COMPSCI 223: Computational Microeconomics

Instructor: Vincent Conitzer
(Kimberly J. Jenkins University Professor of New Technologies
Professor of Computer Science, Professor of Economics, and Professor of Philosophy)
conitzer@cs.duke.edu
https://www2.cs.duke.edu/courses/spring18/compsci223/

TAs: Harsh Parikh and Hanrui Zhang

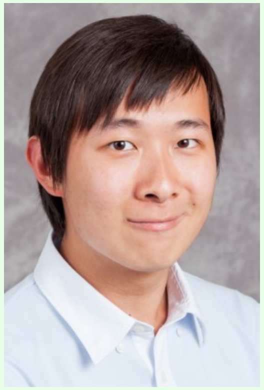

CS-ECON

http://econ.cs.duke.edu

CS-ECON@DUKE
Exploring the Intersection of Computer Science and Economics

Home | Schedule | Past Talks | People | Mini-retreat | Reading Group

Who Are We?

We are a group of Duke University faculty, postdocs, and students interested in the intersection of computer science and economics (and the social sciences more broadly) and the impact of this interplay on decisions in information technology and digital business. This includes applying techniques from computer science and optimization to economics -- for example, using computation to design market clearing mechanisms and to implement efficient allocation and pricing in them -- as well as applying techniques from economics to computer science -- for example, designing incentives for users of networked computer systems and social networks.

Contacts

For organizational questions about the seminar series:

- Yuan Deng
- Catherine Moon

For other matters, contact the relevant faculty member(s):

- Atila Abdulkadiroglu (Econ)
- Vincent Conitzer (CS)

CS-Econ Talks

- Upcoming Talks
- Past Talks

Related Seminars

- AI Group (CS)
- Algorithms Seminar (CS)
- Decision Sciences Seminar (Fuqua)
- Duke Robotics, Intelligence, and Vision (DRIV) Seminar (CS)

- For Prospective Students
- Degree Programs
  - M.A. Economics
  - M.A. Analytical Political Economy
  - M.S. Economics & Computation
  - M.S. Quantitative Financial Economics
- For Current Students

EcoTeach: Student Services Center

M.S. Economics & Computation

The joint field of economics and computer science has emerged from two converging intellectual needs: Computer science has become increasingly important for economists working with big data to address complex questions. Students interested in learning about computational mechanism design with applications to economics are ideal candidates for this program. Students whose interest is more generally focused on data analytics across a broad range of fields may also be interested in Duke's Master of Quantitative Management (MQM) program, offered at the Fuqua School of Business, and/or Duke's new Master in Interdisciplinary Data Science (MIDS) program, which is accepting its first class in Fall 2018.

The MSEC program combines the strengths of the Departments of Economics and Computer Science to educate students in these important computational skills linked to economics, and to prepare them for Ph.D. studies or careers in economics, finance, government, and business. Reflecting this strong interdisciplinary relationship, Duke University ranks No. 5 for research in economics and computation, according to CSRankings.org.

This program is designed to meet the needs of students with varied levels of exposure to either field, but a strong quantitative background is recommended.

History

computer architecture
(von Neumann
architecture)

***Computer Science
& Engineering***

game theory
(minimax theorem)

Economic Theory

?

linear programming
(duality)

***Mathematical
Optimization &
Operations
Research***

*John von
Neumann*

1900

1950

2000

What is Economics?

- "the social science that studies the production, distribution, and consumption of goods and services." [Wikipedia, Jan. 2018]
- Some key concepts:
  - Economic agents or players (individuals, households, firms, bots, …)
  - Agents' current endowments of goods, money, skills, …
  - Possible outcomes ((re)allocations of resources, tasks, …)
  - Agents' preferences or utility functions over outcomes
  - Agents' beliefs (over other agents' utility functions, endowments, production possibilities, …)
  - Agents' possible decisions/actions
  - Mechanism that maps decisions/actions to outcomes

An economic picture

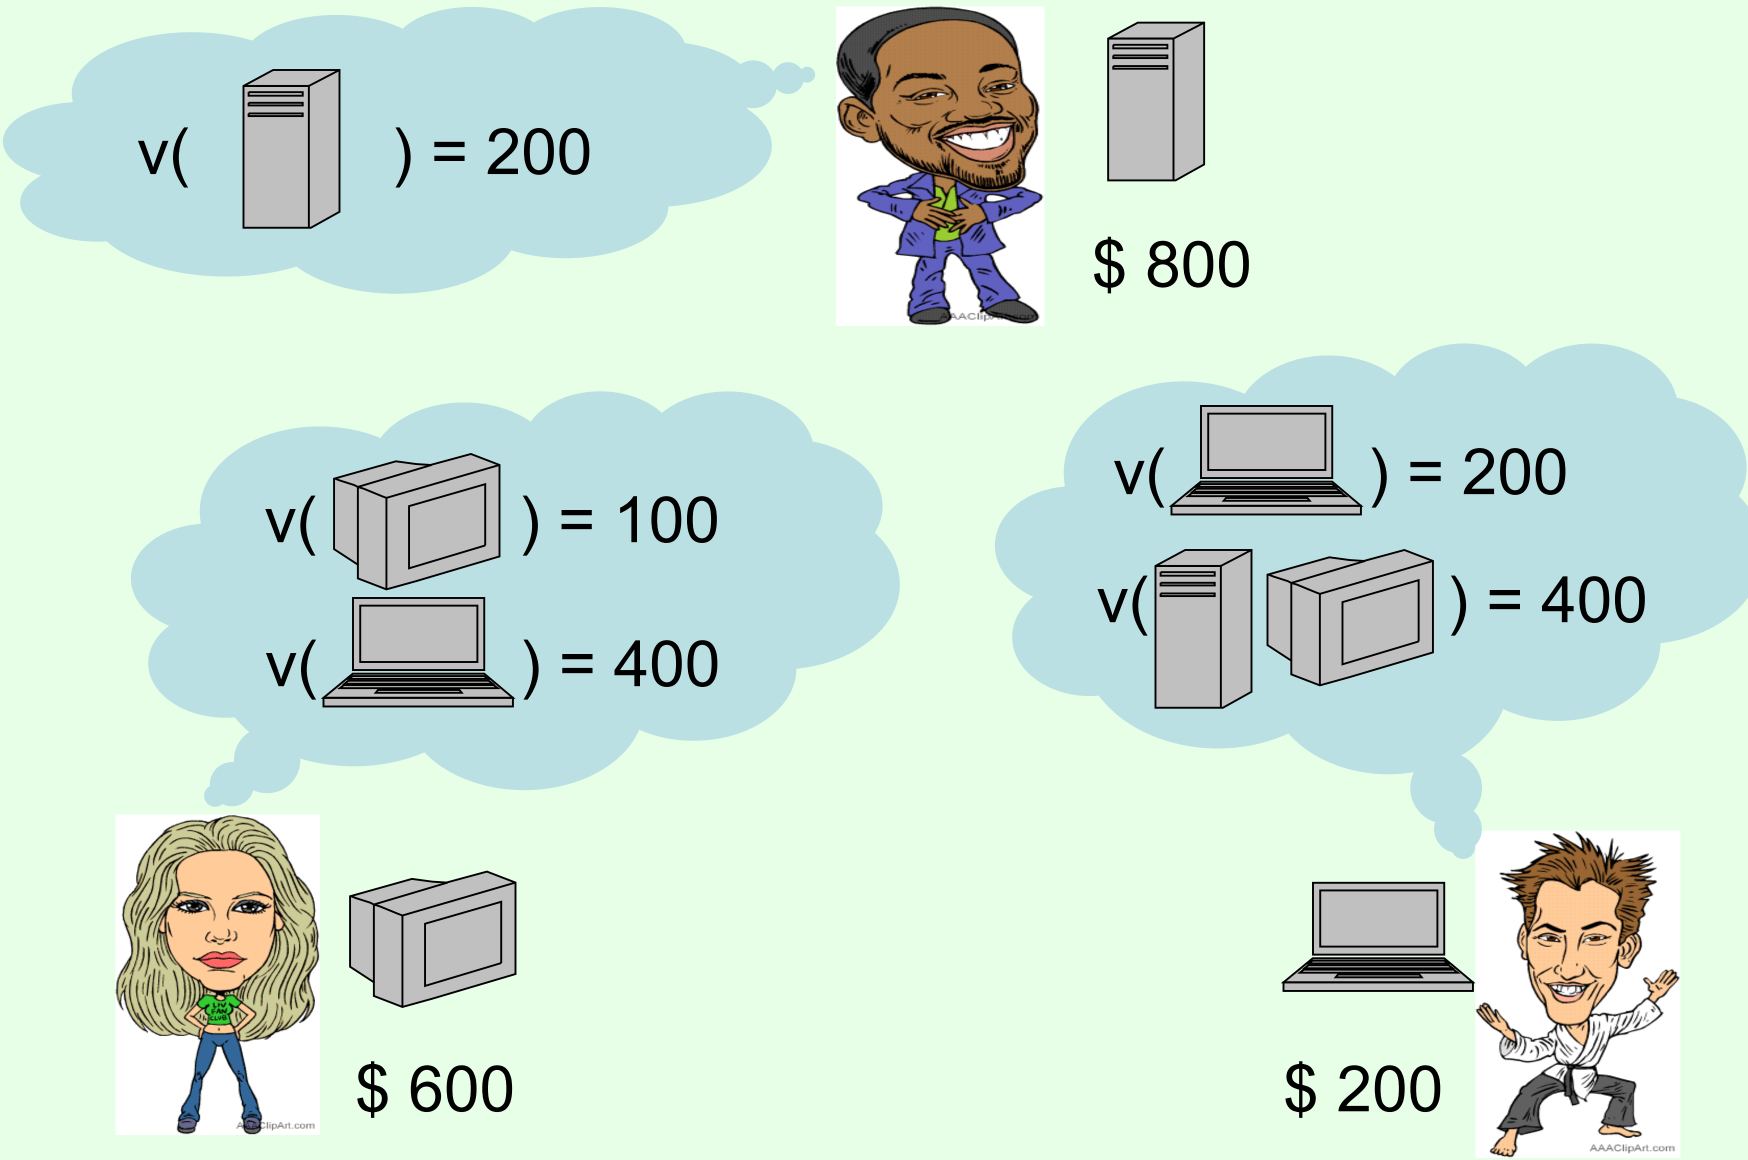

After trade (a more efficient outcome)

… but how do we get here?
Unstructured trade?
Auctions?
Exchanges?

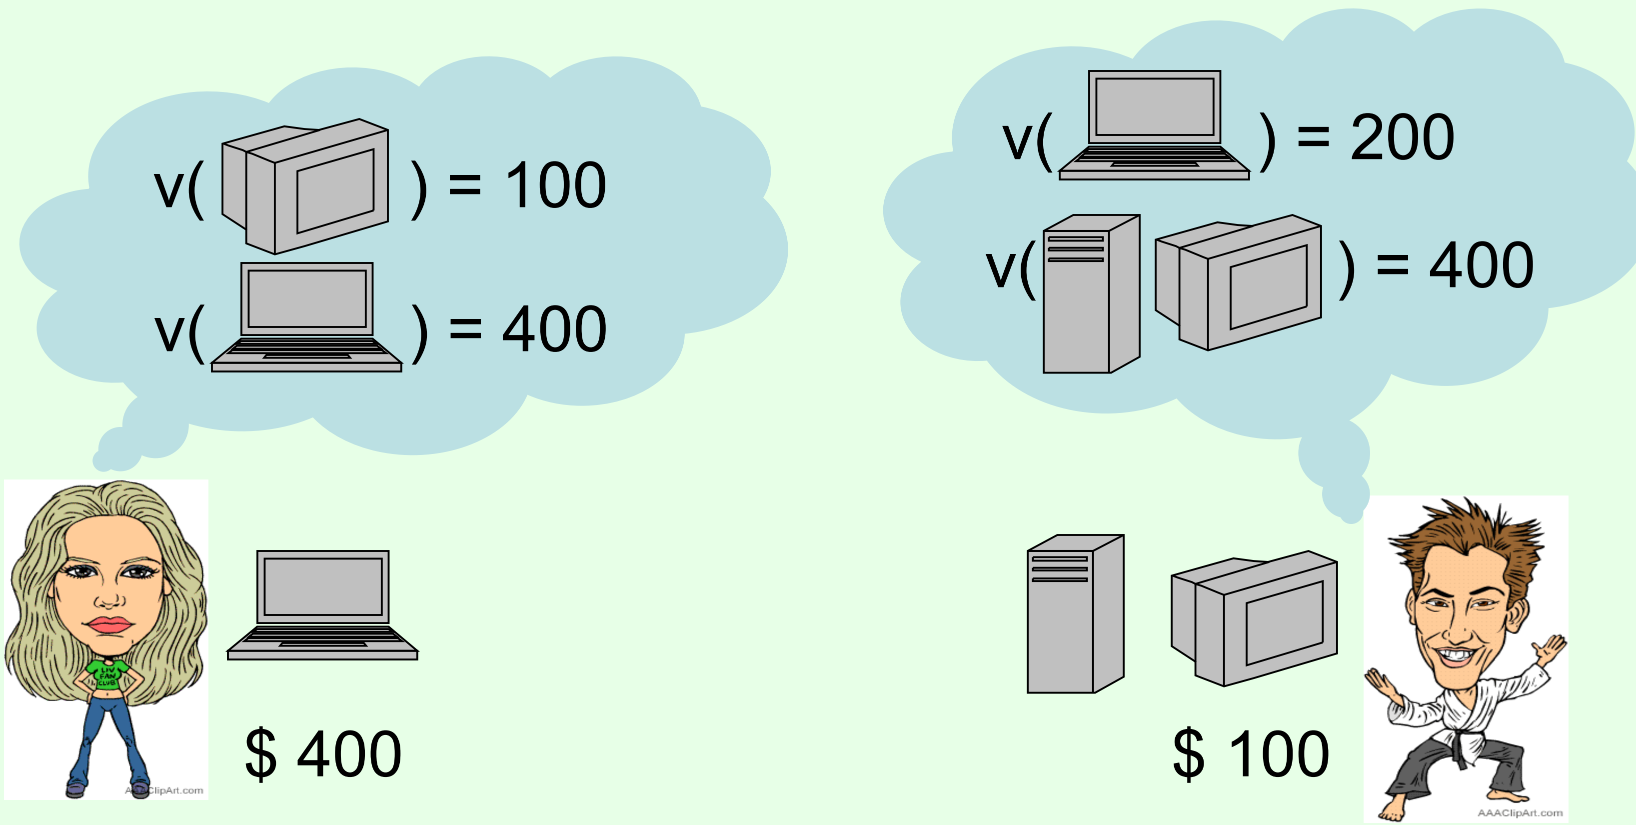

Some distinctions in economics

- Descriptive vs. normative economics
  - Descriptive:
    - seeks only to describe real-world economic phenomena
    - does not care if this is in any sense the “right” outcome
  - Normative:
    - studies how people “should” behave, what the “right” or “best” outcome is
- Microeconomics vs. macroeconomics
  - Microeconomics: analyzes decisions at the level of individual agents
    - deciding which goods to produce/consume, setting prices, …
    - “bottom-up” approach
  - Macroeconomics: analyzes “the sum” of economic activity
    - interest rates, inflation, growth, unemployment, government spending, taxation, …
    - “big picture”

What is Computer Science?

- "the study of automating algorithmic processes that scale. A computer scientist specializes in the theory of computation and the design of computational systems." [Wikipedia, Jan. 2018]
- A computational problem is given by a function f mapping inputs to outputs
  - For integer x, let f(x) = 0 if x is prime, 1 otherwise
  - For initial allocation of resources + agent utilities x, let f(x) be the (re)allocation that maximizes the sum of utilities
- An algorithm is a fully specified procedure for computing f
  - E.g., sieve of Eratosthenes
  - A correct algorithm always returns the right answer
  - An efficient algorithm returns the answer fast
- Computer science is also concerned with building larger artifacts out of these building blocks (e.g., personal computers, spreadsheets, the Internet, the Web, search engines, artificial intelligence, …)

Resource allocation as a computational problem *(Part 1 of the course)*

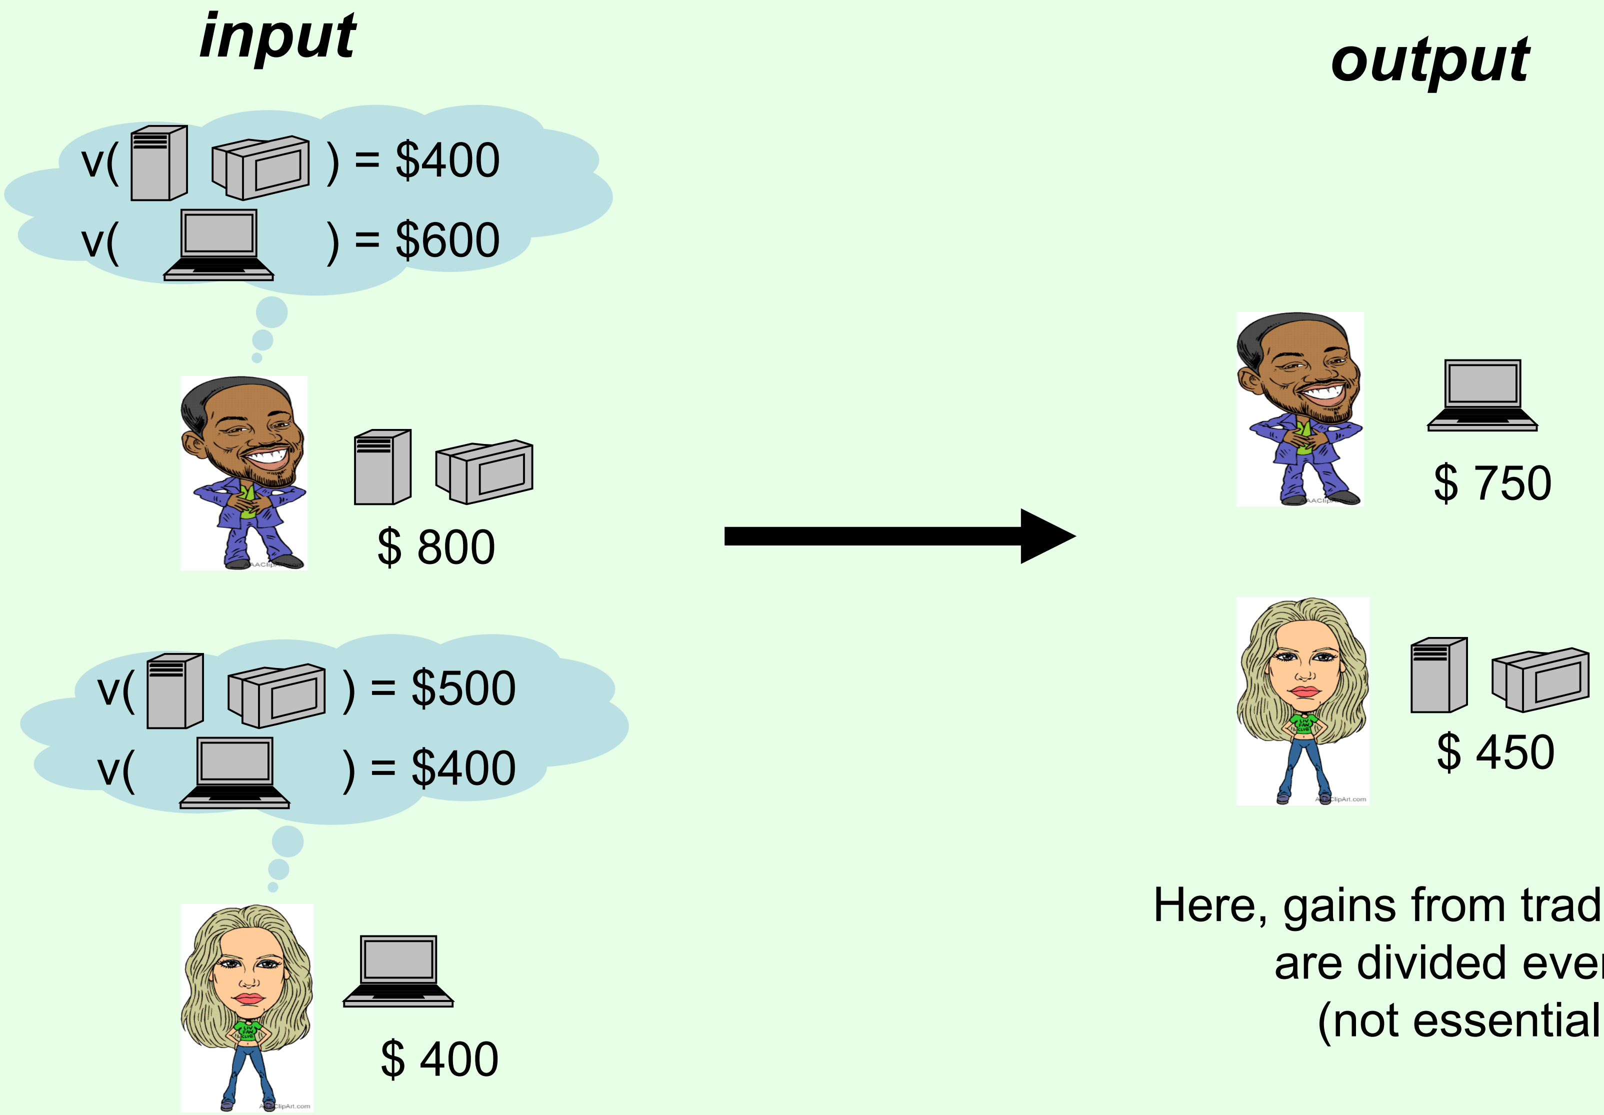

Here, gains from trade ($300)
are divided evenly
(not essential)

Economic mechanisms

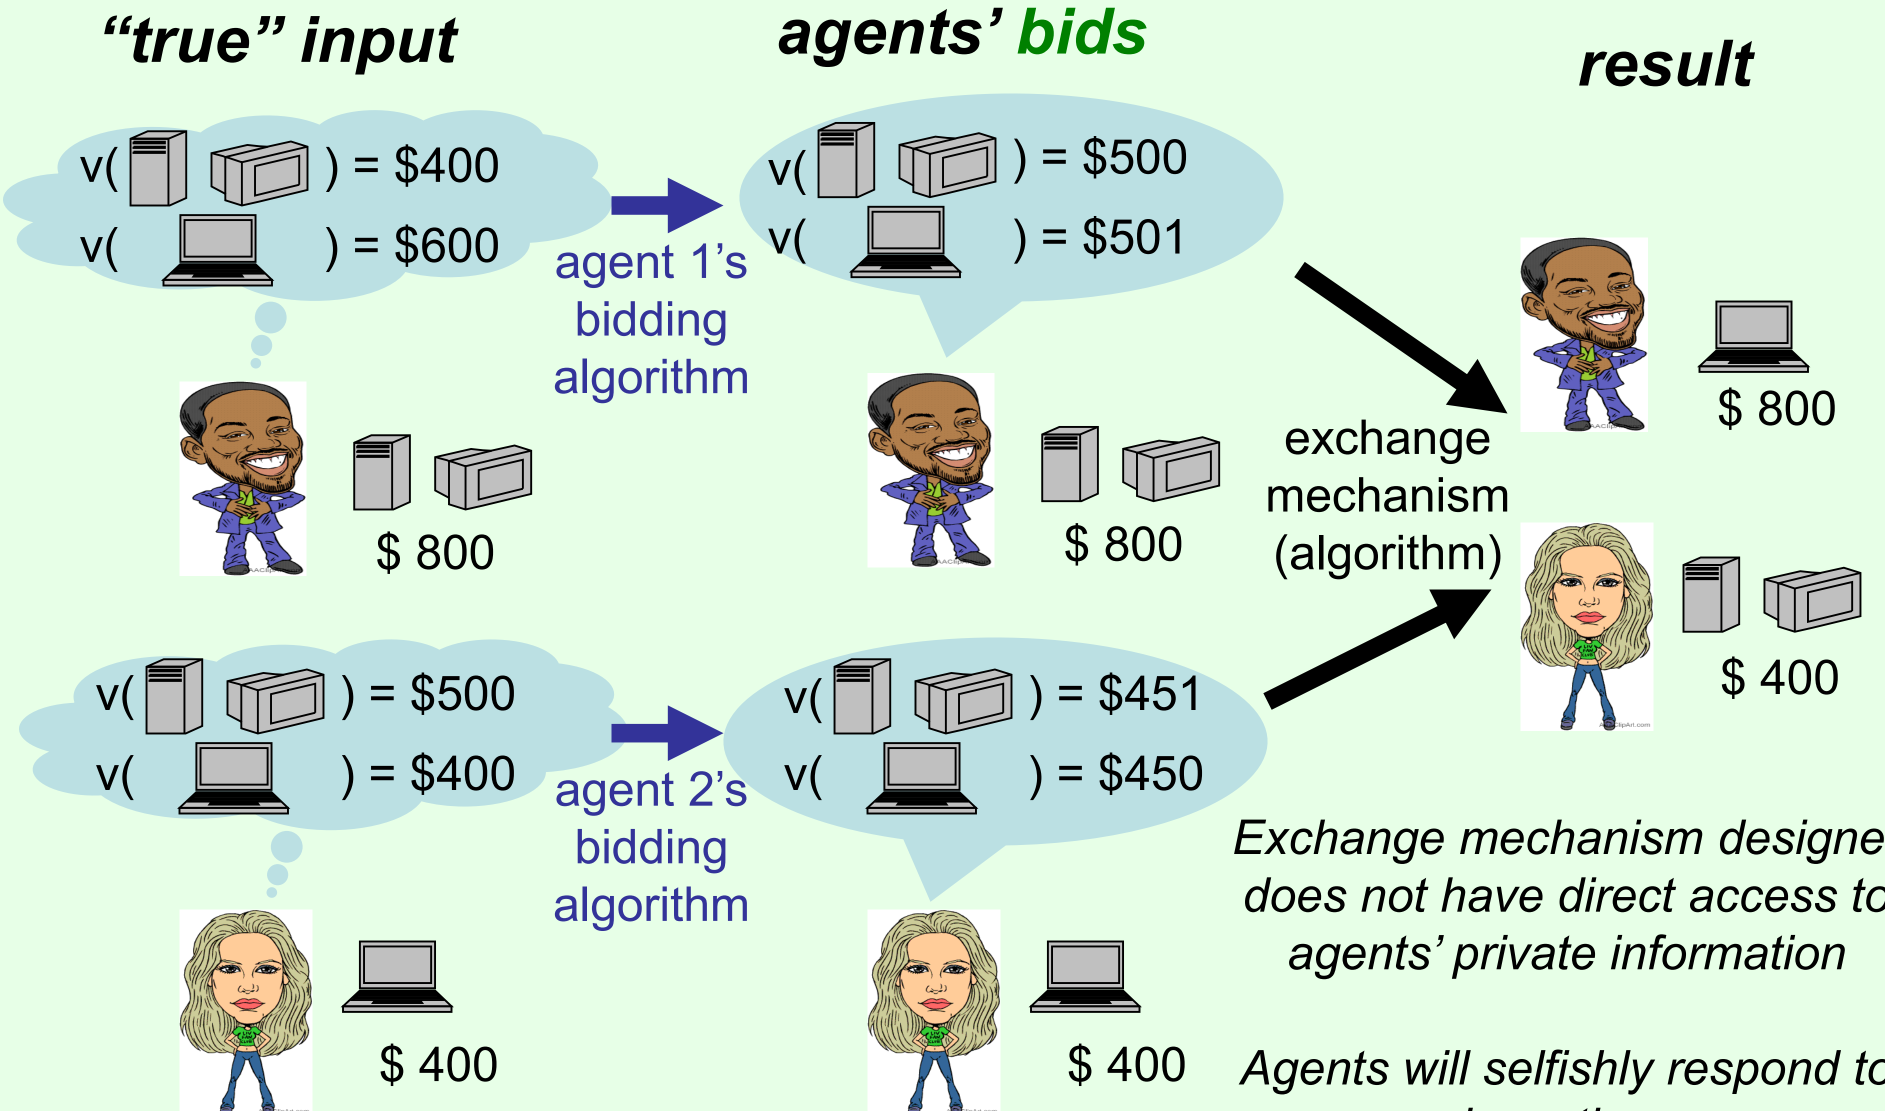

Exchange mechanism designer does not have direct access to agents' private information

Agents will selfishly respond to incentives

Game theory
(Part 2 of the course)

- Game theory studies settings where agents each have
  - different preferences (utility functions),
  - different actions that they can take
- Each agent's utility (potentially) depends on all agents' actions
  - What is optimal for one agent depends on what other agents do
    - Very circular!
- Game theory studies how agents can rationally form beliefs over what other agents will do, and (hence) how agents should act
  - Useful for acting as well as predicting behavior of others

Penalty kick example

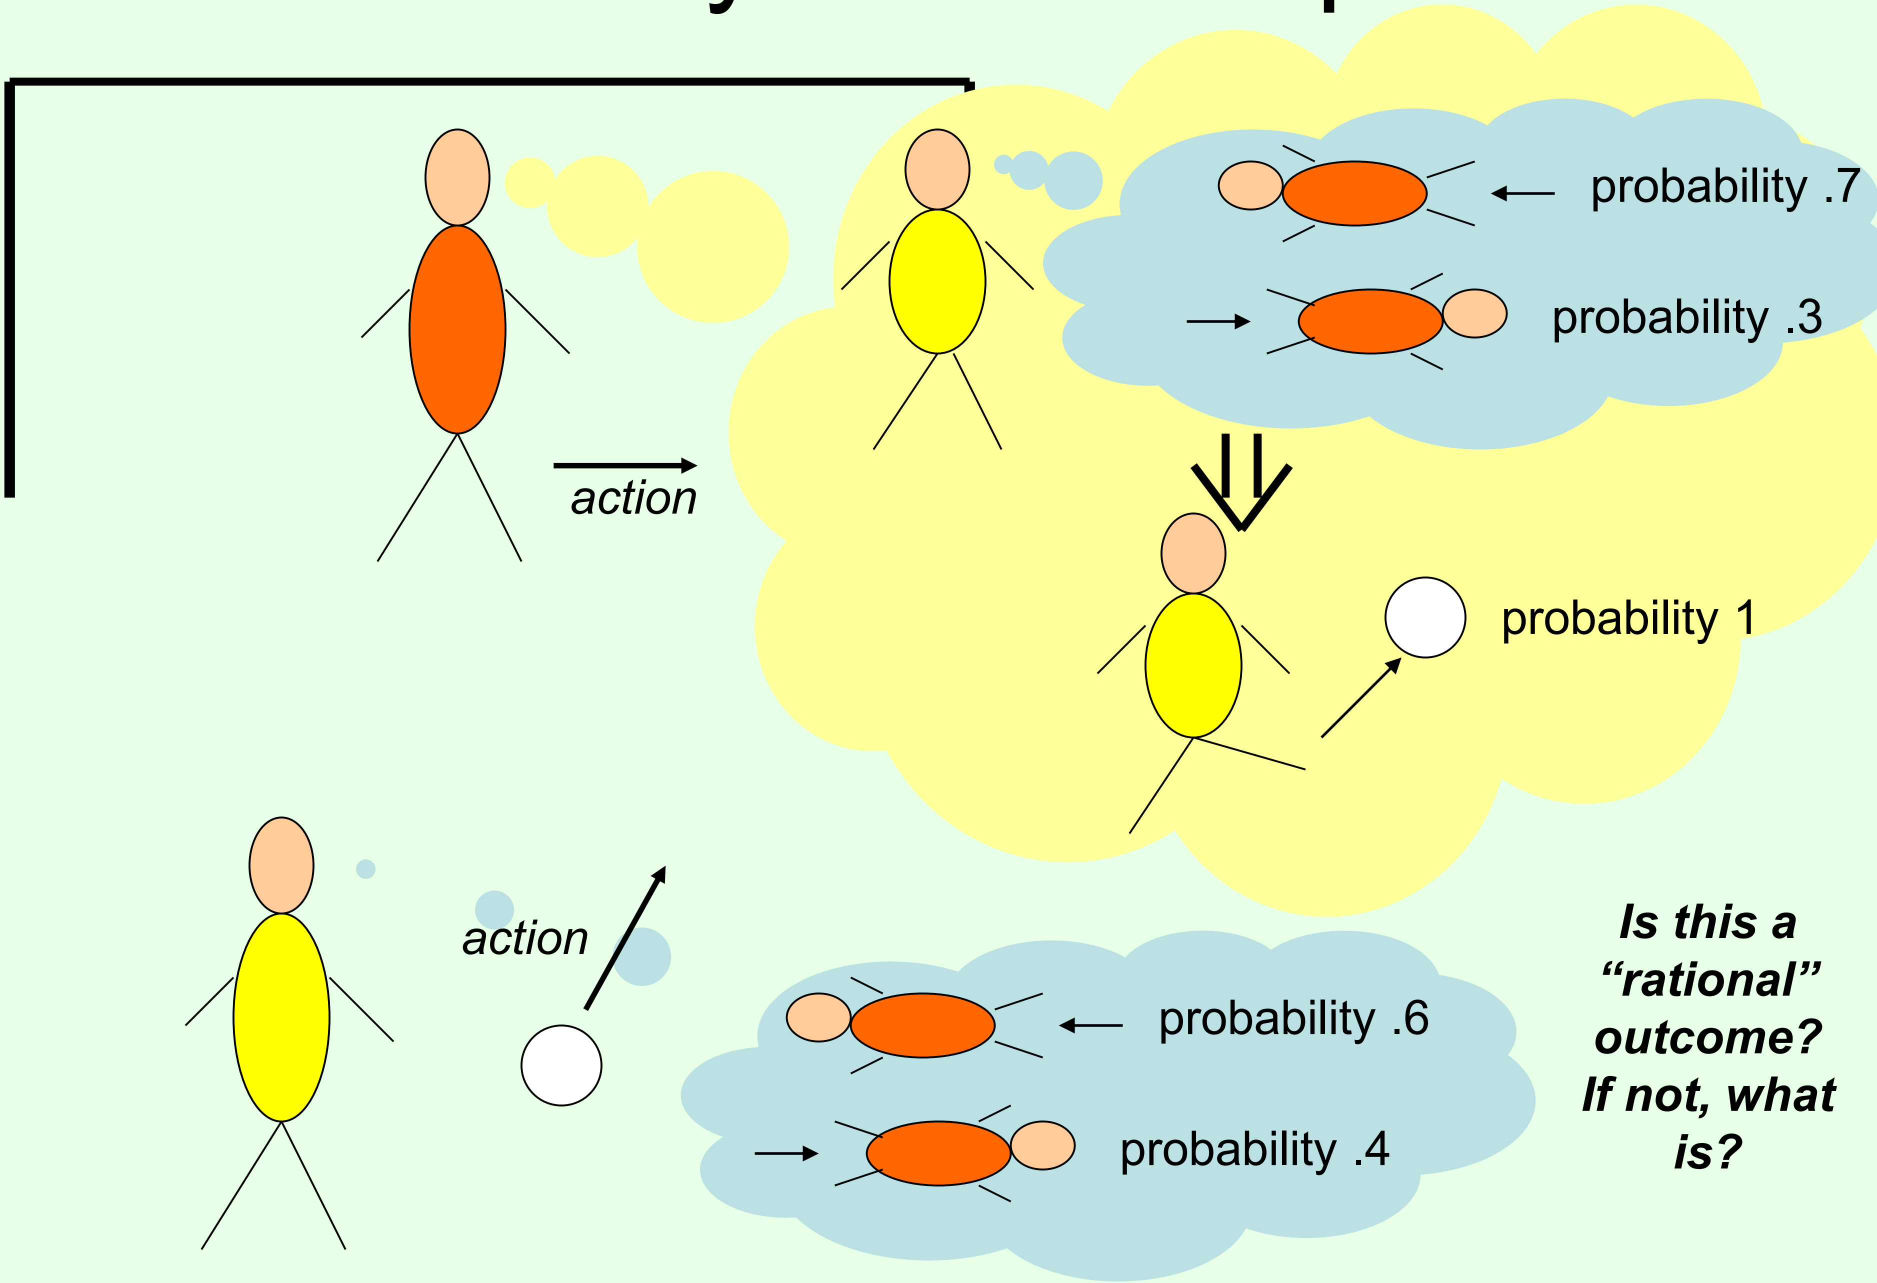

Is this a "rational" outcome? If not, what is?

Mechanism design
(Part 3 of the course)

- Mechanism = rules of auction, exchange, …
- A function that takes reported preferences (bids) as input, and produces outcome (allocation, payments to be made) as output

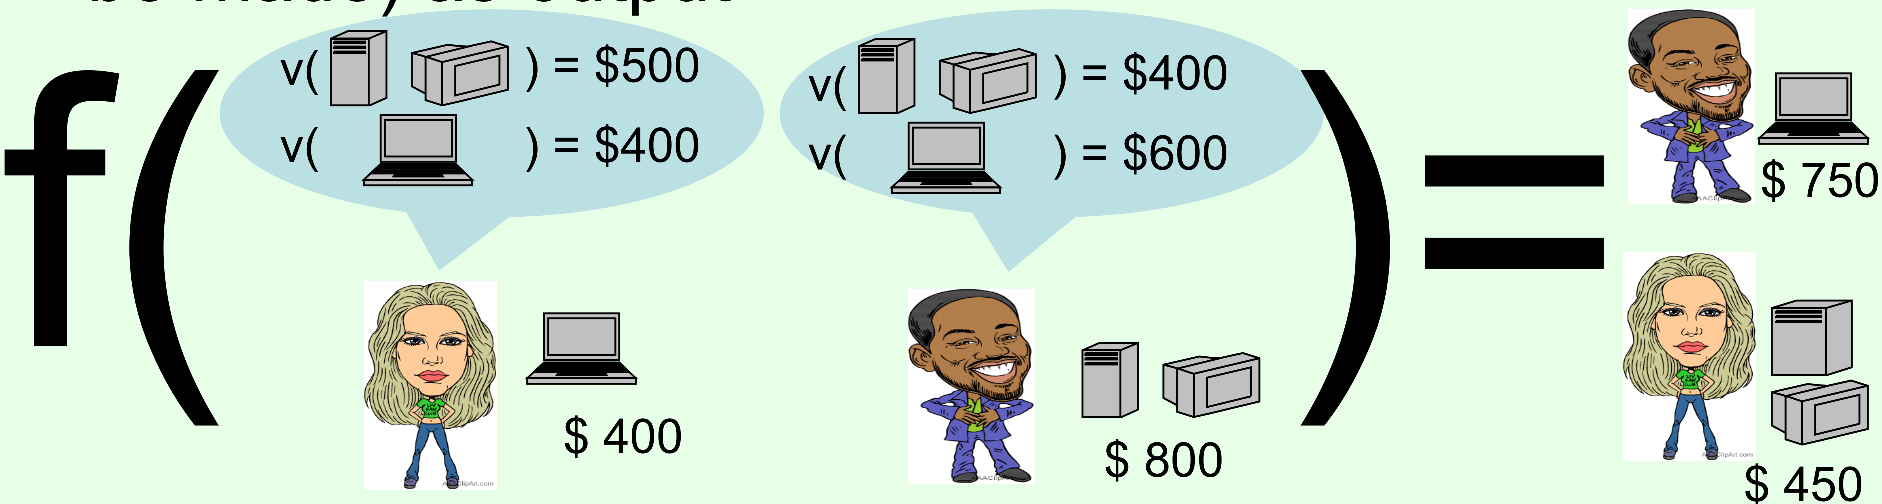

- The entire function f is one mechanism
- E.g., the mechanism from part 1: find allocation that maximizes (reported) utilities, distribute (reported) gains evenly
- Other mechanisms choose different allocations, payments

Example: (single-item) auctions

- Sealed-bid auction: every bidder submits bid in a sealed envelope
- First-price sealed-bid auction: highest bid wins, pays amount of own bid
- Second-price sealed-bid auction: highest bid wins, pays amount of second-highest bid

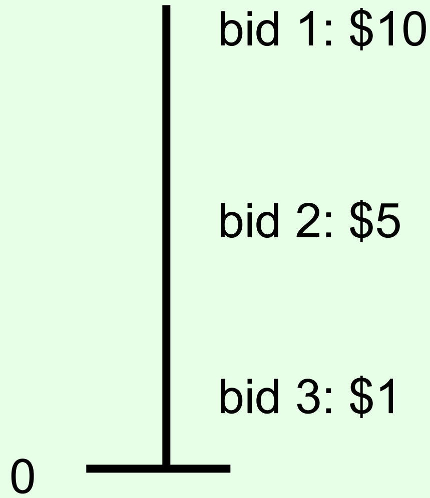

first-price: bid 1 wins, pays $10
second-price: bid 1 wins, pays $5

Which auction generates more revenue?

- Each bid depends on
  - bidder's true valuation for the item (utility = valuation - payment),
  - bidder's beliefs over what others will bid (→ game theory),
  - and... the auction mechanism used
- In a first-price auction, it does not make sense to bid your true valuation
  - Even if you win, your utility will be 0…
- In a second-price auction, (we will see later that) it always makes sense to bid your true valuation

a likely outcome for the first-price mechanism

bid 1: $5

bid 2: $4

bid 3: $1

0

a likely outcome for the second-price mechanism

bid 1: $10

bid 2: $5

bid 3: $1

0

Are there other auctions that perform better? How do we know when we have found the best one?

Mechanism design…

- Mechanism = game
- → we can use game theory to predict what will happen under a mechanism
  - if agents act strategically
- When is a mechanism "good"?
  - Should it result in outcomes that are good for the reported preferences, or for the true preferences?
  - Should agents ever end up lying about their preferences (in the game-theoretic solution)?
  - Should it always generate the best allocation?
  - Should agents ever burn money?(!?)
- Can we solve for the optimal mechanism?

How are we going to solve these problems? *(Part 0)*

- This is not a programming course
- Will use optimization software
  - GNU Linear Programming Kit (GLPK)
  - Linear programming, mixed integer linear programming

Uses of LP, MIP in this course

	Linear programming	**Mixed integer linear programming**
Part 1 (expressive marketplaces)	Winner determination in auctions, exchanges, … with partially acceptable bids	Winner determination in auctions, exchanges, … without partially acceptable bids
Part 2 (game theory)	Dominated strategies Minimax strategies Correlated equilibrium Optimal mixed strategies to commit to	Nash equilibrium
Part 3 (mechanism design)	Automatically designing optimal mechanisms that use randomization	Automatically designing optimal mechanisms that do not use randomization

Other settings/applications

Combinatorial auctions (in Part 1)

Simultaneously for sale:

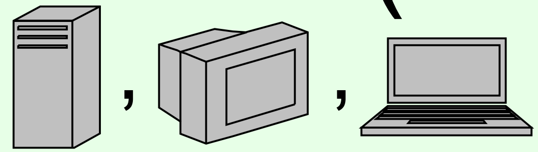

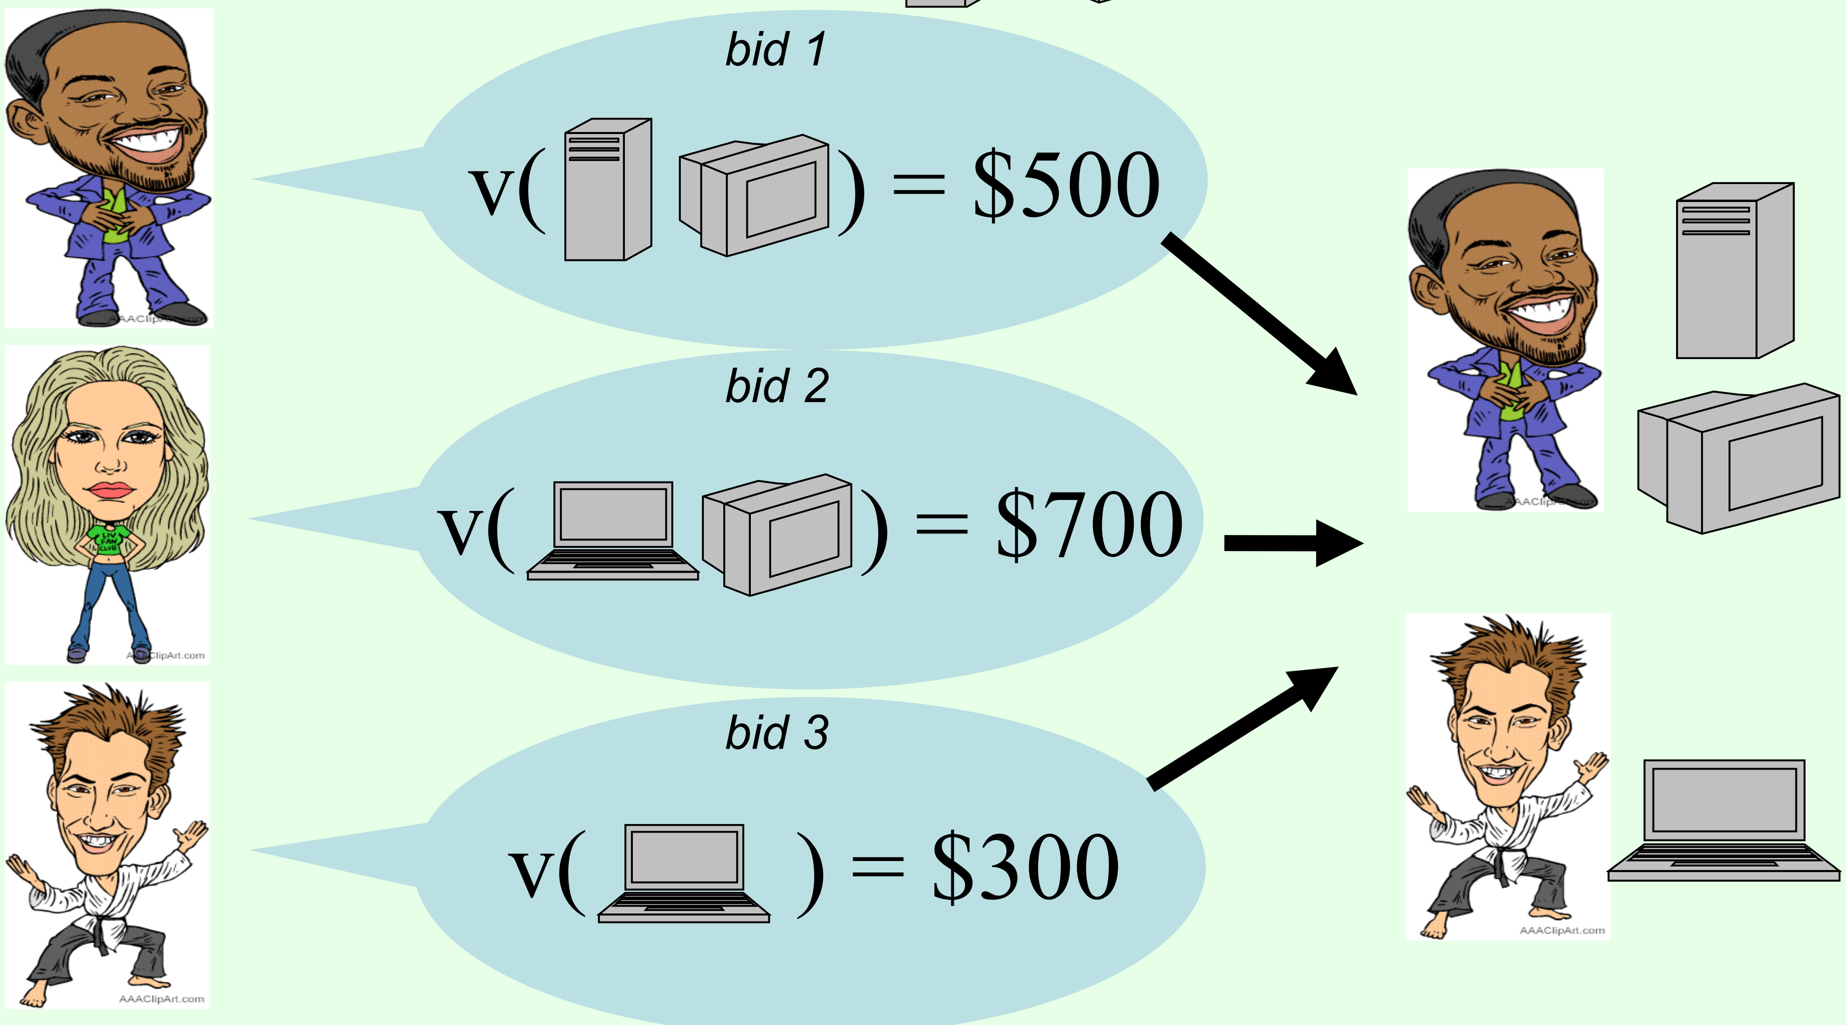

used in truckload transportation, industrial procurement, radio spectrum allocation, …

Voting (in Part 1)

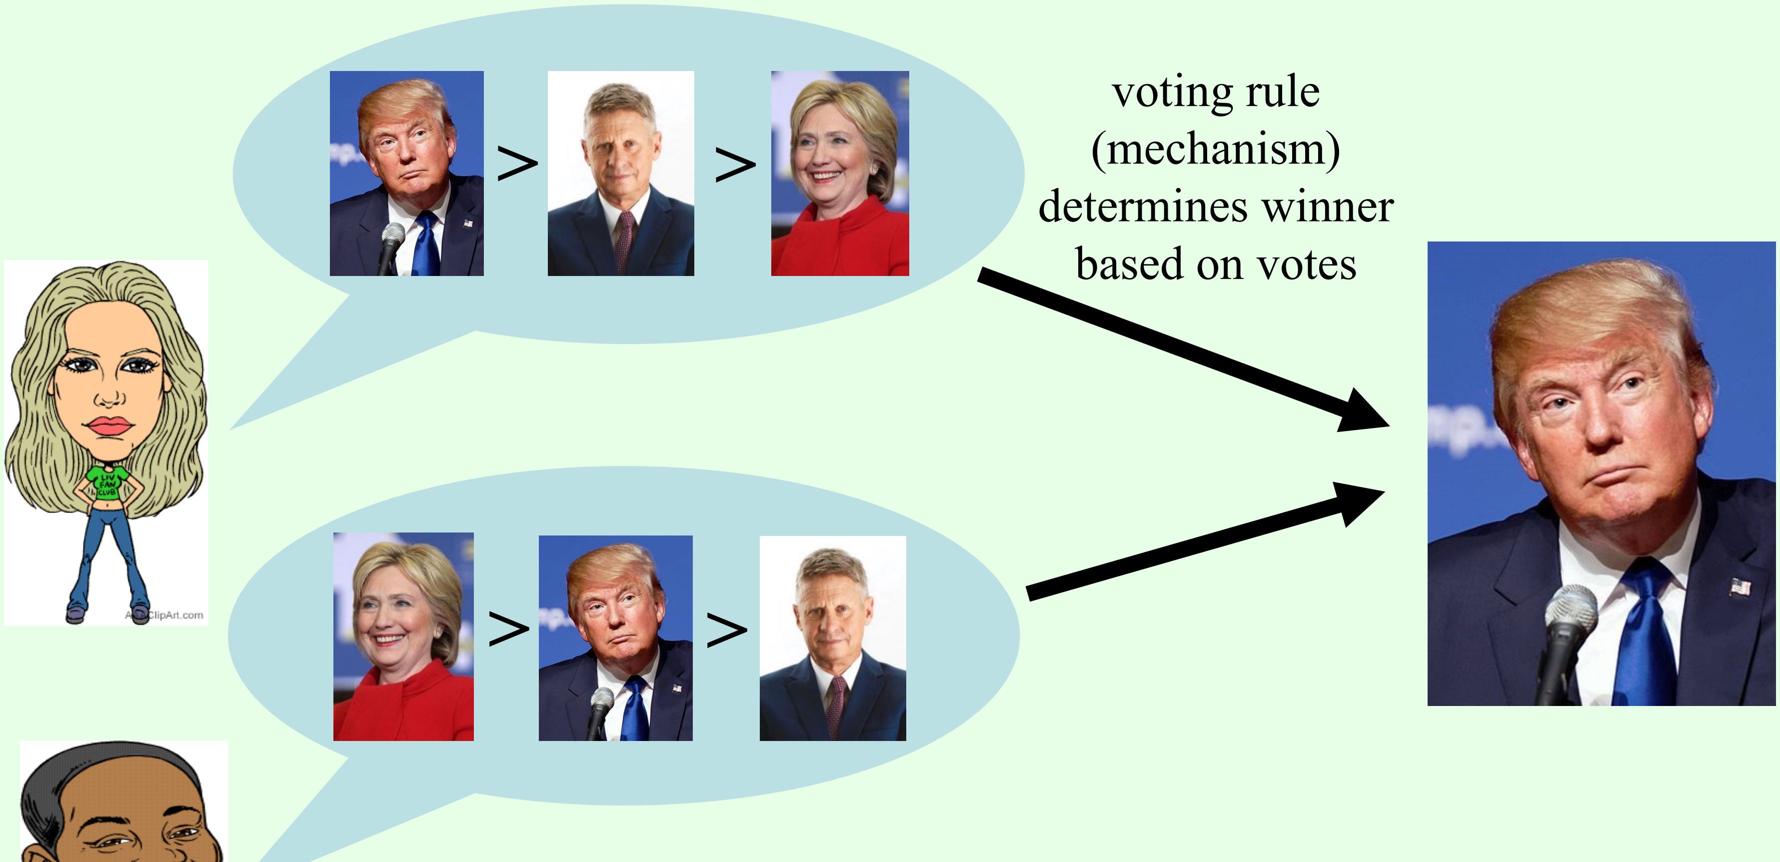

- Can vote over other things too
  - Where to go for dinner tonight, other joint plans, …
- Many different rules exist for selecting the winner

Kidney exchange (in Part 1)

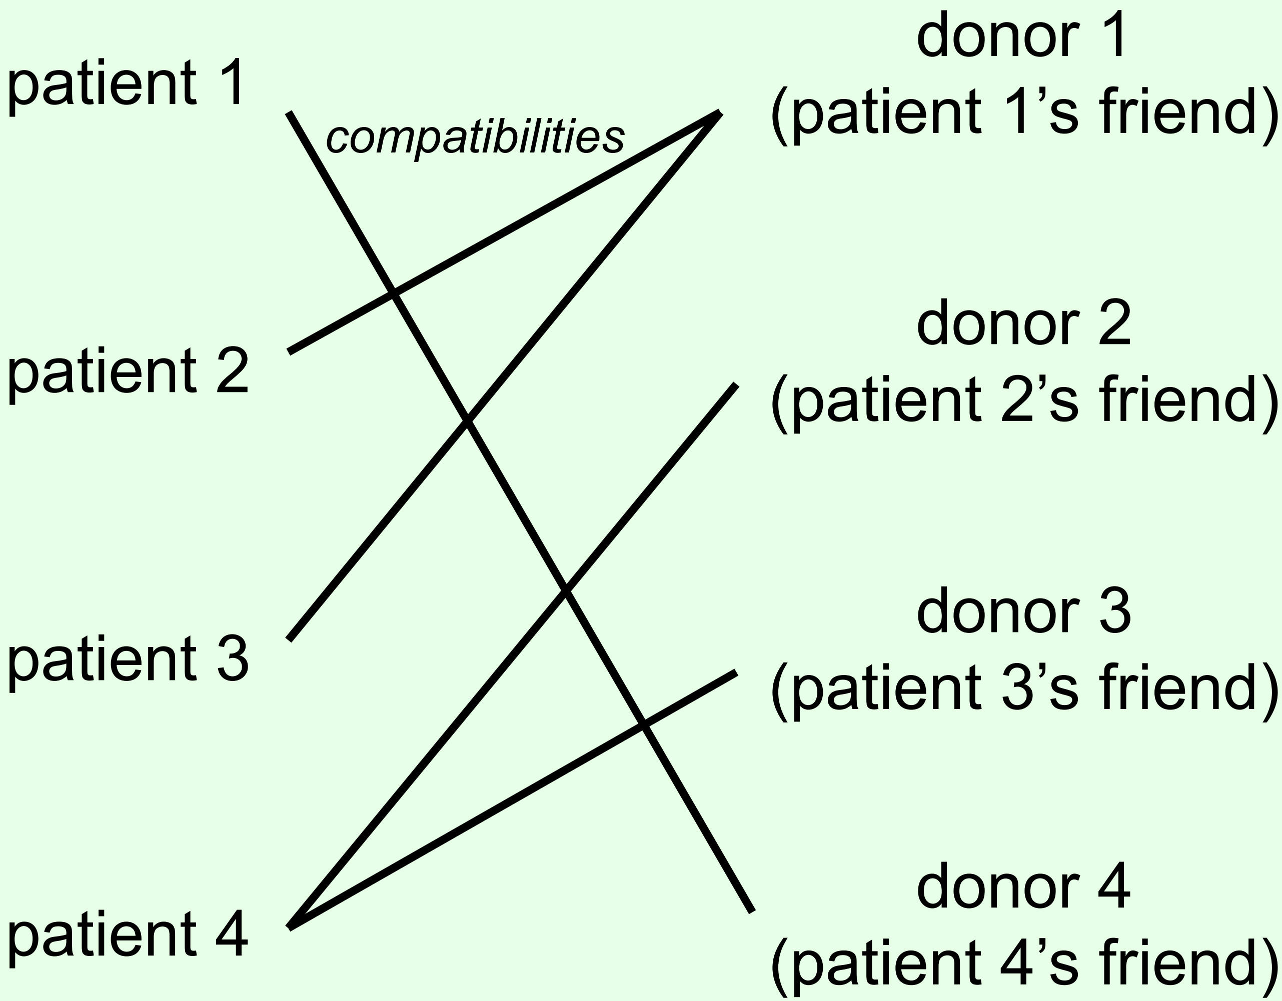

Game playing & AI (in Part 2)

perfect information games: no uncertainty about the state of the game (e.g. tic-tac-toe, chess, Go)

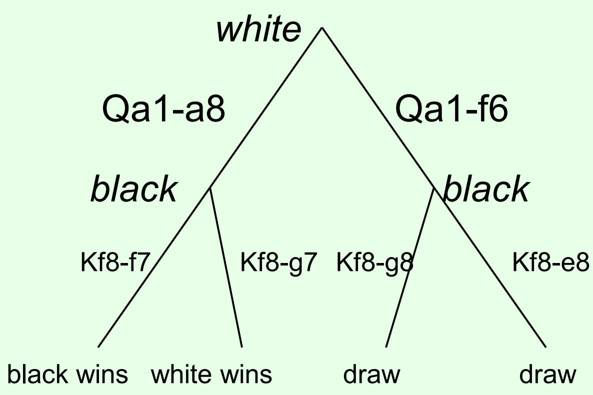

- Optimal play: value of each node = value of optimal child for current player (backward induction, minimax)
- For chess and Go, tree is too large
  - Use other techniques (heuristics, limited-depth search, alpha-beta, deep learning, …)
- Top computer programs better than humans in chess, ~~not yet in Go~~

imperfect information games: uncertainty about the state of the game (e.g., poker)

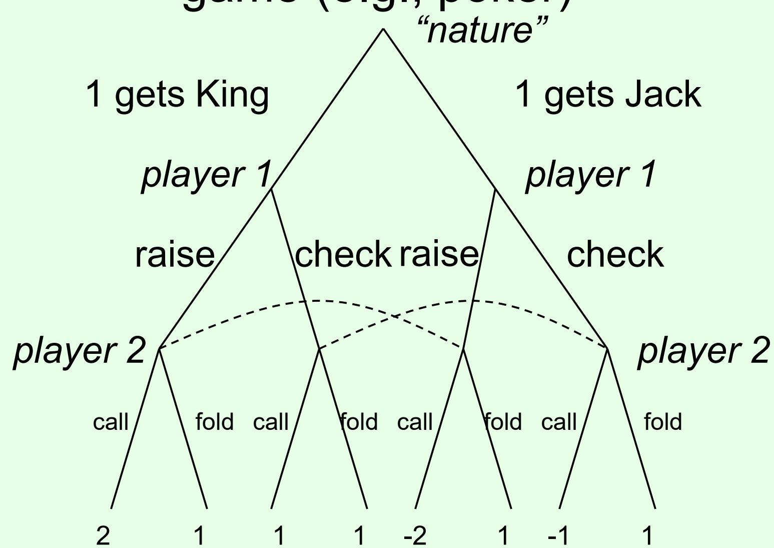

- Player 2 cannot distinguish nodes connected by dotted lines
  - Backward induction fails; need more sophisticated game-theoretic techniques for optimal play
- Small poker variants can be solved optimally
- Humans still better than top computer programs at full-scale poker (at least most versions) ?
- Top computer (heads-up) poker players are based on techniques for game theory

Real-world security applications (in Part 2)

Milind Tambe's TEAMCORE group (USC)

Airport security

Where should checkpoints, canine units, etc. be deployed?

Federal Air Marshals

Which flights get a FAM?

US Coast Guard

Which patrol routes should be followed?

Wildlife Protection

Where to patrol to catch poachers or find their snares?

Prediction markets

PredictIt

Markets | Analysis | About | Login

Closing Soon

U.S. Elections

U.S. Politics

World

Who will win the 2020 U.S. presidential election?

Market Type: Linked

End Date: N/A

Status: Open

Contracts | Rules | Chart

Trade shares from this page by clicking any price in bold. For more information on an individual prediction, click on the name or image.

PREZ.2020	Latest	Buy Yes	Sell Yes	Buy No	Sell No
Donald Trump TRUM.PREZ.2020	31¢ ↓ 1¢	33¢	31¢	69¢	67¢
Bernie Sanders SAND.PREZ.2020	13¢ NC	13¢	12¢	88¢	87¢
Joe Biden BIDE.PREZ.2020	10¢ NC	11¢	10¢	90¢	89¢
Kamala Harris HARR.PREZ.2020	9¢ NC	10¢	8¢	92¢	90¢
Kirsten Gillibrand GILL.PREZ.2020	9¢ ↑ 1¢	9¢	8¢	92¢	91¢
Mike Pence PENC.PREZ.2020	8¢ NC	8¢	7¢	93¢	92¢
Elizabeth Warren WARR.PREZ.2020	7¢ ↑ 1¢	7¢	6¢	94¢	93¢
Cory Booker BOOK.PREZ.2020	7¢ NC	8¢	7¢	93¢	92¢
John Kasich KASI.PREZ.2020	4¢ ↑ 1¢	4¢	3¢	97¢	96¢
Amy Klobuchar KLOB.PREZ.2020	4¢ NC	5¢	4¢	96¢	95¢
Andrew Cuomo CUOM.PREZ.2020	3¢ NC	4¢	3¢	97¢	96¢
Nikki Haley HALE.PREZ.2020	2¢ NC	3¢	2¢	98¢	97¢
Paul Ryan RYAN.PREZ.2020	2¢ NC	3¢	2¢	98¢	97¢
Mark Cuban CUBA.PREZ.2020	1¢ NC	2¢	1¢	99¢	98¢
Mark Zuckerberg	1¢ NC	2¢	1¢	99¢	98¢

Financial securities (in Part 1)

- Tomorrow there must be one of   
- Agent 1 offers \$5 for a security that pays off \$10 if  or 
- Agent 2 offers \$8 for a security that pays off \$10 if  or 
- Agent 3 offers \$6 for a security that pays off \$10 if 
- Can we accept some of these at offers at no risk?

How to incentivize a weather forecaster (in Part 3)

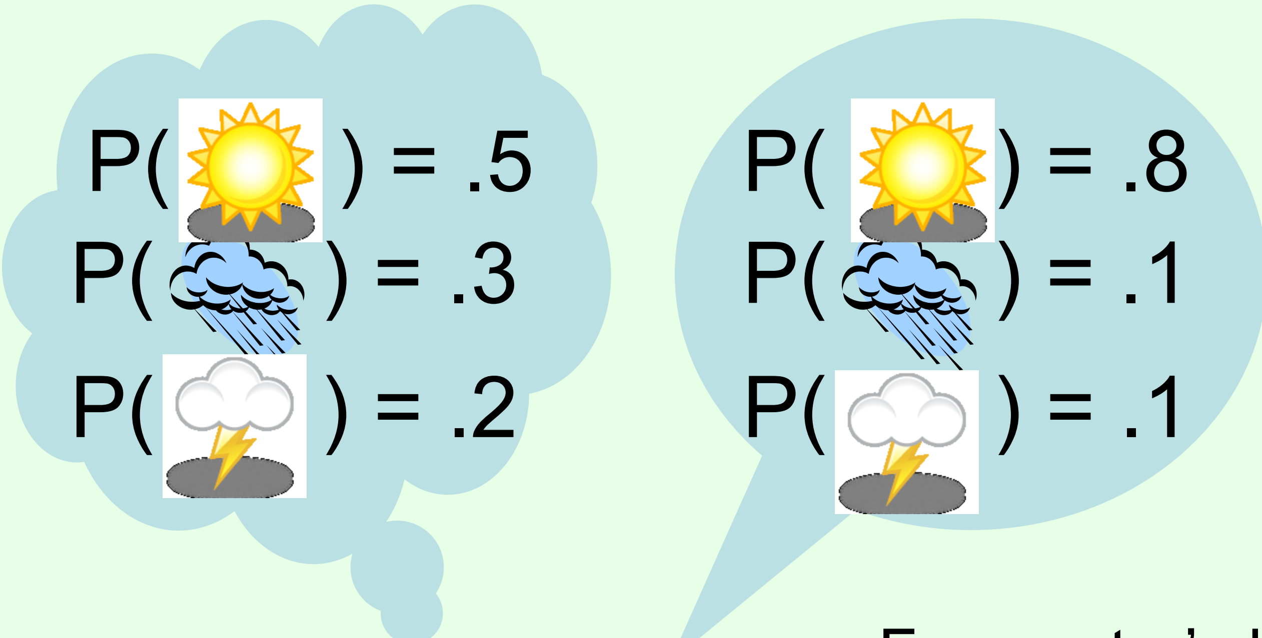

- Forecaster's bonus can depend on
  - Prediction
  - Actual weather on predicted day
- Reporting true beliefs should maximize expected bonus

Sponsored search / ad auctions (in Part 3)

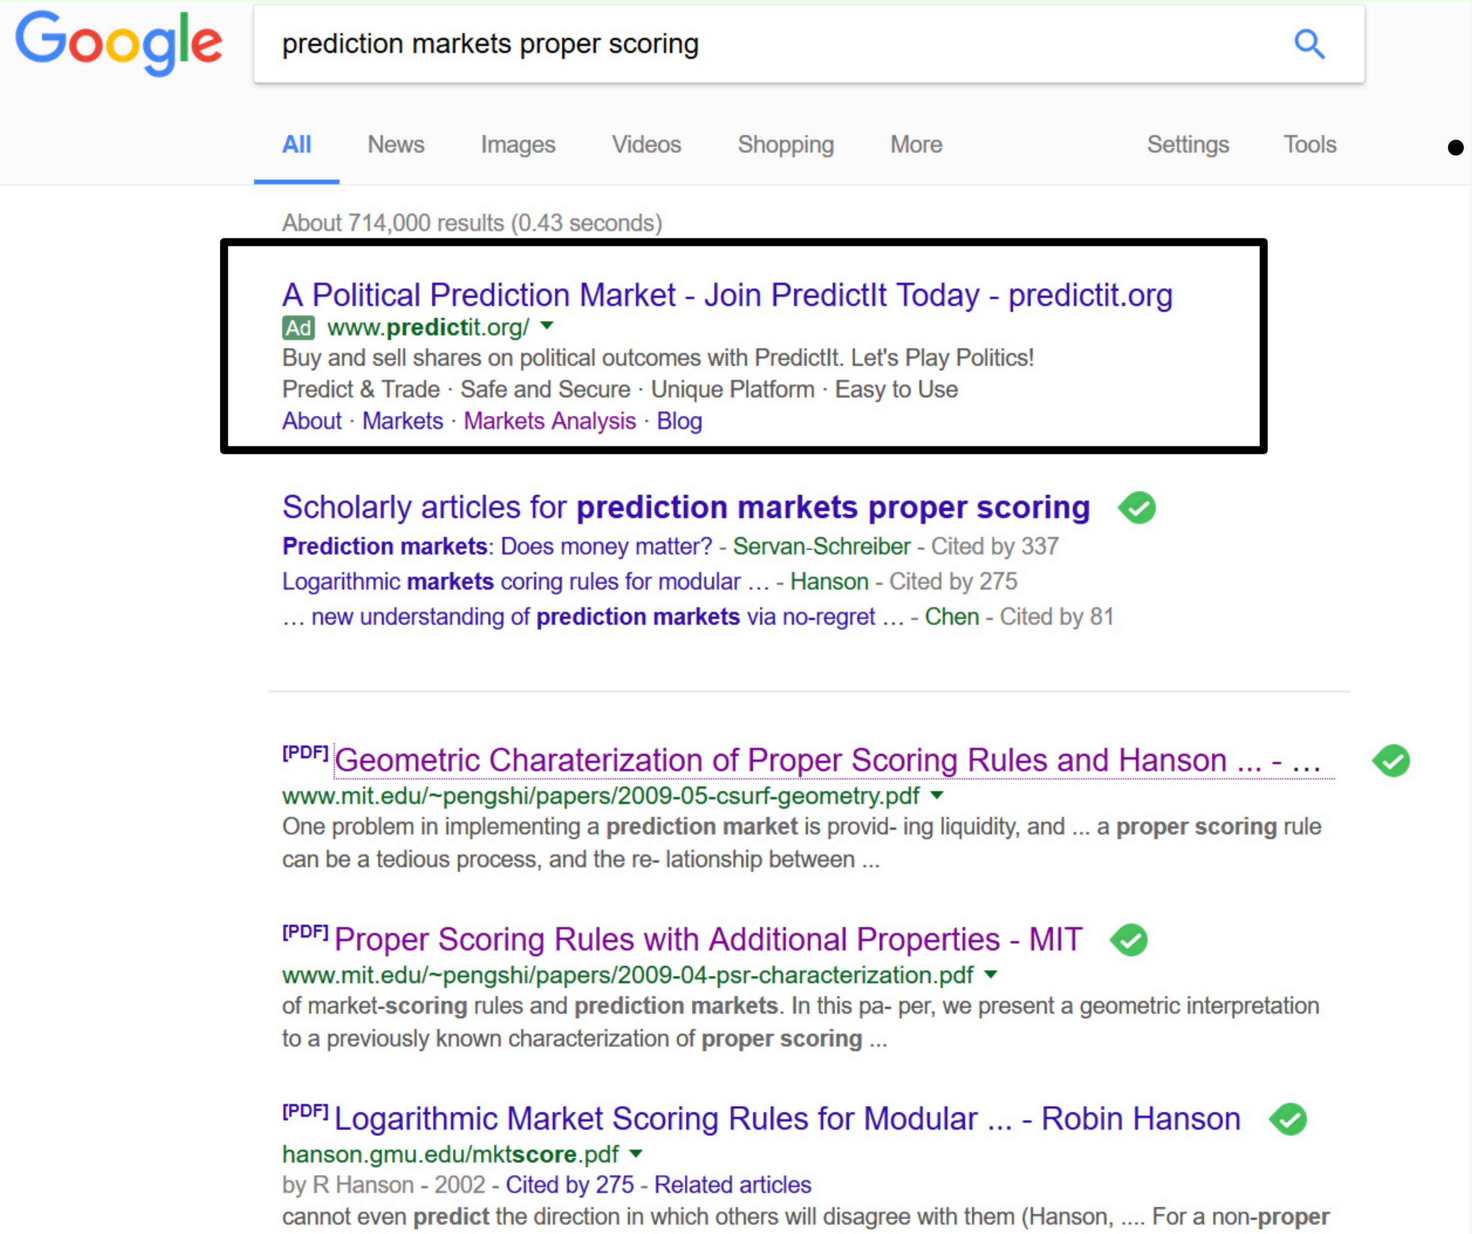

- Choice of ads (if any) to show determined by:
  - Advertiser bid
  - Predicted likelihood of click