Establishing ‘perfect price collusion’ Bayesian Nash equilibria for Kyle’s (1989) market microstructure model when assets have non-negative values*

Alexander Zimper[†] Youcheng Lou[‡]

August 23, 2023

Abstract

Under the assumption of normally distributed asset values, Kyle (1989) describes Bayesian Nash equilibria which partially reveal the speculators’ private information. We adopt Kyle’s market microstructure model to economies with non-negative asset values. We argue that the most plausible equilibria for such economies are Bayesian Nash equilibria in which all speculators push the asset price to zero irrespective of their private information. For assets with non-negative values, Kyle’s (1989) model of ‘imperfect competition’ with partially revealed information thus becomes a model of ‘perfect price collusion’ which does not reveal any private information.

Keywords: Asymmetric information; Market microstructure; Bayesian Nash equilibrium

JEL Classification Numbers: C71; D82; G12

*Youcheng Lou gratefully acknowledges the financial support of the National Natural Science Foundation of China 71971208 and 72192804.

†Department of Economics, University of Pretoria. South Africa. Email address: alexander.zimper@up.ac.za

‡Academy of Mathematics and Systems Science, Chinese Academy of Sciences. Email address: louyoucheng@amss.ac.cn

1 Introduction

Competitive equilibria have a measurability problem when it comes to asymmetric information: if traders' demand functions are measurable with respect to their different information types, then there typically do not exist equilibrium prices which can clear markets in every state of the world. One prominent attempt to circumvent this problem is the introduction of *rational expectations equilibria* in the specific sense that the inverse of the equilibrium price function is supposed to reveal information to all traders which they can use in their ex ante utility maximization problem.[1] Radner (1979) and Allen (1981) have shown that the information revealed by prices is generically one-to-one to the full communication information, i.e., the information which would obtain if all traders truthfully shared their private information. As a consequence, all traders' demand functions become in a rational expectations equilibrium measurable with respect to the same full communication information. The problem with this 'solution' to the measurability problem is that rational expectations equilibria ignore the traders' incentives to hide their private information through the manipulation of prices.[2] For example, why should the owner of a bad car in Akerlof's (1970) lemon market reveal the low quality of his car by accepting a low price?

Game-theoretic models of asset trade can offer solutions to the question of how incentive compatibility can be combined with the competitive equilibrium requirement of market-clearing in every state of the world. On the one hand, the strategic solution of a Bayesian Nash equilibrium ensures, by its very definition, incentive compatibility for all information types. On the other hand, the consideration of strategies in terms of demand-schedule correspondences—through which speculators' information types submit their different demands at different prices to an auctioneer—might allow for prices to clear markets in every state of the world as long as these demand-schedules contain sufficiently rich combinations of demands and prices.

[1]Early examples of "fulfilled expectations" (=rational expectations) equilibria appear in Green (1975), Grossman (1976, 1977, 1978), and Kreps (1977). Radner (1979)—for the case of a finite state space—and Allen (1981)—for the case of a general probability space—prove the generic existence of fully revealing rational expectations equilibria.

[2]Hellwig (1980) basically argues that rational expectations equilibria are only incentive compatible for large markets in which each atomless trader has no impact on the equilibrium price.

A pioneering article in this literature is Kyle (1989) whose strategic model of 'imperfect competition' links different realizations of speculator's private information to different market outcomes through equilibrium demand-schedule correspondences that have a linear structure.[3] Although already published in 1989, Kyle's model is still highly popular. The *www.dimensions.ai* database website writes: "This publication in The Review of Economic Studies has been cited 627 times. 9% of its citations have been received in the past two years, which is higher than you might expect, suggesting that it is currently receiving a lot of interest."[4] Unlike in competitive equilibrium models, competition in Kyle (1989) is 'imperfect' because traders are not merely price-takers but can influence the market price through their choice of demand-schedule correspondences:

> "After private observations are realized, each speculator chooses a demand function, taking as given the strategies other speculators use to choose theirs. Since the market clearing price is determined after the demand functions are chosen, each speculator realizes that his choice of demand functions influences the market clearing price. In other respects, the equilibrium looks Walrasian, i.e. an "auctioneer" calculates the equilibrium price by aggregating demand curves." (Kyle 1989, p.318)

This paper adopts Kyle's (1989) market microstructure model to situations in which asset values are always non-negative. Kyle's (1989) expected utility specification—which combines a normally distributed asset value with a CARA (=constant absolute risk aversion) Bernoulli utility function—comes with the analytically convenient feature that demand functions are linear in the price. But as a severe drawback, Kyle's utility specification makes his model very restrictive: all information types of the speculators will only demand some limited amount of the asset even if its price is zero. More generally, the assumption that asset values can be negative excludes relevant asset classes traded at stock exchanges such as shares or bonds whose minimal value to a limited-liability investor cannot be worse than zero. Our own model applies Kyle's (1989) original game-theoretic structure to speculators who have (i) arbitrary Bernoulli utility functions, (ii) arbitrary

[3] For a discussion of Kyle's (1989) model within the wider context of the market microstructure literature see, e.g., Easley and O'Hara (1995).

[4] Accessed 03 August 2023 at https://badge.dimensions.ai/details/id/pub.1069868950

(heterogeneous) beliefs, and (iii) arbitrary private information types whereby we restrict attention to assets that can only have non-negative values.

We establish the existence of a symmetric Bayesian Nash equilibrium (=BNE) in which the information types of all speculators collude to push the asset price towards zero in every state of the world. We argue that this 'perfect price collusion' BNE is highly plausible because the corresponding BNE strategy of each speculator is a best response against any demand-schedule through which some opponent would like to buy as much as possible of the valuable asset at a zero price. Our findings rely heavily on Kyle's structural assumption that speculators have to choose between convex-valued demand-schedule correspondences and not just between demand-schedule functions.[5] As a consequence, our game-theoretic analysis of best responses cannot be reduced to first-order conditions but has to take into account 'jumps' that arise from the auctioneer's market-clearing algorithm.

The 'perfect price collusion' BNE of our model is in stark contrast to the partially-revealing Bayesian Nash equilibria described in Kyle (1989) in which market prices depend on the speculators' private information. We construct an illustrative example for which partially information revealing BNE that Pareto-dominate the 'perfect price collusion' BNE can only be established under non-generic conditions imposed on the speculators' beliefs. Kyle's (1989) model of 'imperfect competition' thus becomes for non-negative asset values a model of 'perfect price collusion' in which the corresponding zero market price does not reveal any private information.

The remainder of our analysis proceeds as follows. Section 2 formally translates Kyle's (1989) structural assumptions into a class of Bayesian games. Section 3 proves the existence of 'perfect price collusion' BNE. Section 4 discusses our findings and Section 5 concludes. Formal proofs are relegated to the Appendix.

[5] To restrict attention to demand-schedule functions only would mean that the speculators are only allowed to submit "fill-or-kill" orders, which is unrealistic for most markets.

2 The Kyle (1989) market microstructure model with non-negative asset values

2.1 Kyle's structural assumptions

Whereas our own approach will make different distributional assumptions than Kyle's (1989) model, we want to stay as closely as possible to the details of his original assumptions concerning the model's market microstructure. The following quotes from Kyle (1989, p.321) recall these structural assumptions (which we label for later reference).

> K1: "A "demand schedule" $X_n(\cdot; i_n)$ is allowed to be any convex valued, upper-hemicontinuous correspondence mapping prices p into non-empty subsets of the closed infinite interval $[-\infty, +\infty]$."
>
> K2: "So-called "all-or-nothing" orders (orders to purchase all of the desired quantity or none, with partial executions not accepted) are excluded [....] : if the market-clearing price equals the limit price in a trader's limit order, he must accept as a satisfactory execution all of the quantity he requested, or *any fraction thereof.*"
>
> K3: "These correspondences are submitted to an auctioneer (i.e. a broker or non-trading specialist), who proceeds as follows. First, the set of market-clearing prices and quantity allocations is calculated. [...] If a market-clearing price exists, the auctioneer chooses from the set of all such prices (which is closed by upper hemicontinuity) that price with minimum absolute value (or the positive one if p and $-p$ both have minimal absolute value)."
>
> K4: "He then chooses the market-clearing quantity allocation which minimizes the sum of squared quantities traded by speculators."
>
> K5: "If a market-clearing price does not exist, the fact that correspondences are convex-valued implies that there is either positive excess demand at all prices or negative excess demand at all prices. In the former case, the auctioneer announces a price $p = +\infty$, and all buyers of bounded quantities receive negatively infinite utility. Similarly, in the latter case, the auctioneer announces a price $p = -\infty$, and all sellers of bounded quantities receive negatively infinite utility."

In contrast to Kyle, we will neither consider any infinite quantities or prices nor the possibility of *negative* prices. In this regard, the original formulation of K5 is somewhat unclear.[6] The game-theoretic intention of K5, however, is very clear: "Infinite quantities and prices are a theoretical possibility, but they do not occur in the equilibrium defined below because they imply infinitely negative utility." (Kyle 1989, p.321). That is, the purpose of K5 is to enforce market-clearing strategy choices in a Bayesian Nash equilibrium by punishing the possible non-existence of a market-clearing outcome through an infinitely negative utility. We therefore reformulate K5 more directly as follows:

> K5': If markets do not clear in every state of the world, each speculator receives an infinitely negative utility.

In what follows, we will translate Kyle's structural assumptions K1,...,K5' into a class of Bayesian games. We will thereby distinguish between two different scenarios in order to address an ambiguity concerning Kyle's definition of admissible demand-schedule correspondences.

2.2 Translating Kyle's structural assumptions into a class of Bayesian games

We consider an asset trade economy in which N speculators can buy (but not sell) quantities of some risky asset from a noise trader who—randomly and inelastically—supplies the asset. Given the submitted demand-schedules for the speculators' realized information types and the noise trader's supply, markets are cleared by an auctioneer who uses a specific algorithm to determine the market price and market allocation (provided that at least one market-clearing price exists).

2.2.1 Information types and beliefs

The set of possible values of the asset—which are always non-negative by assumption—is given by $V \subseteq \mathbb{R}_{\geq 0}$. Denote by Θ_i an arbitrary index set of possible information types for

[6] If an auctioneer can announce a price $p = -\infty$, all speculators might collude to demand positive but sufficiently small amounts of the asset in order to generate 'negative excess demand' which would transfer via $p = -\infty$ infinitely positive utility to them.

any given speculator $i \in \{1, ..., N\}$ with the general interpretation that any given type $\theta_i \in \Theta_i$ impacts somehow on speculator i's belief about the asset's true value. Finally, denote by $Z \subseteq \mathbb{R}_{\geq 0}$ the possible amounts of the asset supplied by the noise trader. This random amount is always non-negative because we restrict attention to situations in which speculators compete about some exogenously given supply of the asset.

Define the relevant state space as

$$\Omega = V \times \Theta_1 \times \cdots \times \Theta_N \times Z$$

with $(v, \theta_1, ..., \theta_N, z) \in \Omega$ as generic element. Each speculator's prior beliefs are given by some common additive probability measure π_i defined on the measurable space $(\Omega, \mathcal{B})$ with $\mathcal{B}$ denoting the Borel-sigma algebra on the state space Ω.[7] When speculator i learns his private information type θ_i–which happens before any demand-schedule is submitted–he updates his prior to the conditional probability measure $\pi_i(\cdot \mid \theta_i)$. In case we want to model the information type of some *uninformed* speculator i (as, e.g., considered in Kyle (1989)), we simply set $\Theta_i = \{\theta_i\}$ to obtain $\pi_i(\cdot \mid \theta_i) = \pi_i(\cdot)$.

Introduce the following coordinate random variables on the probability space $(\Omega, \mathcal{B}, \pi_i)$

$$\begin{aligned} \mathbf{v}(v, \theta_i, \theta_{-i}, z) &= v, \\ \boldsymbol{\theta}_i(v, \theta_i, \theta_{-i}, z) &= \theta_i, \\ \mathbf{z}(v, \theta_i, \theta_{-i}, z) &= z \end{aligned}$$

for all $(v, \theta_1, ..., \theta_N, z) \in \Omega$, which stand for the random asset value $\mathbf{v}$, information type $\boldsymbol{\theta}_i$, and supplied amount $\mathbf{z}$, respectively. Kyle (1989) stipulates that an *informed* speculator i's information type is given as the true value of the asset plus some random error term, i.e.,

$$\boldsymbol{\theta}_i = \mathbf{v} + \mathbf{e}_i,$$

such that $\mathbf{v}$ and the error terms $\mathbf{e}_1, ..., \mathbf{e}_N$ are independently and normally distributed with zero means under the probability measure π_i. Kyle (1989) also assumes that the random supply $\mathbf{z}$ is independently and normally distributed with zero mean so that negative values

[7] Recall that a Borel-sigma algebra is generated by the *open* sets of Ω. We endow Ω with the standard product topology such that the V and Z are endowed with the Euclidean topology on $\mathbb{R}_{\geq 0}$. The same holds for Θ_i if $\Theta_i \subseteq \mathbb{R}_{\geq 0}$. If Θ_i is just some finite non-numerical set, we use the discrete topology.

of $\mathbf{z}$ would stand for a positive demand of the (aggregate) noise trader. At this point we deviate from Kyle's specific distributional assumptions: we only consider non-negative, but arbitrarily distributed, values of the random variables $\mathbf{v}$ and $\mathbf{z}$ whereby we allow for heterogeneous π_i.

2.2.2 Actions and strategies

A demand-schedule correspondence of information type θ_i is any non-empty set-valued mapping

$$\varphi_i [\theta_i] : P \rightrightarrows \mathbb{R}_{\geq 0}$$

for some set of admissible prices $P \subseteq \mathbb{R}_{\geq 0}$ satisfying $0 \in P$. By Kyle's Assumption K1 (i) $\varphi_i [\theta_i] (p)$ has to be convex-valued for every p and (ii) $\varphi [\theta_i]$ must be upper-hemicontinuous on P, i.e., for any open set $\mathcal{U}$ around the set $\varphi_i [\theta_i] (p)$, $p \in P$, the pre-image of $\mathcal{U}$ under $\varphi_i [\theta_i]$

$$\mathcal{U}^{-1} = \{p' \in P \mid \varphi_i [\theta_i] (p') \in \mathcal{U}\}$$

is also open in the Euclidean topology on P. While we adopt the assumption of convex-valuedness of $\varphi_i [\theta_i]$, we use a weaker assumption than upper-hemicontinuity in that we only require that the values $\varphi_i [\theta_i] (p)$, $p \in P$, are closed sets.

By Kyle's Assumption K2, the speculator "must accept as a satisfactory execution all of the quantity he requested, or *any fraction thereof*." Kyle (1989, p.321) writes that K2, and in particular the exclusion of "all-or-nothing" orders, is already implied "by the assumption that demand correspondences be convex-valued". But this is technically not correct: to formalize Assumption K2, one needs to combine convex-valuedness with the additional assumption that the zero demand-offer is always included. To distinguish between these two different definitions of admissible demand-schedule correspondences—one imposing K2, the other not—, we consider two different scenarios.

Definition 1.

(i) *We speak of the 'Convexity' scenario iff, for any given $p \in P$, the value $\varphi_i [\theta_i] (p)$ of an admissible demand-schedule correspondence is some closed and convex interval* *on* $\mathbb{R}_{\geq 0}$*, i.e.,*

$$\text{either } \varphi_i[\theta_i](p) = [\underline{x}, \overline{x}] \text{ or } \varphi_i[\theta_i](p) = [\underline{x}, \infty) \tag{1}$$

such that $0 \leq \underline{x} \leq \overline{x}$ *whereby* $\underline{x}$ *and* $\overline{x}$ *may depend on the price* p*.*

(ii) *The 'Any-fraction' scenario is the special case of the 'Convexity' scenario such that* $\underline{x} = 0$*. That is, we speak of the 'Any-fraction' scenario iff, for any given* $p \in P$*, the value* $\varphi_i[\theta_i](p)$ *of an admissible demand-schedule correspondence is some closed and convex interval with zero as lower bound, i.e.,*

$$\text{either } \varphi_i[\theta_i](p) = [0, \overline{x}] \text{ or } \varphi_i[\theta_i](p) = [0, \infty) \tag{2}$$

where $\overline{x} \geq 0$ *denotes the "quantity requested" at price* p*.*

Denote by A^k with $k \in \{conv, anyf\}$ the set of all admissible demand-schedule correspondences, i.e., admissible actions, of any given information type θ_i. The index $k = conv$ refers to the 'Convexity scenario'—with admissible demand-schedule correspondences given by (1)—whereas $k = anyf$ refers to the 'Any-fraction scenario'—with admissible demand-schedule correspondences given by (2). A *strategy* s_i of speculator i is any mapping that assigns to every information type of i some action, i.e., $s_i : \Theta_i \to A^k$ whereby we also write $s_i = (\varphi_i[\theta_i])_{\theta_i \in \Theta}$. The interpretation of a strategy is as follows. Suppose that speculator i chooses strategy $s_i = (\varphi_i[\theta_i])_{\theta_i \in \Theta}$. After learning his information type θ_i, he submits through the action $\varphi_i[\theta_i]$ the following demand-schedule to the auctioneer

$$\left\{(x_i, p) \in \mathbb{R}^2_{\geq 0} \mid x_i \in \varphi_i[\theta_i](p), \, p \in P\right\}.$$

Through this demand-schedule speculator i commits himself to buy any amount $x_i \in \varphi_i[\theta_i](p)$ provided that the auctioneer announces p as the market price. For example, by submitting some bounded demand-schedule $\varphi_i[\theta_i](p) = [0, \bar{x}]$ in accordance with the 'Any-fraction' scenario, the speculator commits himself to buy at the market price p "all of the quantity he requested", i.e., $x_i = \bar{x}$, or "*any fraction thereof*", i.e., $x_i \in [0, \bar{x})$. If he submits, e.g., $\varphi_i[\theta_i](p) = [\bar{x}, \bar{x}] = \{\bar{x}\}$ in accordance with the 'Convexity' scenario, the speculator commits himself to buy $\bar{x}$ but rejects any other amount at price p, which stands for an ‘all-or-nothing’ (i.e., ‘fill-or-kill’) order that would be excluded under Assumption K2.

Denote by S_i the set of all strategies of speculator i and denote by S the set of all *strategy-profiles*, i.e., $S = \times_{i=1}^{N} S_i$.

Remark 1. The ‘Convexity’ scenario covers the possibility of single-valued demand-schedule correspondences (i.e., functions), which are convex-valued but would not be admissible under the ‘Any-fraction’ scenario (unless these functions are constantly zero). On the one hand, Kyle’s analysis of BNE deals exclusively with (linear) demand-schedule functions, which would suggest that he either works under the ‘Convexity’ scenario or that he—implicitly—makes the even stronger assumption that only demand-functions are admissible. On the other hand, Kyle’s own formulation of Assumption K2 strongly suggests that he intends to work within the ‘Any-fraction’ scenario. To deal with this ambiguity in Kyle’s definition of admissible demand-schedule correspondences, we will distinguish between the two scenarios $k \in \{conv, anyf\}$ when we establish existence of a ‘perfect price collusion’ BNE.[8]

2.2.3 Market-clearing algorithm

Fix some strategy profile $(s_1, ..., s_N) \in S$ with $s_i = (\varphi_i [\theta_i])_{\theta_i \in \Theta}$ for all i. The set of all potentially market clearing prices in state $(v, \theta_1, ..., \theta_N, z) \in \Omega$ is given as

$$P^* [s] (v, \theta_1, ..., \theta_N, z) = \left\{ p \in P \mid \sum_{i=1}^{N} x_i = z \text{ such that } x_i \in \varphi_i [\theta_i] (p) \right\}.$$

By Kyle’s Assumption K3, the market price chosen by the auctioneer in state $(v, \theta_1, ..., \theta_N, z)$ is pinned down as

$$\mathbf{p}^* [s] (v, \theta_1, ..., \theta_N, z) = \min P^* [s] (v, \theta_1, ..., \theta_N, z) \tag{3}$$

whenever this minimal market-clearing price exists. In what follows, we write $\mathbf{p}^* [s] : \Omega \to \mathbb{R}_{\geq 0}$ for the random variable $\mathbf{p}^* [s] = \min P^* [s]$ whenever (3) exists for all $(v, \theta_1, ..., \theta_N, z) \in$

[8] Our Propositions 1 and 2, and their respective proofs, show that these different scenarios give rise to different strategic considerations. Simply put, the possibility to commit to ‘all-or-nothing’ offers gives speculators under the ‘Convexity’ scenario a greater power to influence market-clearing outcomes than under the ‘Any-fraction’ scenario (also cf. Remark 2 in the following subsection).

Ω. Else, we say that $\mathbf{p}^*[s]$ is not well-defined. We write $\mathbf{p}^*[\varphi_i[\theta_i], s_{-i}]$ for the random market price that obtains in states $(\cdot, \theta_i)$ if information type θ_i submits the demand-schedule correspondence $\varphi_i[\theta_i]$.

The set of all potentially market-clearing demand offers at the market price $\mathbf{p}^*[s](v, \theta_1, ..., \theta_N, z)$ is given as

$$X(\mathbf{p}^*[s](\cdot)) = \left\{(x_1, ..., x_N) \in \times_{i=1}^N \varphi_i[\theta_i](\mathbf{p}^*[s](\cdot)) \mid \sum_{i=1}^N x_i = z\right\}. \tag{4}$$

By Kyle's Assumption K4, the market-clearing demand offers that are chosen by the auctioneer in state $(v, \theta_1, ..., \theta_N, z)$ is some allocation $(x_1^*, ..., x_N^*)$ that minimizes the 'sum of squared market clearing demands', i.e.,

$$(x_1^*, ..., x_N^*) \in \arg \min_{(x_1,...,x_N) \in X(\mathbf{p}^*[s](\cdot))} \sum_{i=1}^N x_i^2. \tag{5}$$

For any given state $(v, \theta_1, ..., \theta_N, z)$ select some minimizing $(x_1^*, ..., x_N^*)$. In what follows, we write

$$\mathbf{x}^*[s] = (\mathbf{x}_1^*[s], ..., \mathbf{x}_N^*[s]) : \Omega \to \mathbb{R}_{\geq 0}^N$$

for the corresponding random vector, i.e., for all $(v, \theta_1, ..., \theta_N, z) \in \Omega$

$$\mathbf{x}_i^*[s](v, \theta_1, ..., \theta_N, z) = x_i^* \text{ for all } i,$$

provided that such minimizers exist for all states. Else, we say that $\mathbf{x}^*[s]$ is not well-defined. We write $\mathbf{x}^*[\varphi_i[\theta_i], s_{-i}]$ for the random market allocation if θ_i submits the demand-schedule correspondence $\varphi_i[\theta_i]$.

To emphasize our terminology: $\mathbf{p}^*[s]$ is not well-defined iff there exists some state of the world in which the above market-clearing algorithm cannot pin down any market price. Analogously, $\mathbf{x}^*[s]$ is not well-defined iff there exists some state of the world in which the above market-clearing algorithm cannot pin down any market allocation.

Remark 2. Let us briefly illustrate the different implications of the 'Any-fraction' and the 'Convexity' scenario, respectively, on market-clearing outcomes. Consider two speculators $i \in \{A, B\}$ such that $\varphi_A[\theta_A](p) = [0, \infty)$. If θ_B chooses under the 'Any-fraction' scenario $\varphi_B[\theta_B](p) = [0, z]$, we obtain in any state $(v, \theta_A, \theta_B, z)$ at the market price $\mathbf{p}^*[s](v,\theta_A,\theta_B,z) = p$ the following set of market clearing allocations

$$X(\mathbf{p}^*[s](\cdot)) = \{(x_A, x_B) \mid x_A \in [0,z], \, x_B = z - x_A\}.$$

In accordance with (5) the auctioneer pins down $(x_A, x_B) = \left(\frac{z}{2}, \frac{z}{2}\right)$ as market allocation in any state (v,θ_A,θ_B,z). Suppose now that θ_B chooses under the 'Convexity' scenario only a single amount such that $\varphi_B[\theta_B](p) = \{z\}$. Then the set of market clearing allocations becomes in any state (v,θ_A,θ_B,z) at market price $\mathbf{p}^*[s](v,\theta_A,\theta_B,z) = p$ the singleton set

$$X(\mathbf{p}^*[s](\cdot)) = \{(x_A, x_B) = (0,z)\}.$$

By (5), the auctioneer (trivially) pins down $(x_A, x_B) = (0,z)$ as market allocation in any state (v,θ_A,θ_B,z). Under the 'Convexity' scenario a speculator, here information type θ_B, can thus reject through an 'all-or-nothing' offer to share the supplied amount with other speculators. This is not possible under the 'Any-fraction' scenario.

2.2.4 Utility functions

If $\mathbf{p}^*[s]$ and $\mathbf{x}^*[s]$ are well-defined, the expected utility of speculator i from the strategy profile $s = (s_i, s_{-i})$ is given as

$$U_i(s_i, s_{-i}) = \mathbb{E}_{\pi_i} u_i\left((\mathbf{v} - \mathbf{p}^*[s_i, s_{-i}])\,\mathbf{x}_i^*[s_i, s_{-i}]\right)$$

where u_i denotes some increasing Bernoulli utility function. Whereas Kyle (1989) only considers CARA Bernoulli utility functions, our speculators can have arbitrary Bernoulli utility functions as long as they are integrable so that the (finite) Lebesgue integral

$$\mathbb{E}_{\pi_i} u_i(\cdot) = \int_\Omega u_i(\cdot)\, d\pi_i$$

exists.

If $\mathbf{p}^*[s]$ or $\mathbf{x}^*[s]$ are not well-defined, we set, in accordance with Assumption K5',

$$U_i(s_i, s_{-i}) = -\infty. \tag{6}$$

3 'Perfect price collusion' Bayesian Nash equilibria

Denote by $\Gamma^k = \langle S_i, U_i, i \in \{1, ..., N\} \rangle$ with $k \in \{conv, anyf\}$ any Bayesian game whose components satisfy the definitions of the previous section.[9] The strategy profile

$$(s_1^*, ..., s_N^*) = \left((\varphi_1^* [\theta_1])_{\theta_1 \in \Theta_1}, ..., (\varphi_N^* [\theta_N])_{\theta_N \in \Theta} \right)$$

is a Bayesian Nash equilibrium (=BNE) of Γ^k iff, for all $i \in \{1, ..., N\}$,

$$\begin{aligned} U_i \left(s_i^*, s_{-i}^* \right) &\geq U_i \left(s_i, s_{-i}^* \right) \\ &\Leftrightarrow \\ \mathbb{E}_{\pi_i} u_i \left(\left(\mathbf{v} - \mathbf{p}^* \left[s_i^*, s_{-i}^* \right] \right) \mathbf{x}_i^* \left[s_i^*, s_{-i}^* \right] \right) &\geq \mathbb{E}_{\pi_i} u_i \left(\left(\mathbf{v} - \mathbf{p}^* \left[s_i, s_{-i}^* \right] \right) \mathbf{x}_i^* \left[s_i, s_{-i}^* \right] \right) \end{aligned} \tag{7}$$

for all $s_i \in S_i$. By the towering property of conditional expectations (cf. Theorem 34.4 in Billingsley 1995), (7) is with π_i-probability one equivalent to

$$\begin{aligned} & \mathbb{E}_{\pi_i(\theta_i)} \mathbb{E}_{\pi_i(\cdot|\theta_i)} u_i \left(\left(\mathbf{v} - \mathbf{p}^* \left[\varphi_i^* [\theta_i], s_{-i}^* \right] \right) \mathbf{x}_i^* \left[\varphi_i^* [\theta_i], s_{-i}^* \right] \right) \\ \geq\ & \mathbb{E}_{\pi_i(\theta_i)} \mathbb{E}_{\pi_i(\cdot|\theta_i)} u_i \left(\left(\mathbf{v} - \mathbf{p}^* \left[\varphi_i [\theta_i], s_{-i}^* \right] \right) \mathbf{x}_i^* \left[\varphi_i [\theta_i], s_{-i}^* \right] \right) \end{aligned}$$

whenever $\mathbf{p}^* [s]$ and $\mathbf{x}^* [s]$ are well-defined. A sufficient condition for a market-clearing strategy profile s^* being a BNE is thus

$$\begin{aligned} & \mathbb{E}_{\pi_i(\cdot|\theta_i)} u_i \left(\left(\mathbf{v} - \mathbf{p}^* \left[\varphi_i^* [\theta_i], s_{-i}^* \right] \right) \mathbf{x}_i^* \left[\varphi_i^* [\theta_i], s_{-i}^* \right] \right) \\ \geq\ & \mathbb{E}_{\pi_i(\cdot|\theta_i)} u_i \left(\left(\mathbf{v} - \mathbf{p}^* \left[\varphi_i [\theta_i], s_{-i}^* \right] \right) \mathbf{x}_i^* \left[\varphi_i [\theta_i], s_{-i}^* \right] \right) \end{aligned} \tag{8}$$

for all $\theta_i \in \Theta$, $i \in \{1, ..., N\}$.[10]

Proposition 1. *Consider the 'Any-fraction' scenario according to which the values of the admissible demand-schedule correspondences of Γ^{anyf} are given by (2).*

[9]Some authors stipulate that the players in a Bayesian game must share a common prior. In contrast to this view, our definition of a Bayesian game is more general in that it allows for the possibility of heterogeneous priors. We would like to refer any reader who is interested in the plausibility of the common prior assumption to the discussion between Gul (1998) and Aumann (1998).

[10]For a finite number of types in Θ such that $\pi(\theta_i) > 0$ for all $\theta_i \in \Theta$, the sufficiency condition (8) is also necessary for s^* being a BNE.

(i) *There exists a symmetric BNE*

$$s^* = \left((\varphi_1^* [\theta_1])_{\theta_1 \in \Theta}, ..., (\varphi_N^* [\theta_N])_{\theta_N \in \Theta} \right) \tag{9}$$

of Γ^{anyf} *such that, for all* $\theta_i \in \Theta$, $i \in \{1, ..., N\}$,

$$\varphi_i^* [\theta_i] (p) = [0, \infty) \text{ for all } p \in P. \tag{10}$$

(ii) *The market price and allocations corresponding to* s^* *are independent of the speculators' information types and given as follows*

$$\begin{aligned} \mathbf{p}^* [s^*] (v, \theta_1, ..., \theta_N, z) &= 0, \\ \mathbf{x}_i^* [s^*] (v, \theta_1, ..., \theta_N, z) &= \frac{z}{N} \text{ for all } i \end{aligned} \tag{11}$$

for all states $(v, \theta_1, ..., \theta_N, z) \in \Omega$.

The basic intuition for the formal proof of Proposition 1 is straightforward. In the 'Any-fraction' scenario it is impossible for any given information type θ_i to move the market price away from zero provided that all other speculators stick to their strategy profile s_{-i}^*. As the asset—having non-negative values—is always valuable at price zero, every information type θ_i would therefore like to buy as many units as possible at price zero. But this can be achieved through the BNE action (10).

Proposition 2. *Now consider the weaker 'Convexity' scenario according to which the values of the admissible demand-schedule correspondences of* Γ^{conv} *are given by (1).*

If the zero supply is a possibility, i.e., if $0 \in Z$, *then the strategy profile (9)—with corresponding market price and allocations (11)—is also a BNE of* Γ^{conv}.

In contrast to the 'Any-fraction' scenario of Proposition 1, information type θ_i has—at the fixed profile s_{-i}^*—now the power to change the market price from zero to some strictly positive price. The reason is that he can commit under the 'Convexity' scenario to some strictly positive demand at price zero. Combined with the assumption that a zero supply is a possibility, markets would then not clear at price zero in states in which the supply is zero. In spite of this power to impact on the market-price, information type θ_i has no strict incentive to move the market price away from zero: given his opponents' strategy profile s^*_{-i} he can never achieve a greater market share of the valuable asset than the amount $\frac{z}{N}$ which he already obtains in the BNE (9) at a zero price.

Observe that the BNE action (10) is always a best response of information type θ_i against all strategy profiles s_{-i} which include at least one $j \neq i$ who offers to purchase any amount of the asset at price zero, i.e.,

$$\varphi_j^* [\theta_j] (0) = [0, \infty) \text{ for all } \theta_j \in \Theta_j. \tag{12}$$

In other words, under the perfectly natural expectation that some speculator would always buy as much as possible of the asset at price zero, any other speculator could not do better than adopting the BNE action (10). For this reason, we regard the 'perfect price collusion' BNE (9) as the most plausible BNE for the class of games considered in this paper.

Remark 3. Without the additional requirement (12), the BNE action (10) is not necessarily a best response against all s_{-i}. That is, the 'perfect price collusion' BNE is not necessarily an equilibrium in *weakly dominating strategies*. To see this, let $i \in \{A, B\}$, $Z = \{z\}$, and consider s_A such that for all $\theta_A \in \Theta_A$

$$\varphi_A [\theta_A] (p) = \begin{cases} [0, \frac{z}{2}] & p = 0 \\ \{0\} & p = \varepsilon \end{cases}$$

Then there exists—for a strictly increasing Bernoulli utility function and $z > 0$—some sufficiently small $\varepsilon > 0$ such that the action

$$\varphi_B [\theta_B] (p) = \begin{cases} \{0\} & p = 0 \\ [0, \infty) & p = \varepsilon \end{cases}$$

is strictly better against s_A than the BNE action (10) because θ_B would receive the full amount z of the asset at the small price ε rather than $\frac{z}{2}$ at price zero.

4 Discussion

4.1 'Perfect price collusion' BNE versus Pareto-superior strategy profiles: an illustrative example

Although the 'perfect price collusion' BNE fully exploits the noise-trader, there typically exist Pareto-superior strategy profiles that would result for each speculator in a greater utility. The reason is that the egalitarian allocation $\frac{z}{N}$ (whenever the supply is z) for all information types under the 'perfect price collusion' BNE does not take account of possibly different beliefs (or/and risk-appetites) of the speculators. The more relevant question is whether such Pareto-superior strategy profiles could also be established as BNE.

The 'perfect price collusion' BNE of Propositions 1 and 2 have been obtained for general Bernoulli utility functions and beliefs as long as asset values are non-negative. To establish the existence of any information revealing BNE—and check for their Pareto-superiority—is impossible at this level of generality. In what follows we construct an illustrative example which shows the difficulty of transforming Pareto-superior strategy profiles into BNE.

Consider the game Γ^{anyf} under the 'Any-fraction' scenario such that $i \in \{A, B\}$ and

$$\begin{aligned} V &= \{0, \bar{v}\} \text{ with } \bar{v} > 0, \\ \Theta_i &= \{L, H\}, \\ Z &= \{0, \bar{z}\} \text{ with } \bar{z} > 0. \end{aligned}$$

Both speculators are risk-neutral, i.e., $u_i(x) = x$. We restrict attention to a countable set of admissible prices

$$P = \{\varepsilon_0, \varepsilon_1, ...\} \text{ with } \varepsilon_0 = 0 \text{ and } \varepsilon_k < \varepsilon_{k+1}.$$

Next define the symmetric strategy profile

$$s^{**} = ((\varphi_A^{**}[L], \varphi_A^{**}[H]), (\varphi_B^{**}[L], \varphi_B^{**}[H])) \quad (13)$$

such that

$$\begin{aligned}
\varphi_A^{**}[L](p) &= \begin{cases} \{0\} & \text{if } p = \varepsilon_0 \\ [0,\infty) & \text{if } p \geq \varepsilon_1 \end{cases} \\
\varphi_A^{**}[H](p) &= [0,\infty) \text{ for all } p,
\end{aligned}$$

$$\begin{aligned}
\varphi_B^{**}[L](p) &= \begin{cases} \{0\} & \text{if } p = \varepsilon_0 \\ [0,\infty) & \text{if } p \geq \varepsilon_1 \end{cases} \\
\varphi_B^{**}[H](p) &= [0,\infty) \text{ for all } p.
\end{aligned}$$

The market price and allocations that correspond to s^{**} in state $(v, \theta_A, \theta_B, z) \in \Omega$ are given as

$$\begin{aligned}
\mathbf{p}^*[s^{**}](v,\theta_A,\theta_B,z) &= \begin{cases} \varepsilon_0 & \text{if } (\theta_A,\theta_B) \in \{(H,L),(L,H),(H,H)\} \\ \varepsilon_0 & \text{if } (\theta_A,\theta_B) = (L,L) \text{ and } z = 0 \\ \varepsilon_1 & \text{if } (\theta_A,\theta_B) = (L,L) \text{ and } z = \bar{z} \end{cases} \\
\mathbf{x}_A^*[s^{**}](v,\theta_A,\theta_B,z) &= \begin{cases} z & \text{if } (\theta_A,\theta_B) = (H,L) \\ 0 & \text{if } (\theta_A,\theta_B) = (L,H) \\ \frac{z}{2} & \text{else} \end{cases} \\
\mathbf{x}_B^*[s^{**}](v,\theta_A,\theta_B,z) &= \begin{cases} 0 & \text{if } (\theta_A,\theta_B) = (H,L) \\ z & \text{if } (\theta_A,\theta_B) = (L,H) \\ \frac{z}{2} & \text{else} \end{cases}
\end{aligned}$$

Note that the price function $\mathbf{p}^*[s^{**}]$ is partially revealing: whenever we observe price ε_1, we know that both speculators must be of type L.

The strategy profile (13) represents a form of cooperation between both speculators such that whenever an L and an H-type meet the L-type would leave the whole amount of the asset to the H-type. If the H-type appreciates the asset more than the L-type and the partially revealing price system (here $\varepsilon_1 > 0$) is not too costly, there is scope for welfare improvement for both speculators compared to the 'perfect price collusion' BNE.

Proposition 3. *Fix the game Γ^{anyf} as defined above with s^{**} given by (13) and s^* given by (9).*

(i) *We have* $U_A(s^{**}) \geq U_A(s^*)$ *iff*

$$\begin{aligned} & \bar{v}\left(\pi_A(\bar{v}, H, L, \bar{z}) - \pi_A(\bar{v}, L, H, \bar{z})\right) \\ \geq\ & \varepsilon_1 \pi_A(\boldsymbol{\theta}_A = L, \boldsymbol{\theta}_B = L, \mathbf{z} = \bar{z}). \end{aligned}$$

(ii) *We have* $U_B(s^{**}) \geq U_B(s^*)$ *iff*

$$\begin{aligned} & \bar{v}\left(\pi_B(\bar{v}, L, H, \bar{z}) - \pi_B(\bar{v}, H, L, \bar{z})\right) \\ \geq\ & \varepsilon_1 \pi_B(\boldsymbol{\theta}_A = L, \boldsymbol{\theta}_B = L, \mathbf{z} = \bar{z}). \end{aligned}$$

(iii) *Suppose that the heterogeneous priors of both speculators satisfy*

$$\pi_A(\bar{v}, H, L, \bar{z}) > \pi_A(\bar{v}, L, H, \bar{z}) \text{ and } \pi_B(\bar{v}, H, L, \bar{z}) < \pi_B(\bar{v}, L, H, \bar{z}). \tag{14}$$

Then one can always find a sufficiently small price $\varepsilon_1 > 0$ *such that the strategy profile* s^{**} *is Pareto-superior over the 'perfect price collusion' BNE* s^**.*

Proposition 3 shows that risk-neutral speculators who agree to disagree through their heterogeneous beliefs (14) might want to coordinate on a Pareto-superior strategy profile s^{**} whenever the price ε_1 is sufficiently small. The problem is that in the absence of any commitment mechanism such coordination might not be incentive compatible, that is, a Pareto-superior strategy profile s^{**} might fail to be a BNE. In the above example, the speculators' information types $\theta_i = L$ might have a strict incentive to deviate from $\varphi_i^{**}[L](\varepsilon_0) = \{0\}$ given s_{-i}^{**}. To be precise, the L-type of speculator A would experience an upward-jump in his utility from demanding more than the zero amount at price ε_0 either when he finds himself in state $(\bar{v}, L, H, \bar{z})$ or when he encounters the L-type of his opponent B and there is strictly positive supply $\bar{z}$. In the former case, the L-type of A would obtain a greater amount, i.e., $\frac{\bar{z}}{2}$ instead of 0, at a better price, i.e., zero instead of ε_1. In the latter case, the L-type of A would pay zero for the amount $\bar{z}$ instead of paying ε_1 for the amount $\frac{\bar{z}}{2}$. This gives the L-type very strong incentives to deviate from s_A^{**}, which can only be counteracted by attaching zero probabilities to relevant events.

More generally, it is easy to see that the 'perfect price collusion' BNE action

$$\varphi_i^*[\theta_i](p) = [0, \infty) \text{ for all } p \in P$$

is for every information type θ_i a best response against any strategy profile s^{**}_{-i}. Namely, by choosing $\varphi^*_i[\theta_i]$ over $\varphi^{**}_i[\theta_i]$, information type θ_i will always achieve a zero price for a (weakly) larger amount of the asset.

Proposition 4. *Any Pareto-superior strategy profile s^{**}—given by (13)—is also a BNE of Γ^{anyf} iff the speculators' heterogeneous priors satisfy, in addition to (14),*

$$\begin{aligned}\pi_A(\bar{v}, L, H, \bar{z}) &= \pi_A(\bar{v}, L, L, \bar{z}) = \pi_A(0, L, L, \bar{z}) = 0 \text{ and}\\ \pi_B(\bar{v}, H, L, \bar{z}) &= \pi_B(\bar{v}, L, L, \bar{z}) = \pi_B(0, L, L, \bar{z}) = 0.\end{aligned}$$

For the above example, any Pareto-superior strategy profile s^{**} could thus only be established as a BNE of Γ^{anyf} under very strong (non-generic) assumptions on the speculators' beliefs. Moreover, even if the Pareto-superior strategy profile s^{**} is a BNE, the 'perfect price collusion' BNE strategy s^*_i would be a best response against s^{**}_{-i}. We take these findings as further evidence that 'perfect price collusion' BNE are a rather robust solution for the class of games considered in this paper.

4.2 Relationship to Kyle's (1989) original analysis

One of Kyle's (1989) main contribution is to identify conditions which ensure the existence of symmetric BNE given as linear demand-schedule functions[11]

$$\begin{aligned}s^K &= \left(\left(f^K_1[\theta_1]\right)_{\theta_1\in\Theta}, ..., \left(f^K_N[\theta_N]\right)_{\theta_N\in\Theta}\right)\\ &\quad \text{such that}\\ f^K_i[\theta_i](p) &= \mu + \beta[\theta_i] - \gamma p \text{ for } p \in \mathbb{R}.\end{aligned}$$

Such a linear-functions BNE has very attractive features: the information type θ_i who is more optimistic about the asset's value than the information type θ'_i shifts the intercept of his linear demand-schedule function upwards by $\beta[\theta_i] - \beta[\theta'_i] > 0$. In this specific sense,

[11] The other main contribution is the comparison of his 'partially revealing', 'imperfect competition' BNE with the fully revealing benchmark case of the rational expectations equilibrium in Hellwig (1980) which obtains if the number of (informed) speculators goes to infinity.

different information types move the equilibrium demand into the 'right' direction so that more positive signals come with a higher demand. Key to this celebrated result is Kyle's combination of a CARA Bernoulli function with normally distributed asset values. The resulting (conditional) expected utility of information type θ_i at any given price p

$$\begin{aligned} & \mathbb{E}_{\pi(\cdot|\theta_i)} u\left((\mathbf{v}-p)\, x_i\left[\theta_i\right]\right) \\ = \quad & -\frac{1}{\rho} \exp\left(\left(\mathbb{E}\left(\mathbf{v} \mid \theta_i\right)-p\right) x_i\left[\theta_i\right]-\frac{\rho}{2} var\left(\mathbf{v} \mid \theta_i\right) x_i^2\left[\theta_i\right]\right) \end{aligned}$$

is then maximized at the demand

$$x_i^*\left[\theta_i\right]=\frac{\mathbb{E}\left(\mathbf{v} \mid \theta_i\right)-p}{\rho \cdot var\left(\mathbf{v} \mid \theta_i\right)} \tag{15}$$

where ρ denotes the absolute risk aversion coefficient.

The problem with the linear demand-function $x_i^*\left[\theta_i\right](p)$ given by (15), however, is that it excludes all assets which can be owned by speculators with limited liability. According to (15), the information types in Kyle (1989) would only like to purchase some finite amount of the asset at a zero price because they do not like the risk of greater losses that would come with a greater amount of the asset. In contrast, the driving-force behind our 'perfect price collusion' BNE is the fact that all information types would like to purchase as much as possible of the asset at a zero price. This fact is an immediate implication of our assumption of non-negative asset values. We consider this desire to purchase as much of a valuable asset at a zero price as the *natural scenario* for any assets (e.g., shares, bonds, etc.) that come with limited liability for speculators.

Remark 4. Although the purpose of this paper is not an in-depth investigation of Kyle's (1989) original analysis within his normal distribution framework, we must admit that we are not sure in how far his first-order analysis adequately captures the possibility that speculators can choose arbitrary demand-schedule correspondences rather than just demand-schedule functions. For example, the characterization of unique market-clearing prices through equation (39) in Kyle (1989) seems to ignore the 'Any fraction' scenario as proclaimed by Kyle's (1989) own structural assumption K2. Under the 'Any fraction' scenario, market-clearing prices are often not unique so that Kyle's structural assumption K3—according to which the auctioneer chooses the minimal market-clearing price—would be needed to determine the resulting market price.[12] We did not find the possibility of multiple market-clearing prices reflected in Kyle's formal derivations. In a way it looks as if we are applying Kyle's own market microstructure assumptions more rigorously than he did himself.

5 Concluding remarks

We have translated Kyle's (1989) original market microstructure assumptions into a class of Bayesian games in which different information types of speculators can purchase a risky asset with non-negative values from a noise-trader. In contrast to the original analysis in Kyle (1989)—which combines CARA Bernoulli functions with normally distributed asset values—there is an unlimited demand for the asset in our model whenever the asset price is zero. As our main finding we have established the existence of 'perfect price collusion' BNE for two different scenarios of admissible demand schedules. Whereas all convex-valued demand-schedule correspondences are admissible under the 'Convexity' scenario, the 'Any fraction' scenario comes with the additional requirement that these convex-valued demand-schedule correspondences must also include zero as a possible demand.

We have argued that 'perfect price collusion' BNE are more plausible than any other BNE for the class of games considered in this paper. The main reason is that a 'perfect price collusion' BNE strategy is a best response against any strategy profile for which at least one opponent always wants to buy as much as possible of the asset at price zero. To expect other speculators to demand as much as possible of the asset at price zero, is a very natural assumption given that the asset has non-negative value for all information types. Since the 'perfect price collusion' BNE does not reveal any information, Kyle's (1989) market microstructure model does not seem to be well suited for generating (partially) information revealing Bayesian Nash equilibria if the asset can only have non-negative values.

[12] In our 'perfect price collusion' BNE all prices clear markets in every state of the world. Even for the Pareto-superior strategy (13) of our illustrative example all prices $p \geq \varepsilon_1$ clear markets in every state of the world.

Appendix: Formal proofs

Proof of Proposition 1. Fix s_{-i}^*. If there is no information type θ_i who has a strict incentive to deviate from the supposed BNE action $\varphi_i^*[\theta_i]$, the sufficiency condition (8) holds and s^* must be a BNE. Suppose that θ_i chooses action $\varphi_i'[\theta_i] \neq \varphi_i^*[\theta_i]$.

Step 1. Observe at first that under the 'Any-fraction' scenario the market price will always be zero irrespective of the choice of $\varphi_i'[\theta_i]$. To see this, note that the 'Any-fraction' scenario implies, by (2), for any action $\varphi_i'[\theta_i]$ that

$$\text{either } \varphi_i'[\theta_i](0) = [0, \overline{x}] \text{ or } \varphi_i[\theta_i](0) = [0, \infty)$$

for some $\overline{x} \geq 0$. Given s_{-i}^* with $\varphi_j^*[\theta_i](0) = [0, \infty)$ the markets will thus clear in every state at price 0 because the set (4) of potentially market-clearing demand offers at zero, given as

$$X(p^* = 0) = \left\{ (x_i, x_{-i}) \in \varphi_i'[\theta_i](0) \times \mathbb{R}_{\geq 0}^{N-1} \mid \sum_{j=1}^{N} x_j = z \right\},$$

is non-empty for every $z \geq 0$. Irrespective of any other existing market-clearing prices, zero is always the minimal market-clearing price so that we obtain for all states

$$\mathbf{p}^*\left[\varphi_i'[\theta_i], s_{-i}^*\right](v, \theta_1, ..., \theta_N, z) = 0$$

for any choice $\varphi_i'[\theta_i] \neq \varphi_i^*[\theta_i]$.

Step 2. Because we can, by Step 1, fix zero as market price, it is sufficient to consider some action $\varphi_i'[\theta_i] \neq \varphi_i^*[\theta_i]$ such that $\varphi_i'[\theta_i](0) = [0, \bar{x}] \neq [0, \infty)$ for some $\bar{x} \geq 0$. What would be the optimal interval $[0, \bar{x}]$ from the perspective of information type θ_i? Since the Bernoulli utility

$$u_i\left((v - 0)\, \mathbf{x}_i^*\left[[0, \bar{x}], s_{-i}^*\right](\cdot)\right)$$

increases in $\mathbf{x}_i^*\left[[0, \bar{x}], s_{-i}^*\right]$ for all states with $v, z > 0$ (and is constant else), θ_i wants to choose some $[0, \bar{x}]$ that maximizes $\mathbf{x}_i^*\left[[0, \bar{x}], s_{-i}^*\right](\cdot)$ for all states. But this can be achieved by choosing

$$\varphi_i'[\theta_i](0) = [0, \infty). \tag{16}$$

The reason is as follows. Suppose that θ_i picks some $\varphi_i'[\theta_i](0) = [0, \bar{x}]$ such that $\bar{x} < \frac{z}{N}$ in some state $(v, \theta_1, ..., \theta_N, z)$. In this state his market allocation becomes

$$\mathbf{x}_i^*\left[\varphi_i'[\theta_i], s_{-i}^*\right](v, \theta_1, ..., \theta_N, z) = \bar{x}$$

which is strictly less than if he had chosen $[0, \infty)$ because of

$$\mathbf{x}_i^* \left[[0, \infty), s_{-i}^* \right] (v, \theta_1, ..., \theta_N, z) = \frac{z}{N}.$$

If he picks instead some $\varphi_i' [\theta_i] (0) = [0, \bar{x}]$ with $\bar{x} \geq \frac{z}{N}$, he will obtain the same allocation of $\frac{z}{N}$ as if he had picked $\varphi_i' [\theta_i] (0) = [0, \infty)$. In other words, choosing (16) weakly dominates any bounded interval $[0, \bar{x}]$ as demand choice at price 0 for all states. This proves that θ_i has no strict incentive to deviate from the BNE action (10) provided that the other speculators stick to their BNE strategies s_{-i}^*.□□

Proof of Proposition 2. In contrast to the 'Any-fraction' scenario of Proposition 1, the information type θ_i has now the power to shift the market price away from zero even if the other speculators stick with s_{-i}^*. To see this, note that under the 'Convexity' scenario, any actions $\varphi_i' [\theta_i]$ such that

$$\varphi_i' [\theta_i] (0) = [\underline{x}, \overline{x}] \text{ or } \varphi_i' [\theta_i] (0) = [\underline{x}, \infty)$$

for some $\underline{x} > 0$ becomes now admissible. In that case, markets would no longer clear at $p = 0$ in any state $(v, \theta_1, ..., \theta_N, z)$ such that $\underline{x} > z$. The reason is that θ_i alone would already demand a strictly higher quantity through $\underline{x}$ than the quantity z supplied by the noise-trader.

To investigate whether θ_i has a strict incentive to choose $\varphi_i' [\theta_i]$ over $\varphi_i^* [\theta_i]$, we have thus to distinguish between two cases. In the first case, θ_i chooses some $\varphi_i' [\theta_i]$ such that the price $p = 0$ remains the market price in all states. In the second case, θ_i chooses some $\varphi_i' [\theta_i]$ such that the price $p = 0$ does no longer clear the market in all states.

Case 1. Suppose that θ_i chooses some $\varphi_i' [\theta_i] \neq [0, \infty)$ such that the price $p = 0$ remains the market price in all states $(v, \theta_1, ..., \theta_N, z)$. In this case, the proof becomes identical to the proof of Proposition 1.

To conclude Case 1: whenever θ_i does not choose any action which moves the market price away from zero, he cannot strictly improve upon the supposed BNE action $\varphi_i^* [\theta_i]$.

Case 2. Suppose now that θ_i chooses some $\varphi_i' [\theta_i] \neq [0, \infty)$ such that the price $p = 0$ is no longer a market-clearing price in all states $(v, \theta_1, ..., \theta_N, z)$. Since θ_i wants the markets to clear in every state in order to avoid infinitely negative utility, he will determine a new minimal market clearing price $p' > 0$ through his choice of $\varphi_i' [\theta_i]$. Such p' would be the smallest price at which markets clear in every state given the fixed profile s^*_{-i} which implies $\varphi^*_{-i}[\theta_{-i}](p') = \mathbb{R}^{N-1}_{\geq 0}$. Under our assumption that $0 \in Z$, this p' is uniquely pinned down as the smallest admissible price p' such that the demand-schedule $\varphi'_i[\theta_i](p')$ contains 0 as possible demand-offer. Formally, we have that

$$p' = \min\{p \in P \mid \min \varphi'_i[\theta_i](p) = 0\}. \tag{17}$$

To see this, note that

$$\varphi'_i[\theta_i](p) = [\underline{x}, \overline{x}] \text{ or } \varphi'_i[\theta_i](p) = [\underline{x}, \infty)$$

with $\underline{x} > 0$ implies that markets cannot clear at price p in any state $(v, \theta_1, ..., \theta_N, z)$ such that $z = 0$. On the other hand, given $\varphi^*_{-i}[\theta_{-i}](p) = \mathbb{R}^{N-1}_{\geq 0}$ for all $p \in P$ markets would clear in every state at any price p satisfying

$$\text{either } \varphi'_i[\theta_i](p) = [0, \overline{x}] \text{ or } \varphi'_i[\theta_i](p) = [0, \infty).$$

The price p' defined by (17) is the smallest of these market-clearing prices and therefore the market price that will be picked by the auctioneer in accordance with Assumption K3.

Choosing any new market price $p' \in P$ is thus in the power of θ_i under the 'Convexity' scenario. The question is, does θ_i have any strict incentive to move the market price away from 0 to some $p' > 0$? The answer is no. Whereas he received in the supposed BNE s^* the Bernoulli utility

$$u_i\left((v - 0)\frac{z}{N}\right) \tag{18}$$

in every state $(\cdot, z)$ of the world, he now receives

$$u_i\left((v - p')\,\mathbf{x}^*_i\left[\varphi'_i[\theta_i], s^*_{-i}\right](\cdot, z)\right). \tag{19}$$

The definition of the new market price p' through (17) implies, by convex-valuedness, that

$$\text{either } \varphi'_i[\theta_i](p') = [0, \overline{x}] \text{ or } \varphi'_i[\theta_i](p') = [0, \infty)$$

for some $\overline{x} \geq 0$. Consequently, (18) clearly dominates (19) in every state of the world because of

$$(v - p')\,\mathbf{x}^*_i\left[\varphi'_i[\theta_i], s^*_{-i}\right] \leq 0 \leq (v - 0)\frac{z}{N} \text{ if } v \leq p' \text{ or } z = 0$$

and

$$(v - p')\,\mathbf{x}_i^*\left[\varphi_i'\left[\theta_i\right], s_{-i}^*\right] < (v - 0)\,\frac{z}{N} \text{ if } v > p' \text{ and } z > 0$$

whereby the last inequality follows from $\mathbf{x}_i^*\left[\varphi_i'\left[\theta_i\right], s_{-i}^*\right] \leq \frac{z}{N}$ since

$$\begin{aligned}
\mathbf{x}_i^*\left[[0,\overline{x}], s_{-i}^*\right] &< \frac{z}{N} \text{ if } \overline{x} < \frac{z}{N}, \\
\mathbf{x}_i^*\left[[0,\overline{x}], s_{-i}^*\right] &= \frac{z}{N} \text{ if } \overline{x} \geq \frac{z}{N}.
\end{aligned}$$

To conclude Case 2: whenever θ_i chooses some action against s_{-i}^* that moves the market price away from zero, he cannot strictly improve upon the supposed BNE action $\varphi_i^*\left[\theta_i\right]$.□□

Proof of Proposition 3. Note that

$$\begin{aligned}
U_i\left(s^{**}\right) &\geq U_i\left(s^*\right) \\
&\Leftrightarrow
\end{aligned}$$

$$\begin{aligned}
&\mathbb{E}_{\pi_i}\left(\left(\mathbf{v} - \mathbf{p}^*\left[s_i^{**}, s_{-i}^{**}\right]\right)\mathbf{x}_i^*\left[s_i^{**}, s_{-i}^{**}\right]\right) \\
\geq\; &\mathbb{E}_{\pi_i}\left(\left(\mathbf{v} - \mathbf{p}^*\left[s_i^{*}, s_{-i}^{*}\right]\right)\mathbf{x}_i^*\left[s_i^{*}, s_{-i}^{*}\right]\right).
\end{aligned}$$

For $i = A$, we have

$$\begin{aligned}
&\mathbb{E}_{\pi_A}\left(\left(\mathbf{v} - \mathbf{p}^*\left[s_A^{**}, s_B^{**}\right]\right)\mathbf{x}_A^*\left[s_A^{**}, s_B^{**}\right]\right) \\
=\; &(\bar{v} - \varepsilon_0)\,\bar{z}\pi_A\left(\bar{v}, H, L, \bar{z}\right) + (0 - \varepsilon_0)\,\bar{z}\pi_A\left(0, H, L, \bar{z}\right) \\
&+ (\bar{v} - \varepsilon_0)\,\frac{\bar{z}}{2}\pi_A\left(\bar{v}, H, H, \bar{z}\right) + (0 - \varepsilon_0)\,\frac{\bar{z}}{2}\pi_A\left(0, H, H, \bar{z}\right) \\
&+ (\bar{v} - \varepsilon_1)\,\frac{\bar{z}}{2}\pi_A\left(\bar{v}, L, L, \bar{z}\right) + (0 - \varepsilon_1)\,\frac{\bar{z}}{2}\pi_A\left(0, L, L, \bar{z}\right)
\end{aligned}$$

as well as

$$\begin{aligned}
&\mathbb{E}_{\pi_A}\left(\left(\mathbf{v} - \mathbf{p}^*\left[s_A^{*}, s_B^{*}\right]\right)\mathbf{x}_A^*\left[s_A^{*}, s_B^{*}\right]\right) \\
=\; &\bar{v}\frac{\bar{z}}{2}\left(\pi_A\left(\bar{v}, H, L, \bar{z}\right) + \pi_A\left(\bar{v}, L, H, \bar{z}\right) + \pi_A\left(\bar{v}, H, H, \bar{z}\right) + \pi_A\left(\bar{v}, L, L, \bar{z}\right)\right).
\end{aligned}$$

Using $\varepsilon_0 = 0$ gives after straightforward calculations

$$\begin{aligned}
U_A\left(s^{**}\right) &\geq U_A\left(s^*\right) \\
&\Leftrightarrow
\end{aligned}$$

$$
\begin{aligned}
& \bar{v}\left(\pi_A\left(\bar{v}, H, L, \bar{z}\right)-\pi_A\left(\bar{v}, L, H, \bar{z}\right)\right) \\
\geq\ & \varepsilon_1\left(\pi_A\left(\bar{v}, L, L, \bar{z}\right)+\pi_A\left(0, L, L, \bar{z}\right)\right) \\
=\ & \varepsilon_1 \pi_A\left(\boldsymbol{\theta}_A=L, \boldsymbol{\theta}_B=L, \mathbf{z}=\bar{z}\right)
\end{aligned}
$$

For $i=B$, we obtain analogously

$$
U_B\left(s^{**}\right) \quad \geq \quad U_B\left(s^*\right)
$$

$$
\Leftrightarrow
$$

$$
\begin{aligned}
& \bar{v}\left(\pi_B\left(\bar{v}, L, H, \bar{z}\right)-\pi_B\left(\bar{v}, H, L, \bar{z}\right)\right) \\
\geq\ & \varepsilon_1\left(\pi_B\left(\bar{v}, L, L, \bar{z}\right)+\pi_B\left(0, L, L, \bar{z}\right)\right) \\
=\ & \varepsilon_1 \pi_B\left(\boldsymbol{\theta}_A=L, \boldsymbol{\theta}_B=L, \mathbf{z}=\bar{z}\right).
\end{aligned}
$$

□□

Proof of Proposition 4. Since the 'perfect price collusion' strategy s_i^* is a best response against s_{-i}^{**}, we have that s^{**} is a BNE iff, for $i \in \{A, B\}$,

$$
\begin{aligned}
& \mathbb{E}_{\pi_i}\left(\left(\mathbf{v}-\mathbf{p}^*\left[s_i^{**}, s_{-i}^{**}\right]\right) \mathbf{x}_i^*\left[s_i^{**}, s_{-i}^{**}\right]\right) \\
=\ & \mathbb{E}_{\pi_i}\left(\left(\mathbf{v}-\mathbf{p}^*\left[s_i^{*}, s_{-i}^{**}\right]\right) \mathbf{x}_i^*\left[s_i^{*}, s_{-i}^{**}\right]\right).
\end{aligned} \tag{20}
$$

Observe that the market price and allocations that correspond to (s_A^*, s_B^{**}) in state $(v, \theta_A, \theta_B, z) \in \Omega$ are given as

$$
\begin{aligned}
\mathbf{p}^*\left[s_A^*, s_B^{**}\right](v, \theta_A, \theta_B, z) &= \varepsilon_0 \\
\mathbf{x}_A^*\left[s_A^*, s_B^{**}\right](v, \theta_A, \theta_B, z) &= \begin{cases} z & \text{if } \theta_B = L \\ \frac{z}{2} & \text{else} \end{cases} \\
\mathbf{x}_B^*\left[s_A^*, s_B^{**}\right](v, \theta_A, \theta_B, z) &= \begin{cases} \frac{z}{2} & \text{if } \theta_B = H \\ 0 & \text{else} \end{cases}
\end{aligned}
$$

For $i=A$, we therefore have

$$
\begin{aligned}
& \mathbb{E}_{\pi_A}\left(\left(\mathbf{v}-\mathbf{p}^*\left[s_A^{**}, s_B^{**}\right]\right) \mathbf{x}_A^*\left[s_A^{**}, s_B^{**}\right]\right) \\
=\ & \left(\bar{v}-\varepsilon_0\right) \bar{z} \pi_A\left(\bar{v}, H, L, \bar{z}\right)+\left(0-\varepsilon_0\right) \bar{z} \pi_A\left(0, H, L, \bar{z}\right) \\
& +\left(\bar{v}-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(\bar{v}, H, H, \bar{z}\right)+\left(0-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(0, H, H, \bar{z}\right) \\
& +\left(\bar{v}-\varepsilon_1\right) \frac{\bar{z}}{2} \pi_A\left(\bar{v}, L, L, \bar{z}\right)+\left(0-\varepsilon_1\right) \frac{\bar{z}}{2} \pi_A\left(0, L, L, \bar{z}\right)
\end{aligned}
$$

as well as

$$\begin{aligned}
& \mathbb{E}_{\pi_A}\left(\left(\mathbf{v}-\mathbf{p}^*\left[s_A^*, s_B^{**}\right]\right) \mathbf{x}_A^*\left[s_A^*, s_B^{**}\right]\right) \\
= & \left(\bar{v}-\varepsilon_0\right) \bar{z} \pi_A\left(\bar{v}, H, L, \bar{z}\right)+\left(0-\varepsilon_0\right) \bar{z} \pi_A\left(0, H, L, \bar{z}\right) \\
& +\left(\bar{v}-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(\bar{v}, L, H, \bar{z}\right)+\left(0-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(0, L, H, \bar{z}\right) \\
& +\left(\bar{v}-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(\bar{v}, H, H, \bar{z}\right)+\left(0-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(0, H, H, \bar{z}\right) \\
& +\left(\bar{v}-\varepsilon_0\right) \bar{z} \pi_A\left(\bar{v}, L, L, \bar{z}\right)+\left(0-\varepsilon_0\right) \bar{z} \pi_A\left(0, L, L, \bar{z}\right).
\end{aligned}$$

Straightforward calculations show that equality (20) is equivalent to

$$\begin{aligned}
& \left(\bar{v}-\varepsilon_1\right) \frac{\bar{z}}{2} \pi_A\left(\bar{v}, L, L, \bar{z}\right)+\left(0-\varepsilon_1\right) \frac{\bar{z}}{2} \pi_A\left(0, L, L, \bar{z}\right) \\
= & \left(\bar{v}-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(\bar{v}, L, H, \bar{z}\right)+\left(0-\varepsilon_0\right) \frac{\bar{z}}{2} \pi_A\left(0, L, H, \bar{z}\right) \\
& +\left(\bar{v}-\varepsilon_0\right) \bar{z} \pi_A\left(\bar{v}, L, L, \bar{z}\right)+\left(0-\varepsilon_0\right) \bar{z} \pi_A\left(0, L, L, \bar{z}\right)
\end{aligned}$$

$$\Leftrightarrow$$

$$\begin{aligned}
& -\varepsilon_1\left(\pi_A\left(\bar{v}, L, L, \bar{z}\right)+\pi_A\left(0, L, L, \bar{z}\right)\right) \\
= & \bar{v}\left(\pi_A\left(\bar{v}, L, H, \bar{z}\right)+\pi_A\left(\bar{v}, L, L, \bar{z}\right)\right)
\end{aligned}$$

which requires

$$\pi_A\left(\bar{v}, L, H, \bar{z}\right)=\pi_A\left(\bar{v}, L, L, \bar{z}\right)=\pi_A\left(0, L, L, \bar{z}\right)=0$$

to hold because of $\varepsilon_1, \bar{v}>0$.

For $i=B$, we obtain analogously

$$\pi_B\left(\bar{v}, H, L, \bar{z}\right)=\pi_B\left(\bar{v}, L, L, \bar{z}\right)=\pi_B\left(0, L, L, \bar{z}\right)=0.$$

□□

References

Akerlof, G.A. (1970) "The market for "lemons": quality uncertainty and the market mechanism," *Quarterly Journal of Economics* **84**, 88-100.

Allen, B. (1981) "Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information," *Econometrica* **49**, 1173-1199.

Aumann, R.J. (1998) "Common priors: a reply to Gul," *Econometrica* **66**, 929-933.

Billingsley, P. (1995) *Probability and measure*, John Wiley and & Sons, Inc.: New York.

Easley, D. and M. O'Hara (1995) "Market microstructure," Chapter 12 in R. Jarrow et al. (Eds.), *Handbooks in Operations Research and Management Science*, vol. **9**, Elsevier, 357-383.

Green, J.R. (1975) "Information, equilibrium, efficiency," mimeo.

Grossman, S.J. (1976) "On the efficiency of competitive stock markets where traders have diverse information," *Journal of Financial* **31**, 573-585.

Grossman, S.J. (1977) "Existence of futures markets, noisy rational expectations and informational externalities," *Review of Economic Studies* **44**, 431-449.

Grossman, S.J. (1978) "Further results on the informational efficiency of competitive stock markets," *Journal of Economic Theory* **18**, 81-101.

Gul, F. (1998) "A comment on Aumann's Bayesian view," *Econometrica* **66**, 923-927.

Hellwig, M.R. (1980) "On the aggregation of information in competitive markets," *Journal of Economic Theory* **22**, 477-498.

Kreps, D. (1977) "A note on 'fulfilled expectations' equilibria," *Journal of Economic Theory* **14**, 32-43.

Kyle, A.S. (1989) "Informed speculation with imperfect competition," *Review of Economic Studies* **56**, 317-355.

Radner, R. (1979) "Rational expectations equilibrium: Generic existence and the information revealed by prices," *Econometrica* **47**, 655-678.