia pairs (Q^n, T^n) and loss realizations $L^n \in \{0, 1\}$. Given this data, a natural two-step empirical procedure may run as follows.

The first step would be to construct an estimate of the market tariff T or, more precisely, of the marginal price schedule T'.[26] Although data on firms' offers are typically not available,

[25] This is actually a prediction of the regulated game G^{CSR} under free-entry (AMS (2019b, Theorem 3)): in equilibrium, basic contracts, traded by both low- and high-risk consumers, offer more coverage than complementary contracts, traded by high-risk consumers only. Thus, with data originating from a single firm, we could well observe a negative correlation between risk and coverage.

[26] This is in line with the classical problem of estimating a firm's production frontier.

we could, to this end, use the data on individual aggregate coverage and premia, assuming that each consumer strives to minimize the price she pays for her aggregate coverage. For instance, we could perform a nonparametric regression

$$T^n = T(Q^n) + \varepsilon^n,$$

with one-sided error terms ε^n capturing the idea that consumers may fail to combine the firms' offers optimally.

The second step would be to test whether the estimator $\hat{T}'$ of T' satisfies the property that each marginal quantity is priced at the expected cost of serving the consumers who purchase it. This would involve comparing, for each aggregate coverage level Q, the estimated marginal price $\hat{T}'(Q)$ with the empirical loss frequency

$$\hat{\bar{c}}\,(Q) \equiv \frac{\sum_n 1_{\{Q^n \geq Q, L^n = 1\}}}{\sum_n 1_{\{Q^n \geq Q\}}}$$

of the consumers whose aggregate coverage is at least Q.

Estimates of prices and costs play a crucial role in this procedure. This contrasts with tests of the positive-correlation property, which only rely on aggregate coverage amounts and loss realizations. The procedure is thus closer to that proposed by Einav, Finkelstein, and Cullen (2010) in a setting where consumers have a zero-one demand for coverage: evidence of adverse selection is obtained if the average cost of serving the consumers choosing to buy an additional layer of insurance is affected by the price of that layer. Our analysis suggests that the upper-tail conditional expectation function is the generalization of the firms' cost function in Einav, Finkelstein, and Cullen (2010) to richer environments where firms offer nonexclusive insurance contracts and consumers can choose different levels of coverage. An attractive feature of this approach is that it is fully nonparametric: there is no need to make assumptions about consumers' underlying utility functions nor about the distribution of their private information.

Appendix

Proof of Theorem 6. The proof consists of three steps.

Step 1 For each i in the support of the distribution of types, let $Q_i^{-k}(0,0)$ be a solution to the maximization problem that defines $z_i^{-k}(0,0)$. Then

$$z_i^{-k}(0,0) = u_i(Q_i^{-k}(0,0), T^{-k}(Q_i^{-k}(0,0))) \leq u_i(Q_i^{-k}(0,0), T(Q_i^{-k}(0,0))) \leq U_i, \tag{20}$$

where the first inequality follows from $T \leq T^{-k}$, and the second inequality follows from the fact that U_i is type i's equilibrium utility. Now, if $U_i = z_i^{-k}(0,0)$ for some type i and some seller k, all the inequalities in (20) are in fact equalities. This has two fundamental consequences. First, we have

$$u_i(Q_i^{-k}(0,0), T^{-k}(Q_i^{-k}(0,0))) = u_i(Q_i^{-k}(0,0), T(Q_i^{-k}(0,0))),$$

which implies

$$T^{-k}(Q_i^{-k}(0,0)) = T(Q_i^{-k}(0,0))$$

and, hence, by P2,

$$Q_i^{-k}(0,0) = 0. \tag{21}$$

Second, we have

$$U_i = u_i(Q_i^{-k}(0,0), T^{-k}(Q_i^{-k}(0,0)))$$

and, hence, by (21),

$$U_i = u_i(0,0). \tag{22}$$

Because $U_i \geq z_i^{-l}(0,0) \geq u_i(0,0)$ for all $l \neq k$, the upshot from this reasoning is that, if $U_i = z_i^{-k}(0,0)$ for some type i and some seller k, then, for this type i, $U_i = z_i^{-k}(0,0)$ for any seller k, and hence (21) must hold for all k. In other terms, the indispensability property P2 implies that, if, for some type, the individual-rationality constraint binds for some seller, then it must bind for all sellers.

Step 2 By P1 and Step 1, there exists some i such that $U_i = z_i^{-k}(0,0)$ for all k. Let (q_i^k, t_i^k) be the contract traded by such a type i with seller k in equilibrium, and let $Q_i^{-k} = \sum_{l \neq k} q_i^l$ be the quantity purchased by type i from the sellers other than k, so that

$$U_i = u_i(q_i^k + Q_i^{-k}, t_i^k + T^{-k}(Q_i^{-k})). \tag{23}$$

We claim that $Q_i^{-k} = 0$. Indeed, suppose, by way of contradiction, that $Q_i^{-k} > 0$. Because type i could abstain from trading with the sellers other than k, it must be that

$$u_i(q_i^k + Q_i^{-k}, t_i^k + T^{-k}(Q_i^{-k})) \geq u_i(q_i^k, t_i^k). \tag{24}$$

Similarly, consider the maximization problem that defines $z_i^{-k}(0,0)$, and whose unique solution, by (21), is $Q_i^{-k}(0,0) = 0$. As type i could instead purchase Q_i^{-k} from the sellers other than k, it must be that

$$u_i(0,0) > u_i(Q_i^{-k}, T^{-k}(Q_i^{-k})). \tag{25}$$

We know from (22)–(23) that the left-hand sides of (24)–(25) are both equal to U_i. Hence

$$u_i(0,0) \geq u_i(q_i^k, t_i^k). \tag{26}$$

Representing by $T = \phi(Q)$ the equilibrium indifference curve of type i, (25)–(26) amount to

$$T^{-k}(Q_i^{-k}) > \phi(Q_i^{-k}) \tag{27}$$

and

$$t_i^k \geq \phi(q_i^k). \tag{28}$$

Because ϕ is concave and $\phi(0) = 0$, ϕ is subadditive. Thus, by (27)–(28),

$$t_i^k + T^{-k}(Q_i^{-k}) > \phi(q_i^k + Q_i^{-k}),$$

which amounts to

$$u_i(0,0) > u_i(q_i^k + Q_i^{-k}, t_i^k + T^{-k}(Q_i^{-k})),$$

in contradiction to (22)–(23). Hence $Q_i^{-k} = 0$, as claimed. Because this is true for all k, we obtain that

$$\sum_k q_i^k = \frac{1}{K-1} \sum_k Q_i^{-k} = 0,$$

and thus that $q_i^k = 0$ for all k.

Step 3 We now verify that there exists some type i_1 such that $Q_i \equiv \sum_k q_i^k$ vanishes if and only if $i \leq i_1$, which concludes the proof. By P1 and Steps 1–2, we know that there exists at least one type i for which $Q_i = 0$. This implies in particular that

$$\text{For each } Q > 0,\ u_i(0,0) \geq u_i(Q, T(Q)),$$

and, hence, by strict single-crossing, that

$$\text{For all } j < i \text{ and } Q > 0,\ u_j(0,0) > u_j(Q, T(Q)),$$

so that $Q_j = 0$ for all $j < i$. Thus the set of types who are excluded from trade is an interval $\mathcal{I}^0$ at the bottom of the type distribution. We just need to check that $\mathcal{I}^0$, if it is not reduced to a single type, contains its least upper bound i_1. (This is obvious if the distribution of types is discrete.) By assumption, the mapping $i \mapsto u_i(0,0)$ is continuous, and so is the mapping $i \mapsto U_i$ by Berge's maximum theorem. Because $U_j = u_j(0,0)$ for all $j \in \mathcal{I}_0$, it follows that $U_{i_1} = u_{i_1}(0,0)$ as well, and hence that $U_{i_1} = z_{i_1}^{-k}(0,0)$ for all k. Proceeding as in Step 2 then shows that $Q_{i_1} = 0$, as desired. Hence the result. ■

References

[1] Akerlof, G.A. (1970): "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84(3), 488–500.

[2] Allen, F. (1985): "Repeated Principal-Agent Relationships with Lending and Borrowing," *Economics Letters*, 17(1-2), 27–31.

[3] Arnott, R.J., and J.E. Stiglitz (1991): "Equilibrium in Competitive Insurance Markets with Moral Hazard," NBER Working Paper No. 3588.

[4] Attar, A., and A. Chassagnon (2009): "On Moral Hazard and Nonexclusive Contracts," *Journal of Mathematical Economics*, 45(9–10), 511–525.

[5] Attar, A., V. Farinha Luz, T. Mariotti, and F. Salanié (2021): "Mixed-Strategy Equilibria in Limit-Order Markets," Mimeo, Toulouse School of Economics.

[6] Attar, A., T. Mariotti, and F. Salanié (2011): "Nonexclusive Competition in the Market for Lemons," *Econometrica*, 79(6), 1869–1918.

[7] Attar, A., T. Mariotti, and F. Salanié (2014): "Nonexclusive Competition under Adverse Selection," *Theoretical Economics*, 9(1), 1–40.

[8] Attar, A., T. Mariotti, and F. Salanié (2018): "Concurrence Non Exclusive et Sélection Adverse," *Revue Economique*, 69(6), 1009–1023.

[9] Attar, A., T. Mariotti, and F. Salanié (2019a): "On Competitive Nonlinear Pricing," *Theoretical Economics*, 14(1), 297–343.

[10] Attar, A., T. Mariotti, and F. Salanié (2019b): "Regulating Insurance Markets: Multiple Contracting and Adverse Selection," Working Paper No. 19-1033, Toulouse School of Economics.

[11] Attar, A., T. Mariotti, and F. Salanié (2020): "The Social Costs of Side Trading," *Economic Journal*, 130(630), 1608–1622.

[12] Attar, A., T. Mariotti, and F. Salanié (2021): "Entry-Proofness and Discriminatory Pricing under Adverse Selection," Working Paper No. 17-788, Toulouse School of Economics.

[13] Ausubel, L.M. (2004): "An Efficient Ascending-Bid Auction for Multiple Objects,"

American Economic Review, 94(5), 1452–1475.

[14] Back, K., and S. Baruch (2013): "Strategic Liquidity Provision in Limit Order Markets," *Econometrica*, 81(1), 363–392.

[15] Baruch, S., and L.R. Glosten (2019): "Tail Expectation and Imperfect Competition in Limit Order Book Markets," *Journal of Economic Theory*, 183, 661–697.

[16] Beaudry, P., and M. Poitevin (1995): "Competitive Screening in Financial Markets when Borrowers Can Recontract," *Review of Economic Studies*, 62(3), 401–423.

[17] Biais, B., T. Mariotti (2005): "Strategic Liquidity Supply and Security Design," *Review of Economic Studies*, 72(3), 615–649.

[18] Biais, B., D. Martimort, and J.-C. Rochet (2000): "Competing Mechanisms in a Common Value Environment," *Econometrica*, 68(4), 799–837.

[19] Biais, B., D. Martimort, and J.-C. Rochet (2013): "Corrigendum to "Competing Mechanisms in a Common Value Environment"," *Econometrica*, 81(1), 393–406.

[20] Bisin, A., and P. Gottardi (2006): "Efficient Competitive Equilibria with Adverse Selection.," *Journal of Political Economy*, 114(3), 485–516.

[21] Bisin, A., and D. Guaitoli (2004): "Moral Hazard and Nonexclusive Contracts," *RAND Journal of Economics*, 35(2), 306–328.

[22] Budish, E., P. Cramton, and J. Shim (2015): "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130(4), 1547–1621.

[23] Carmona, G., and J. Fajardo (2009): "Existence of Equilibrium in Common Agency Games with Adverse Selection," *Games and Economic Behavior*, 66(2), 749–760.

[24] Cawley, J., and T. Philipson (1999): "An Empirical Examination of Information Barriers to Trade in Insurance," *American Economic Review*, 89(4), 827–846.

[25] Chiappori P.-A., B. Jullien, B. Salanié, and F. Salanié (2006): "Asymmetric Information in Insurance: Some Testable Implications," *RAND Journal of Economics*, 37(4), 783–798.

[26] Chiappori P.-A., and B. Salanié (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108(1), 56–78.

[27] Chiappori, P.-A., and B. Salanié (2003): "Testing Contract Theory: A Survey of Some Recent Work," in *Advances in Economics and Econometrics*: *Theory and Applications, Eighth World Congress*, *Volume 1*, ed. by M. Dewatripont, L.P. Hansen, and S.J. Turnovsky. Cambridge, UK: Cambridge University Press, 115–149.

[28] Cole, H.L., and N.R. Kocherlakota (2001): "Efficient Allocations with Hidden Income and Hidden Storage," *Review of Economic Studies*, 68(3), 523–542.

[29] Crocker, K.J., and A. Snow (1985): "The Efficiency of Competitive Equilibria in Insurance Markets with Asymmetric Information," *Journal of Public Economics*, 26(2), 207–219.

[30] Dasgupta, P., and E. Maskin (1986): "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications," *Review of Economic Studies*, 53(1), 27–41.

[31] DeMarzo, P., and D. Duffie (1999): "A Liquidity-Based Model of Security Design," *Econometrica*, 67(1), 65–99.

[32] Diamond, D.W., and P.H. Dybvig (1983): "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91(3), 401–419.

[33] Einav, L., A. Finkelstein, and M.R. Cullen (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 125(2), 877–921.

[34] Engers, M., and L. Fernandez (1987): "Market Equilibrium with Hidden Knowledge and Self-Selection," *Econometrica*, 55(2), 425–439.

[35] Farhi, E., M. Golosov, and A. Tsyvinski (2009): "A Theory of Liquidity and Regulation of Financial Intermediation," *Review of Economic Studies*, 76(3), 973–992.

[36] Farinha Luz, V. (2017): "Characterization and Uniqueness of Equilibrium in Competitive Insurance," *Theoretical Economics*, 12(3), 1349–1391.

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

[38] Finkelstein, A., and J. Poterba (2004): "Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market," *Journal of Political Economy*, 112(1), 183–208.

[39] Glosten, L.R. (1989): "Insider Trading, Liquidity, and the Role of the Monopolist Specialist," *Journal of Business*, 62(2), 211–235.

[40] Glosten, L.R. (1994): "Is the Electronic Open Limit Order Book Inevitable?" *Journal of Finance*, 49(4), 1127–1161.

[41] Golosov, M., and A. Tsyvinski (2007): "Optimal Taxation with Endogenous Insurance Markets," *Quarterly Journal of Economics*, 122(2), 487–534.

[42] Hammond, P.J. (1979): "Straightforward Individual Incentive Compatibility in Large Economies," *Review of Economic Studies*, 46(2), 263–282.

[43] Hammond, P.J. (1987): "Markets as Constraints: Multilateral Incentive Compatibility in Continuum Economies," *Review of Economic Studies*, 54(3), 399–412.

[44] Harris, M., and R.M. Townsend (1981): "Resource Allocation under Asymmetric Information," *Econometrica*, 49(1), 33–64.

[45] Hellwig, M.F. (1983): "On Moral Hazard and Non-Price Equilibria in Competitive Insurance Markets," Discussion Paper No. 109, Sonderforschungsbereich 21, Universität Bonn.

[46] Hellwig, M.F. (1988): "A Note on the Specification of Interfirm Communication in Insurance Markets with Adverse Selection," *Journal of Economic Theory*, 46(1), 154–163.

[47] Hendren, N. (2013): "Private Information and Insurance Rejections," *Econometrica*, 81(5), 1713–1762.

[48] Hosseini, R. (2015): "Adverse Selection in the Annuity Market and the Role for Social Security," *Journal of Political Economy*, 123(4), 941–984.

[49] Jacklin, C.J. (1987): "Demand Deposits, Trading Restrictions, and Risk Sharing," in *Contractual Arrangements for Intertemporal Trade*, ed. by E.C. Prescott and N. Wallace. Minneapolis, MI: University of Minnesota Press, 26–47.

[50] Jaynes, G.D. (1978): "Equilibria in Monopolistically Competitive Insurance Markets," *Journal of Economic Theory*, 19(2), 394–422.

[51] Jaynes, G.D. (2011): "Equilibrium and Strategic Communication in the Adverse Selection Insurance Model," Economics Department Working Paper No. 91, Yale University.

[52] Kahn, C.M., and D. Mookherjee (1998): "Competition and Incentives with Nonexclusive Contracts," *RAND Journal of Economics*, 29(3), 443-465.

[53] Kehoe, T.J., D.K. Levine, and E.C. Prescott (2002): "Lotteries, Sunspots, and Incentive Constraints," *Journal of Economic Theory*, 107(1), 39–69.

[54] Leland, H.E., and D.H. Pyle (1977): "Informational Asymmetries, Financial Structure, and Financial Intermediation," *Journal of Finance*, 32(2), 371–387.

[55] Mailath, G., and G. Nöldeke (2008): "Does Competitive Pricing Cause Market Breakdown Under Extreme Adverse Selection?" *Journal of Economic Theory*, 140(1), 97–125.

[56] Malinvaud, E. (1972): "The Allocation of Individual Risks in Large Markets," *Journal of Economic Theory*, 4(2), 312–328.

[57] Martimort, D. (2006): "Multi-Contracting Mechanism Design," in *Advances in Economics and Econometrics*: *Theory and Applications*, *Ninth World Congress*, *Volume 1*, ed. by R. Blundell, W.K. Newey, and T. Persson. Cambridge, UK: Cambridge University Press, 57–101.

[58] Martimort, D., and L. Stole (2002): "The Revelation and Delegation Principles in Common Agency Games," *Econometrica*, 70(4), 1659–1673.

[59] Milgrom, P. (2000): "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108(2), 245–272.

[60] Milgrom, P., and C. Shannon (1994): "Monotone Comparative Statics," *Econometrica*, 62(1), 157–180.

[61] Mimra, W., and A. Wambach (2014): "New Developments in the Theory of Adverse Selection in Competitive Insurance," *Geneva Risk and Insurance Review*, 39(2), 136–152.

[62] Mimra, W., and A. Wambach (2019): "Contract Withdrawals and Equilibrium in Competitive Markets with Adverse Selection," *Economic Theory*, 67(4), 875–907.

[63] Mirrlees, J.A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38(2), 175–208.

[64] Miyazaki, H. (1977): "The Rat Race and Internal Labor Markets," *Bell Journal of Economics*, 8(2), 394–418.

[65] Myers, S.C., and N.S. Majluf (1984): "Corporate Financing and Investment Decisions when Firms Have Information that Investors Do not Have," *Journal of Financial Economics*, 13(2), 187–221.

[66] Myerson, R.B. (1983): "Mechanism Design by an Informed Principal," *Econometrica*, 51(6), 1767–1797.

[67] Netzer, N., and F. Scheuer (2014): "A Game Theoretic Foundation of Competitive Equilibria with Adverse Selection," *International Economic Review*, 55(2), 399–422.

[68] Page, F.H. Jr, and P.K. Monteiro (2003): "Three Principles of Competitive Nonlinear Pricing," *Journal of Mathematical Economics*, 39(1–2), 63–109.

[69] Pauly, M.V. (1974): "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection," *Quarterly Journal of Economics*, 88(1), 44–62.

[70] Peters, M. (2001): "Common Agency and the Revelation Principle," *Econometrica*, 69(5), 1349–1372.

[71] Philippon T., and V. Skreta (2012): "Optimal Interventions in Markets with Adverse Selection," *American Economic Review*, 102(1), 1–28.

[72] Prescott, E.C., and R.M. Townsend (1984): "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, 52(1), 21–45.

[73] Riley, J.G. (1979): "Informational Equilibrium," *Econometrica*, 47(2), 331–359.

[74] Riley, J.G. (2001): "Silver Signals: Twenty-Five Years of Screening and Signaling," *Journal of Economic Literature*, 39(2), 432–478.

[75] Rockafellar, R.T. (1970): *Convex Analysis*, Princeton, NJ: Princeton University Press.

[76] Rosenthal, R.W., and A. Weiss (1984): "Mixed-Strategy Equilibrium in a Market with Asymmetric Information," *Review of Economic Studies*, 51(2), 333–342.

[77] Rothschild, C. (2015): "Nonexclusivity, Linear Pricing, and Annuity Market Screening," *Journal of Risk and Insurance*, 82(1), 1–32.

[78] Rothschild, M., and J.E. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90(4), 629–649.

[79] Rustichini, A., P. Siconolfi, P. (2008): "General Equilibrium in Economies with Adverse Selection, *Economic Theory*, 37(1), 1–29.

[80] Salanié, B. (2017): "Equilibrium in Insurance Markets: An Empiricist's View," *Geneva Risk and Insurance Review*, 42(1), 1–14.

[81] Sheshinski, E. (2008): *The Economic Theory of Annuities*, Princeton, NJ: Princeton University Press.

[82] Spence, M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87(3), 355–374.

[83] Spence, M. (1978): "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics*, 10(3), 427–447.

[84] Stiglitz, J.E., J. Yun, and A. Kosenko (2020): "Bilateral Information Disclosure in Adverse Selection Markets with Nonexclusive Competition," NBER Working Paper No. 27041.

[85] Tirole, J. (2012): "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," *American Economic Review*, 102(1), 29–59.

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

[87] Wilson, R.B. (1989): "Game-Theoretic Analyses of Trading Processes," in *Advances in Economic Theory*, *Fifth World Congress*, ed. by T.F. Bewley. Cambridge, UK: Cambridge University Press, 33–70.

[88] Wilson, R.B. (1993): *Nonlinear Pricing*, Oxford: Oxford University Press.

[89] Yamashita, T. (2010): "Mechanism Games With Multiple Principals and Three or More Agents," *Econometrica*, 78(2), 791–801.