15-326 Computational Microeconomics

Instructor: Vincent Conitzer
conitzer@cs.cmu.edu
http://www.cs.cmu.edu/~15326-f24/

TA: Jiayuan Liu

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?

- "a social science that studies the production, distribution, and consumption of goods and services." [Wikipedia, Aug. 2024]
- 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

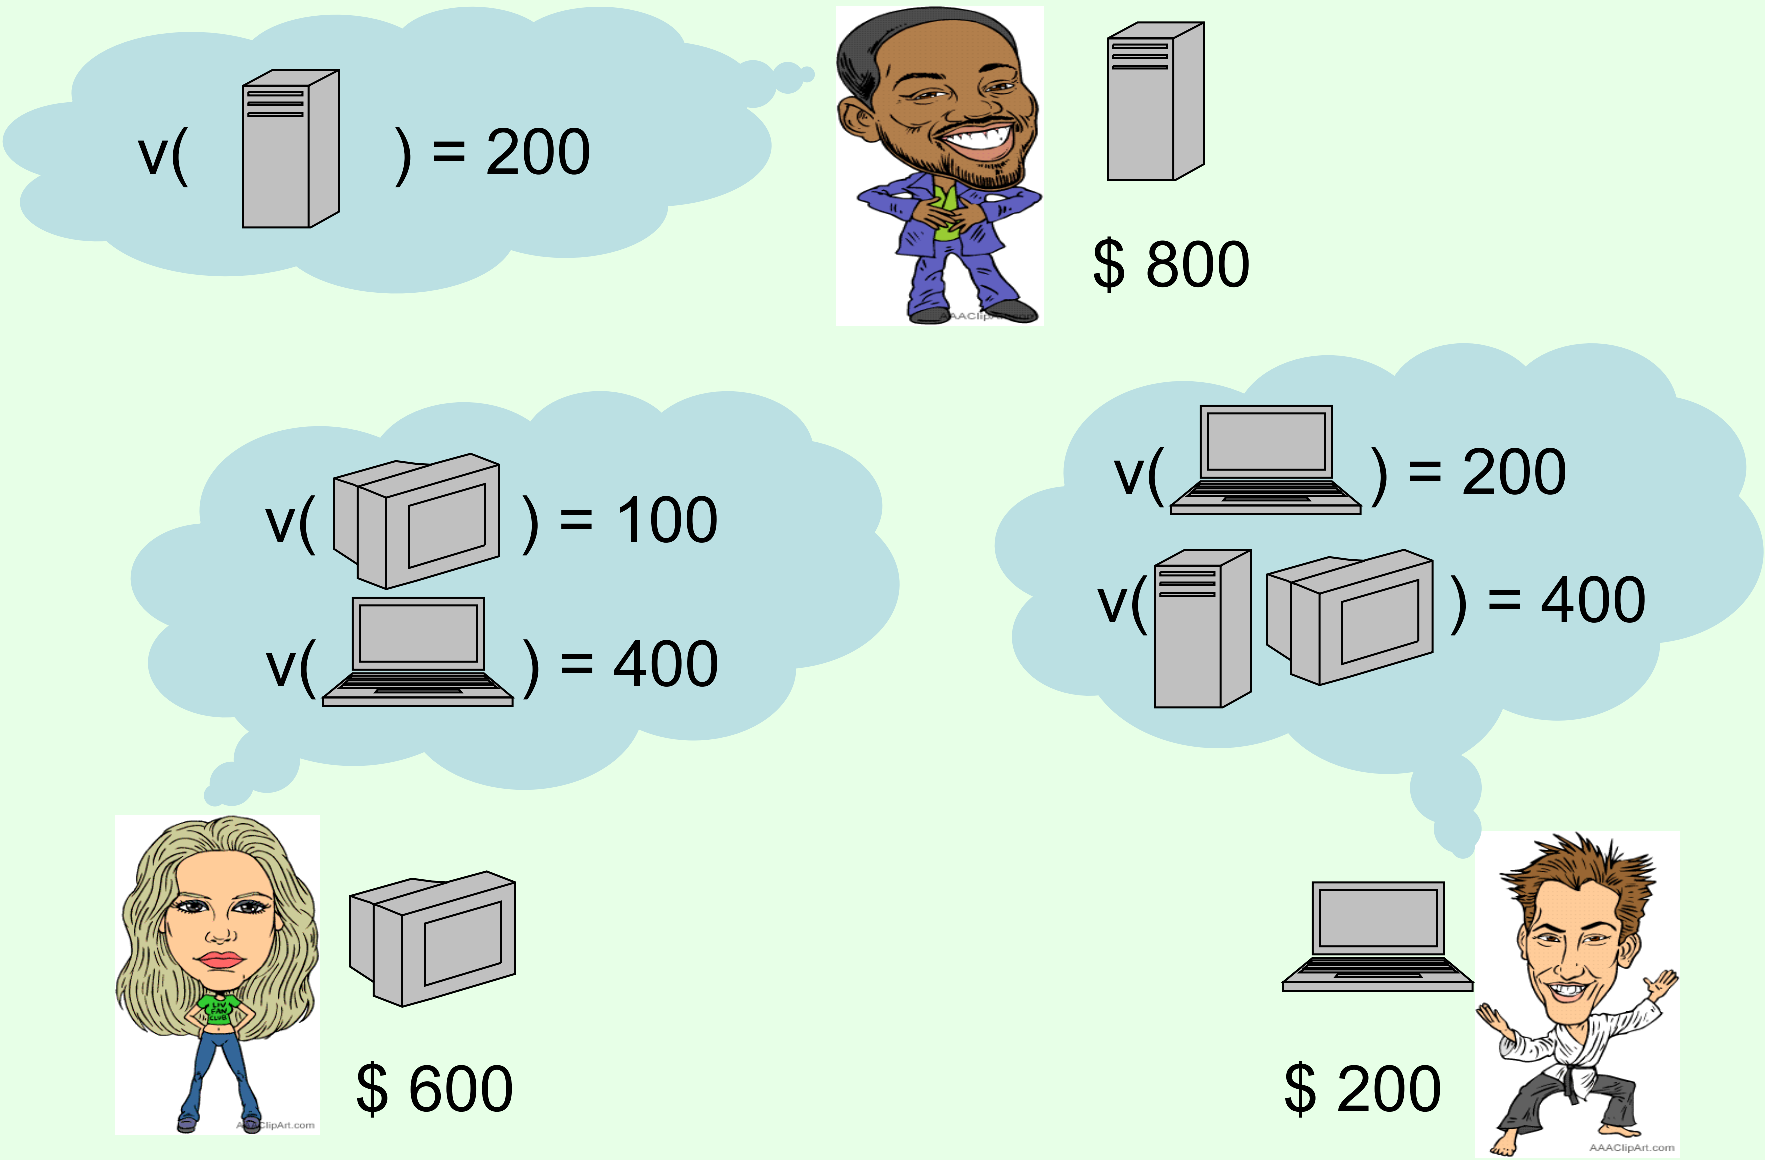

After trade (a more efficient outcome)

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

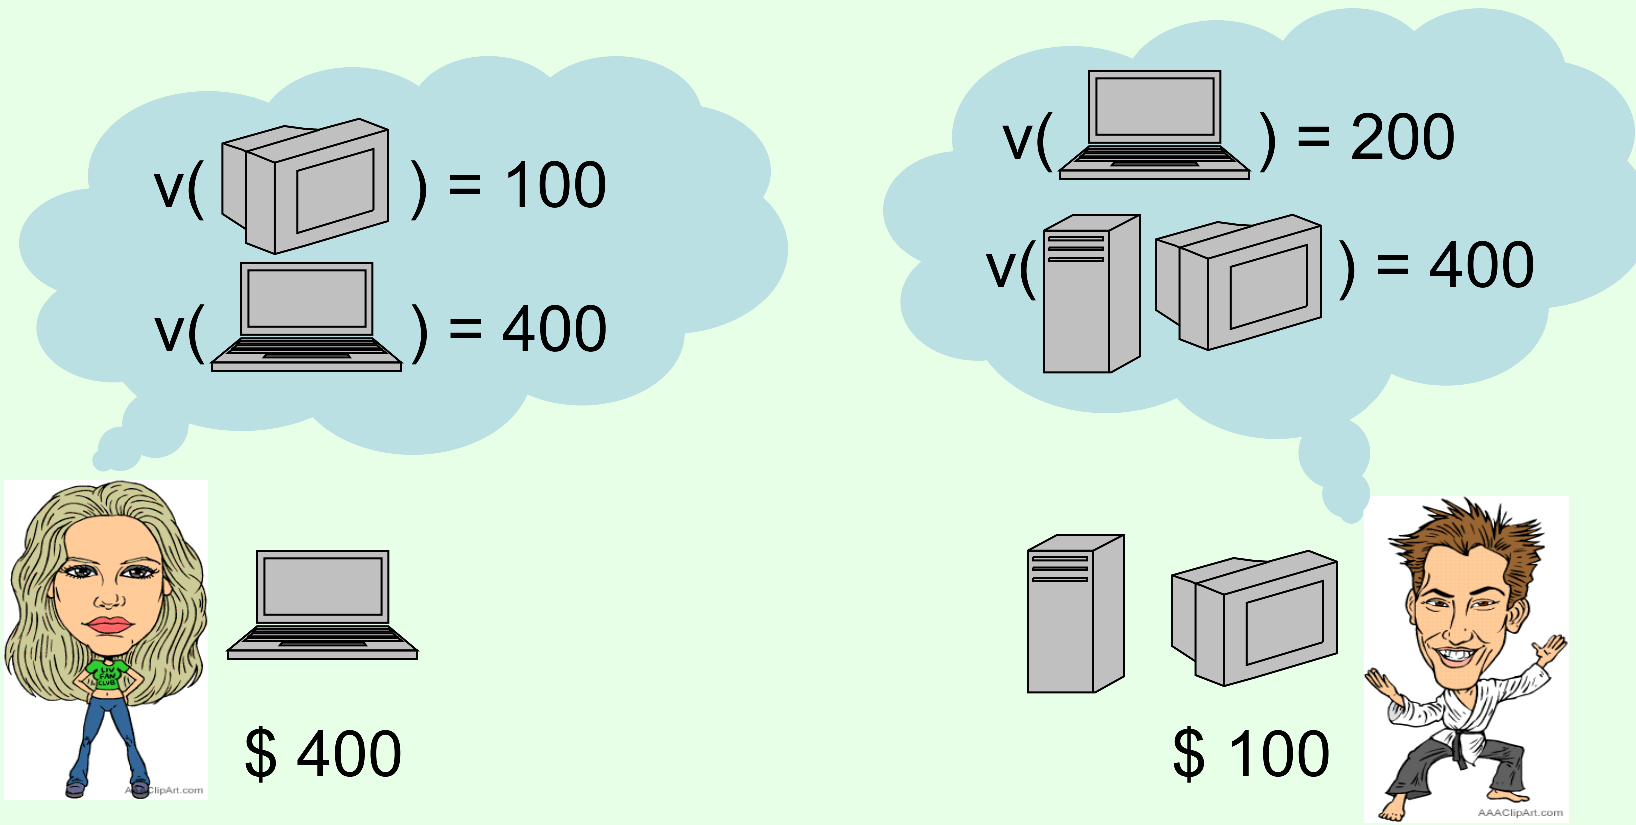

Some distinctions in economics

- 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?

- "Computer science is the study of computation, information, and automation. Computer science spans theoretical disciplines (such as algorithms, theory of computation, and information theory) to applied disciplines (including the design and implementation of hardware and software)." [Wikipedia, Aug. 2024]
- 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)*

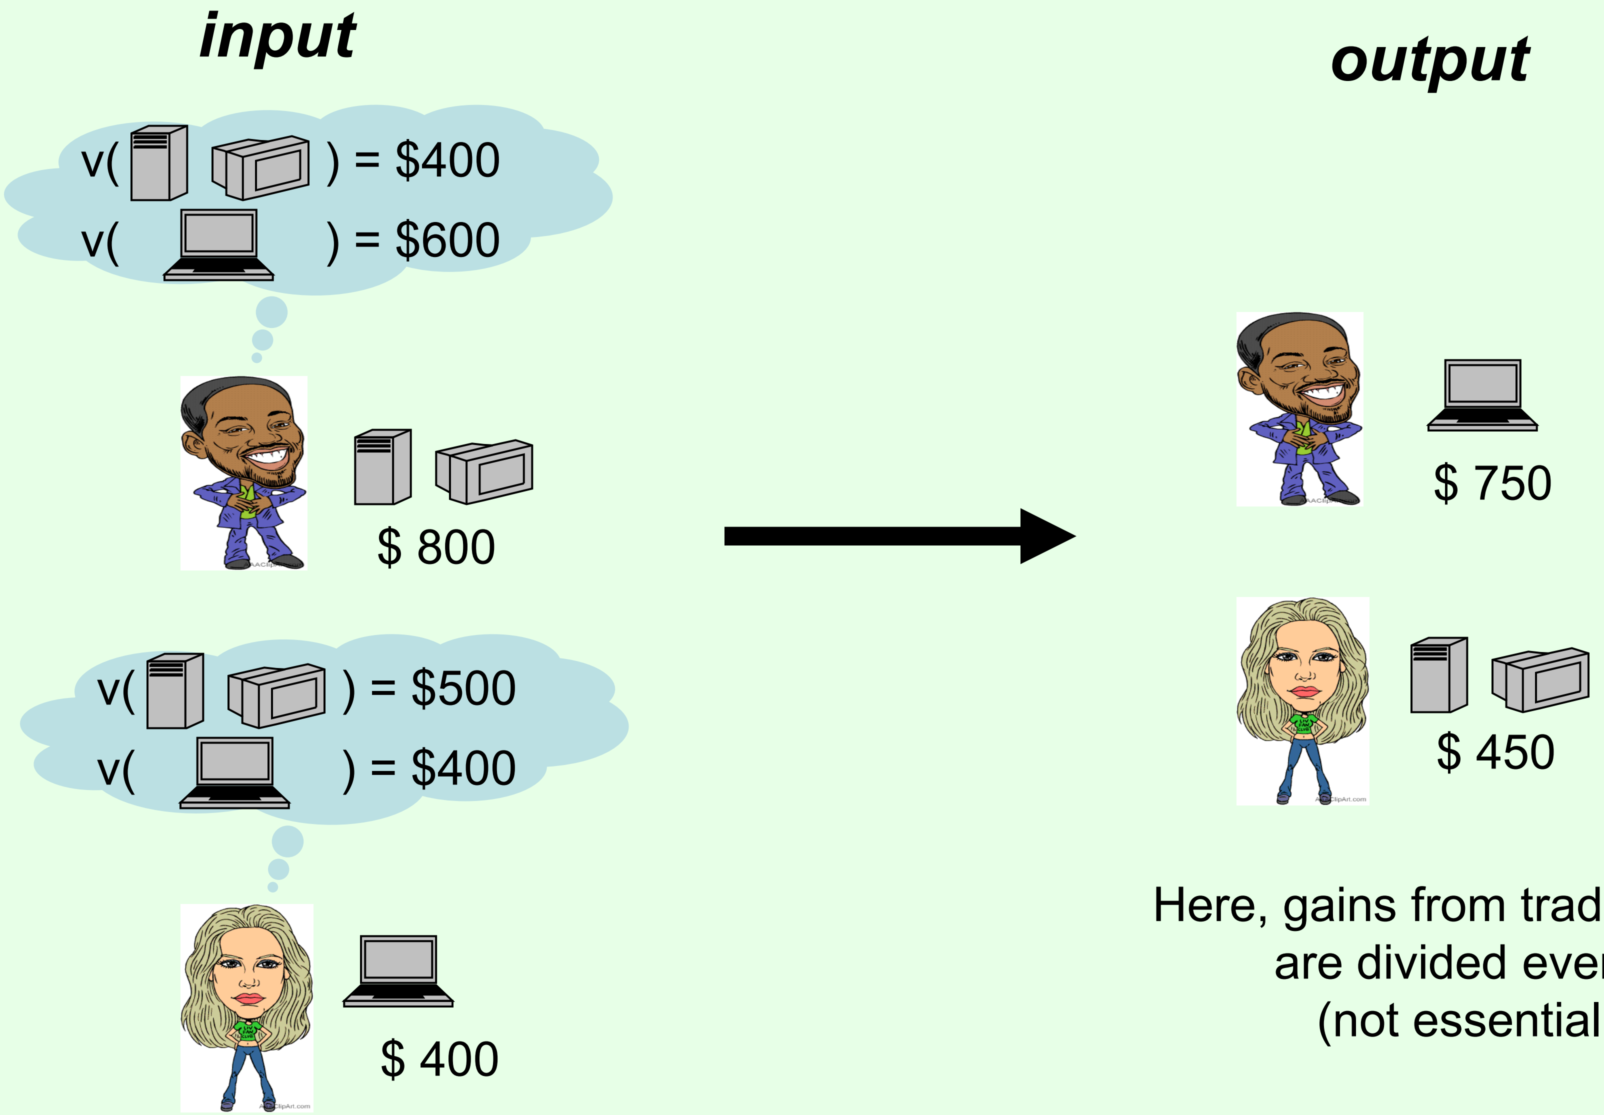

Here, gains from trade ($300) are divided evenly (not essential)

Economic mechanisms

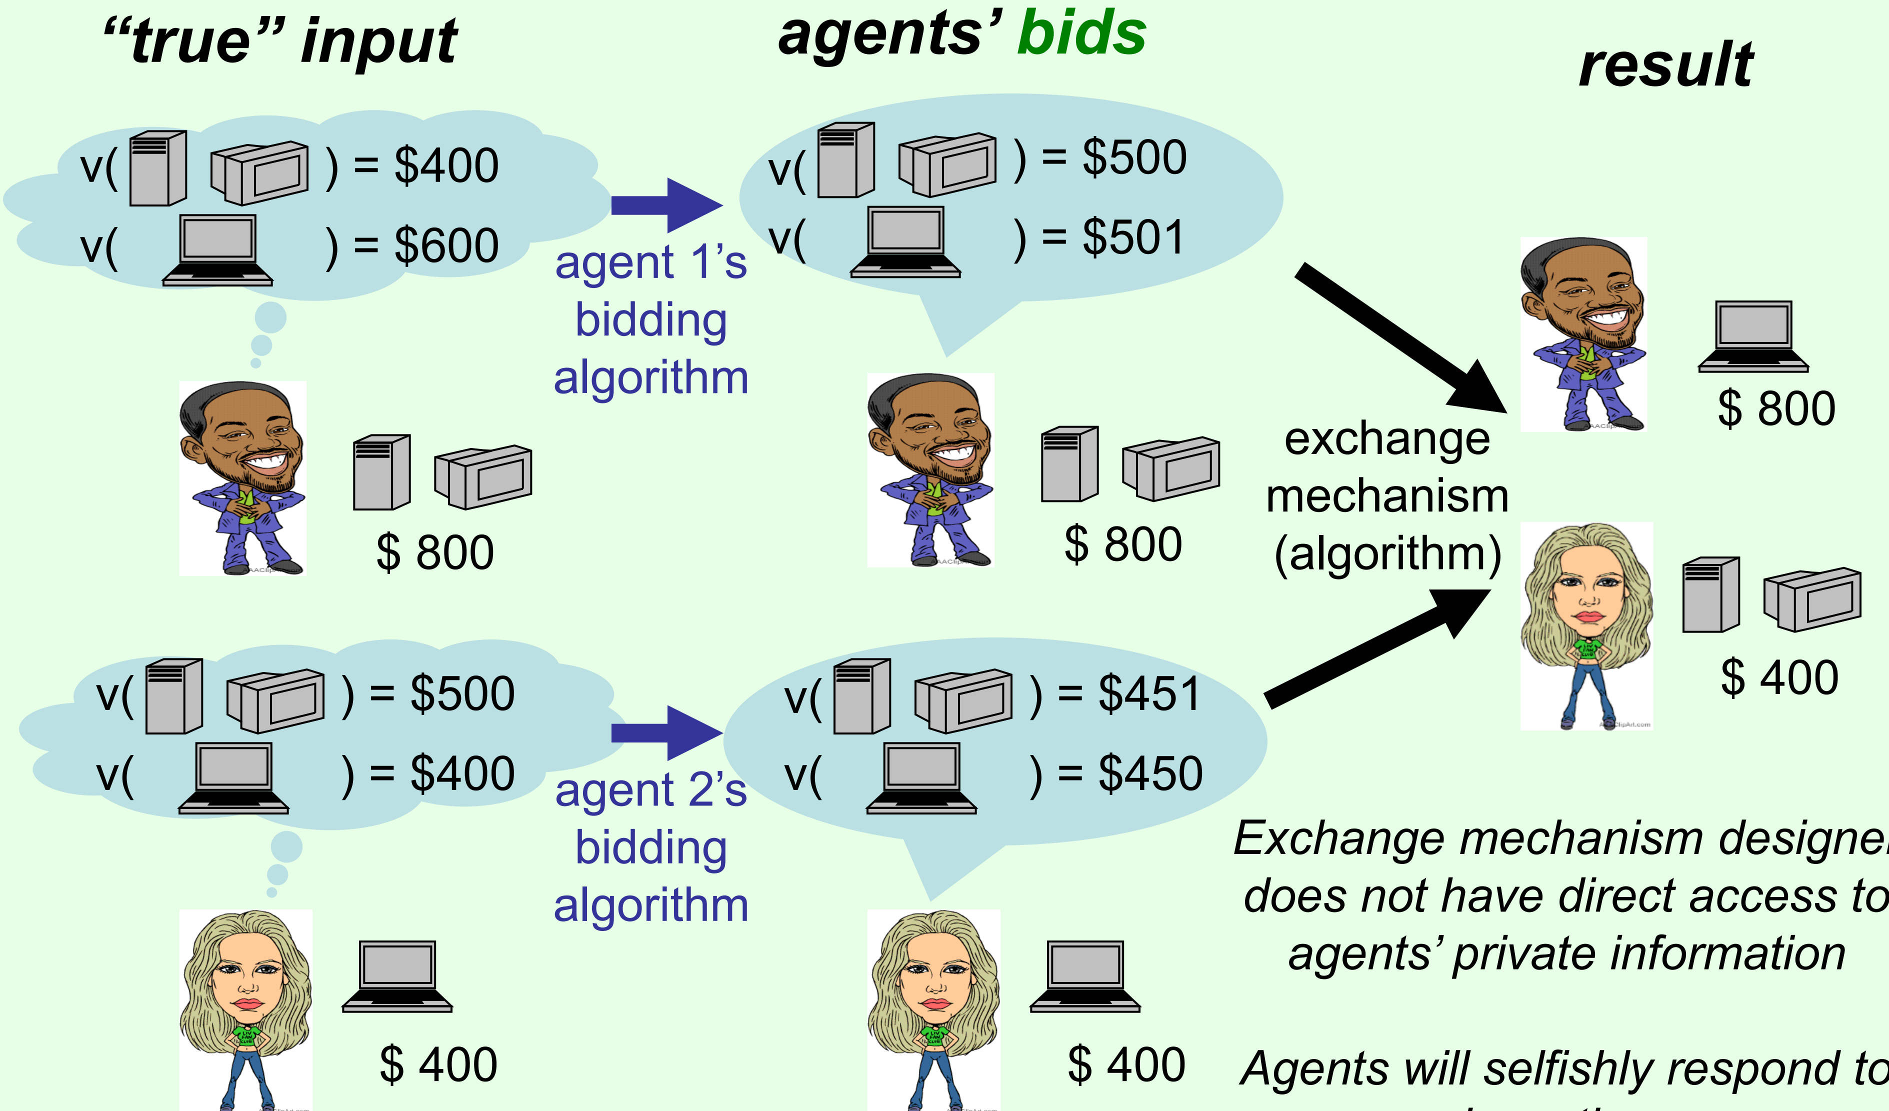

Exchange mechanism designer does not have direct access to agents' private information

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

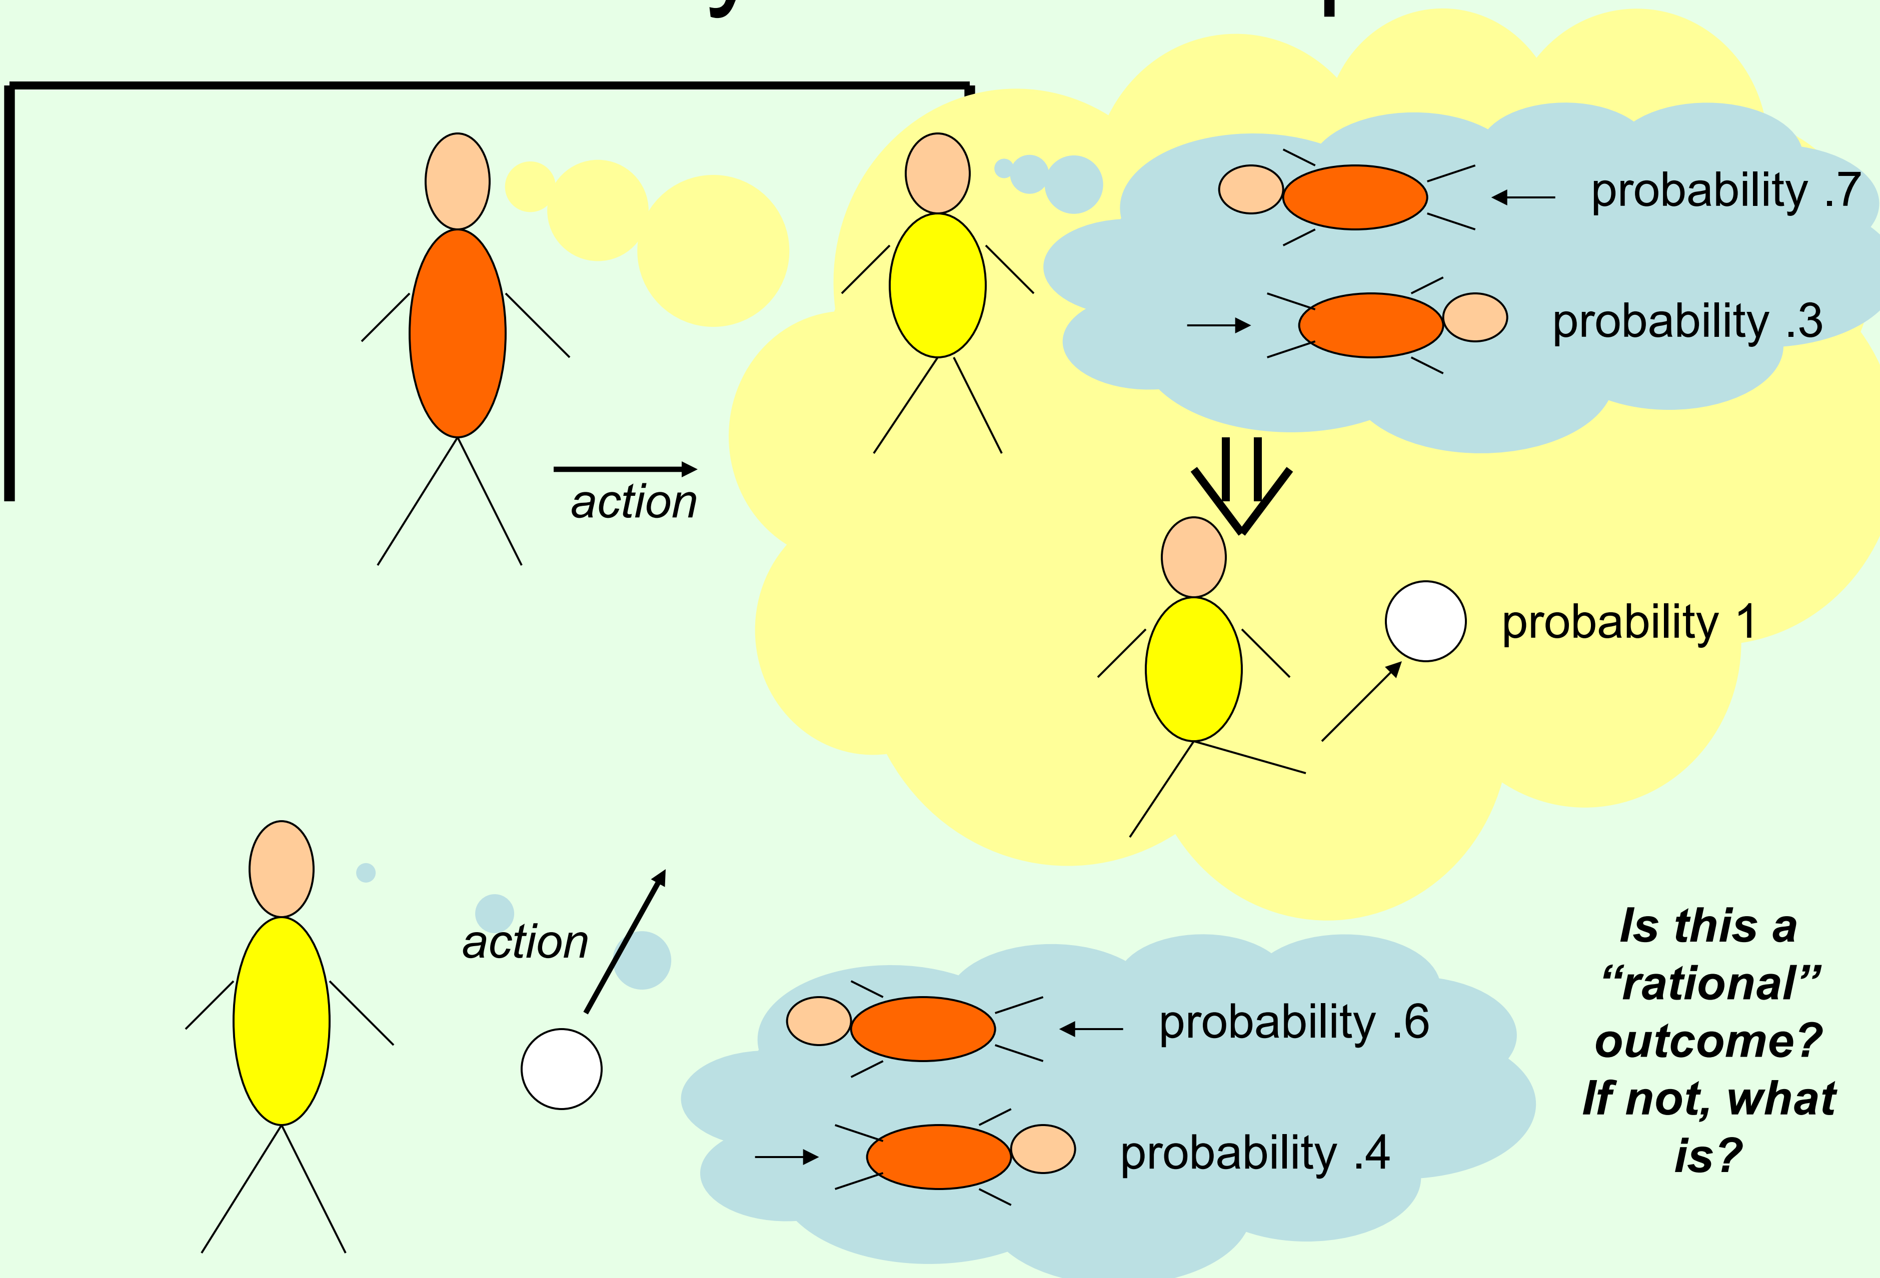

Is this a "rational" outcome? If not, what is?

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

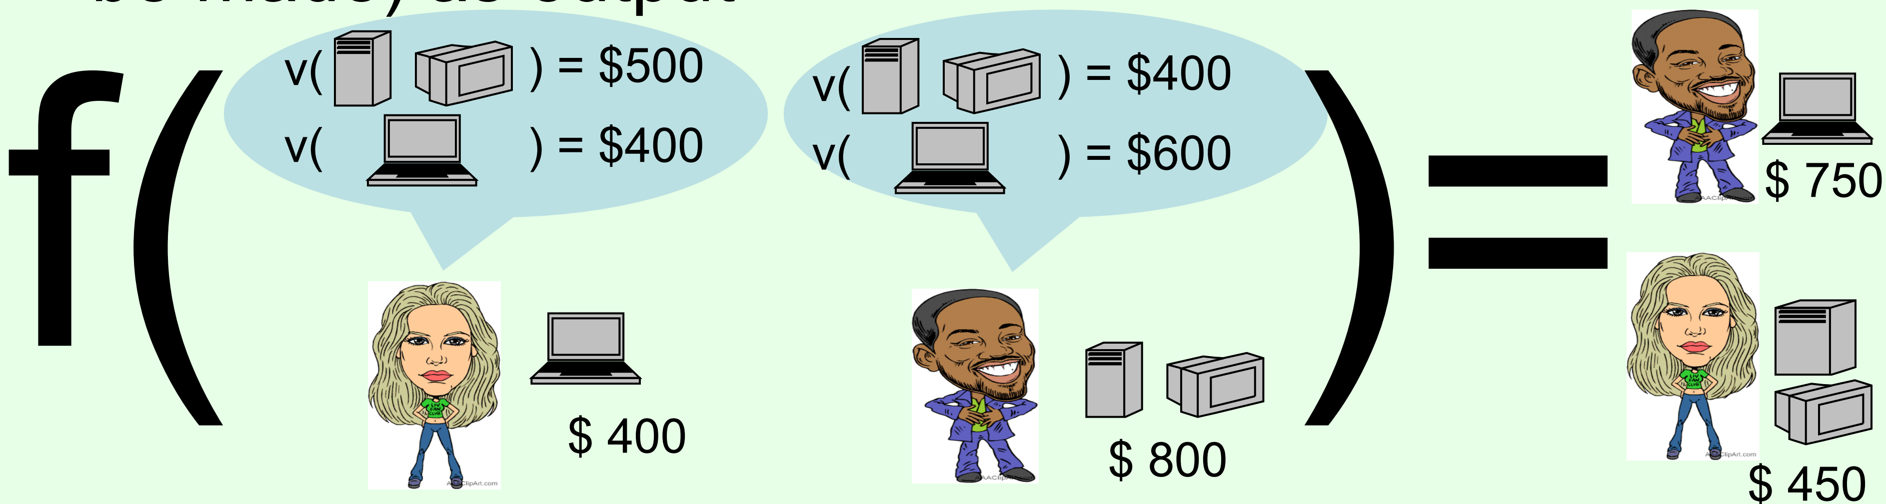

- The entire function f is one mechanism
- E.g., the mechanism from part 1: find allocation that maximizes (reported) utilities, distribute (reported) gains evenly
- Other mechanisms choose different allocations, payments

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

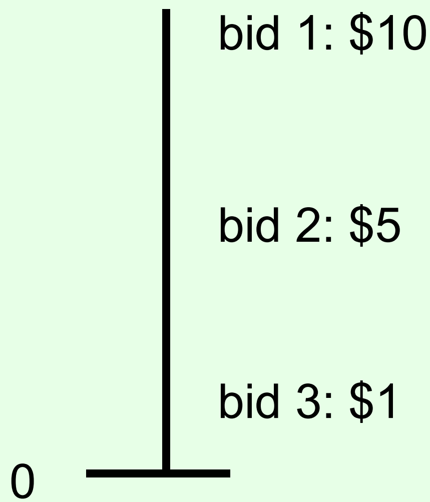

first-price: bid 1 wins, pays $10
second-price: bid 1 wins, pays $5

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 ($\rightarrow$ 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

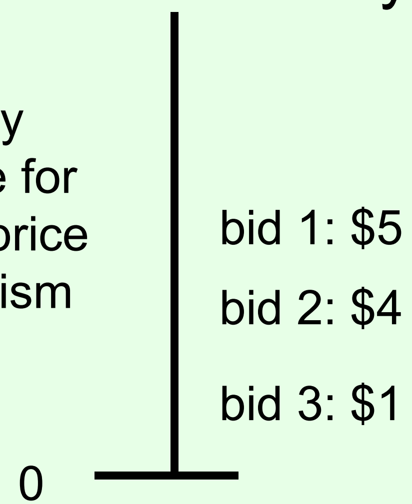

a likely outcome for the second-price mechanism

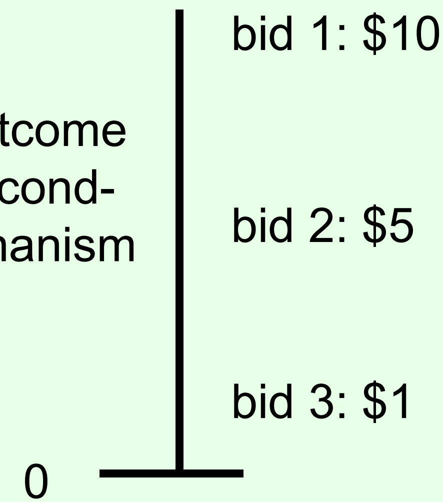

Are there other auctions that perform better? How do we know when we have found the best one?

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: 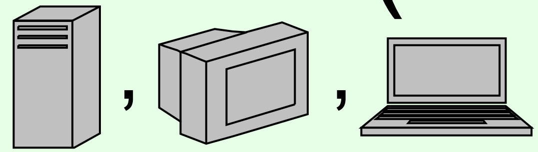

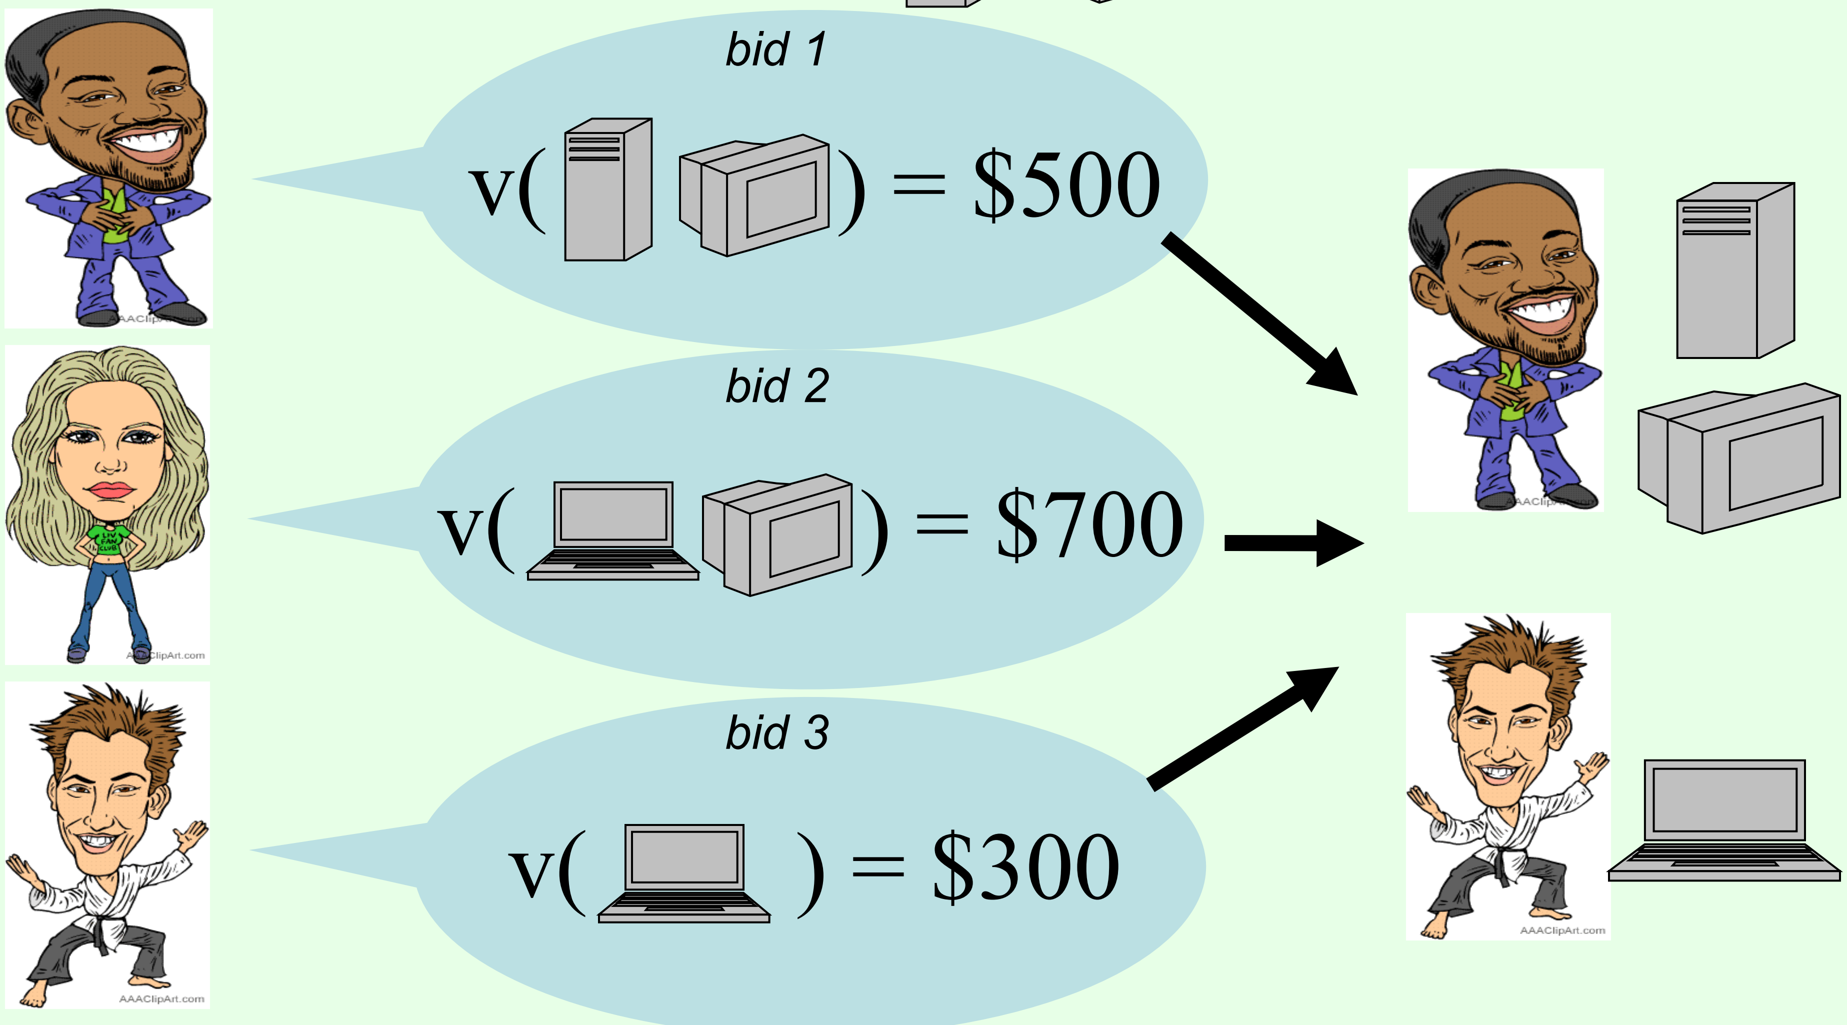

used in truckload transportation, industrial procurement, radio spectrum allocation, …

Voting (in Part 1)

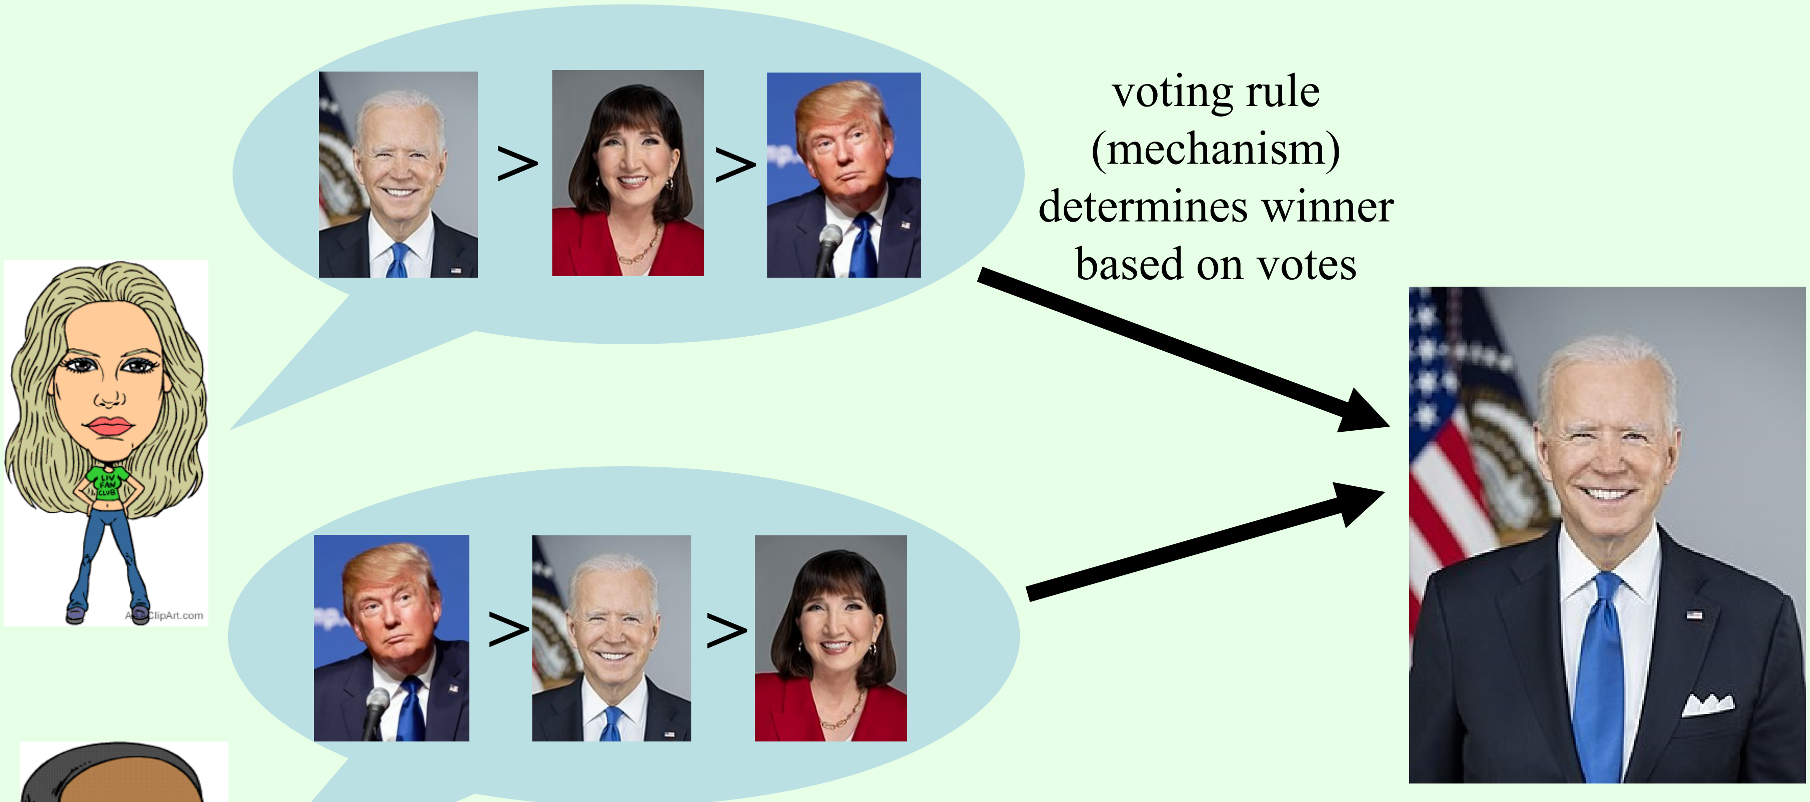

- Can vote over other things too
  - Where to go for dinner tonight, other joint plans, …
- Many different rules exist for selecting the winner

Kidney exchange (in Part 1)

- Kidney exchanges allow patients with willing but incompatible live donors to swap donors

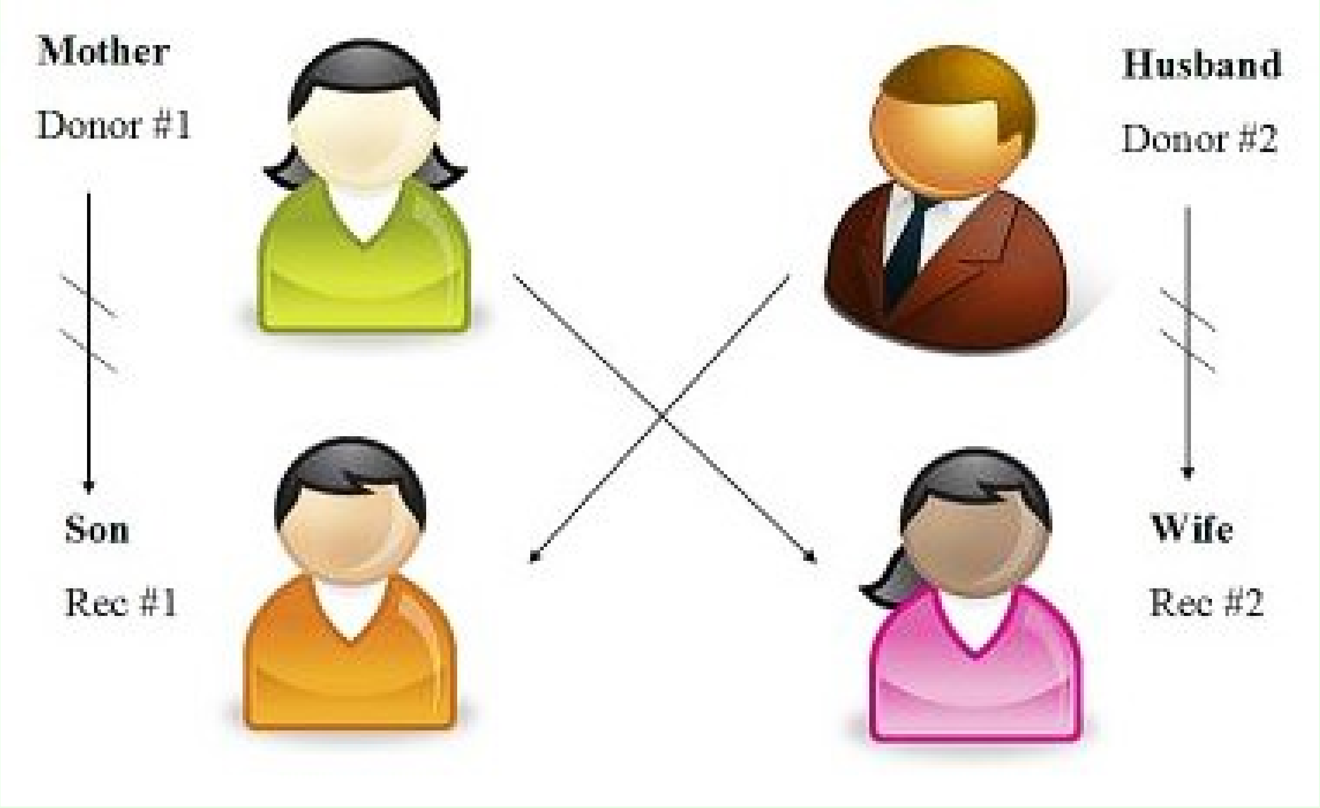

Kidney exchange (in Part 1)

POPULAR | LATEST | FEATURED QUARTZ OBSESSIONS | EMAILS | EDITIONS

Prescription AI

This series explores the promise of AI to personalize, democratize, and advance medicine—and the dangers of letting machines make decisions.

THE BOTPERATING TABLE

How AI changed organ donation in the US

By Corinne Purtill • September 10, 2018

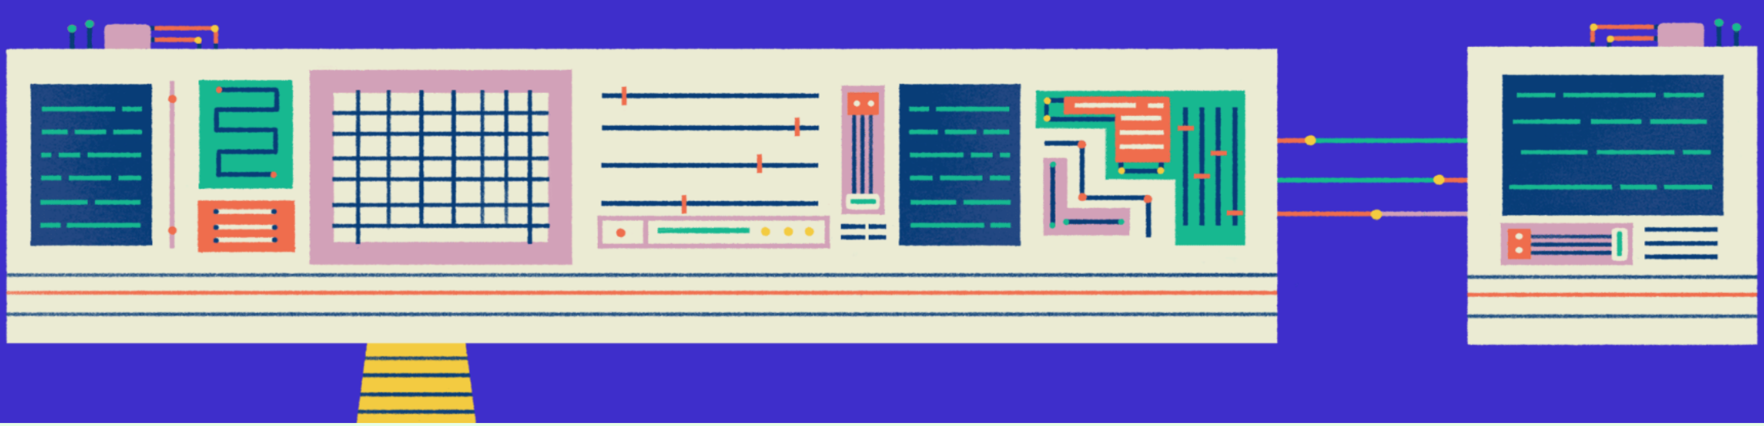

Kidney exchange (in Part 1)

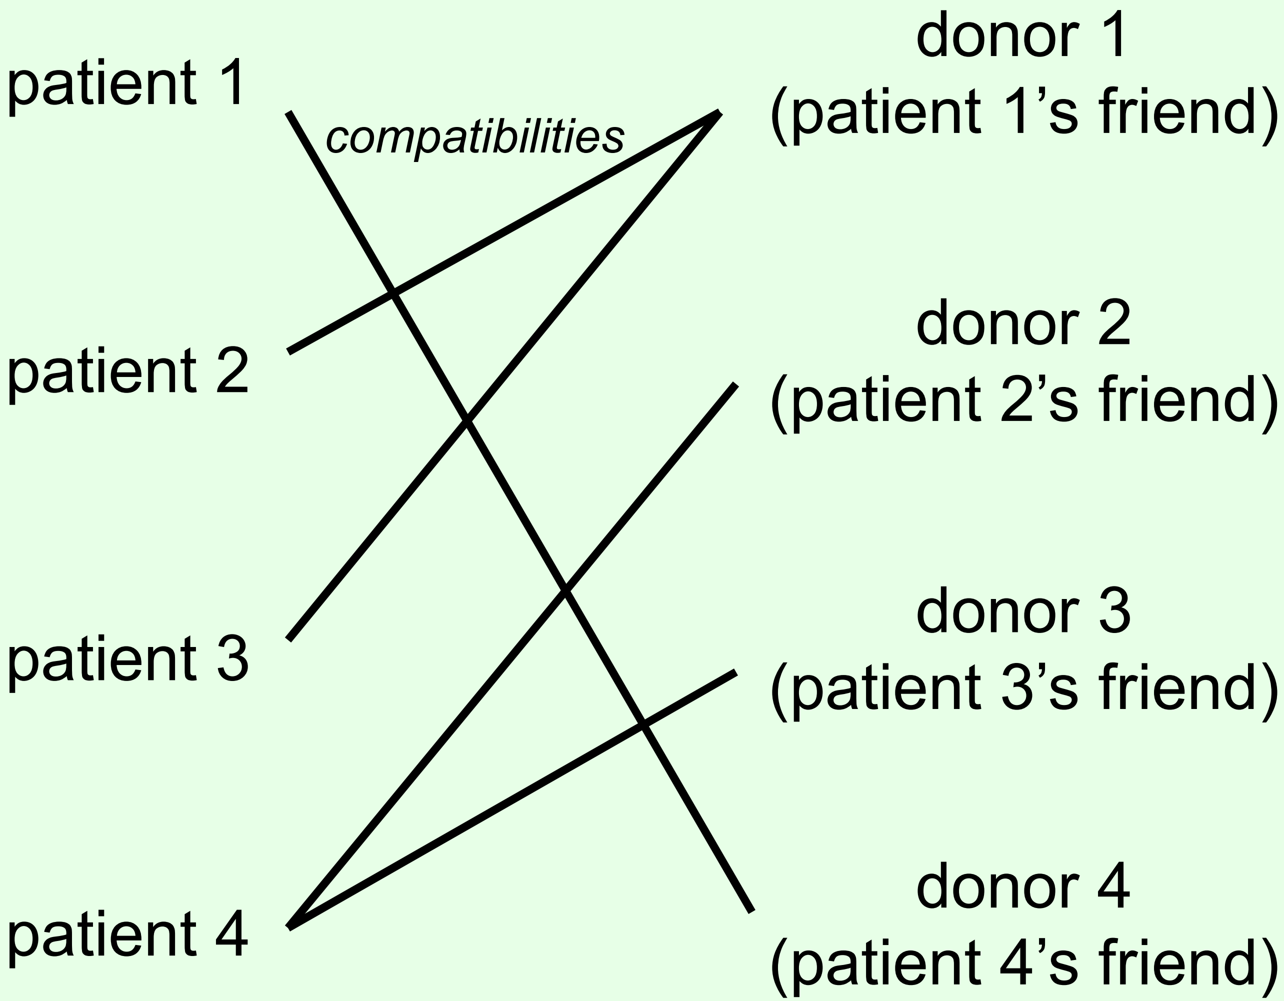

Game playing & AI (in Part 2)

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

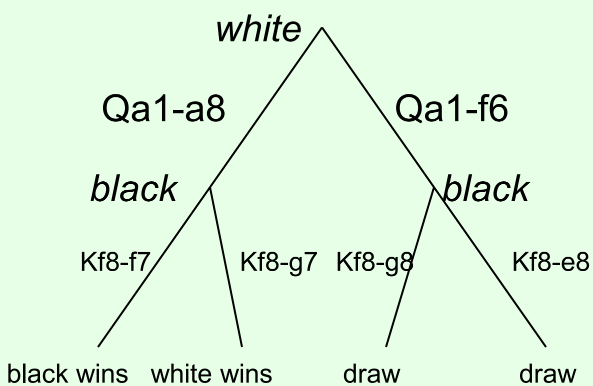

- Optimal play: value of each node = value of optimal child for current player (backward induction, minimax)
- For chess and Go, tree is too large
  - Use other techniques (heuristics, limited-depth search, alpha-beta, deep learning, …)
- Top computer programs better than humans in chess, ~~not yet in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

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

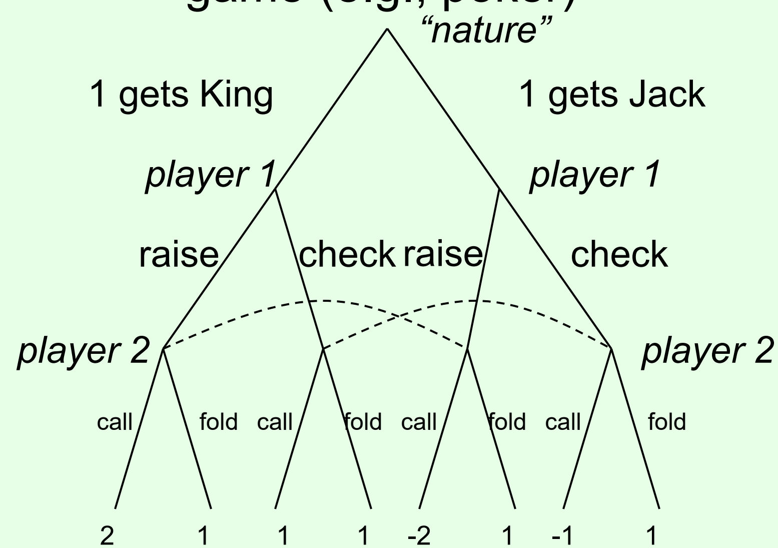

- Player 2 cannot distinguish nodes connected by dotted lines
  - Backward induction fails; need more sophisticated game-theoretic techniques for optimal play
- Small poker variants can be solved optimally
- ~~Humans still better than top computer programs at full-scale poker (at least most versions)~~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

SIGN UP FOR OUR DAILY NEWSLETTER
Country *
Email Address * Sign Up
Click to view the privacy policy.
Required fields are indicated by an asterisk (*)
306

Artificial intelligence masters multiplayer poker

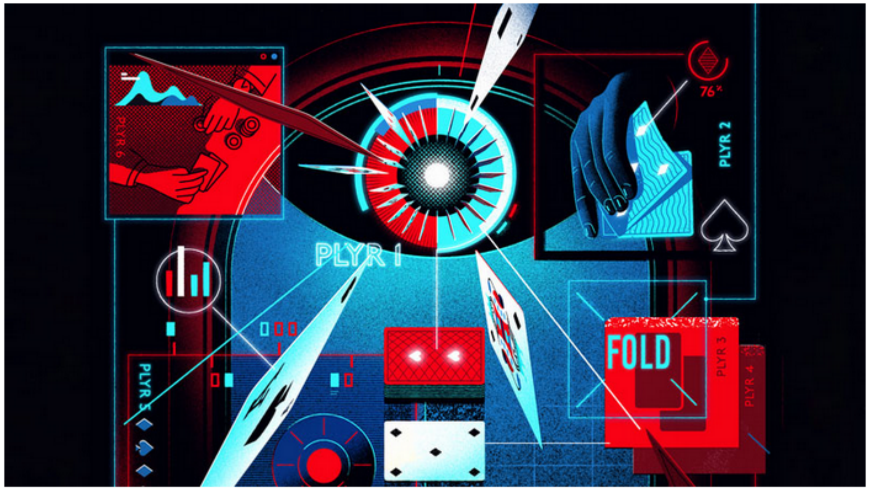

JASON SOLO/THE JACKY WINTER GROUP

This year, an artificial intelligence (AI) program beat some of the world's best players in the most popular version of poker, no-limit Texas Hold 'em. The landmark result marks the first time AI has prevailed in a multiplayer contest in which players have only imperfect information about the state of the game.

AI has been trouncing humans in games at a spectacular rate. In 2007, computer scientists developed a program guaranteed not to lose at checkers. In 2016, another team developed an AI program that defeated the best humans at Go, a board game with vastly more configurations than checkers.

Poker presents a stiffer challenge, as players cannot see their opponents' cards and thus have limited information. In 2017, computer scientists developed an AI program unbeatable at a two-player version of Hold 'em—in which each player forms a hand from five cards laid face up on the table and two more each holds privately.

Now, AI has bested world-class players in the full multiplayer game, as computer scientists at Carnegie Mellon University in Pittsburgh, Pennsylvania, announced in August. By playing 1 trillion games against itself, their program, Pluribus, developed a basic strategy for various kinds of situations—say, playing for an inside straight. For each specific hand, it could also think through how the cards would likely play out. In 20,000 hands with six players it outperformed 15 top-level players, as measured by average winnings per hand.


LEE SEDOL

Real-world security applications (in Part 2)

Milind Tambe's TEAMCORE group (USC → Harvard)

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?

Global Presence of Security Games Efforts

Prediction markets

(Aug. 24, 2024)

PredictIt Markets Leaderboards Support Login Sign Up

Presidency Congress Donald Trump Kamala Harris

Who will win the 2024 US presidential election?

	Contract	Latest Yes Price		Best Offer			Best Offer
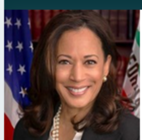	Kamala Harris	55¢	NC	55¢	Buy Yes	Buy No	46¢
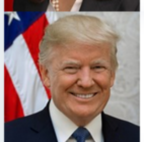	Donald Trump	49¢	NC	49¢	Buy Yes	Buy No	52¢
	Nikki Haley	2¢	NC	2¢	Buy Yes	Buy No	N/A
	Ron DeSantis	1¢	NC	1¢	Buy Yes	Buy No	N/A
	Joe Biden	1¢	NC	1¢	Buy Yes	Buy No	N/A

5 More Contracts

Prediction markets

(Aug. 24, 2024)

PredictIt Markets Leaderboards Support Login Sign Up

Presidency Congress Donald Trump Kamala Harris

PredictIt Announcements

Visit the PredictIt Status Page for real-time updates

Notice to Traders

Victoria University of Wellington ("VUW") has responded to a March 2, 2023 preliminary decision from the Commodity Futures Trading Commission ("CFTC") to withdraw the No Action Letter under which PredictIt operates. The March 2 letter provided an opportunity for VUW, PredictIt's operator to contest the Commission's arguments.

VUW issued the following statement in connection with their response to the CFTC:

"While Te Herenga Waka—Victoria University of Wellington respects the CFTC's authority as a regulator, it strongly disagrees with the version of events put forward by the Commission in its recent letter proposing to withdraw the "No Action" letter. The University has provided a detailed and robust response to the Commission in relation to all of the points raised. In particular, the University believes it has been transparent and has engaged in good faith regarding the operation of PredictIt and how this has changed over time, including the role of Aristotle International, Inc. The University has made no money from PredictIt, with the only payment being $US2,000 per month to the University's subsidiary Wellington UniVentures, to offset costs. The University's goal in relation to the PredictIt platform has been purely to support its development as a research and educational tool for the international research community of which we are a part."

The CFTC's March 2 letter withdrew an earlier August 4, 2022 withdrawal decision that has been temporarily enjoined. PredictIt traders, university educators, and market servicer, Aristotle International Inc. are challenging the CFTC's efforts to close PredictIt in a case that is currently pending before that court. VUW is not a party to that legal action.

PredictIt traders may continue to hold and trade existing contracts pending further consideration by the Fifth Circuit Court of Appeals and the CFTC. There remains the possibility that a judicial or administrative decision may require early termination of those contracts. We have no certainty as to the timing of any such decision.

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)

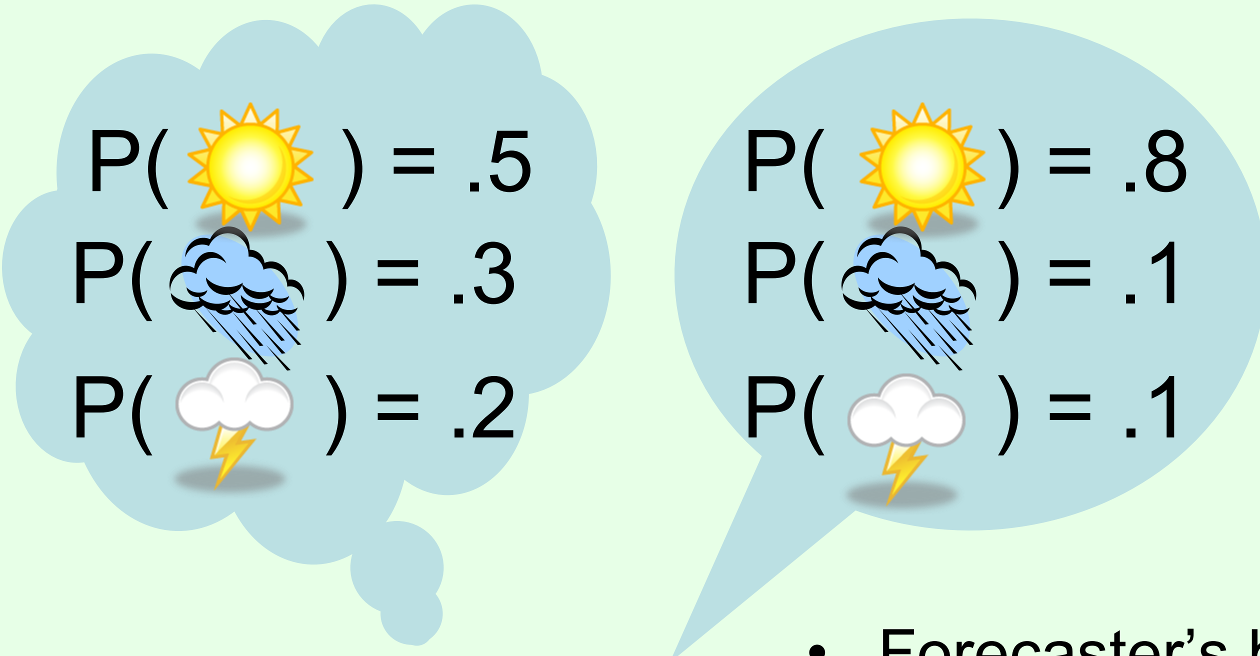

- Forecaster's bonus can depend on
  - Prediction
  - Actual weather on predicted day
- Reporting true beliefs should maximize expected bonus

Sponsored search / ad auctions (in Part 3)

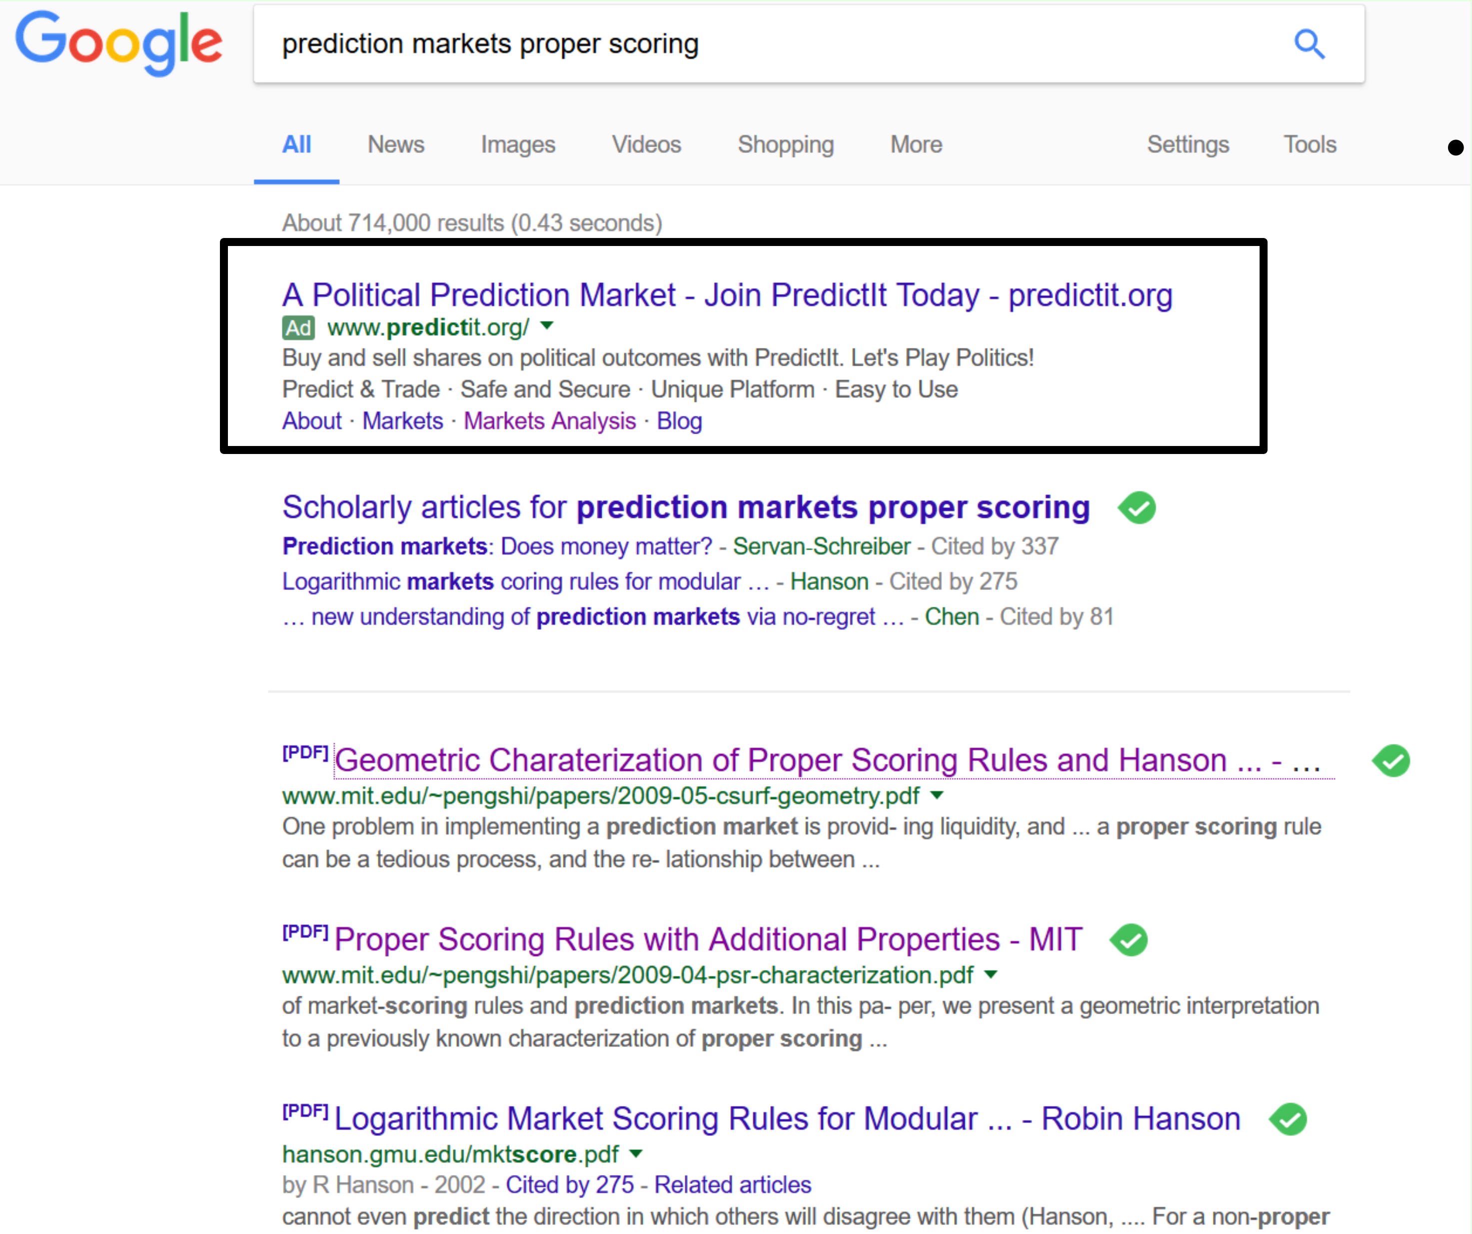

- Choice of ads (if any) to show determined by:
  - Advertiser bid
  - Predicted likelihood of click