Market Communication in Production Economies

Christopher A. Wilkens[⋆]

Computer Science Division,
University of California at Berkeley, CA, 94720 USA
`cwilkens@cs.berkeley.edu`

Abstract. We study the information content of equilibrium prices using the market communication model of Deng, Papadimitriou, and Safra [4]. We show that, in the worst case, communicating an exact equilibrium in a production economy requires a number of bits that is a quadratic polynomial in the number of goods, the number of agents, the number of firms, and the number of bits used to represent an endowment.

1 Introduction

In the European Union, prices are typically expressed in whole-Euro amounts (or as "nice" decimals when they are small). In contrast, buyers and sellers in the United states cling to every penny and advertise prices to the $\frac{1}{100}$-th of a dollar. Does such accuracy serve a computational purpose? We study this question in the case of market equilibrium: how many bits of information must prices express in order to ensure that the economy achieves equilibrium?

The *market communication* model of Deng et al. [4] highlights the unusual properties of communication in markets. In standard market models, communication often comes from central authority, such as a market maker or Walras's fictitious auctioneer [14]. This omniscient authority must broadcast enough information (e.g. prices) for each agent to decide his own behavior without further communication — because each agent has private information (e.g. an endowment), it may be that agents are ignorant of others' equilibrium allocations. By comparison, in Yao's basic two-party model [15], two players follow a protocol (where both may send information) to communicate enough information that both players know the answer to the problem. Here, we study the communication requirements of reaching equilibrium in the market communication model.

Classical economic treatment of communication costs studies the *dimensionality of the message space* required to communicate a Pareto-efficient outcome. In standard convex economies, the seminal work of Arrow and Debreu [1] may be interpreted as a proof that $(m-1)$ real numbers — i.e. normalized prices — are sufficient. Subsequent work [6, 10] shows that normalized prices are optimal. A priori, the results for convex economies are powerful because the amount of

---

[⋆] Supported by a fellowship from the University of California at Berkeley and NSF grant CCF-0635319.

communication is independent of the number of agents and firms. Many subsequent works have sharpened and extended these results [7, 2, 13]. Of particular relevance, Calsamiglia's introduction of parametric communication [2] precisely captures the notion that communication may leverage private information to reduce communication.

Our work focuses on the bit-wise communication requirements for reaching equilibrium — while $(m-1)$ real numbers may be dimensionally optimal, they may hide many bits. Since most real-world applications communicate a price with fixed precision, we follow Deng et al. [4] in believing that bit-wise communication bounds are important. Related communication complexity results [12, 11] consider the problem of communicating preferences or complete allocations, while most research on market equilibria has focused on developing efficient algorithms (e.g. [5, 8, 3]). To the best of our knowledge, Deng et al. give the only result specifically applicable to this model.

Our main result gives a lower bound on the number of bits of information that must be communicated in an Arrow-Debreu market with production. We show that the number of bits depends polynomially on the number of agents, the number of firms, and the amount of private information they hold.

Our bound is significantly stronger than the bound of Deng at al. [4]. First, Deng et al. need $\Theta(\frac{n}{m})$-bit numbers to show a $poly(n)$ lower bound, i.e. they give each agent polynomially many bits of private information. We achieve the same lower bound with a logarithmic number of such bits. Second, Deng et al. must relax the standard non-satiation requirement on utility functions;[1] we do not. Thirdly, our bound is more general because it considers a production economy.

The main shortcoming of our bound is that it critically exploits the fact that real numbers rarely sum to the same value, even if they are very close. Thus, it is unlikely to extend to approximate equilibria.

2 Markets and Market Communication

Market communication complexity aims to study the amount of information that prices must encode to induce equilibrium in an Arrow-Debreu economy [1].

2.1 Arrow-Debreu Markets

An Arrow-Debreu market with production consists of n agents, m goods, and l production firms (indexed by i, j, and k respectively). A bundle of goods is a vector $\boldsymbol{x} \in \mathbb{R}^m$ where x_j represents a quantity of good j.

Each agent has a utility function and an endowment. The utility function $u_i(\boldsymbol{x}_i) : \mathbb{R}^M \to \mathbb{R}$ maps bundles of goods to utilities, and the endowment $\boldsymbol{e}_i \in \mathbb{R}^m$ is a bundle of goods. In order to guarantee the existence of an equilibrium, it is sufficient to assume that u_i is strictly concave in x.

---

[1] They call it "strict concavity." Nonsatiation is required for Arrow and Debreu's proof of the existence of equilibrium [1].

A production firm is specified by a set of net production possibilities $Y_k \subset \mathbb{R}^M$. A vector $\boldsymbol{y}_k \in Y_k$ represents the net quantities of goods produced: a positive value $y_{j,k}$ represents an output of good j, and a negative value $y_{j,k}$ represents an input. Notice that at prices $\boldsymbol{\pi}$, the profit of firm k may be written as $\boldsymbol{\pi} \cdot y_k$. Again, the sets Y_k must satisfy convexity requirements. In particular, it is sufficient to assume the following: Y_k is closed, convex, and contains the 0 vector, and if $y \in \bigcup_k Y_k$, then $-y \in \bigcup_k Y_k$ if and only if $y = 0$.

To link production to consumption, a firm is owned by agents. Agent i may own a share $\alpha_{i,k} \in [0, 1]$ of the profits of firm k, i.e. at prices $\boldsymbol{\pi}$, agent i's budget will be the value of his endowment plus the profit derived from firms he owns, i.e.

$$M_i = \boldsymbol{\pi} \cdot \boldsymbol{e}_i + \sum_{k \in [l]} \sigma_{i,k} \boldsymbol{\pi} \cdot \boldsymbol{y}_k \quad . \tag{1}$$

Since $\sigma_{i,k}$ denotes a share of firm k, it must be that $\sum_{i \in [n]} \sigma_{i,k} = 1$. We omit the precise restrictions on production sets and utility functions for brevity.

The following economic definitions are standard [9]:

Definition 1. *An* economic allocation *is a tuple* $(\{\boldsymbol{x}_i\}, \{\boldsymbol{y}_k\})$ *specifying the bundle* $\boldsymbol{x}_i$ *consumed by each agent and the production vector* $\boldsymbol{y}_k$ *chosen by each firm.*

Definition 2. *An economic allocation is* feasible *if* $\boldsymbol{x}_i \geq 0$, $\boldsymbol{y}_k \in Y_k$, *and the total demand is less than or equal to the total supply, i.e.*

$$\sum_{i \in [n]} \boldsymbol{x}_i \leq \sum_{i \in [n]} \boldsymbol{e}_i + \sum_{j \in [m]} \boldsymbol{y}_k \tag{2}$$

Definition 3. *A* competitive equilibrium *(hereafter equilibrium) in an Arrow-Debreu market is a set of prices* $\boldsymbol{\pi} \in \mathbb{R}^M$ *and a feasible allocation* $(\{\boldsymbol{x}_i\}, \{\boldsymbol{y}_k\})$ *such that agents maximize their utilities and firms maximize their profits at current prices, i.e.*

$$\boldsymbol{x}_i \in \operatorname*{arg\,max}_{\boldsymbol{x} \in \{\boldsymbol{x} | \boldsymbol{x} \cdot \boldsymbol{\pi} \leq \boldsymbol{e}_i \cdot \boldsymbol{\pi}\}} u_i(\boldsymbol{x}) \tag{3}$$

$$\boldsymbol{y}_k \in \operatorname*{arg\,max}_{\boldsymbol{y} \in Y_k} \boldsymbol{\pi} \cdot \boldsymbol{y} \quad . \tag{4}$$

2.2 Market Communication

Deng et al. [4] define the market communication model as follows:

Definition 4. Market Communication*: n agents $[n]$ have private information $x_i \in X_i$ (the sets X_i are common knowledge). Agent i wishes to compute the function $f_i(x_1, \ldots x_n)$. Another agent, agent 0 (the "invisible hand"), knows $(x_1, \ldots x_n)$.*

A protocol is a set of functions $(g_0(\cdot), g_1(\cdot), \ldots g_n(\cdot))$ where $g_0 : X_1 \times \ldots X_n \to X_0$, $g_{i \in [n]} : X_0 \times X_i \to \mathbb{R}$, and $g_i(g_0(x_1, \ldots x_n), x_i) = f_i(x_1, \ldots x_n)$. The amount of market communication is the number of bits in $x_0 = g_0(x_1, \ldots x_n)$.

In essence, the omniscient agent 0 computes $x_0 = g_0(x_1, \ldots x_n)$ and broadcasts x_0 to agents $i \in [n]$. Next, each agent privately uses x_i to compute $g_i(x_0, x_i) = f_i(x_1, \ldots x_n)$.

The Power of Market Communication The addition of an omniscient agent substantially increases the model's power: it is as powerful as standard nondeterministic communication.

Theorem 1. *Assume communication costs are measured in bits. Then any problem $f(x_1, \ldots x_n)$ in NP^{CC} has an efficient market communication protocol.*

Proof. By assumption, there is a communication sequence σ of poly-logarithmic length that solves the problem. Let $T = \{(i_t, \sigma_t)\}$ be a transcript of the communication, i.e. agent i_t sent σ_t at time t.

Note that agent 0 may compute T because she is omniscient. Thus, in the market communication protocol, agent 0 computes T and broadcasts it to the agents. Each agent then simulates his behavior based on T to solve the problem. The size of i_t is $\log n$, so $|T| = \Theta(|\sigma| \log n)$, thus giving an efficient market communication protocol. □

Market Communication in Arrow-Debreu Markets We wish to discuss the number of bits of private information an agent or firm receives; however, such private information is often given in terms of real numbers or functions. To generate a meaningful measure of each agent's private information, we assume that utility functions and production sets are drawn from finite sets.

Specifically, an agent's utility function u_i is drawn from a finite set $\mathbb{U}$. Also, an agent's endowment is a vector of dimension m in which each coordinate is represented in β bits. Similarly, a firm's production set Y_k is drawn from a finite set $\mathbb{Y}$. Our bound will be a function of the number of possible utility functions $|\mathbb{U}|$ and the number of possible production sets $|\mathbb{Y}|$.

The goal of an agent or firm is to compute its consumption vector $\boldsymbol{x}_i$ or production vector $\boldsymbol{y}_k$. Thus, if E represents all private information in the economy, we have $\boldsymbol{g}_i = \boldsymbol{x}_i(E)$ for the agents and $\boldsymbol{g}_k = \boldsymbol{y}_k(E)$ for the firms. (While the definition of an equilibrium includes prices, we take the position that prices are merely a communication tool and that, at the end of the day, we only care if each agent chooses the correct allocation. Thus, we do not explicitly require agents to compute prices as part of g_i.)

For example, a trivial protocol might broadcast each agent's utility function and each firm's production set, using $O(n \log |\mathbb{U}| + l \log |\mathbb{Y}|)$ bits of communication. Agents would then know everything and, therefore, could compute an equilibrium.

3 A Lower Bound for the Arrow-Debreu Model

Our main result shows that the number of bits is polynomial in the number of goods, agents, firms, and bits of private information.

Theorem 2. *In the worst case, communicating a market equilibrium in the market communication model requires at least*

$$\frac{m}{2} \left(\beta + \lg(n-1)\right) + n + l - O(1) \tag{5}$$

bits of communication to reach equilibrium, where n is the number of agents, m is the number of goods, l is the number of firms, and β is the number of bits used to represent a value in an agent's endowment.

The $(n+l)$ term is the most significant — it implies that the number of bits is linear in the number of agents and production firms. From a practical perspective, it is unrealistic to believe that prices contain $(n+l)$ bits of information.

The $\frac{m}{2}(\beta + \lg(n-1))$ term corresponds to communicating the total global endowment of resources. The total endowment of each resource is, in general, a $(\beta + \lg n)$-bit number. Thus, this term roughly corresponds to communicating the total endowment of $\frac{m}{2}$ goods.

Instead of proving the theorem directly, we prove a more general lemma that allows us to compare our bound to Deng et al. [4]. They achieve an $\Omega(n \log(m+n))$ lower bound using an exponentially large set of utility functions that require $poly(m,n)$-bit numbers, i.e. $|\mathbb{U}| = \Omega\left(2^{poly(m,n)}\right)$ and $\beta = poly(m,n)$. By comparison, taking the number of utility functions and production functions to be polynomial in our construction (i.e. $|\mathbb{U}| = |\mathbb{Y}| = poly(m+n)$) gives the same $\Omega(n \log(m+n))$ bound.

Moreover, if we allow players and firms to value arbitrary bundles of goods, then we get $|\mathbb{U}| = |\mathbb{Y}| = \Omega(2^m)$ and thus a lower bound of

$$\Omega\left(\frac{m}{2}(\beta + \lg(n-1)) + m \cdot (n+l) - m\right) \quad . \tag{6}$$

In this case, each item requires $(n+l)$ bits of information in its price. (While attributing an arbitrary bundle value to an item is slightly unrealistic, it is a common worst case setting in areas such as combinatorial auctions.)

Lemma 1. *Communicating a market equilibrium in the market communication model requires at least*

$$\frac{m}{2}(\beta + \lg(n-1)) + (n - 1 - |\mathbb{U}|^{\frac{4}{m}}) \lg |\mathbb{U}| + (l - 1 - |\mathbb{Y}|^{\frac{4}{m}}) \lg |\mathbb{Y}| \tag{7}$$

bits of communication in the worst case, where $|\mathbb{U}|$ is the number of possible utility functions u_i, and $|\mathbb{Y}|$ is the number of possible production sets Y_k.

Moreover, there are specific sets $\mathbb{U}$ and $\mathbb{Y}$ such that the economy requires

$$\frac{m}{2}(\beta + \lg(n-1)) + n + l - O(1) \tag{8}$$

bits of communication to reach equilibrium.

Proof. We construct an economy with m goods, n agents, and l firms. The main trick is to make each combination of utility functions (or production functions) correspond to a unique prime factorization. Thus, no two combinations of utility functions (or production functions) will have the same optimal allocation.

The second trick is to leave one agent (and firm) without any private information, so the number of communication sequences is trivially lower-bounded by the number of possible equilibrium choices she may make.

The Economy. Assume m is divisible by 4, $n \geq 2$, and $l \geq 2$.

Partition the goods into four groups modulo 4, i.e.

$$M_a = \{j | j \in [m] \text{ and } j \equiv a \mod 4\} \tag{9}$$

The sets will serve the following purposes:

- Goods in M_0 are production inputs and goods in M_1 are outputs. Nobody wants goods in M_0. Consequently, the entire supply of goods M_0 is converted to goods in M_1. Goods in M_1 are indistinguishable to the agents, so Pareto-optimality will imply that producers maximize the total output of goods in M_1.
- Goods in M_2 and M_3 are traded among agents. Goods are paired such that an agent balances the quantity of a good $m_2 \in M_2$ with some good $m_3 \in M_3$ to match marginal utilities.

Agents $i > 1$ have utility functions of the form

$$u_i(\boldsymbol{x}_i) = \sum_{j \in M_3} \left(2\sqrt{x_{i,j} \lg c_{i,j}} + x_{i,j-1}\right) + \sum_{j \in M_1} x_{1,j} \tag{10}$$

where $c_{i,j} \in C$, and the set C will be determined later. Agents $i > 1$ are endowed with goods from M_0, M_2 and M_3 only, i.e.

$$e_{i,j} = \begin{cases} e_{i,j}, & j \in M_0 \cup M_3 \\ \bar{e}, & j \in M_2 \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

For endowed goods, $e_{i,j} \in [2^\beta]$ is a β-bit integer and $\bar{e} = n \cdot 2^\beta$ is a large number (large enough that, in equilibrium, an agent will always keep a positive quantity of each good in M_2).

Agent 1 has the utility function

$$u_1(\boldsymbol{x}_1) = \sum_{j \in M_3} \left(2\sqrt{x_{1,j}} + x_{1,j-1}\right) \tag{12}$$

(the first term is equivalent to setting $c_{1,j} = 2$). She is endowed with 1 unit each of goods in M_2, and M_3, i.e.

$$e_{1,j} = \begin{cases} 1, & j \in M_2 \cup M_3 \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

Note that agent 1 has no private information.

The firms have technology to convert goods in M_0 to goods in M_1. Like agents' utilities, the production functions are parameterized by coefficients $c_{j,k} \in C$. (We will take C to be the same for both parts of the economy, but this is certainly not necessary.) We define the production of firm k in terms of a production function, i.e. firm k may transform $y_{j-1,k}$ units of good $(j-1)$ into $y_{j,k}$ units of good j according to the following function $f_{j,k}$:

$$y_{j,k} = f_{j,k}(y_{j-1,k}) = 2\sqrt{y_{j-1,k} \cdot \lg c_{j,k}} \tag{14}$$

This function is translated to a set of vectors to match the model.[2] In order to create a firm with no private information, we require that $c_{j,1} = 2$. For simplicity, we also specify that all firms are owned by agents $i > 1$.

Analysis. First, we show that agent 1 must be able to select

$$\left((n-1)2^{\beta}\right)^{\frac{m}{4}} |\mathbb{U}|^{n-1-|\mathbb{U}|^{\frac{4}{m}}} \tag{15}$$

distinct consumption vectors. A similar proof gives a lower bound for the production side of the economy.

Consider a good $j \in M_3$. (Note that good $(j-1)$ is in M_2.) Let $p_j = \frac{\pi_j}{\pi_{j-1}}$ be the relative price of good j compared to good $(j-1)$. Note that each agent has a term of the form $2\sqrt{x_{i,j} \lg c_{i,j}} + x_{i,j-1}$ in u_i. In equilibrium, we know that agent i does not wish to sell good $j-1$ to get good j (or vice-versa). Thus, agent i must balance her marginal utilities from the $2\sqrt{x_{i,j} \lg c_{i,j}}$ and $x_{i,j-1}$ terms. This gives the relation

$$\frac{\partial\left(2\sqrt{x_{i,j} \lg c_{i,j}}\right)}{\partial x_{i,j}} = \frac{\partial\left(p_j x_{i,j-1}\right)}{\partial x_{i,j-1}} \tag{16}$$

$$\sqrt{\frac{\lg c_{i,j}}{x_{i,j}}} = p_j \tag{17}$$

$$x_{i,j} = \frac{\lg c_{i,j}}{(p_j)^2} \tag{18}$$

(Note that by construction, i.e. by choice of $\bar{e}$, this is always possible.) Since goods in M_2 and M_3 do not involve production, we know that

$$\sum_{i \in [n]} x_{i,j} = \sum_{i \in [n]} e_{i,j} \quad . \tag{19}$$

Let $\alpha_j = \sum_{i \in [n]} e_{i,j}$. Using this constraint and the equations $x_{i,j} = \frac{\lg c_{i,j}}{(p_j)^2}$, it follows that

$$p_j^2 = \frac{1}{\alpha_j} \sum_{j \in M_3} \lg c_{i,j} = \frac{1}{\alpha_j} \lg\left(\prod_{j \in M_3} c_{i,j}\right) \quad , \tag{20}$$

and thus

$$x_{i,j} = \frac{\alpha_j \lg c_{i,j}}{\lg\left(\prod_{j \in M_3} c_{i,j}\right)} \quad . \tag{21}$$

---

[2] The only trick to converting $f_{j,k}$ to a set is to allow firm k to produce any amount of good j between 0 and $f_{j,k}$. In equilibrium, production will always occur on the boundary defined by $f_{j,k}$, so this change is inconsequential.

For agent 1, we get

$$x_{1,j} = \frac{1}{\lg\left(\left(\prod_{j \in M_3 \setminus 1} c_{i,j}\right)^{\frac{1}{\alpha_j}}\right)} \quad . \tag{22}$$

To show a lower bound, we want to show that we can choose the set C such that the number of possible values for $x_{1,j}$ is large. We take C to be the $|C|$ smallest primes. (Note that this implies the total number of possible utility functions is $|\mathbb{U}| = |C|^{\frac{m}{4}}$.) To count the number of possible values for $x_{1,j}$, consider the value

$$\prod_{j \in M_3 \setminus 1} (c_{i,j})^{\frac{1}{\alpha_j}} \quad . \tag{23}$$

This represents a fractional prime factorization of a number where the prime factors are in C. Thus, the number of distinct values is the number of distinct sets of the form

$$\left\{ \frac{k_{c,j}}{\alpha_j} \right\} \tag{24}$$

where $k_{c,j}$ is the number of times a prime c occurs in the factorization and $\sum_i k_{c,j} = n$. Note that if two sets $\left\{\frac{k_{c,j}}{\alpha_j}\right\}$ and $\left\{\frac{k'_{c,j}}{\alpha'_j}\right\}$ are the same, then $\sum \frac{k_{c,j}}{\alpha_j} = \sum \frac{k'_{c,j}}{\alpha'_j}$. Since $\sum_i k_{c,j} = n$ is fixed, the only way for two sets to be the same is if $\alpha_j = \alpha'_j$. Thus, since there are $(n-1)2^\beta$ possible values for α_j, there are

$$(n-1)2^\beta \frac{|C|^{n-1}}{|C|!} \geq (n-1)2^\beta |C|^{n-|C|-1} \tag{25}$$

possible values for this quantity, and, therefore, the same number of possible values for $x_{1,j}$. Counting over all $\frac{m}{4}$ goods in M_3 and assuming that C is the same for all j (thus $|\mathbb{U}| = |C|^{\frac{m}{4}}$), it follows that agent 1 has at least

$$\left((n-1)2^\beta |C|^{n-|C|-1}\right)^{\frac{m}{4}} = \left((n-1)2^\beta\right)^{\frac{m}{4}} |\mathbb{U}|^{n-1-|\mathbb{U}|^{\frac{4}{m}}} \tag{26}$$

possible choices. Moreover, note that we may derive this bound for other sizes $|\mathbb{U}|$ by fixing all c_{ij} for some j. Assume that we fix the values of c_{ij} for $\left(\frac{m}{4} - k\right)$ goods (i.e. k goods still have c_{ij} drawn from C), then we get

$$\left((n-1)2^\beta\right)^{\frac{m}{4}} |\mathbb{U}|^{n-1-|\mathbb{U}|^{\frac{1}{k}}} \tag{27}$$

where $|\mathbb{U}| = |C|^k$. (This will be useful when we wish to set $|\mathbb{U}| = 2$.)

The analysis for the firms is similar: firm 1 must be able to select

$$\left((n-1)2^\beta\right)^{\frac{m}{4}} |\mathbb{Y}|^{l-1-|\mathbb{Y}|^{\frac{4}{m}}} \tag{28}$$

distinct production vectors. First, we characterize optimal production. Consider a single good $j \in M_1$ and observe that

$$\sum_{k \in [l]} y_{j-1,k} = \sum_{i \in [n]} e_{i,j-1} = \alpha_j \quad . \tag{29}$$

Let $p_{j-1} = \frac{\pi_{j-1}}{\pi_j}$ be the equilibrium price of good $(j-1)$ relative to good j. Then we know that firm k maximizes

$$y_{j,k} - p_{j-1} y_{j-1,k} = 2\sqrt{y_{j-1,k} \cdot \lg c_{j,k}} - p_{j-1} y_{j-1,k} \tag{30}$$

Taking the first derivative with respect to $y_{j-1,k}$ implies that $y_{j,k} = \frac{\lg c_{j,k}}{(p_{j-1})^2}$, so we repeat the analysis used for agent 1. This shows that firm 1 must be able to make at least

$$\left((n-1)2^{\beta}\right)^{\frac{m}{4}} |\mathbb{Y}|^{l-1-|\mathbb{Y}|^{\frac{4}{m}}} \tag{31}$$

different choices.

Because the choices of agent 1 and firm 1 are independent, all combinations of choices are possible. Thus, the total number of communication sequences must be at least

$$\left((n-1)2^{\beta}\right)^{\frac{m}{2}} |\mathbb{U}|^{n-1-|\mathbb{U}|^{\frac{4}{m}}} |\mathbb{Y}|^{l-1-|\mathbb{Y}|^{\frac{4}{m}}} \tag{32}$$

and the total number of bits of communication is at least

$$\frac{m}{2}\left(\beta + \lg(n-1)\right) + (n - 1 - |\mathbb{U}|^{\frac{4}{m}}) \lg |\mathbb{U}| + (l - 1 - |\mathbb{Y}|^{\frac{4}{m}}) \lg |\mathbb{Y}| \tag{33}$$

If we take $|\mathbb{U}| = |\mathbb{Y}| = 2$ (i.e. the set $C = \{2, 3\}$ for one good, fixed at $c_{ij} = 2$ otherwise), then we get a lower bound of

$$\frac{m}{2}\left(\beta + \lg(n-1)\right) + n + l - O(1) \tag{34}$$

bits of communication. □

4 Conclusion

Our main theorem diminishes the power of prices, with the caveat that we demand an exact equilibrium. While $(m-1)$ prices are sufficient, the amount of information they contain may be highly dependent on the parameters of the market.

Most significantly, the number of bits of information they must communicate is linear in the number of agents and firms in the worst case (the $(n+l)$ term). This implies that even though a price is supposed to be “universal,”prices must contain a unique bit of information for every agent in the economy. In the context of decimal prices, this roughly translates to one digit for every four buyers of a good. This is quite impractical.

It remains an open problem to give tight bounds. For example, we currently do not have any nontrivial upper bounds. Also, there are a few reasons why our lower bound may not be tight. First, instead of a multiplicative factor of $\frac{m}{2}$, one might expect a multiplicative factor of $(m-1)$ since that is the number of prices that must be communicated. Second, the multiplicative $\log(m+n)$ factor shown by Deng et al. [4] arises from an effect not present in our construction.

A more significant open problem is to give lower bounds for communicating approximate equilibria. Since our construction is highly dependent on the fact that two sets of irrational numbers rarely sum to the same value, it is unlikely to survive when an approximate equilibrium is sufficient. In particular, it is plausible that the polynomial dependence on n and l fundamentally requires $poly(n, l)$-bit numbers.

Furthermore, lower bounds for approximate equilibria would be more realistic. Because market clearing is also measured to finite precision, a lower bound approximate equilibria would give a stronger result on the amount of precision required in prices.

References

1. Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
2. X. Calsamiglia. Informational requirements of parametric resource allocation processes. In T. Groves, R. Radner, and S. Reiter, editors, *Information, Incentives, and Economic Mechanisms*. Basil Blackwell, 1988.
3. Richard Cole and Lisa Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 315–324, New York, NY, USA, 2008. ACM.
4. Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. On the complexity of equilibria. In *STOC '02: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 67–71, New York, NY, USA, 2002. ACM.
5. Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *FOCS '02: Proceedings of the 43rd Symposium on Foundations of Computer Science*, pages 389–395, Washington, DC, USA, 2002. IEEE Computer Society.
6. Leonard Hurwicz. Optimality and informational efficiency in resource allocation processes. In K.J. Arrow, S. Karlin, and P. Suppes, editors, *Mathematical Methods in the Social Sciences*, pages 27–46. Stanford University Press, Stanford, CA, 1960.
7. J. S. Jordan. The competitive allocation process is informationally efficient uniquely. *Journal of Economic Theory*, 28(1):1 – 18, 1982.
8. Sanjiv Kapoor, Aranyak Mehta, and Vijay Vazirani. An auction-based market equilibrium algorithm for a production model. *Theor. Comput. Sci.*, 378(2):153–164, 2007.
9. Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory.* Oxford University Press, June 1995.
10. Kenneth Mount and Stanley Reiter. The informational size of message spaces. *Journal of Economic Theory*, 8(2):161–192, June 1974.
11. Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA, 2007.

12. Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, July 2006.
13. Guoqiang Tian. On the informational requirements of decentralized pareto-satisfactory mechanisms in economies with increasing returns. October 2006.
14. Leon Walras. *Elements of pure economics, or, The theory of social wealth / Leon Walras ; translated by William Jaffe*. Published for the American Economic Association and the Royal Economic Society by Allen and Unwin, London :, 1954.
15. Andrew Chi-Chih Yao. Some complexity questions related to distributive computing(preliminary report). In *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213, New York, NY, USA, 1979. ACM.