The Computational Complexity of Games, Equilibria, and Fixed Points

Kousha Etessami

LFCS, School of Informatics
University of Edinburgh

Appetizer

- What is the computational complexity of computing a Nash Equilibrium in a finite game?

Appetizer

- What is the computational complexity of computing a Nash Equilibrium in a finite game?
- What is the complexity of computing a Market price equilibrium, in a finite market consisting of consumers, (producers,) and commodities?

Appetizer

- What is the computational complexity of computing a Nash Equilibrium in a finite game?
- What is the complexity of computing a Market price equilibrium, in a finite market consisting of consumers, (producers,) and commodities?
- What is the complexity of computing a fixed point for an algebraically defined Brouwer function, $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$?

Appetizer

- What is the computational complexity of computing a Nash Equilibrium in a finite game?
- What is the complexity of computing a Market price equilibrium, in a finite market consisting of consumers, (producers,) and commodities?
- What is the complexity of computing a fixed point for an algebraically defined Brouwer function, $F : \mathbb{R}^n \longrightarrow \mathbb{R}^n$?
- How are these and other fundamental computational problems related?

Appetizer

- What is the computational complexity of computing a Nash Equilibrium in a finite game?
- What is the complexity of computing a Market price equilibrium, in a finite market consisting of consumers, (producers,) and commodities?
- What is the complexity of computing a fixed point for an algebraically defined Brouwer function, $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$?
- How are these and other fundamental computational problems related?
- Why are these problems important?

Outline of lecture

- Background: Games, Markets, Equilibria, Brouwer Fixed Points.
- Weak vs. Strong approximation of Fixed Points.
- Scarf's classic algorithm, and its complexity implications.
- The complexity class PPAD, and weak approximation.
- PPAD-completeness results for ϵ-Nash, and 2-player Nash.
- Hardness of strong approximation: square-root-sum & arithmetic circuits.
- A new complexity class: **FIXP**. Nash is FIXP-complete.
- linear-FIXP $=$ PPAD.
- Other FIXP problems:
  market equilibria, *stochastic games*, *branching processes*...
- Conclusions and future challenges.

some background and motivation

- Game theory is a key foundation for mathematical economics. Nash equilibria are the central "solution concept" of game theory: they are the basic predicted/prescribed "outcomes".
- General equilibrium theory extends the equilibrium concept to a vast variety of economies and markets. (Far reaching generalization of "supply = demand".) Again, market (price) equilibria are the basic predicted/prescribed "outcomes" of market interactions.
- Many phenomena on the internet (online markets, ad word auctions, selfish routing,...) are best studied in a game-theoretic/economic framework, as interactions between collections of self-interested agents.
- Because of the huge scales of the internet, algorithmic and computational complexity issues can not be ignored:

  If equilibria are the outcomes predicted by economic theory, how hard is it to find one? (I.e., what is the complexity?)

two quotes

Kamal Jain (MSR), 2006: "If your laptop can't find it, then neither can the market."

Mas-Colell,Whinston, & Green, 1995 (standard graduate text in Microeconomic Theory): "A characteristic feature [of] economics is that for us the equations of equilibrium constitute the center of our discipline. By contrast, other sciences put more emphasis on the dynamic laws of change. The reason... is that economists are good at recognizing a state of equilibrium, but are poor at predicting precisely how an economy in disequilibrium will evolve..."

Or, to paraphrase: "Modern economic theory is largely non-algorithmic." By contrast, computer science is very good at "dynamics" (algorithmics).

Algorithmic Game Theory is an active research field at the intersection of CS and economics that, broadly speaking, aims to remedy this difficiency of modern economic theory.

Reminder: standard Complexity Classes ("the complexity zoo")

How complexity theorists classify "difficulty" of computational problems:

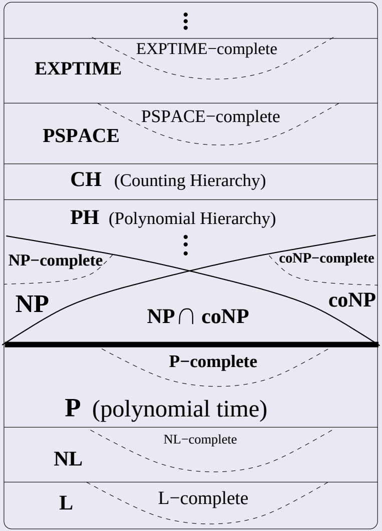

- The most important dichotomy in complexity theory is:

  NP-hard ("intractable")
  vs.
  in **P** ("tractable")

  (The **P** $\neq$ **NP** conjecture says **NP**-hard problems are not in **P**.)
- But a variety of important game-theoretic/economic problems have resisted such a classification:

  they are neither known to be in **P**, nor known to be **NP**-hard.

- The most important dichotomy in complexity theory is:

  NP-hard ("intractable")

  vs.

  in **P** ("tractable")

  (The **P** $\neq$ **NP** conjecture says **NP**-hard problems are not in **P**.)
- But a variety of important game-theoretic/economic problems have resisted such a classification:

  they are neither known to be in **P**, nor known to be **NP**-hard.
- This hasn't stopped us from trying to "classify" them.

Finite Games

A finite (normal form) game, Γ, consists of:

- A set $N = \{1, \ldots, n\}$ of players.
- Each player $i \in N$ has a finite set $S_i = \{1, \ldots, m_i\}$ of (pure) strategies. Let $S = S_1 \times S_2 \times \ldots \times S_n$.
- Each player $i \in N$, has a payoff (utility) function:

$$u_i : S \mapsto \mathbb{Q}$$

- A mixed strategy, $x_i = (x_{i,1}, \ldots, x_{i,m_i})$, for player i is a probability distribution over S_i.
  A profile of mixed strategies: $x = (x_1, \ldots, x_n)$
  Let X denote the set of all profiles.
- The expected payoff for player i:

$$U_i(x) = \sum_{s=(s_1,\ldots,s_n)\in S} \left(\prod_{k=1}^{n} x_{k,s_k}\right) u_i(s)$$

- Let x_{-i} denote everybody's strategy in x except player i's.
  Let $(x_{-i}; y_i)$ denote the new profile: $(x_1, \ldots, x_{i-1}, y_i, x_{i+1}, \ldots, x_n)$.

Nash Equilibria

A mixed strategy profile x is called:

- a Nash Equilibrium if:
  $\forall\, i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i)$

  In other words: *No player can increase its own payoff by unilaterally switching its strategy.*

- a ϵ-Nash Equilibrium, for $\epsilon > 0$, if:
  $\forall\, i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i) - \epsilon$

  In other words: *No player can increase its own payoff by more than ϵ by unilaterally switching its strategy.*

Nash Equilibria

A mixed strategy profile x is called:

- a Nash Equilibrium if:
  $\forall\ i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i)$

  In other words: *No player can increase its own payoff by unilaterally switching its strategy.*

- a ϵ-Nash Equilibrium, for $\epsilon > 0$, if:
  $\forall\ i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i) - \epsilon$

  In other words: *No player can increase its own payoff by more than ϵ by unilaterally switching its strategy.*

Theorem (Nash, 1950)

Every finite game has a Nash Equilibrium.

Market Equilibria

A simple Exchange Economy

- n agents and m commodities.
- Each agent j starts off with an initial endowment of commodities $w_j = (w_{j,1}, \ldots, w_{j,m})$.
- Each agent has a utility function, $u_i(x)$, that defines how much it likes different bundles $x \in \mathbb{R}^m_{\geq 0}$ of commodities.
- For a given price vector, $p \geq 0$, each agent j has an optimal demand vector $d^j(p) \in \mathbb{R}^m_{\geq 0}$ for commodities.
  This demand maximizes its utility using the budget obtained by selling all its endowment w_j at the prevailing at prices p.
  Under certain conditions on utility functions (e.g., continuity and strict quasi-concavity, and local non-satiation) demands $d^j(p)$ are uniquely determined continuous functions of prices.

The Arrow-Debreu Theorem

Price Equilibrium for an exchange economy

A vector of prices $p^* \geq 0$ is a price equilibrium if at prices p^* everybody can sell all of their goods into the market and buy precisely their optimal demand bundle of goods with the money earned, such that the market clears (i.e., no goods are left in the marketplace: supply $=$ demand).

The Arrow-Debreu Theorem

Price Equilibrium for an exchange economy

A vector of prices $p^* \geq 0$ is a price equilibrium if at prices p^* everybody can sell all of their goods into the market and buy precisely their optimal demand bundle of goods with the money earned, such that the market clears (i.e., no goods are left in the marketplace: supply $=$ demand).

Remarkably....

The Arrow-Debreu Theorem

Price Equilibrium for an exchange economy

A vector of prices $p^* \geq 0$ is a price equilibrium if at prices p^* everybody can sell all of their goods into the market and buy precisely their optimal demand bundle of goods with the money earned, such that the market clears (i.e., no goods are left in the marketplace: supply $=$ demand).

Remarkably....

Theorem (Arrow-Debreu, 1954])

Every exchange economy has a market price equilibrium.

(They proved a much much more general theorem.)

The Arrow-Debreu Theorem

Price Equilibrium for an exchange economy

A vector of prices $p^* \geq 0$ is a price equilibrium if at prices p^* everybody can sell all of their goods into the market and buy precisely their optimal demand bundle of goods with the money earned, such that the market clears (i.e., no goods are left in the marketplace: supply $=$ demand).

Remarkably....

Theorem (Arrow-Debreu, 1954])

Every exchange economy has a market price equilibrium.

(They proved a much much more general theorem.)

Unfortunately, the theorem's proof is completely non-algorithmic.
Both Nash's and Arrow-Debreu's proofs crucially use fixed point theorems.

Nash's proof

Brouwer's fixed point theorem

Every continuous function $F : D \mapsto D$ from a compact convex set $D \subseteq \mathbb{R}^m$ to itself has a fixed point: $x^* \in D$, such that $F(x^*) = x^*$.

- The NEs of a finite game, Γ, are precisely the fixed points of the following Brouwer function $F_\Gamma : X \mapsto X$:

$$F_\Gamma(x)_{(i,j)} = \frac{x_{i,j} + \max\{0, g_{i,j}(x)\}}{1 + \sum_{k=1}^{m_i} \max\{0, g_{i,k}(x)\}}$$

where $g_{i,j}(x) \doteq U_i(x_{-i}; j) - U_i(x)$.

Note: $g_{i,j}(x)$ are polynomials in the variables in x, and they measure:

"how much better off would player i be if it switched to pure strategy j?"

A basic computational question

Question

What is the complexity of the following search problem:

> *("Strong") ϵ-approximation of a Nash Equilibrium:*
> *Given a finite (normal form) game, Γ, with 3 or more players, and given $\epsilon > 0$, compute a rational vector x' such that there is some (exact!) Nash Equilibrium x^* of Γ so that:*

$$\|x^* - x'\|_\infty < \epsilon$$

Note:

This is NOT the same thing as asking for an ϵ-Nash Equilibrium.

Weak vs. Strong approximation of Fixed Points

- 2-player finite games always have rational NEs, and there are algorithms for computing an exact rational NE in a 2-player game (Lemke-Howson'64).
- For games with ≥ 3 players, all NEs can be irrational (Nash,1951). So we can't hope to compute one "*exactly*".

Two different notions of ϵ-approximation of fixed points:

- (**Weak**) Given $F : \Delta_n \mapsto \Delta_n$, compute x' such that:

$$\|F(x') - x'\| < \epsilon$$

- (**Strong**) Given $F : \Delta_n \mapsto \Delta_n$, compute x' s.t. there exists x^* where $F(x^*) = x^*$ and:

$$\|x^* - x'\| < \epsilon$$

Scarf's classic algorithm

Scarf (1967) gave a beautiful algorithm (refined by Kuhn and others) for computing (weak!) ϵ-fixed points of a given Brouwer function $F : \Delta_n \mapsto \Delta_n$:

1. Subdivide the simplex Δ_n into "small" subsimplices of diameter $\delta > 0$ (δ depending on ϵ and on the "modulus of continuity" of F).
2. Color every *vertex*, $\mathbf{z}$, of every subsimplex with a color i such that $\mathbf{z}_i > 0$ & $F(\mathbf{z})_i \leq \mathbf{z}_i$.
3. By **Sperner's Lemma** there must exist a panchromatic subsimplex. (And the proof provides a way to "navigate" toward such a simplex.)
4. **Fact:** If $\delta > 0$ is chosen such that $\delta \leq \epsilon/2n$ and $\forall x, y \in \Delta_n$, $||x - y||_\infty < \delta \ \Rightarrow \ ||F(x) - F(y)||_\infty < \epsilon/2n$, then all points in a panchromatic subsimplex are weak ϵ-fixed point.
5. They need NOT in general be anywhere near an actual fixed point.

Sperner's Lemma

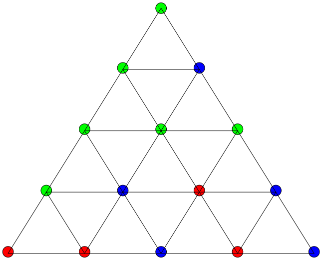

“Proof” of Sperner’s lemma

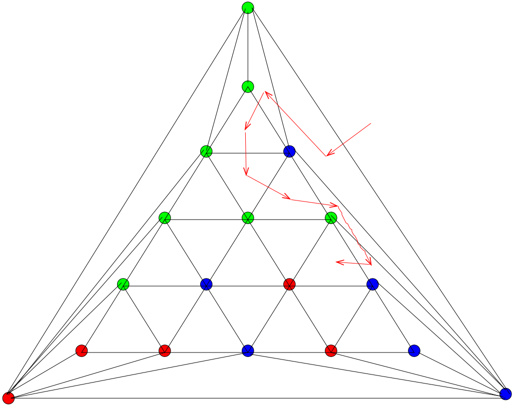

(Things are more involved in higher dimensions.)

(The same combinatorial argument was also used by (Lemke-Howson'64) for an algorithm for computing a 2-player Nash Equilibrium.)

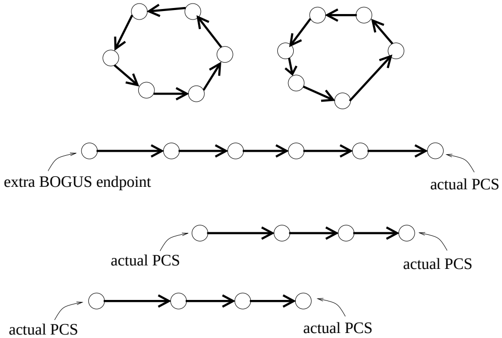

ϵ-NEs are weak ϵ-fixed points

Proposition

For finite games, Γ, computing an ϵ-NE is P-time equivalent to computing a weak ϵ-fixed point of Nash's function F_Γ.

ϵ-NEs are weak ϵ-fixed points

Proposition

For finite games, Γ, computing an ϵ-NE is P-time equivalent to computing a weak ϵ-fixed point of Nash's function F_Γ.

Thus, to compute an ϵ-NE, we can simply apply Scarf's algorithm to F_Γ.

ϵ-NEs are weak ϵ-fixed points

Proposition

For finite games, Γ, computing an ϵ-NE is P-time equivalent to computing a weak ϵ-fixed point of Nash's function F_Γ.

Thus, to compute an ϵ-NE, we can simply apply Scarf's algorithm to F_Γ.

Question

What does this tell us about the complexity of computing an ϵ-NE?

The complexity class PPAD

Papadimitriou (1992) defined **PPAD**, based on the "directed line" parity argument, to capture (approximate) Nash and Brouwer, etc...

Definition

PPAD is the class of search problems polynomial-time reducible to:
Directed line endpoint problem: Given two boolean circuits, S ("Successor") and P ("Predecessor"), each with n input bits and n output bits, such that $P(0^n) = 0^n$, and $S(0^n) \neq 0^n$, find a n-bit vector, $\mathbf{z}$, such that either: $P(S(\mathbf{z})) \neq \mathbf{z}$ or $S(P(\mathbf{z})) \neq \mathbf{z} \neq 0^n$.
(By the directed line parity argument such a $\mathbf{z}$ exists.)

PPAD lies somewhere between (the search versions of) **P** and **NP**.

By Scarf's algorithm, computing a ϵ-NE is in PPAD.
Can we do better??

By Scarf's algorithm, computing a ϵ-NE is in PPAD.
Can we do better?? No.

By Scarf's algorithm, computing a ϵ-NE is in PPAD.
Can we do better?? No.

Theorem

1. *[Daskalakis-Goldberg-Papadimitriou'06][Chen-Deng'06]:*
   Computing a ϵ-NE for a 3 player game is PPAD-complete.
2. *[Chen-Deng'06]:*
   Computing an exact (rational) NE for a 2 player game is PPAD-complete.

By Scarf's algorithm, computing a ϵ-NE is in PPAD.
Can we do better?? No.

Theorem

1. *[Daskalakis-Goldberg-Papadimitriou'06][Chen-Deng'06]:*
   Computing a ϵ-NE for a 3 player game is PPAD-complete.
2. *[Chen-Deng'06]:*
   Computing an exact (rational) NE for a 2 player game is PPAD-complete.

But what if we want to approximate actual NEs for games with ≥ 3 players and to approximate actual fixed points?
I.e., what if we want to do strong approximation of fixed points?
Warning: Scarf's algorithm does not in general yield strong ϵ-fixed points.

Why care about strong approximation of fixed points?

- It can be argued (Scarf (1973) implicitly did) that for many applications in economics weak ϵ-fixed points of Brouwer functions are sufficient.
- However, many important problems boil down to a fixed point computation for which weak ϵ-FPs are useless, unless they also happen to be strong ϵ-FPs.
  Examples:
  –Stochastic Games (Shapley, 1953; Condon, 1992);
  –(multi-type) Branching Processes (Kolmogorov, 1947);
- Moreover, a highly prized property in economics is uniqueness of equilibrium (unique predicted/prescribed "outcome"). In settings where there is a unique equilibrium, we naturally want to compute that unique outcome!

Proposition

Given game Γ and $\epsilon > 0$, we can Strong ϵ-approximate a NE in **PSPACE**.

Proof.

For Nash's functions, F_Γ, the expression

$$\exists \mathbf{x}(\mathbf{x} = F_\Gamma(\mathbf{x}) \wedge \mathbf{a} \leq \mathbf{x} \leq \mathbf{b})$$

can be expressed as a formula in the Existential Theory of Reals (ETR). So we can Strong ϵ-approximate an NE, $x^* \in \Delta_n$, in **PSPACE**, using $\log(1/\epsilon)n$ queries to a PSPACE decision procedure for ETR ([Canny'89],[Renegar'92]).
(These are deep, but thusfar impractical algorithms.) □

Can we do better than **PSPACE**?

two hard problems

Sqrt-Sum

The square-root sum problem is the following decision problem:

Given $(d_1, \ldots, d_n) \in \mathbb{N}^n$ and $k \in \mathbb{N}$, decide whether $\sum_{i=1}^{n} \sqrt{d_i} \leq k$.

It is solvable in PSPACE.

Open problem ([GareyGrahamJohnson'76]) whether it is solvable even in NP (or even the polynomial time hierarchy).

PosSLP (Allender, Bürgisser,Kjeldgaard-Pedersen, Miltersen,2006)

Given an arithmetic circuit (Straight Line Program) over basis $\{+, *, -\}$ with integer inputs, decide whether the output is > 0.

[Allender et. al.'06] Gave a (Turing) reduction from Sqrt-Sum to PosSLP and showed both can be decided in the Counting Hierarchy:

$$P^{PP^{PP^{PP}}}$$

why isn't PosSLP easy??

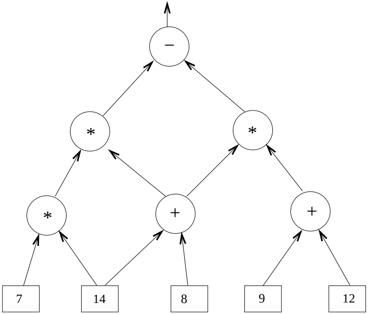

Sqrt-Sum & PosSLP $\leq_p$ approximation of actual NE

Theorem [E.-Yannakakis,2007]

Any non-trivial approximation of an actual NE is both **Sqrt-Sum**-hard and **PosSLP**-hard.

More precisely: for every $\epsilon > 0$, both **Sqrt-Sum** and **PosSLP** are P-time reducible to the following problem:

Given a 3-player (normal form) game, Γ, with the property that:

1. It has a unique NE, x^*, and
2. the probability, $x^*_{1,1}$, with which player 1 plays it first pure strategy is either:

$$(a.) \quad = 0 \quad , \qquad\qquad \text{or} \qquad\qquad (b.) \quad \geq (1-\epsilon)$$

Decide which of (a.) or (b.) is the case.

ϵ-NEs can be very far from actual NEs

Theorem [E.-Yannakakis,2007]

- For every n, there exists a 4-player game Γ_n of size $O(n)$ with an ϵ-NE, x', where $\epsilon = \frac{1}{2^{2^{\Omega(n)}}}$, and yet x' has distance 1 in l_∞ to any actual NE. (Thus worst possible distance in l_∞.)
- The same holds for 3 players, but with distance 1 replaced by distance $(1-\delta)$, for any fixed constant $\delta > 0$ (and even for $\delta = 2^{-poly(n)}$).

A new complexity class: FIXP

Consider the following class of fixed point problems:

FIXP [E.-Yannakakis,2007]

- **Input:** algebraic circuit (straight-line program) over basis $\{+, *, -, /, \max, \min\}$ with rational constants, having n input variables and n outputs, such that the circuit represents a continuous function $F : [0,1]^n \mapsto [0,1]^n$.
  (The domain can be much more general than $[0,1]^n$.)
- **Output:** Compute, or strong ϵ-approximate, a fixed point of F.

Close these problems under suitable P-time reductions.
Call the resulting class **FIXP**.

We shall see that many interesting problems besides Nash are in **FIXP**.

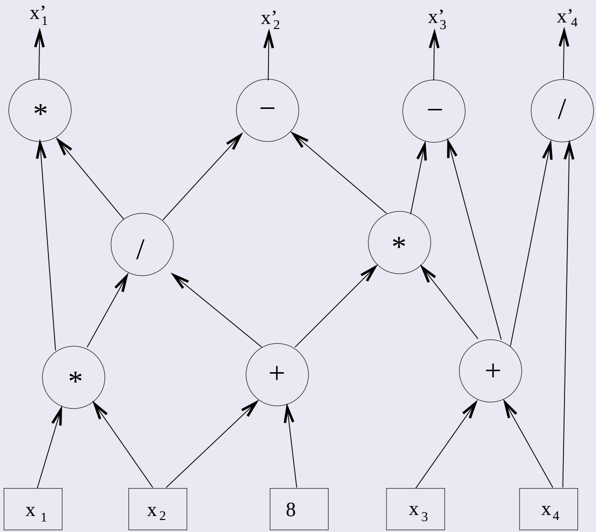
x'1
x'2
x'3
x'4
*
−
−
/
/
*
*
+
+
x 1
x 2
8
x 3
x 4

Nash is FIXP-complete

Theorem [E.-Yannakakis,2007]

Computing a 3-player Nash Equilibrium is **FIXP**-complete.

It is complete in several senses:

- In terms of "exact" (real valued) computation;
- In terms of strong ϵ-approximation,
- An appropriate "decision" version of the problem: Given a game,Γ, rational value $q \in \mathbb{Q}$, and coordinate i: if for all NEs x^*, $x_i^* \geq q$, then *"Yes"*; if for all NEs x^*, $x_i^* < q$, then *"No"*. Otherwise, any answer is fine.

A new characterization of PPAD

Let **linear-FIXP** denote the subclass of FIXP where the algebraic circuits are restricted to basis $\{+, \max\}$ and multiplication by rational constants only.

Theorem [E.-Yannakakis,2007]

The following are all equivalent:

1. PPAD
2. linear-FIXP
3. exact fixed point problems for "polynomial piecewise-linear functions"

Corollary

`Simple-Stochastic-Games` *(and* `Parity Games`*, etc.) are in PPAD.*

Market equilibria, excess demands, Walras's law, etc.

- From the demand functions we directly get excess demand functions: $g_i^j(p) = d_i^j(p) - w_{j,i}$, for agent j and commodity i.
- The total excess demand for commodity i is $g_i(p) = \sum_j g_i^j(p)$.
- Excess demands are continuous and satisfy economically justified axioms:
  - (Homogeneous of degree 0): For all $\alpha > 0$, $p \geq 0$, $g_i^j(\alpha p) = g_i^j(p)$. (So, we can w.l.o.g. consider only "normalized" price vectors in Δ_m.)
  - (Walras's law): $\sum_i p_i g_i(p) = 0$.

Excess demand functions can be essentially arbitrary continuous functions (Sonnenschein-Mantel-Debreu,1973-74).

Price Equilibrium

A vector of prices $p^* \geq 0$ such that $g_i(p^*) \leq 0$ for all i ($= 0$ if $p_i^* > 0$).

Theorem ((Arrow-Debreu'54) (proved a much more general fact)

Every exchange economy has a price equilibrium.

The proof is via Brouwer's fixed point theorem, and for more general market equilibrium results, via the closely related Kakutani fixed point theorem.

We can use Scarf's algorithm to compute a so called ϵ-Price equilibrium (where the excess demand for all commodities is $\geq \epsilon$).

Theorem [E.-Yannakakis,2007]

Any non-trivial approximation of an actual Price Equilibrium in an arbitrary exchange exchange economy (even one with a unique PE) whose excess demands are given by algebraic circuits over $\{+, *, -, /, \max, \min\}$ is **SQRT-SUM**-hard and **PosSLP**-hard.

Proposition [E.-Yannakakis,2007]

Computing price equilibria in exchange economies with excess demands given by algebraic circuits over $\{+, *, -, /, \max, \min\}$ is FIXP-complete.

Proof.

One direction of proof is via the following variant of Nash's function:

$$H(p)_i = \frac{p_i + \max\{0, g_i(p)\}}{1 + \sum_{j=1}^{m} \max\{0, g_j(p)\}}$$

where $g_i(x)$ is the total excess demand for commodity i.
The (Brouwer) fixed points of $H(p)$ are the price equilibria of the economy.
The other direction (Uzawa (1962)): given Brouwer function $F : \Delta_n \mapsto \Delta_n$, define total excess demand function $g : \Delta_n \mapsto \mathbb{R}^n$ by

$$g(p) = F(p) - \left(\frac{\langle p, F(p)\rangle}{\langle p, p\rangle}\right)p$$

$g(p)$ satisfies excess demand axioms. The price equilibria of $g(p)$ are the fixed points of $F(p)$. $\square$

Special restricted classes of Markets

Theorem

- [Eisenberg-Game'59],[Devanur,et.al.,2002][Jain,2008] For markets with linear, separable, utility functions, we can compute a Arrow-Debreu (or Fischer) price equilbrium in polynomial time.

- [Chen et. al.,2009],[Vazirani-Yannakakis,2010] Unfortunately, already for markets with piecewise-linear, separable (concave) utility functions, computing an equilibrium is PPAD-hard (and in PPAD).

Application: price equilibrium in Google's TV ad auction mechanism

- Currently, Google runs a TV advertisement slot combinatorial auction.
- The algorithm it uses to determines prices, and allocates TV ad slots, to the various bidders (devised by N. Nisan, a computer scientist), is via a Walrasian market price equilibrium computation for a very special class of markets where the equilibrium is polynomial time computable. [Demange-Gale-Sotomayor,1986], [Nisan et. al., 2009].

Conclusions

- The convergence of algorithmics and economics is likely to continue apace in the coming decade, because a lot of Economic theory (e.g., auction theory) is now becoming economic practice on the internet.
- At the same time, ecomomics and game theory have generated deep and difficult algorithmic questions (fixed points, etc.) that will continue to challenge researchers in algorithms and complexity for many years.
- Can Nash/market equilibria be computed in polynomial time?
- More modestly, can strong approximation of Nash/market equilibria be done in anything better than PSPACE?

K. Etessami and M. Yannakakis, "On the complexity of Nash Equilibria and Other Fixed Points", FOCS'07.
(Journal version in *SIAM Journal on Computing*, 2010 (66pp), see:
`http://homepages.inf.ed.ac.uk/kousha`)