

Optimality in a large two-stage market game*

Tai-Wei Hu[†] and Neil Wallace[‡]

June 21, 2012

Abstract

We study a market-game mechanism in a setting with potential information aggregation. The mechanism has two stages and is inspired by pari-mutuel betting. The second stage resembles the determination of final odds and payoffs, while the first stage is a version of the announced running bet totals and odds. In a model that is essentially a finite-state, divisible-goods version of the private/common values model of Reny-Perry (2006), we show existence of an equilibrium that is almost ex post efficient when the number of players is finite, but large.

1 Introduction

We study strategic trade in a setting with potential information aggregation: there is an unobserved state-of-the-world and there is dispersed and incomplete information about that state in the form of private signals. Our model of strategic trade resembles a market. There are two main approaches to modeling trade in that way. One uses a double-auction mechanism (see, in particular, Reny and Perry [5], Cripps and Swinkels [1], and Vives [8]). The other uses a Cournot quantity mechanism (see Palfrey [3] and Vives [7]) or some version of a Shapley-Shubik [6] market game. We use a market-game here. Our setting is essentially a divisible-good generalization of that in Reny and Perry [5]. In particular, we follow them in assuming that utility depends on the unobserved state-of-the-world and on the private signal that each person receives.

Our mechanism has two stages and is inspired by pari-mutuel betting—perhaps, the most significant actual use of a market game. The second stage resembles the determination of final odds and payoffs, while the first stage is a version of the announced running bet totals and odds. The mechanism also resembles some of those used in experiments that are devoted to information aggregation (see, for example, Axelrod *et al* [4])—mechanisms

*We are indebted to Jim Jordan for helpful discussions when we started this work. This is a preliminary version. Please do not cite without permission of the authors.

[†]MEDS, Northwestern, <t-hu@kellogg.northwestern.edu>

[‡]Penn State, <neilw@psu.edu>

which have not been analyzed theoretically. Prior theoretical work in private-information settings does not use such mechanisms and does not achieve ex post optimality in anything like our structure. Palfrey [3] uses a Cournot mechanism and obtains ex post optimality, but, as Vives [7] points out, only because marginal cost is common and constant so that it does not matter how production is allocated among the firms in the model.

A crucial aspect of our mechanism is the way it elicits information at stage-1. At stage-1, each agent names an offer. Then, in a random way, the agents are divided into two groups: a small *inactive* group and a large *active* group. For the inactive group, there is no further participation; their stage-1 offers are executed at an exogenous price. For the active group, after the histogram of their stage-1 offers is announced, they participate in a stage-2 market game—a version that uses the refinement in Dubey and Shubik [2] in order to eliminate no-trade as an equilibrium. The threat of being in the inactive group induces agents to be serious about their stage-1 offers.

Our results are about existence and ex post efficiency. Under a mild genericity condition, if the number of agents is sufficiently large, then there exists a symmetric equilibrium in pure strategies in which stage-1 actions reveal the private-information held by all active agents. Moreover, the stage-2 equilibrium outcome is almost ex post efficient. And, if there exist other equilibria, then they, too, have almost efficient ex post stage-2 outcomes.

A single-stage, double-auction mechanism could possibly achieve something like our asymptotic existence and efficiency results. Vives [8] gets such results and explicit rates of convergence in a model with linear-quadratic payoffs and normally distributed shocks. However, in order to get ex post efficiency, both in his setting and, therefore, in more general settings, actions have to be (supply) functions. In contrast, actions in our mechanism are simpler; they are quantities. Relative to a direct mechanism, a virtue of our mechanism is that it works in any environment in our class; the mechanism does not rely on the detailed structure of the economy.

2 The model and the mechanism

We describe, in turn, the environment, our mechanism, and the equilibrium concept.

2.1 Environment

Our economy is an endowment economy with two goods and N agents. (The set of agents is denoted \mathcal{N} .) Each agent is assigned a type, denoted x , where $x \in X$, a finite set. An agent of type $x \in X$ maximizes expected utility with ex post utility function, $u(q, r; x, z)$, where $(q, r) \in \mathbb{R}_+^2$ is the vector of quantities of the two goods consumed and $z \in Z$, a finite set, is a state-of-the-world. The function $u(\cdot, \cdot; x, z)$ is strictly increasing, strictly concave, continuously differentiable with partial derivatives u_q and u_r , and satisfies Inada conditions. For simplicity, each agent is endowed with the per capita endowment of each

good, denoted \bar{q} and \bar{r} , respectively.¹ Finally, we assume that $u(\cdot, \cdot; x, z)$ is such that the implied complete-information competitive demands are monotone.² After presenting our analysis, we comment in detail on the role of this monotonicity assumption.

The sequence of events is as follows. First, nature draws a state-of-the-world $z \in Z$ with probability $\kappa(z)$, a state which no one observes. Then each agent gets a type realization, $x \in X$, which is private to the agent. Conditional on the realization z , these realizations are *i.i.d.* across people. We denote the conditional probability $\mathbb{P}[x | z]$ by $\mu_z(x)$, and denote the implied posterior probability $\mathbb{P}[z | x]$ by $\tau_x(z)$. We assume that x is informative in the sense that $z \neq z'$ implies $\mu_z(x) \neq \mu_{z'}(x)$ for some $x \in X$. This informativeness assumption is without loss of generality: If $\mu_z(x) = \mu_{z'}(x)$ for all $x \in X$, then we treat z and z' as a single state z'' with utility $u(q, r; x, z'') = P[z]u(q, r; x, z) + P[z']u(q, r; x, z')$. Our interpretation is that x is an idiosyncratic taste shock and that z is a common taste shock. Notice that the realized type, x , plays two roles: it serves as private information about z and it is private information about preferences.³

2.2 The mechanism

After the type realizations, each agent n chooses an offer $a^n = (a_q^n, a_r^n) \in \mathcal{O}$, where $\mathcal{O} = \{(o_q, o_r) \in [0, \bar{q}] \times [0, \bar{r}] : o_q o_r = 0\}$. Then agents are randomly divided into two groups in the following way. Let $\eta \in (0, 1)$ be a (small) rational number and let $\lceil (1 - \eta)N \rceil = M$ denote the smallest integer that is no less than $(1 - \eta)N$. An assignment $\pi : \mathcal{N} \rightarrow \{1, \dots, N\}$ is drawn from the uniform distribution over the set of all such assignments, and agent n is called *active* if $\pi(n) \leq M$ and is called *inactive* if $\pi(n) > M$. The payoff for each inactive agent is given by trade at the fixed price, $p_1 = \bar{r}/\bar{q}$. That is,

$$(q^n, r^n) = \left(\bar{q} - a_q^n + \frac{a_r^n}{p_1}, \bar{r} - a_r^n + p_1 a_q^n \right) \text{ for } n \notin \mathcal{M}, \quad (1)$$

where \mathcal{M} is the set of active agents. Next, the mechanism announces the histogram of the stage-1 offers of the active agents, denoted $\nu : \mathcal{O} \rightarrow \{1, 2, \dots, M\}$.⁴ Then, given that information, stage-2 has active agents participating in a market game. Each active agent n makes an offer $b^n = (b_q^n, b_r^n) \in \mathcal{O}$ and gets payoff

$$(q^n, r^n) = \left(\bar{q} - b_q^n + \frac{b_r^n}{p_2}, \bar{r} - b_r^n + p_2 b_q^n \right) \text{ for } n \in \mathcal{M}, \quad (2)$$

where $p_2 = R/Q$ and

¹What really matters is common knowledge about individual endowments.

²A gross-substitutes assumption about $u(\cdot, \cdot; x, z)$ is sufficient for such monotonicity.

³Of course, we could have formulated the types x as $x = (x^t, x^s)$, where x^t affects the utility function and x^s is a signal about z . However, this formulation is equivalent to ours and only complicates the notation.

⁴We could let the mechanism announce two histograms, one for active agents and one for inactive agents. However, that would complicate the notation and would not change the results.

$$(Q, R) = M\varepsilon + \sum_{n \in \mathcal{M}} b^n. \quad (3)$$

Here, $\varepsilon = (\varepsilon_q, \varepsilon_r)$, where $\varepsilon_q > 0$ and $\varepsilon_r > 0$ are exogenous (small) quantities that prevent no-trade from being an equilibrium, a formulation borrowed from Dubey and Shubik [2]. Notice that p_2 functions as a “price,” but a price that depends on the offer of each active agent.

As it stands, this mechanism violates feasibility. The trades of inactive agents at the fixed price do not clear that market. In addition, resources are required for the positive ε_q and ε_r . We proceed as if the mechanism designer has the resources required to support this mechanism. At the end, we suggest that small entry fees could be used to provide those resources. In any case, the departure from feasibility (and *balancedness*) can be made arbitrarily small in percapita terms.

The restriction in \mathcal{O} that agents can only make offers on one side of the market plays a significant role in our analysis. It is used to obtain the uniqueness of best responses. The following lemma shows that the restriction is not binding on the agent.⁵

Lemma 1. Fix the stage-2 offers of all other agents. If $b' \in [0, \bar{q}] \times [0, \bar{r}]$ with payoff (q', r') is such that $b'_q b'_r > 0$, then there exists $b'' \in \mathcal{O}$ with payoff (q', r') .

Obviously, the restriction is not binding in the same sense on payoffs for inactive agents.

2.3 Strategies, beliefs, and equilibrium

Now we formulate strategies and beliefs. A stage-1 strategy is $s_1^n(x) \in \mathcal{O}$, while a stage-2 strategy is $s_2^n(x, a, \nu^{-a}) \in \mathcal{O}$, where the second component in the domain is the agent’s stage-1 action, and the third is the announced histogram of offers of active agents *net of the agent’s own action*. (That is, for any $a' \in \mathcal{O}$, $\nu^{-a}(a') = \nu(a')$ if $a \neq a'$ and $\nu^{-a}(a) = \nu(a) - 1$.) A strategy profile $\{(s_1^n, s_2^n) : n \in \mathcal{N}\}$ is a Perfect Bayesian Equilibrium (PBE) if for each $n \in \mathcal{N}$, s_1^n is a best response to $\{(s_1^{n'}, s_2^{n'}) : n' \neq n\}$ and s_2^n is a best response to $\{s_2^{n'} : n' \neq n\}$ with respect to a belief φ^n that is consistent with Bayes’ rule whenever possible.

Throughout the paper, we focus exclusively on symmetric equilibrium in pure strategies. A PBE $\{(s_1^n, s_2^n) : n \in \mathcal{N}\}$ is a *symmetric equilibrium* if for all $n \in \mathcal{N}$, $(s_1^n, s_2^n) = (s_1, s_2)$ and $\varphi^n = \varphi$. In a symmetric equilibrium, an agent’s expected payoff at stage-2 depends only on his private history (x, a) and the configuration of other active agents’ private histories $\theta : X \times \mathcal{O} \rightarrow \{0, 1, 2, \dots, M - 1\}$. Thus, we may formulate the belief $\varphi(x, a, \nu^{-a})$ as an element of $\Delta(Z \times \Theta)$, where Θ is the set of all configurations θ of type/stage-1-action of the other active agents. Then, a symmetric equilibrium is a triple (s_1, s_2, φ) such that (a) $s_1(x)$ is a best response to s_1 and s_2 ; (b) $s_2(x, a, \nu^{-a})$ is a best response to s_2 and $\varphi(x, a, \nu^{-a})$; (c) $\varphi(x, a, \nu^{-a})$ is derived from equilibrium behavior using Bayes’ rule whenever possible.

⁵All the proofs appear in section 4.

It is also useful to define a *separating* equilibrium. A separating equilibrium is a symmetric equilibrium in which $x \neq x'$ implies $s_1(x) \neq s_1(x')$. In a separating equilibrium, the belief on the equilibrium path does not depend on an agent's private history. In particular, in such an equilibrium, ν^{-a} and the agent's own type reveal the true configuration of types of the active agents and each agent uses that true configuration and Bayes' rule to form a common posterior over Z . Thus, on the equilibrium path in a separating equilibrium, while the agent's private type matters for the agent's preferences at stage 2, all active agents share the same information at that stage regardless of their private types.

We show that a separating equilibrium exists generically for sufficiently large N and that the equilibrium outcome is almost ex post efficient. We establish existence by demonstrating that it is optimal for an agent at stage 1 to choose an action that is best contingent on being inactive when others do so. The main ingredient in that argument is the off-equilibrium belief formulation. The genericity qualification is very simple: it requires that the ratios of marginal utilities at the optimal consumption levels under the fixed price p_1 differ across types. As regards ex post efficiency, we show that the limit of stage-2 actions in a separating equilibrium is close to the unique complete-information competitive equilibrium. We get almost ex post efficiency instead of efficiency for two reasons. First, we get a competitive outcome only in the limit as $N \rightarrow \infty$. Second, that competitive limit is for the economy of active agents that has exogenous per capita trades ε . Because, we can let η and ε approach zero, we have almost efficiency. Finally, under a stronger version of informativeness and a minor modification of the mechanism, we show that any equilibrium outcome in symmetric pure strategies is also almost ex post efficient.

3 Separating equilibrium: existence, characterization, and uniqueness

Our existence proof is partly constructive. Therefore, it is helpful to begin with existence and characterization of stage-2 when stage-1 is separating.

3.1 Second-stage equilibrium when stage-1 is separating

In a separating equilibrium (s_1, s_2) , the belief φ about the type/stage-1-action configuration is degenerate on the configuration θ given by $\theta(x, s_1(x)) = \nu^{-a}(s_1(x))$. This implies that there is common knowledge at stage 2 about the type-configuration of active agents, a configuration we denote $\sigma : X \rightarrow \{0, 1, \dots, M\}$, where M is the number of active agents.⁶ It also implies a common posterior over Z , denoted ϕ , which is derived from the type-configuration σ via Bayes' rule.

Therefore, stage-2 in a separating equilibrium only depends on the type-configuration σ . In fact, the stage-2 game along such an equilibrium path can be regarded as a Bayesian

⁶In what follows, and only to simplify the notation, we assume that $\sigma(x) > 0$ for all $x \in X$. Obviously, this holds with probability one for all sufficiently large N .

game defined as follows. Given a type-configuration σ , the stage-2 game on such an equilibrium path is a Bayesian game in which: (a) the players are the active agents; (b) the action set for each player consists of offers $(b_q, b_r) \in \mathcal{O}$; (c) the payoffs are determined by the market game and u ; (d) the number of players of type x is $\sigma(x)$, which is common knowledge among the players; (e) the common prior over Z is given by $\phi^\sigma \in \Delta(Z)$ that is derived from σ .⁷ We denote an equilibrium for this game by $\beta^\sigma : X \rightarrow \mathcal{O}$.

Then we have

Proposition 1. For any type-configuration $\sigma : X \rightarrow \{0, 1, \dots, M\}$, β^σ exists.

The proof is a routine application of Brouwer's fixed point theorem. It does, however, depend crucially on the constraint $b_q b_r = 0$. With it, the best response, which is the mapping studied in order to get a fixed point, is a function; without that constraint, the mapping is a nonconvex correspondence. It also makes use of $\varepsilon > 0$, which gives both upper and lower bounds on p_2 .

Although there is no uniqueness claim in Proposition 1, we can characterize the limit of β^σ as $N \rightarrow \infty$ for any suitable sequence of type-configurations σ^N . Our characterization result is based on the following implication of our informativeness assumption. Fix $z \in Z$ and let σ^N be the type configuration of active agents for an economy of size N . If the sequence $\{\sigma^N\}_{N=1}^\infty$ is such that $\lim_{N \rightarrow \infty} \sigma^N(x)/[(1-\eta)N] = \mu_z(x)$ for each $x \in X$ (which holds *almost surely* conditional on z), then $\lim_{N \rightarrow \infty} \phi^{\sigma^N}(z) = 1$. It also makes use of the following candidate limit for a stage-2 equilibrium. Let $\mathcal{L}^z(\varepsilon)$ denote the economy with z known, with a nonatomic measure of agents, with fraction of type- x agents equal to $\mu_z(x)$, and with exogenous per capita trades ε . A competitive equilibrium in $\mathcal{L}^z(\varepsilon)$ is a tuple $\{p^z, (q^z(x), r^z(x))_{x \in X}\}$ such that $(q^z(x), r^z(x))$ maximizes $u(q, r; x, z)$ subject to $p q + r = p \bar{q} + \bar{r}$ for each $x \in X$ and

$$\frac{\varepsilon_r}{p} + \sum_{x \in X} \mu_z(x) q(x) = \bar{q} + \varepsilon_q.$$

The following lemma shows the equivalence between the competitive equilibrium in $\mathcal{L}^z(\varepsilon)$ and the Nash equilibrium of the market game in $\mathcal{L}^z(\varepsilon)$.

Lemma 2. The economy $\mathcal{L}^z(\varepsilon)$ has a unique competitive equilibrium and it is continuous in ε . Moreover, $(q^z(x), r^z(x))_{x \in X}$ is a competitive allocation for $\mathcal{L}^z(\varepsilon)$ if and only if

$$\beta_q^z(x) = \max\{\bar{q} - q^z(x), 0\} \text{ and } \beta_r^z(x) = \max\{\bar{r} - r^z(x), 0\}$$

is a Nash equilibrium for the market game for $\mathcal{L}^z(\varepsilon)$.

This is one of the places where positive ε plays a role. Without it, no-trade could be another Nash equilibrium of the market game. In what follows, we denote $\{\beta_q^z(x), \beta_r^z(x)\}_{x \in X}$ by β^z .

Proposition 2. Fix $z \in Z$. If the sequence $\{\sigma^N\}_{N=1}^\infty$ is such that $\lim_{N \rightarrow \infty} \sigma^N(x)/[(1-\eta)N] = \mu_z(x)$ for each $x \in X$, then $\lim_{N \rightarrow \infty} \beta^{\sigma^N} = \beta^z$.

⁷Indeed, it can be regarded as a one-shot, complete-information game in which ϕ is a preference parameter.

The proof uses the fact that the sequence $\{\beta^{\sigma^N}\}_{N=1}^{\infty}$ is bounded and applies the Theorem of the Maximum to any subsequence of such equilibria.

3.2 Existence of separating equilibrium

Contingent on being inactive, an agent at stage-1 chooses $a \in \mathcal{O}$ to maximize

$$G_x(a) = \sum_{z \in Z} \tau_x(z) u(\bar{q} - a_q + \frac{a_r}{p_1}, \bar{r} - a_r + p_1 a_q; x, z). \quad (4)$$

By the argument in the proof of Proposition 1, a unique maximum of $G_x(a)$ exists. We denote it $\alpha^* = \{\alpha^*(x)\}_{x \in X}$. Generically, α^* is separating in the sense that $x \neq y$ implies $\alpha^*(x) \neq \alpha^*(y)$. That is, if, instead, $\alpha^*(x) = \alpha^*(y) = (\check{a}_q, \check{a}_r)$ for $x \neq y$, then

$$\frac{\sum_{z \in Z} \tau_x(z) u_q(\check{q}, \check{r}; x, z)}{\sum_{z \in Z} \tau_x(z) u_r(\check{q}, \check{r}; x, z)} = p_1 = \frac{\sum_{z \in Z} \tau_y(z) u_q(\check{q}, \check{r}; y, z)}{\sum_{z \in Z} \tau_y(z) u_r(\check{q}, \check{r}; y, z)},$$

where $\check{q} = \bar{q} - \check{a}_q + \frac{\check{a}_r}{p_1}$ and $\check{r} = \bar{r} - \check{a}_r + p_1 \check{a}_q$. But this restriction holds only for knife-edge cases for two distinct aspects of the environment: the probabilities, $\tau_x(\cdot)$ and $\tau_y(\cdot)$, and the utilities, $u(\check{q}, \check{r}; x, \cdot)$ and $u(\check{q}, \check{r}; y, \cdot)$.

Now we describe beliefs under the assumption that α^* is separating. For any $a \in \mathcal{O}$, let


$$q_1(a) = \bar{q} - a_q + \frac{a_r}{p_1} \in [0, 2\bar{q}]. \quad (5)$$

By separation, $x \neq y$ implies $q_1(\alpha^*(x)) \neq q_1(\alpha^*(y))$. Let $k = \#X$ and, without loss of generality, order the elements of X so that $q_1(\alpha^*(x_i)) < q_1(\alpha^*(x_{i+1}))$ for $i \in \{1, 2, \dots, k-1\}$. Next, partition the interval $[0, 2\bar{q}]$ into k subintervals indexed by that ordering as follows:

$$I(x_i) = \begin{cases} \left[0, \frac{q_1(\alpha^*(x_2)) + q_1(\alpha^*(x_1))}{2} \right) & \text{for } i = 1 \\ \left[\frac{q_1(\alpha^*(x_i)) + q_1(\alpha^*(x_{i-1}))}{2}, \frac{q_1(\alpha^*(x_{i+1})) + q_1(\alpha^*(x_i))}{2} \right) & \text{for } i = 2, 3, \dots, k-1 \\ \left[\frac{q_1(\alpha^*(x_i)) + q_1(\alpha^*(x_{i-1}))}{2}, 2\bar{q} \right] & \text{for } i = k \end{cases} \quad (6)$$

For $k = 2$, $I(x_1)$ and $I(x_2)$ are depicted in Figure 1.

Figure 1. Intervals for beliefs: $k = 2$



An agent's belief is a joint distribution over the type/stage-1-action configuration of the other active agents and the state-of-the-world z . It is derived from the observed histogram, ν , and from knowledge of the agent's own private information. In particular, each active agent sees the stage-1 offers of all other active agents, which imply a configuration of quantities via the mapping q_1 . Here, then, is our candidate for beliefs, which is defined for arbitrary stage-1 outcomes.

Candidate for equilibrium beliefs, φ^ : $\varphi^*(x, a, \nu^{-a})$ puts probability 1 on the configuration $\theta_{\nu^{-a}}$ defined by*

$$\theta_{\nu^{-a}}(y, a') = \begin{cases} \nu^{-a}(a') & \text{if } q_1(a') \in I(y) \\ 0 & \text{otherwise} \end{cases} . \quad (7)$$

Its marginal distribution over Z is given by the posterior derived from Bayes' rule using the type-configuration of all active agents $\sigma^* : X \rightarrow \{0, 1, \dots, M\}$ defined by

$$\sigma^*(y) = \begin{cases} \theta_{\nu^{-a}}(y) & \text{if } y \neq x \\ \theta_{\nu^{-a}}(x) + 1 & \text{if } y = x \end{cases} . \quad (8)$$

That is, (7) says that each agent forms a degenerate distribution over the type/stage-1-action configuration of the other active agents by treating an observed stage-1 action in $I(x_i)$ as coming from an agent of type x_i . As for (8), it says that the agent gets a type-configuration over all active agents by using the type-configuration for other active agents implied by (7) and the agent's own true type.

In order to describe the candidate for equilibrium strategies, it is helpful to distinguish between two classes of active agents according to their private histories. We call an agent of type x a *nondefector* if the agent's stage-1 action is in $I(x)$; otherwise, the agent is called a *defector*. Notice that if no one defects, then all agents have beliefs that are symmetric in the sense assumed in proposition 1: all have the same posterior on z and all active agents have the same belief about the type-configuration over other active agents, which happens to be the true configuration. If one agent defects or more than one defect, then all nondefectors have symmetric beliefs; they have the same posterior on z and any such active agent has the same belief about the type-configuration over other active agents, which, however, is not the true configuration. Each defector has a different posterior on z and a different belief about the type-configuration for the active agents.

The belief φ^* has each agent believing that other agents do not defect. Our specification for a candidate equilibrium is consistent with that belief. In particular, when we describe a stage-2 strategy for arbitrary stage-1 actions, we have each agent believing that other agents did not defect at stage-1.

Candidate for equilibrium strategies. For stage-1, our candidate is

$$s_1^*(x) = \alpha^*(x), \quad (9)$$

the maximum of $G_x(a)$ (see (4)).

To describe stage-2 strategies, consider an agent with private history (x, a, ν^{-a}) such that $q_1(a) \in I(x')$. Given the private history, let σ^ be the agent's belief about the type-configuration of all active agents under φ^* (see (8)). Because the agent believes that all*

other agents are nondefectors, he believes that all other agents share the same believed type-configuration, denoted σ' . If $x' = x$ (the agent is a nondefector), then $\sigma'(y) = \sigma^*(y)$ for all y ; otherwise (the agent is a defector), $\sigma'(x) = \sigma^*(x) - 1$, $\sigma'(x') = \sigma^*(x') + 1$, and $\sigma'(y) = \sigma^*(y)$ for all $y \notin \{x, x'\}$. Then $s_2^*(x, a, \nu^{-a})$ satisfies

$$s_2^*(x, a, \nu^{-a}) \in \arg \max_{b \in \mathcal{O}} \sum_{z \in Z} \phi^{\sigma^*}(z) u(\bar{q} + \frac{b_r Q_- - b_q R_-}{R_- + b_r}, \bar{r} + \frac{b_q R_- - b_r Q_-}{Q_- + b_q}; x, z), \quad (10)$$

where $\phi^{\sigma^*}(z)$ is derived from σ^* using Bayes rule and where

$$(Q_-, R_-) = M\varepsilon + \sum_{y \neq x} \sigma^*(y) \beta^{\sigma'}(y) + (\sigma^*(x) - 1) \beta^{\sigma'}(x).$$

When the agent is a nondefector, that is, when $\sigma' = \sigma^*$, we have $s_2^*(x, a, \nu^{-a}) = \beta^{\sigma^*}(x)$. Notice that in the above construction we fix a β^σ for any σ ; that is, agents coordinate on a particular proposition-1 equilibrium for any believed type-configuration.

Theorem 1. Suppose that $\alpha^*(x) \neq \alpha^*(y)$ for any $x \neq y$. There exists some \bar{N} such that if $N \geq \bar{N}$, then the N -agent economy has a separating equilibrium.

The proof shows that the above candidate is an equilibrium. By construction, $s_1^* = \alpha^*$ implies that φ^* is consistent with Bayes' rule. Also, by construction, $s_2^*(x, a, \nu^{-a})$ is a best response to s_2^* with respect to φ^* . That follows because, according to φ^* , the agent believes that every other active agent is a nondefector. And, if they follow s_2^* , then their actions are described by $\beta^{\sigma'}$. Therefore, what remains, and is the focus of the proof, is to show that α^* is optimal given that other agents follows the candidate equilibrium. An agent at stage-1 faces a tradeoff. Conditional on being inactive, playing α^* is optimal for any N . Conditional on being active, a type- x agent could gain by playing something not in $I(x)$. By doing that, the agent influences the beliefs and, thereby, the stage-2 actions of other active agents. The proof shows that any such gain vanishes as N gets large and is, therefore, smaller than the loss implied by playing something that is not in $I(x)$, which, by construction, is bounded away from $\alpha^*(x)$.

Corollary. Any Theorem-1 equilibrium is almost ex post efficient.

This follows from proposition 2.

3.3 Uniqueness of equilibrium outcomes

In the previous section we have established the existence of a separating equilibrium and have shown that it is almost ex post efficient. Here, with a small modification of the mechanism, we show that *any* equilibrium outcome in pure symmetric strategies is almost ex post efficient if a generic condition about the information structure holds. We are *not* showing that nonseparating equilibria exist.

The modification is that the second stage does not exist (the market shuts down) if all agents announce the same offer in the first stage. Under the assumption that

$\alpha^*(x) \neq \alpha^*(y)$ for all $x \neq y$, this modification then rules out any equilibrium (s_1, s_2) such that $s_1(x) = s_1(y)$ for all $x, y \in X$, but does not change any other equilibrium if it exists. (In particular, this modification does not affect the existence of a separating equilibrium.) To see that $s_1(x) = \check{a}$ for all $x \in X$ is not an equilibrium under this modification, it is enough to show that one agent has a profitable deviation. By assumption, there exists a type, say type y , for whom $\check{a} \neq \alpha^*(y)$. Because no trade is always available for the agent at stage 2, playing $\alpha^*(y)$ is a profitable deviation.

To show that any other equilibrium (s_1, s_2) is almost ex post efficient, we need the following stronger informativeness assumption:

A1. Let $\mathcal{Y} = \{Y_1, Y_2\}$ be any bipartition of X and let $\mu_z(Y_i) \equiv \sum_{y \in Y_i} \mu_z(y)$. For any $z \neq z'$, $\mu_z(Y_1) \neq \mu_{z'}(Y_1)$.

Assumption A1 implies that for any partition $\mathcal{Y} = \{Y_1, \dots, Y_K\}$ of X with $K \geq 2$ and for any $z \neq z'$, there exists some k such that $\mu_z(Y_k) \neq \mu_{z'}(Y_k)$. Although this assumption is stronger than our original informativeness assumption, it holds generically.

Because the modification rules out the complete pooling equilibrium, any equilibrium (s_1, s_2) in symmetric pure strategies is a semi-pooling equilibrium in the sense that there is a partition $\mathcal{Y} = \{Y_1, \dots, Y_K\}$ of X with $K \geq 2$ such that $s_1(y) = s_1(y')$ if $y, y' \in Y_k$ and $s_1(y) \neq s_1(y')$ if $y \in Y_k$ and $y' \in Y_{k'}$ with $k \neq k'$. We call each Y_k a *signal*. Let $\lambda : \mathcal{Y} \rightarrow \mathbb{N}$ denote a configuration of signals.

In an equilibrium (s_1, s_2) with associated partition \mathcal{Y} , the signal configuration λ implied by the public announcement along the equilibrium path is a random variable. Moreover, following any public announcement ν and the implied signal configuration λ along the equilibrium path, the stage-2 game can be analyzed as an incomplete-information game: players are the active agents; the action set for each player is \mathcal{O} ; payoffs are determined by the market game and u ; the number of players with types belonging to Y_k is $\lambda(Y_k)$, which is common knowledge among the players; and the common prior over Z and X^M is given by $\kappa(z)$ and μ_z for $z \in Z$. However, in contrast to the stage-2 game in a separating equilibrium, an active agent does not know the type configuration of other active agents.

Fix a state of the world z . Because the efficiency of the equilibrium outcome is mainly determined by the stage-2 outcome, our goal is to show that conditional on z , the stage-2 equilibrium outcome converges to the competitive outcome of $\mathcal{L}^z(\varepsilon)$. Specifically, we show that for any sequence of equilibria, $\{(s_1^N, s_2^N)\}_{N=1}^\infty$, the stage-2 equilibrium outcome converges to the competitive outcome conditional on z as N goes to infinity. Because there are only finitely many nondegenerate partitions of X , it is sufficient to show that for any given partition \mathcal{Y} of X , if the subsequence $\{(s_1^{N_s}, s_2^{N_s})\}_{N=1}^\infty$ consists of equilibria associated with \mathcal{Y} , then the stage-2 equilibrium outcome converges to the competitive outcome conditional on z as s goes to infinity. For each N_s , let λ^{N_s} (which is a random variable) be the signal configuration on the equilibrium path in $\{(s_1^{N_s}, s_2^{N_s})\}_{N=1}^\infty$ and let $\beta^{\lambda^{N_s}}$ be the stage-2 equilibrium actions.

Theorem 2. Let $\mathcal{Y} = \{Y_1, \dots, Y_K\}$ be a partition of X with $K \geq 2$. Assume that there is an increasing sequence of $\{N_s\}_{s=1}^\infty$ such that for each s , there exists an equilibrium

$(s_1^{N_s}, s_2^{N_s})$ associated with \mathcal{Y} . Under assumption A1, for any $z \in Z$, $\lim_{s \rightarrow \infty} \beta^{\lambda^{N_s}} = \beta^z$ almost surely conditional on z .

The proof has two parts. First, under A1, the histogram of stage-1 offers reveals z almost surely as N gets large. Then, the assumption that types are *i.i.d.* conditional on z implies that the type-configuration distribution of active agents converges to $\mu_z(x)$. The second part of the proof follows the logic in the proof of proposition 2. However, here, instead of appealing to the Theorem of the Maximum, an explicit limiting argument is given.

4 Proofs

Lemma 1. Fix the stage-2 offers of all other agents. If $b' \in [0, \bar{q}] \times [0, \bar{r}]$ with payoff (q', r') is such that $b'_q b'_r > 0$, then there exists $b'' \in \mathcal{O}$ with payoff (q', r') .

Proof. Let $(Q_-, R_-) \in \mathbb{R}_{++}^2$ be total offers of other agents (including the exogenous offers). For any $b \in [0, \bar{q}] \times [0, \bar{r}]$, we have from (2),

$$q = \bar{q} + \frac{b_r Q_- - b_q R_-}{R_- + b_r} \text{ and } r = \bar{r} + \frac{b_q R_- - b_r Q_-}{Q_- + b_q}. \quad (11)$$

Case (i): $b'_r Q_- - b'_q R_- > 0$. In this case, let $b''_q = 0$ and let b''_r be the unique solution to

$$\frac{b''_r Q_-}{R_- + b''_r} = \frac{b'_r Q_- - b'_q R_-}{R_- + b'_r} \equiv \gamma, \quad (12)$$

where it follows that $\gamma \in (0, Q_-)$. The solution is $b''_r = R_- \gamma / (Q_- - \gamma)$. It follows by (12) that $q(b''_q, b''_r) = q(b'_q, b'_r)$. Now

$$r(b''_q, b''_r) - \bar{r} = b''_r = R_- \gamma / (Q_- - \gamma) = r(b'_q, b'_r) - \bar{r},$$

where the last equality follows from the definition of γ .

Case (ii): $b'_r Q_- - b'_q R_- < 0$. This is completely analogous, but with $b''_r = 0$.

Case (iii): $b'_r Q_- - b'_q R_- = 0$. Here, of course, we let $b''_q = b''_r = 0$. ■

Proposition 1. For any type-configuration $\sigma : X \rightarrow \{1, \dots, M\}$, β^σ exists.

Proof. Let $k = \#X$ and let $S = \{[0, \bar{q}] \times [0, \bar{r}]\}^k$, which is compact and convex. We let $s = \{s_y\}_{y \in X}$ with $s_y = (s_{yq}, s_{yr})$ denote a generic element of S . For $s \in S$ and $x \in X$, let $F : S \rightarrow S$ be given by

$$F_x(s) = \arg \max_{b \in \mathcal{O}} H_x(b; Q_-, R_-), \quad (13)$$

where

$$H_x(b; Q_-, R_-) = \sum_{z \in Z} \phi(z) u\left(\bar{q} + \frac{b_r Q_- - b_q R_-}{R_- + b_r}, \bar{r} + \frac{b_q R_- - b_r Q_-}{Q_- + b_q}; x, z\right) \quad (14)$$

and

$$(Q_-, R_-) = M\varepsilon + \sum_{y \neq x} \sigma(y) s_y + [\sigma(x) - 1] s_x.$$

Here ϕ is the common posterior on z . We have to show that $F_x(s)$ is unique and is continuous in s . We start with uniqueness. Notice that $(Q_-, R_-) \in \mathbb{R}_{++}^2$ for any $s \in S$.

Because of the $b_q b_r = 0$ constraint in (13), it is helpful to separately consider $H_x(b_q, 0; Q_-, R_-)$ and $H_x(0, b_r; Q_-, R_-)$, where

$$H_x(b_q, 0; Q_-, R_-) = \sum_{z \in Z} \phi(z) u(\bar{q} - b_q, \bar{r} + \frac{b_q R_-}{Q_- + b_q}; x, z) \equiv g(b_q),$$

and

$$H_x(0, b_r; Q_-, R_-) = \sum_{z \in Z} \phi(z) u(\bar{q} + \frac{b_r Q_-}{R_- + b_r}, \bar{r} - b_r; x, z) \equiv h(b_r).$$

For $(c_1, c_2) \in \mathbb{R}_{++}^2$, let $f(y) = c_1 y / (y + c_2)$ for $y \in \mathbb{R}_+$. It follows that f is increasing and strictly concave. Then, because u is strictly concave and because a strictly concave function of an increasing concave function is strictly concave, both g and h are strictly concave. It follows that g has a unique maximum and that h has a unique maximum, denoted \hat{b}_q and \hat{b}_r , respectively. Moreover, by the Inada conditions on u , these maxima are characterized by

$$\hat{b}_q = \begin{cases} 0 & \text{if } g'(0) \leq 0 \\ \text{satisfies } g'(\hat{b}_q) = 0 & \text{if } g'(0) > 0 \end{cases}, \quad (15)$$

and

$$\hat{b}_r = \begin{cases} 0 & \text{if } h'(0) \leq 0 \\ \text{satisfies } h'(\hat{b}_r) = 0 & \text{if } h'(0) > 0 \end{cases}. \quad (16)$$

Therefore, a sufficient condition for uniqueness is $\min\{g'(0), h'(0)\} \leq 0$. But,

$$g'(0) = \sum_{z \in Z} \phi(z) \left[-u_q(\bar{q}; x, z) + u_r(\bar{r}; x, z) \frac{R_-}{Q_-} \right],$$

and

$$h'(0) = \sum_{z \in Z} \phi(z) \left[u_q(\bar{q}; x, z) \frac{Q_-}{R_-} - u_r(\bar{r}; x, z) \right].$$

Therefore,

$$\text{sign}[h'(0)] = \text{sign}\left[\frac{R_-}{Q_-} h'(0)\right] = \text{sign}[-g'(0)] = -\text{sign}[g'(0)], \quad (17)$$

which implies $\min\{g'(0), h'(0)\} \leq 0$.

Now we turn to continuity in s , which follows if (\hat{b}_q, \hat{b}_r) is continuous in (Q_-, R_-) . By (17), $g'(0) = 0$ iff $h'(0) = 0$. That and (15) and (16) imply that $\max\{\hat{b}_q, \hat{b}_r\}$ satisfies a first-order condition with equality. Then, the implicit-function theorem applied to that first-order condition gives the required continuity.

It follows that the mapping F satisfies the hypotheses of Brouwer's fixed-point theorem. Notice that although the domain of the mapping, S , does not satisfy $b_q b_r = 0$, the range does. Therefore, the fixed point satisfies $b_q b_r = 0$. ■

Lemma 2. Let $\mathcal{L}^z(\varepsilon)$ denote the economy with z known, with a nonatomic measure of agents, with fraction of type- x agents equal to $\mu_z(x)$, and with exogenous per capita trades (and endowments) ε . The economy $\mathcal{L}^z(\varepsilon)$ has a unique competitive equilibrium and it is continuous in ε . Moreover, $(q^z(x), r^z(x))_{x \in X}$ is a competitive allocation for $\mathcal{L}^z(\varepsilon)$ if and only if

$$\beta_q^z(x) = \max\{\bar{q} - q^z(x), 0\} \text{ and } \beta_r^z(x) = \max\{\bar{r} - r^z(x), 0\} \quad (18)$$

is a Nash equilibrium for the market game for $\mathcal{L}^z(\varepsilon)$.

Proof. In order to simplify the notation, we suppress the dependence on z . A competitive equilibrium (CE) with exogenous offers ε is a price p and an allocation $(q(x), r(x))_{x \in X}$ such that $(q(x), r(x))$ maximizes $u(q, r; x, z)$ subject to $pq + r = p\bar{q} + \bar{r}$ and

$$\frac{\varepsilon_r}{p} + \sum_{x \in X} \mu_z(x) q(x) = \bar{q} + \varepsilon_q. \quad (19)$$

Given our assumptions about $u(\cdot, \cdot; x, z)$, existence and uniqueness of a CE is entirely standard, as is continuity in ε . Hence, we turn to the coincidence claim.

(i) Assume that $(q(x), r(x))_{x \in X}$ is a competitive allocation. It follows from (18) that $\beta(x) \in \mathcal{O}$. Next, we show that the CE price is

$$p = \frac{\sum_{x \in X} \mu_z(x) \beta_r(x) + \varepsilon_r}{\sum_{x \in X} \mu_z(x) \beta_q(x) + \varepsilon_q}. \quad (20)$$

The market-clearing condition, (19), can be written,

$$\frac{\varepsilon_r}{p} + \sum_{x \in X^+} \mu(x) [q(x) - \bar{q}] - \sum_{x \in X^-} \mu(x) [q(x) - \bar{q}] = \varepsilon_q$$

where X^+ means $x \in X$ such that $q(x) - \bar{q} > 0$ and X^- means $x \in X$ such that $q(x) - \bar{q} < 0$. By the competitive budget, $pq + r = p\bar{q} + \bar{r}$, this can be written

$$\frac{\varepsilon_r}{p} + \frac{1}{p} \sum_{x \in X^+} \mu(x) [\bar{r} - r(x)] + \sum_{x \in X^-} \mu(x) [q(x) - \bar{q}] = \varepsilon_q,$$

where $\bar{r} - r(x) > 0$. Using (18), we have

$$\frac{\varepsilon_r}{p} + \frac{1}{p} \sum_{x \in X^+} \mu(x) \beta_r(x) - \sum_{x \in X^-} \mu(x) \beta_q(x) = \varepsilon_q,$$

which gives (20).

It remains to show that β satisfies the best-response correspondence. This is done by showing that the feasible choices for consumption in the game when everyone else plays β

are the same as the competitive budget set when the price is given by (20). This follows from (2) and the fact that one person's action does not affect the r.h.s. of (20).

(ii) Now assume that β is a stage-2 Nash equilibrium. In the last part of the proof of (i), we showed that the individual choice sets in the game and those implied by a competitive budget are the same for a given price. So all we need is a price that makes (q, r) implied by (18) satisfy market-clearing. That is guaranteed by aggregating the payoffs when the price is given by (20). ■

Proposition 2. Fix $z \in Z$. If the sequence $\{\sigma^N\}_{N=1}^\infty$ is such that $\lim_{N \rightarrow \infty} \sigma^N(x)/[(1 - \eta)N] = \mu_z(x)$ for each $x \in X$, then $\lim_{N \rightarrow \infty} \beta^{\sigma^N} = \beta^z$.

Proof. This proof applies the Theorem of the Maximum to a sequence of proposition 1 equilibria. We can write the best-response objective (see (14)) as

$$H_x(b; Q_-, R_-, \phi) = \sum_{z' \in Z} \phi(z') u(q, r; x, z'), \quad (21)$$

with

$$q = \bar{q} + \frac{b_r}{p(1 + \frac{b_r}{R_-})} - \frac{b_q}{1 + \frac{b_r}{R_-}},$$

$$r = \bar{r} - \frac{b_r}{1 + \frac{b_q}{Q_-}} + \frac{pb_q}{1 + \frac{b_q}{Q_-}},$$

and $p = R_-/Q_-$.

Now, let

$$F_x(b; p, c_1, c_2, \phi) = \sum_{z' \in Z} \phi(z') u(q, r; x, z')$$

with

$$q = \bar{q} + \frac{b_r}{p(1 + c_2 b_r)} - \frac{b_q}{1 + c_2 b_r},$$

and

$$r = \bar{r} - \frac{b_r}{1 + c_1 b_q} + \frac{pb_q}{1 + c_1 b_q},$$

and where the domain for F is $A = \mathcal{O} \times \left[\frac{\varepsilon_r}{\bar{q} + \varepsilon_q}, \frac{\bar{r} + \varepsilon_r}{\varepsilon_q} \right] \times \left[0, \frac{1}{\varepsilon_r} \right] \times \left[0, \frac{1}{\varepsilon_q} \right] \times \Delta(Z)$. It follows that $F_x(b; p, 1/Q_-, 1/R_-, \phi) = H_x(b; Q_-, R_-, \phi)$. Therefore, by the argument used in the proof of proposition 1, $F_x(\cdot; p, c_1, c_2, \phi)$ has a unique maximum, $g_x(p, c_1, c_2, \phi)$. And because F_x is continuous on A , the Theorem of the Maximum implies that $g_x(p, c_1, c_2, \phi)$ is continuous.

Now consider

$$H_x(b; Q_-^N, R_-^N, \phi) = F_x(b; \frac{R_-^N}{Q_-^N}, \frac{1}{Q_-^N}, \frac{1}{R_-^N}, \phi)$$

with

$$(Q_-^N, R_-^N) = \sum_{y \in X} \sigma^N(y) (\beta^N(y) + \varepsilon) - \beta^N(x).$$

Notice that $(\frac{R_-^N}{Q_-^N}, \frac{1}{Q_-^N}, \frac{1}{R_-^N}) \in [\frac{\varepsilon_r}{\bar{q} + \varepsilon_q}, \frac{\bar{r} + \varepsilon_r}{\varepsilon_q}] \times [0, \frac{1}{\varepsilon_r}] \times [0, \frac{1}{\varepsilon_q}]$. Therefore, by the definition of β^N , $\beta^N(x) = g_x(\frac{R_-^N}{Q_-^N}, \frac{1}{Q_-^N}, \frac{1}{R_-^N}, \phi)$. Because $\{\beta^N\}_{N=1}^\infty$ is bounded, it has a convergent subsequence, say $\{\beta^{N_s}\}_{s=1}^\infty$, with limit denoted $\hat{\beta}$. By the continuity of g_x , it follows that

$$\hat{\beta}(x) = \lim_{s \rightarrow \infty} g_x(\frac{R_-^N}{Q_-^N}, \frac{1}{Q_-^N}, \frac{1}{R_-^N}, \phi) = g_x(\lim_{s \rightarrow \infty} \frac{R_-^N}{Q_-^N}, \lim_{s \rightarrow \infty} \frac{1}{Q_-^N}, \lim_{s \rightarrow \infty} \frac{1}{R_-^N}, \lim_{s \rightarrow \infty} \phi) = g_x(\hat{p}, 0, 0, \delta_z)$$

where

$$\hat{p} = \frac{\sum \mu_z(y) \hat{\beta}_r(y) + \varepsilon_r}{\sum \mu_z(y) \hat{\beta}_q(y) + \varepsilon_q}$$

and $\delta_z = 1$ for $z' = z$ (and 0 otherwise). By the definition of F_x , it follows that $\hat{\beta}(x)$ maximizes $u(\bar{q} - b_q + \frac{b_r}{\hat{p}}; x, z) + w(\bar{r} - b_r + \hat{p}b_q)$. Therefore, it is a Nash equilibrium in $\mathcal{L}^z(\varepsilon)$. By lemma 2, it follows that $\hat{\beta} = \beta^z$. ■

Theorem 1. Suppose that $\alpha^*(x) \neq \alpha^*(y)$ for any $x \neq y$. There exists some \bar{N} such that if $N \geq \bar{N}$, then the N -agent economy has a separating equilibrium.

Proof. We show that for large N 's, $((s_1^*, s_2^N), \varphi^N)$ is a PBE, where $s_1^*(x) = \alpha^*(x)$ for all $x \in X$, and s_2^N and φ^N are given by (10) and (7)-(8), respectively. By construction, s_2^N is a best response against s_2^N w.r.t. φ^N and φ^N is consistent with Bayes rule. It remains to show that s_1^* is a best response to (s_1^*, s_2^N) for sufficiently large N .

Let $M^N = \lceil (1 - \eta)N \rceil$. Fix an agent of type x . Because the assignment π is drawn independently from the types, conditional on being active, the agent's belief about other agents' types is such that those types are i.i.d. with marginal probabilities $(\mu_z(x))_{x \in X}$ conditional on each state z . Let γ_z^N be the i.i.d. distribution over $X^{M^N - 1}$ generated by $(\mu_z(x))_{x \in X}$. Given s_2^* , the first stage problem for the agent of type x is $\max_{a \in \mathcal{O}} G_x^N(a)$, where

$$G_x^N(a) = \eta G_x(a) + (1 - \eta) F_x^N(a). \quad (22)$$

Here, G_x is the stage-1 problem contingent on being inactive; while, for $a \in I(\bar{x})$,

$$F_x^N(a) = \sum_{z \in Z} \tau_x(z) \left[\sum_{\xi \in X^{M^N - 1}} \gamma_z^N(\xi) [u(q^N(a; z, \xi), r^N(a; z, \xi); x, z)] \right], \quad (23)$$

where for each z and $\xi = (\xi_1, \dots, \xi_{M^N - 1}) \in X^{M^N - 1}$,

$$q^N(a; z, \xi) = \bar{q} + \frac{s_{2,r}^*(x, a, \nu^{\xi, -a}) Q_-^N - s_{2,q}^*(x, a, \nu^{\xi, -a}) R_-^N}{s_{2,r}^*(x, a, \nu^{\xi, -a}) + R_-^N},$$

$$r^N(a; z, \xi) = \bar{r} + \frac{s_{2,q}^*(x, a, \nu^{\xi, -a}) R_-^N - s_{2,r}^*(x, a, \nu^{\xi, -a}) Q_-^N}{s_{2,q}^*(x, a, \nu^{\xi, -a}) + Q_-^N},$$

$\nu^{\xi, -a}$ is the announced histogram given that other active agents' types are ξ and that other agents follow s_1^* , and Q_-^N and R_-^N are the implied stage-2 offers of other active

agents according to the candidate equilibrium. That is,

$$\nu^{\xi, -a}(s_1^*(y)) = \sum_{i=1}^{M^N-1} \mathbf{1}_y(\xi_i) \text{ for each } y \in X \text{ and } \nu^{\xi, -a}(a') = 0 \text{ otherwise,} \quad (24)$$

and

$$(Q_-^N, R_-^N) = \sum_{y \in X} \sigma^\xi(y)(\beta^{\sigma^\xi}(y) + \varepsilon) - \beta^{\sigma^\xi}(\bar{x}) \quad (25)$$

where σ^ξ is the type-configuration believed by other active agents; namely,

$$\sigma^\xi(y) = \sum_{i=1}^{M^N-1} \mathbf{1}_y(\xi_i) \text{ for each } y \neq \bar{x} \text{ and } \sigma^\xi(\bar{x}) = \sum_{i=1}^{M^N-1} \mathbf{1}_{\bar{x}}(\xi_i) + 1. \quad (26)$$

The theorem is proved using the following two claims.

Claim 1. There exists an $\epsilon > 0$ such that for all $x \in X$, if $q_1(a) \notin I(x)$, then $G_x(a) < G_x(s_1^*(x)) - \epsilon$.

Proof of claim 1. As mentioned before, $\max_{a \in \mathcal{O}} G_x(a)$ is equivalent to $\max_{q \in [0, 2\bar{q}]} L_x(q)$, where

$$L_x(q) = \sum_{z \in Z} \tau_x(z) u(q, p_1 \bar{q} + \bar{r} - p_1 q; x, z).$$

The function L_x is strictly concave in q for each x . Let $\delta = \min\{\frac{|q_1(\alpha^*(x)) - q_1(\alpha^*(y))|}{2} : x, y \in X, x \neq y\}$. Then, for each $x \in X$, if $q \notin I(x)$, then $|q - q_1(\alpha^*(x))| \geq \delta$. Fix some $x \in X$. $L_x(q)$ is strictly concave in q and has maximum at $q_1(\alpha^*(x))$. Let $A_x = \min\{-L'_x(q_1(\alpha^*(x)) + \frac{\delta}{2}), L'_x(q_1(\alpha^*(x)) - \frac{\delta}{2})\} > 0$. Then, for any q such that $|q - q_1(\alpha^*(x))| \geq \delta$, $L_x(q) \leq L_x(q_1(\alpha^*(x))) - \frac{\delta}{2} A_x$. Take $\epsilon = \frac{1}{2} \min\{\frac{\delta}{2} A_x : x \in X\}$. Then, if $q_1(a) \notin I(x)$, then $G_x(a) = L_x(q_1(a)) \leq L_x(q_1(\alpha^*(x))) - 2\epsilon < G_x(s_1^*(x)) - \epsilon$. \square

Claim 2. Let $\xi = (\xi_1, \dots, \xi_n, \dots)$ be an infinite sequence of X -valued random variables that is i.i.d. w.r.t. the marginal distribution $(\mu_z(x))_{x \in X}$ and let $\xi^{M^N-1} = (\xi_1, \dots, \xi_{M^N-1})$, where ξ^{M^N-1} is interpreted as the types of the other active agents when there are M^N active agents. Then

$$\lim_{N \rightarrow \infty} q^N(a; z, \xi^{M^N-1}) = \bar{q} + \frac{\beta_r^z(x)}{p^z} - \beta_q^z(x), \quad \lim_{N \rightarrow \infty} r^N(a; z, \xi^{M^N-1}) = \bar{r} + \beta_q^z(x) p^z - \beta_r^z(x),$$

in probability and

$$\lim_{N \rightarrow \infty} F_x^N(a) = \sum_{z \in Z} \tau_x(z) u(q^z(x), r^z(x); x, z),$$

uniformly in $a \in \mathcal{O}$, where $(q^z(x), r^z(x))$ is the CE allocation of $\mathcal{L}(\varepsilon_q, \varepsilon_r)$ (see Lemma 2).

Proof of claim 2. By definition, $\xi^{M^N-1} = (\xi_1, \dots, \xi_{M^N-1})$ is distributed according to γ_z^N . For each N , let $\sigma^N = \sigma^{\xi^{M^N-1}}$ as defined in (26) and let $\nu^N = \nu^{\xi^{M^N-1}, -a}$ as defined in (24). Then, the sequence $\{\sigma^N\}$ is such that $\sum_{y \in X} \sigma^N(y) = M^N$ and for each $y \in X$, $\lim_{N \rightarrow \infty} (\sigma^N(y)/M^N) = \mu_z(y)$ almost surely. Consider a realization of ξ for which

$\lim_{N \rightarrow \infty} (\sigma^N(y)/M^N) = \mu_z(y)$. Then, by Proposition 2, we have $\lim_{N \rightarrow \infty} \beta^{\sigma^N} = \beta^z$. This implies that

$$\lim_{N \rightarrow \infty} \left(\frac{Q_-^N}{M^N}, \frac{R_-^N}{M^N} \right) = \sum_{y \in X} \mu_z(y) (\beta^z(y) + \varepsilon) \text{ and } \lim_{N \rightarrow \infty} \frac{R_-^N}{Q_-^N} = p^z, \quad (27)$$

where Q_-^N and R_-^N are defined in (25) with $\xi = \xi^{M^N-1}$.

Finally, we show that $\lim_{N \rightarrow \infty} s_2^*(x, a, \nu^N) = \beta^z(x)$. First notice that

$$\lim_{N \rightarrow \infty} \phi^N[z] = \lim_{N \rightarrow \infty} \text{marg}_Z \varphi^N(x, a, \nu^N)[z] = 1.$$

For each N , $s_2^*(x, a, \nu^N)$ solves

$$\max_{b \in \mathcal{O}} H_x^N(b) = \sum_{z' \in Z} \phi^N[z'] u \left(\bar{q} + \frac{b_r Q_-^N - b_q R_-^N}{R_-^N + b_r}, \bar{r} + \frac{b_q R_-^N - b_r Q_-^N}{Q_-^N + b_q}; x, z \right). \quad (28)$$

Now, let

$$J_x(b; p, c_1, c_2, \phi) = \sum_{z' \in Z} \phi[z'] u(q, r; x, z')$$

with $q = \bar{q} + \frac{b_r}{p(1+c_2b_r)} - \frac{b_q}{1+c_2b_r}$ and $r = \bar{r} - \frac{b_r}{1+c_1b_q} + \frac{pb_q}{1+c_1b_q}$, and where the domain for J_x is $\mathcal{O} \times \left[\frac{\varepsilon_r}{\bar{q}+\varepsilon_q}, \frac{\bar{r}+\varepsilon_r}{\varepsilon_q} \right] \times \left[0, \frac{1}{\varepsilon_r} \right] \times \left[0, \frac{1}{\varepsilon_q} \right] \times \Delta(Z)$. It follows that $J_x(b; \frac{R_-^N}{Q_-^N}, 1/Q_-^N, 1/R_-^N, \phi^N) = H_x^N(b)$. Therefore, by the argument used in the proof of Proposition 1, $J_x(\cdot; p, c_1, c_2, \phi)$ has a unique maximum, $j_x(p, c_1, c_2, \phi)$. And because J_x is continuous on its domain, the Theorem of the Maximum implies that $j_x(p, c_1, c_2, \phi)$ is continuous.

Now, for each N , $s_2^*(x, a, \nu^N) = j_x(\frac{R_-^N}{Q_-^N}, 1/Q_-^N, 1/R_-^N, \phi^N)$. By (27) and the continuity of j_x , it follows that

$$b^* = \lim_{N \rightarrow \infty} s_2^*(x, a, \nu^N) = \lim_{N \rightarrow \infty} j_x \left(\frac{R_-^N}{Q_-^N}, 1/Q_-^N, 1/R_-^N, \phi^N \right) = j_x(p^z, 0, 0, \mathbf{1}_z)$$

where $\mathbf{1}_z[z'] = 1$ for $z' = z$ (and 0 otherwise). By the definition of J_x , it follows that b^* maximizes $u(\bar{q} - b_q + \frac{b_r}{p^z}; x, z) + w(\bar{r} - b_r + p^z b_q)$. Therefore, it is a separating stage-2 equilibrium in the limit model. By Lemma 1, it follows that $b^* = \beta^z(x)$. \square

In order to have any effect on $F_x^N(a)$, the agent must choose an offer sufficiently far from s_1^* , the offer that maximizes $G_x(a)$. Claim 1 shows that that the implied loss in terms of $G_x(a)$ is bounded away from zero (and does not depend on N). By claim 2, any effect on $F_x^N(a)$ goes to zero as $N \rightarrow \infty$. Together, they imply that s_1^* is a best response to (s_1^*, s_2^N) for sufficiently large N . \blacksquare

Theorem 2. Let $\mathcal{Y} = \{Y_1, \dots, Y_K\}$ be a partition of X with $K \geq 2$. Assume that there is an increasing sequence of $\{N_s\}_{s=1}^\infty$ such that for each s , there exists an equilibrium $(s_1^{N_s}, s_2^{N_s})$ associated with \mathcal{Y} . Under assumption A1, for any $z \in Z$, $\lim_{s \rightarrow \infty} \beta^{\lambda^{N_s}} = \beta^z$ almost surely conditional on z .

Proof. Fix $\bar{z} \in Z$. In order to simplify the notation, in what follows we assume that $\{N_s\}_{s=1}^\infty = \{N\}_{N=1}^\infty$; that is, we assume that for each N , there exists an equilibrium (s_1^N, s_2^N) associated with \mathcal{Y} and with N agents. The general case is proved in exactly the same way. For any k , we also abuse the notation and write $s_1^N(Y_k) = s_1^N(y)$ for $y \in Y_k$.

Recall that $\lambda^N(Y_k) = \nu^N(s_1^N(Y_k))$ for each k , where ν^N is the public announcement along the equilibrium path with N agents. From the ex ante perspective, both ν^N and λ^N are random variables. Conditional on \bar{z} , we have

$$\lim_{N \rightarrow \infty} \frac{\lambda^N(Y_k)}{\lceil (1 - \eta)N \rceil} = \mu_{\bar{z}}(Y_k)$$

almost surely for each $k = 1, \dots, K$. Thus, consider a sequence of realizations $\{\lambda^N\}_{N=1}^\infty$ for which this holds. We show that $\lim_{N \rightarrow \infty} \beta^{\lambda^N} = \beta^{\bar{z}}$ almost surely.

Fix some type \bar{x} with $\bar{x} \in Y_{\bar{k}}$. As in Theorem 1, because the assignment π is drawn independently from the types, conditional on the agent being active, the agent's belief about other agents' types is the i.i.d. distribution generated by $(\mu_z(x))_{x \in X}$ conditional on each state z . To simplify notation we denote β^{λ^N} by β^N and $\lceil (1 - \eta)N \rceil$ by M^N . Define λ_-^N by $\lambda_-^N(Y_k) = \lambda^N(Y_k)$ for each $k \neq \bar{k}$ and $\lambda_-^N(Y_{\bar{k}}) = \lambda^N(Y_{\bar{k}}) - 1$. That is, λ_-^N is the signal configuration that pertains to other active agents from the perspective of the agent with type \bar{x} . The belief for the agent with type \bar{x} satisfies the following:

(1) The belief over Z is given by

$$\phi_{\bar{x}}^N[z] = \frac{\kappa(z) \mu_z(\bar{x}) \prod_{k=1}^K [\mu_z(Y_k)]^{\lambda_-^N(Y_k)}}{\sum_{z' \in Z} \kappa(z') \mu_{z'}(\bar{x}) \prod_{k=1}^K [\mu_{z'}(Y_k)]^{\lambda_-^N(Y_k)}}. \quad (29)$$

(2) The belief over Z and the types of the other active agents' types $\{\xi^k = (\xi_1^k, \dots, \xi_{\lambda_-^N(Y_k)}^k) : k = 1, \dots, K\}$, where $\xi^k \in Y_k^{\lambda_-^N(Y_k)}$ describes the types for those who make offers in Y_k , is given by

$$\phi^{(\lambda^N, \bar{x})}[z, \xi^1, \dots, \xi^K] = \phi_{\bar{x}}^N[z] \prod_{k=1}^K \left[\prod_{i=1}^{\lambda_-^N(Y_k)} \frac{\mu_z(\xi_i^k)}{\mu_z(Y_k)} \right] \equiv \phi_{\bar{x}}^N[z] \gamma_z^{\lambda^N}[\xi^1, \xi^2, \dots, \xi^K]. \quad (30)$$

Thus, conditional on a state z , the agent's belief about other agents' types is *i.i.d.* with respect to the conditional probabilities given by the information contained in the partition \mathcal{Y} and λ_-^N . By A1, the posterior $\phi_{\bar{x}}^N$ satisfies $\lim_{N \rightarrow \infty} \phi_{\bar{x}}^N[\bar{z}] = 1$.

Notice that along the equilibrium path ν^N reveals the true signal configuration λ^N . However, for each Y_k , the type configuration for those types that belong to the set Y_k are unknown except that their total number is $\lambda_-^N(Y_k)$ agents; the distribution of those types is described by $\gamma_z^{\lambda^N}$. To represent this uncertainty and its asymptotic distribution, for each $z \in Z$, consider K infinite sequences of random variables $(\zeta^1, \zeta^2, \dots, \zeta^K)$ such that ζ_i^k is Y_k -valued for all $i \in \mathbb{N}$ and the K sequences are independent of each other and ζ^k is an i.i.d. sequence with marginal distribution $(\frac{\mu_z(y)}{\mu_z(Y_k)})_{y \in Y_k}$. Let γ_z denote the

joint distribution of $(\zeta^1, \zeta^2, \dots, \zeta^K)$. Thus, for for each N , $(\zeta^{1, \lambda_-^N(Y_1)}, \dots, \zeta^{K, \lambda_-^N(Y_K)})$, where $\zeta^{k, \lambda_-^N(Y_k)} = (\zeta_1^k, \dots, \zeta_{\lambda_-^N(Y_k)}^k)$ for each $k = 1, \dots, K$, is distributed according to $\gamma_z^{\lambda_-^N}$. Each ζ^k represents the type-realizations of other active agents whose stage-1 offers belong to Y_k . Given $(\zeta^1, \zeta^2, \dots, \zeta^K)$, for each $k = 1, \dots, K$ and each $y \in Y_k$, define

$$\rho^N(y) = \#\{\zeta_i^k : \zeta_i^k = y, i = 1, \dots, \lambda_-^N(Y_k)\} / \lambda_-^N(Y_k).$$

By the LLN, for each $y \in Y_k$, $\rho^N(y)$ converges to $\mu_z(y) / \mu_z(Y_k)$ in probability under γ_z for any k and for any z .

Now, let $p^N = R_-^N / Q_-^N$, where

$$(Q_-^N, R_-^N) = \sum_{k=1, \dots, K} \sum_{y \in Y_k} \lambda_-^N(Y_k) \rho^N(y) (\beta^N(y) + \varepsilon) + \varepsilon.$$

Notice that Q_-^N , R_-^N , and p^N only depends on $(\zeta^{1, \lambda_-^N(Y_1)}, \dots, \zeta^{K, \lambda_-^N(Y_K)})$. Then, the equilibrium offer $(\beta_q^N(\bar{x}), \beta_r^N(\bar{x}))$ solves $\max_{b \in \mathcal{O}} H_{\bar{x}}^N(b; \lambda^N, \beta^N)$, where

$$H_{\bar{x}}^N(b; \lambda^N, \beta^N) = \sum_{z \in Z} \phi_{\bar{x}}^N[z] \left[\mathbb{E}_{\gamma_z^{\lambda_-^N}} \left(u(\bar{q} + \frac{b_r Q_-^N - b_q R_-^N}{R_-^N + b_r}, \bar{r} + \frac{b_q R_-^N - b_r Q_-^N}{Q_-^N + b_q}; \bar{x}, z) \right) \right].$$

To show that $\lim_{N \rightarrow \infty} \beta^N = \beta^{\bar{z}}$, we show that any convergent subsequence $\{\beta^{N_t}\}_{t=1}^\infty$ has the same limit $\beta^{\bar{z}}$. The following claim shows that if $\lim_{t \rightarrow \infty} \beta^{N_t} = \beta^*$, then for the agent with type \bar{x} , his stage-2 problem converges to the optimization problem in the competitive equilibrium with known state-of-the-world \bar{z} .

Claim 1. Suppose that $\lim_{t \rightarrow \infty} \beta^{N_t}(x) = \beta^*(x)$ for each $x \in X$. Then,

$$\lim_{t \rightarrow \infty} H_{\bar{x}}^{N_t}(b; \lambda^{N_t}, \beta^{N_t}) = u(\bar{q} + \frac{b_r}{p^*} - b_q, \bar{r} - b_r + p^* b_q; \bar{x}, \bar{z})$$

uniformly in $b \in \mathcal{O}$.

Given Claim 1, $\beta^*(\bar{x})$ solves $\max_{b \in \mathcal{O}} u(\bar{q} + \frac{b_r}{p^*} - b_q, \bar{r} - b_r + p^* b_q; \bar{x}, \bar{z})$. This holds for all $\bar{x} \in X$. Hence, by lemma 2, β^* is a separating stage-2 equilibrium at state \bar{z} in $\mathcal{L}(\varepsilon)$. Thus, $\beta^* = \beta^{\bar{z}}$. Because this holds for any convergent subsequence, it follows that $\lim_{N \rightarrow \infty} \beta^{\lambda^N} = \beta^{\bar{z}}$.

Proof of Claim 1. First notice that

$$\lim_{t \rightarrow \infty} \left(\frac{Q_-^{N_t}}{M^{N_t}}, \frac{R_-^{N_t}}{M^{N_t}} \right) = \sum_{y \in X} \mu_{\bar{z}}(y) (\beta^*(y) + \varepsilon) \quad (31)$$

in probability under $\gamma_{\bar{z}}$. This follows directly from the definition of Q_-^N and R_-^N and the fact that, for all $y \in Y_k$, $\lim_{N \rightarrow \infty} \rho^N(y) = \mu_{\bar{z}}(y) / \mu_{\bar{z}}(Y_k)$ in probability under $\gamma_{\bar{z}}$, and, for all $k = 1, \dots, K$, $\lim_{N \rightarrow \infty} \lambda_-^N(Y_k) / M^N = \mu_{\bar{z}}(Y_k)$. Of course, this limit may not hold for other $z \neq \bar{z}$; we focus on \bar{z} because $\lim_{N \rightarrow \infty} \phi_{\bar{x}}^N[\bar{z}] = 1$. This also implies that

$$\lim_{t \rightarrow \infty} p^{N_t} = p^* = \frac{\sum_{y \in X} \mu_{\bar{z}}(y) (\beta_r^*(y) + \varepsilon_r)}{\sum_{y \in X} \mu_{\bar{z}}(y) (\beta_q^*(y) + \varepsilon_q)}$$

in probability under $\gamma_{\bar{z}}$. Moreover, $\lim_{t \rightarrow \infty} \frac{1}{Q_-^{N_t}} = \lim_{t \rightarrow \infty} \frac{1}{R_-^{N_t}} = 0$ in probability under $\gamma_{\bar{z}}$ as well.

Now,

$$H_{\bar{x}}^{N_t}(b; \lambda^{N_t}, \beta^{N_t}) = \sum_{z \in Z} \phi_{\bar{x}}^{N_t}[z] \left[\mathbb{E}_{\gamma_z^{\lambda^{N_t}}} \left(u \left(\bar{q} + \frac{b_r}{p^N + \frac{b_r}{Q_-^{N_t}}} - \frac{b_q}{1 + \frac{b_r}{R_-^{N_t}}}, \bar{r} + \frac{b_q p^{N_t}}{1 + \frac{b_q}{Q_-^{N_t}}} - \frac{b_r}{1 + \frac{b_q}{Q_-^{N_t}}}; \bar{x}, z \right) \right) \right].$$

Notice that

$$\left(\bar{q} + \frac{b_r}{p^N + \frac{b_r}{Q_-^{N_t}}} - \frac{b_q}{1 + \frac{b_r}{R_-^{N_t}}}, \bar{r} + \frac{b_q p^{N_t}}{1 + \frac{b_q}{Q_-^{N_t}}} - \frac{b_r}{1 + \frac{b_q}{Q_-^{N_t}}} \right) \in \left[0, \bar{q} + \frac{\bar{r}(\bar{q} + \varepsilon_q)}{\varepsilon_r} \right] \times \left[0, \bar{r} + \frac{\bar{q}(\bar{r} + \varepsilon_r)}{\varepsilon_q} \right]$$

for any N_t and any realization. Because the domain is compact and $u(\cdot; \bar{x}, \bar{z})$ is continuous, $u(\cdot; \bar{x}, \bar{z})$ is uniformly continuous over that domain. Thus, for any $\epsilon > 0$, there exists $\delta_1(\epsilon) > 0$ such that $|u(q, r; \bar{x}, \bar{z}) - u(q', r'; \bar{x}, \bar{z})| < \epsilon$ if $|q - q'| \leq \delta_1(\epsilon)$ and $|r - r'| \leq \delta_1(\epsilon)$. For any $\delta_1 > 0$, it is straightforward to check that there exists some $\delta_2(\delta_1) > 0$ such that if $|\frac{Q_-^{N_t}}{M^{N_t}} - \sum_{y \in X} \mu_z(y)(\beta_q^*(y) + \varepsilon_q)| < \delta_2(\delta_1)$ and $|\frac{R_-^{N_t}}{M^{N_t}} - \sum_{y \in X} \mu_z(y)(\beta_r^*(y) + \varepsilon_r)| < \delta_2(\delta_1)$, then

$$\left| \frac{b_r Q_-^{N_t} - b_q R_-^{N_t}}{R_-^{N_t} + b_r} - \left(\frac{b_r}{p^*} - b_q \right) \right| < \delta_1 \quad \text{and} \quad \left| \frac{b_q R_-^{N_t} - b_r Q_-^{N_t}}{Q_-^{N_t} + b_q} - (-b_r + p^* b_q) \right| < \delta_1 \quad (32)$$

for all $b \in \mathcal{O}$.

Now, fix some $\epsilon > 0$. Let t_1 be so large that $t > t_1$ implies that for $\delta = \delta_2(\delta_1(\epsilon/3))$, the event that

$$\left| \frac{Q_-^{N_t}}{M^{N_t}} - \sum_{y \in X} \mu_{\bar{z}}(y)(\beta_q^*(y) + \varepsilon_q) \right| < \delta \quad \text{and} \quad \left| \frac{R_-^{N_t}}{M^{N_t}} - \sum_{y \in X} \mu_{\bar{z}}(y)(\beta_r^*(y) + \varepsilon_r) \right| < \delta$$

has probability less than $1 - \frac{\epsilon}{3\bar{H}}$ under $\gamma_{\bar{z}}$ and $\phi^{N_t}[\bar{z}] > 1 - \frac{\epsilon}{3\bar{H}}$, where \bar{H} is the maximum value for $H_{\bar{x}}^N$ for all N . Such t_1 exists because of (31). Then, for such t ,

$$\begin{aligned} & \left| H_x^{N_t}(b; \lambda^{N_t}, \beta^{N_t}) - u \left(\bar{q} + \frac{b_r}{p^*} - b_q, \bar{r} - b_r + p^* b_q; x, \bar{z} \right) \right| \\ & < \frac{\epsilon}{3\bar{H}} \bar{H} + \phi^{N_t}[\bar{z}] \left[\left(1 - \frac{\epsilon}{3\bar{H}} \right) \frac{\epsilon}{3} + \frac{\epsilon}{3\bar{H}} \bar{H} \right] < \epsilon \end{aligned}$$

for all $b \in \mathcal{O}$. ■

5 Concluding remarks

Several aspects of our model deserve comment. We start with the mechanism, and, in particular, our modelling of payoffs for inactive agents.

As we noted at the outset, our mechanism violates feasibility. The payoffs of inactive agents, which we have assumed are determined by the execution of their stage-1 offers at the exogenous price, \bar{r}/\bar{q} , and the exogenous stage-2 offers, ε , give rise to a net payout of one of the goods. Any such payout could be financed by entry fees levied on all agents before types are realized. In particular, if the entry fee is 2ε , then η can be chosen to insure feasibility. And, provided there is sufficient motivation for trade coming from the appearance of types in the utility function, ε can be chosen to be small enough to induce participation.

A mechanism that would insure feasibility except for ε and would more closely resemble pari-mutuel betting would have the stage-1 offers of the inactive agents be part of the offers that determine the “price” in the second-stage market game and would have their payoffs determined as they are for active agents. However, that would give rise to two-way interaction between the stages. In such a version, if the economy is sufficiently large, agents at stage 1 would, as in our version, make stage-1 offers based on the presumption that they will be chosen to be inactive. Even so, they would want to predict the stage-2 price which, itself, is affected by their offers—both directly and by the information revealed by stage-1 offers. Thus, to get a fixed point, we would have to study a mapping that takes both stages into account. Moreover, the mapping would have to be defined over all feasible stage-1 actions, including stage-1 actions that give rise to asymmetric information at stage-2.

We make one strong assumption about preferences; namely, that complete-information competitive demand is monotone, which assures that there is a unique competitive equilibrium (CE) in the version with no uncertainty. If, instead, there were multiple CE’s, then our existence argument would fail if agents at stage 1 believed that their stage-1 actions would determine the limit to which a sequence of proposition 1 equilibria converges. If that were the case, then the influence of stage-1 actions on payoffs contingent on being active would not disappear as the size of the economy grows. One way to avoid such a belief would be to assume that there is coordination on the sequence of proposition 1 equilibria regarding the limit to which they converge. That would work if there is a sequence of proposition-1 equilibria that converges to any CE. Whether that is true seems not to be known even for complete-information versions of our market game. With a unique CE, that coordination issue does not arise.

We assume a finite number of types and a finite number of states-of-the-world. The former is important for us. Although the realization of types is random, as the size of the economy grows, conditional independence of types gives us something that resembles replication in a deterministic version. Even more important, our existence result, via the specification of beliefs, depends on a finite number of types.⁸ In contrast, a finite support for the state-of-the-world plays no role. Regarding the information structure, two special cases of the model deserve mention. For a specification in which the state-of-the-world does not appear in preferences, all our results apply and there remains an important role

⁸Reny and Perry [5] cannot use a specification with a finite number of types because they have a limit-order mechanism. With such a mechanism and a finite number of types, there can remain an indeterminacy regarding how the gains from trade are distributed. That does not happen for market-game mechanisms.

for two stages and for limits as the number of agents gets large. The same is true for a specification in which types do not appear as arguments of preferences. However, in that case, trade disappears at stage 2 as $\varepsilon \rightarrow 0$ and agents would not want to enter in the presence of an entry fee.

Like all of the previous work on strategic games for accomplishing trade with private information that we cited at the outset and all auction models, we have a two-good model. One well-known way to extend our model to $K + 1$ goods is to treat good $K + 1$ as *cash* and to have K simultaneous markets with market k having trade between cash and good k . However, then, as is well-known, we would want to have multiple rounds of trade because the proceeds of sales in one market cannot be used to make simultaneous purchases in another market. Such an extension with multiple rounds in real time is pertinent for what we see as the main potential application of our model. We can imagine our mechanism being used for spot trades in securities like the common stock of publicly traded companies. However, in order to use it for such trades, two timing questions have to be answered: how frequently should the market operate and what should be the length of time between stages 1 and 2? Any attempt to address those questions requires a dynamic version of our model.

References

- [1] Cripps, M. W. and J. M. Swinkels, Efficiency of large double auctions. *Econometrica*, 74 (2006) 47-92.
- [2] Dubey, P. S. and M. Shubik, A theory of money and financial institutions. 28. The non-cooperative equilibria of a closed trading economy with market supply and bidding strategies. *Journal of Economic Theory*, 17 (1978) 1-20.
- [3] Palfrey, T.R., Uncertainty resolution, private information aggregation and the Cournot competitive limit. *The Review of Economic Studies*, 52 (1985) 69-83.
- [4] Axelrod, B., B. Kulich, C. Plott, and K. Roust, Design improved parimutuel-type information aggregation mechanisms: inaccuracies and the long-shot bias as disequilibrium phenomena. *Journal of Economic Behavior and Organization*, 69 (2009) 170-181.
- [5] Reny, P.J. and M. Perry, Toward a strategic foundation for rational expectations equilibrium. *Econometrica*, 74 (2006) 1231-69.
- [6] Shapley, L.S. and M. Shubik, Trade using one commodity as a means of payment, *Journal of Political Economy*, 85 (1977) 937-68.
- [7] Xavier Vives, Aggregation of information in large cournot markets. *Econometrica*, 56 (1988) 851-76.
- [8] Xavier Vives, Strategic supply function competition with private information. *Econometrica*, 79 (2011) 1119-66.