CS269: Mathematical Frameworks for Social Computing

Jenn Wortman Vaughan

January 10, 2012

What is social computing?

What is social computing?

What is social computing?

"Social computing has to do with supporting 'computations' that are carried out by groups of people."

What is social computing?

"Social computing has to do with supporting 'computations' that are carried out by groups of people."

Usually focuses on types of computations that are hard for computers

Games with a Purpose

Question and Answer Forums

postal

Grease stains?

Is it at all possible to remove grease stains? and if so how?

6 years ago

Report Abuse

smoofus7...

Best Answer - Chosen by Asker

A friend of mine who's a neat freak told me that dish soap works on grease stains (as long as it's a fairly new stain). Just add some to the stain, rub it in and wash it out.

Dish soap is made to get rid of grease.

6 years ago

Report Abuse

Asker's Rating: ★★★★★

cool idea, I'll have to try it

Online Labor Markets

Make Money
by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. Find HITs now.

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work

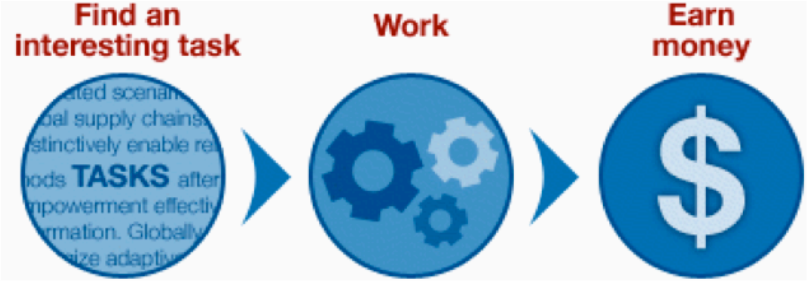

Find HITs Now

Get Results
from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. Register Now

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results

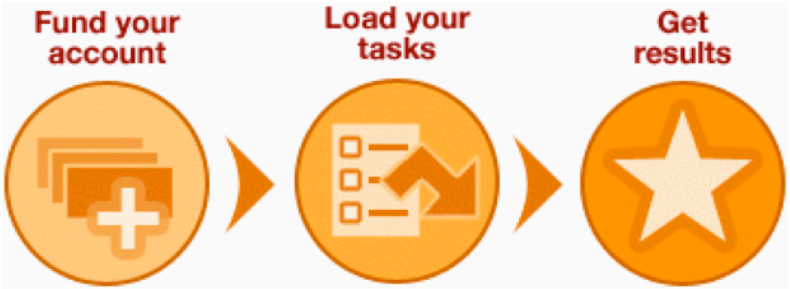

Get Started

What is this class about?

What is this class about?

- We will study the analysis and design of Internet-based systems that allow online communities or crowds to achieve computational goals

What is this class about?

- We will study the analysis and design of Internet-based systems that allow online communities or crowds to achieve computational goals

- We will focus largely on incentives and learning
  - How can we motivate contributions and effort?
  - How can we aggregate information?

What is this class about?

- We will study the analysis and design of Internet-based systems that allow online communities or crowds to achieve computational goals

- We will focus largely on incentives and learning
  - How can we motivate contributions and effort?
  - How can we aggregate information?

- Techniques we study will come from statistics, learning theory, microeconomics / game theory

Structure of the Course

- Introduction and game theory background
- Current research papers
  - Prediction markets
  - Social networks and task routing
  - User generated content and crowdsourcing
  - Reputation mechanisms
  - Games with a purpose
- Project presentations

Course Logistics

Breakdown of Grades

Paper Reviews (35%)

- Submitted the night before each class

Presentations and Discussion Leading (25%)

- Each student will be responsible for leading class discussion twice, possibly in groups of 2-3

Course Project (40%)

- In-class presentation and written report

Paper Reviews

- What are the main contributions of the paper?
- What are the main limitations of the paper?
- What are the most important assumptions made by the authors? Why were these assumptions made?
- Are these assumptions realistic? Are they limiting?

Paper Reviews

- What are the main analytical techniques used?
- Which parts of the paper were unclear to you?
- What are some ways in which this work could be extended?
- Does this relate to other research we have discussed, and if so, how?

Should I really read the *whole* paper?

Should I really read the *whole* paper?

- You are not expected to understand every detail of every proof in the papers that we cover

Should I really read the *whole* paper?

- You are not expected to understand every detail of every proof in the papers that we cover
- However, you should…
  - Be able to answer the paper review questions
  - Thoroughly understand the model used and the statement of the results
  - Attempt to understand the main ideas behind the proofs of major results

Should I really read the *whole* paper?

- You are not expected to understand every detail of every proof in the papers that we cover
- However, you should…
  - Be able to answer the paper review questions
  - Thoroughly understand the model used and the statement of the results
  - Attempt to understand the main ideas behind the proofs of major results
- Plan to spend around three hours per paper

Discussion Leading

- You will be expected to lead the class twice during the quarter, possibly in groups of 2 or 3

Discussion Leading

- You will be expected to lead the class twice during the quarter, possibly in groups of 2 or 3

- Yes, you really do need to read the whole paper when you're the discussion leader…

Discussion Leading

- You will be expected to lead the class twice during the quarter, possibly in groups of 2 or 3
- Yes, you really do need to read the whole paper when you're the discussion leader…
- Prepare a presentation on the paper, but leave lots of time for questions and debate

Discussion Leading

- You will be expected to lead the class twice during the quarter, possibly in groups of 2 or 3
- Yes, you really do need to read the whole paper when you're the discussion leader…
- Prepare a presentation on the paper, but leave lots of time for questions and debate
- Start with the high level motivation, then use the paper review questions to guide you

Course Projects

- Project goals:
  - Develop a deep understanding of a specific research area relevant to the course
  - Make progress towards stating and solving open problems in this area
- The details will be largely up to you, but must relate to the themes of the course

Course Projects

- Project goals:
  - Develop a deep understanding of a specific research area relevant to the course
  - Make progress towards stating and solving open problems in this area
- The details will be largely up to you, but must relate to the themes of the course
- HCOMP 2012 deadline is March 30…

Prerequisites

- Familiarity and comfort with probability theory
- Familiarity and comfort reading and writing mathematical proofs
- Background in algorithms (e.g., big O notation)
- Experience reading mathematical research papers
- Willingness to get involved in class discussions

Prerequisites

- Familiarity and comfort with probability theory
- Familiarity and comfort reading and writing mathematical proofs
- Background in algorithms (e.g., big O notation)
- Experience reading mathematical research papers
- Willingness to get involved in class discussions

A background in learning theory, microeconomics, and game theory is helpful but not required.

Enrollment

- Enrollment will be limited to enable a good environment for active class discussions
- If you would like a PTE and you meet the prerequisites, please
  - Come to class today and Thursday
  - Send me an email after Thursday's class telling me a little about your background and research interests and why you want to take this class

Academic Honesty Policy

Collaboration and discussion outside of class is strongly encouraged, but…

- Each student must write down the paper reviews independently in his or her own words.
- All sources (internet included) must be properly credited in presentations and reports.
- The university academic integrity policy must be followed at all times.

Class Website

The class schedule, announcements, and links to required reading will be posted at:

`http://www.cs.ucla.edu/~jenn/courses/W12.html`

Check this website often!!

Overview of Topics

Prediction Markets

Mitt Romney to be Republican Presidential Nominee in 2012

Last prediction was: **$8.12** / share

Today's Change: ▼ **-$0.17** (-2.1%)

81.2%
CHANCE

Prediction Markets

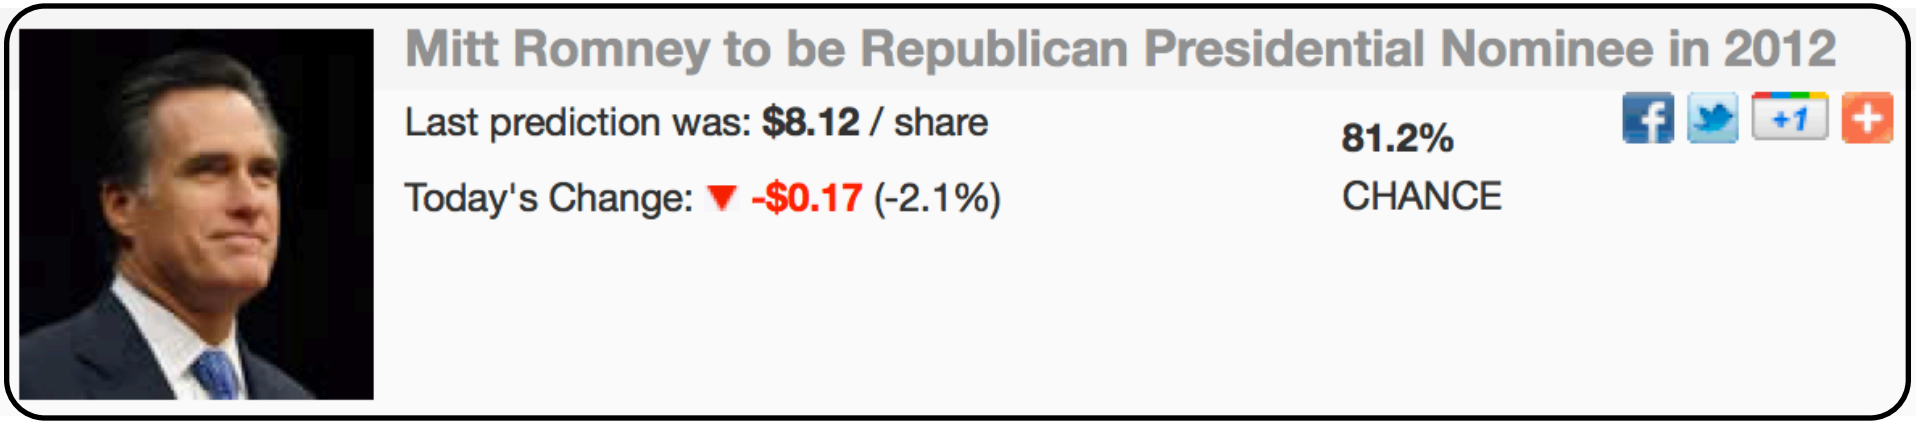

Potential payoff is \$10. If I think that the probability of Romney winning is p, I should

- Buy this security at any price less than \$10p
- Sell this security at any price greater than \$10p

Prediction Markets

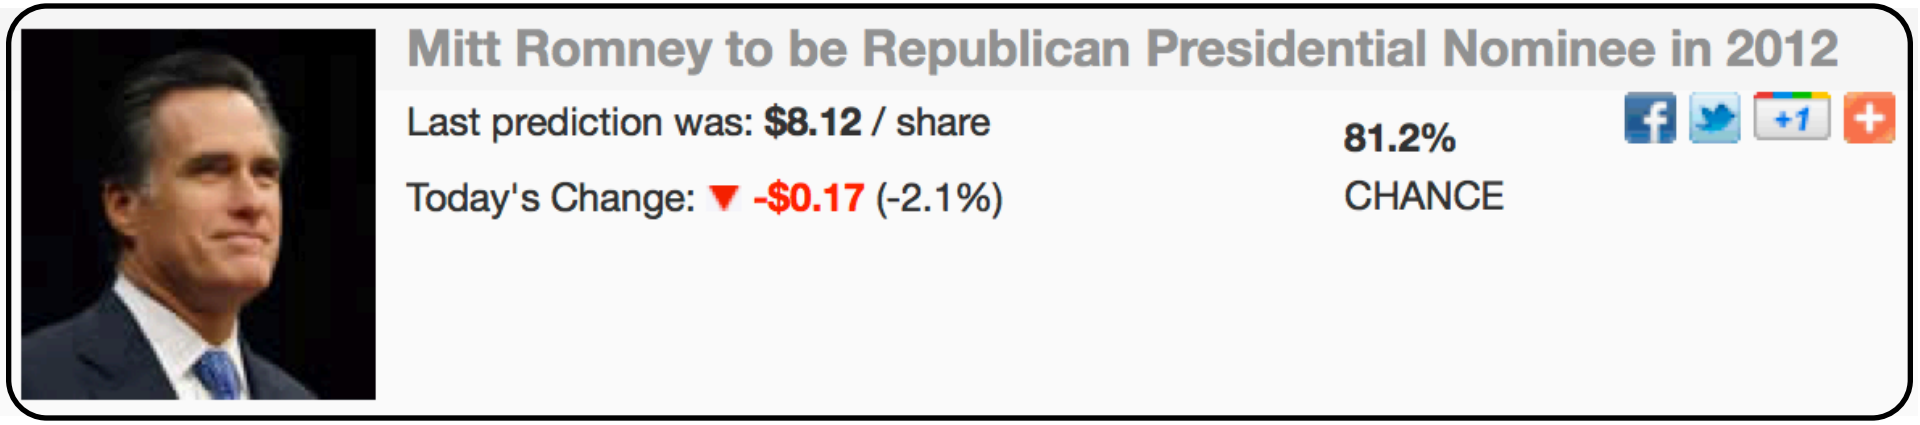

Potential payoff is \$10. If I think that the probability of Romney winning is p, I should

- Buy this security at any price less than \$10p
- Sell this security at any price greater than \$10p

Price measures the population's collective beliefs

Prediction Markets

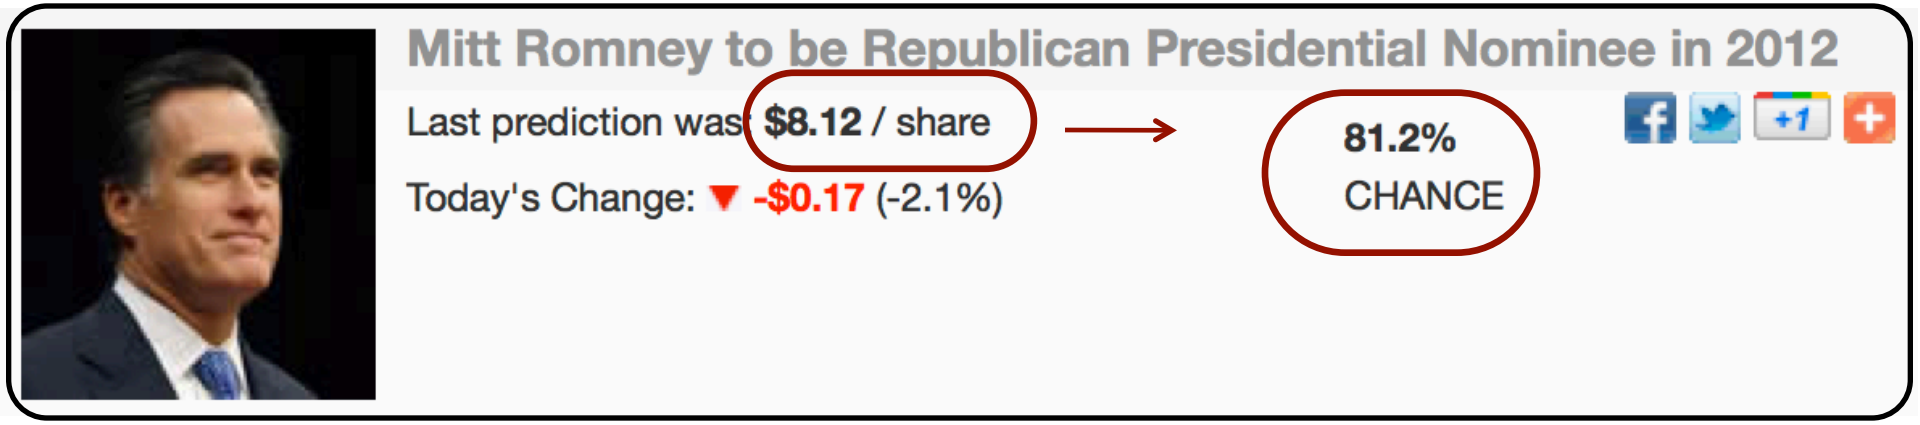

Potential payoff is \$10. If I think that the probability of Romney winning is p, I should

- Buy this security at any price less than \$10p
- Sell this security at any price greater than \$10p

Price measures the population's collective beliefs

How good are these predictions?

- In theory…
  - Under certain assumptions, prices converge and reflect the traders' collective knowledge

How good are these predictions?

- In theory…
  - Under certain assumptions, prices converge and reflect the traders' collective knowledge
- In practice…
  - Election markets beat data from polls (BR02)
  - Oscar markets beat expert columnists (PGN01)
  - HP internal markets beat sales forecasts (P00)
  - Racetrack odds beat experts (F79)
  - and many more

How do they work?

- Securities are offered for every state of the world

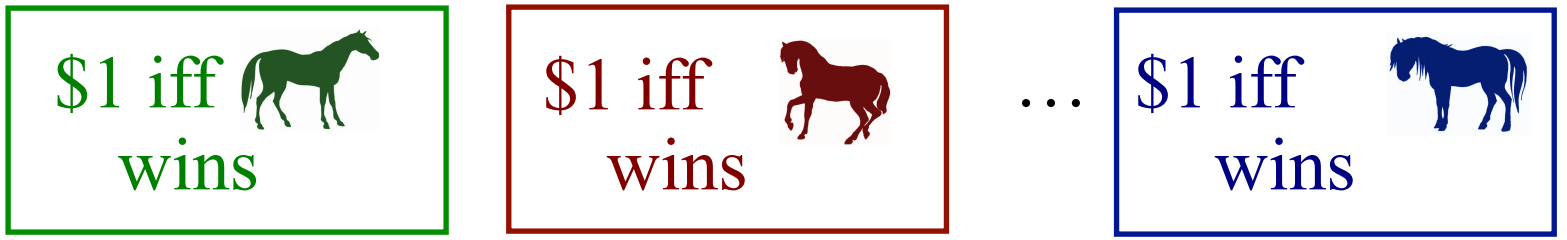

How do they work?

- Securities are offered for every state of the world

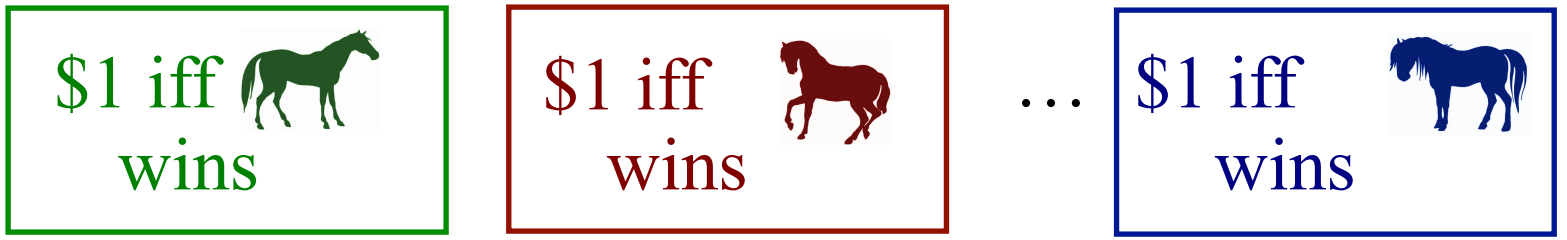

- Mechanism determines prices
  - Traditional stock market style pricing ("continuous double auction") leads to low liquidity, huge spreads

How do they work?

- Securities are offered for every state of the world

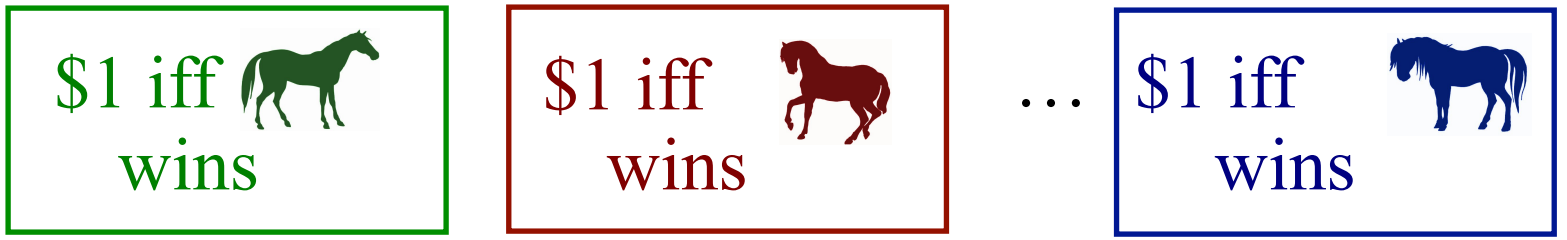

- Mechanism determines prices
  - Traditional stock market style pricing ("continuous double auction") leads to low liquidity, huge spreads
  - What alternatives are there?

Prediction Market Questions

- How can we design a market mechanism to automatically price contracts in a “nice” way?
- How can we design a mechanism that scales well when the number of possible outcomes is large?
- We’ll also examine how markets can be used to “crowdsource” learning problems…

Task Routing in Social Networks

Task Routing Questions

- How can we incentivize individuals to route tasks to appropriate members of their social networks?

- How can individuals in a network find the best/ shortest path?

User Generated Content

Make Money
by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. Find HITs now.

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work

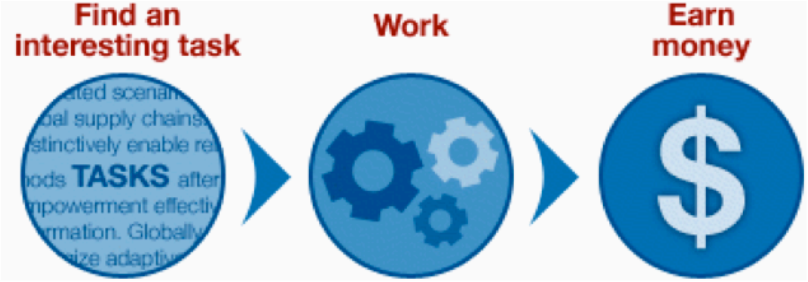

Find HITs Now

Get Results
from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. Register Now

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results

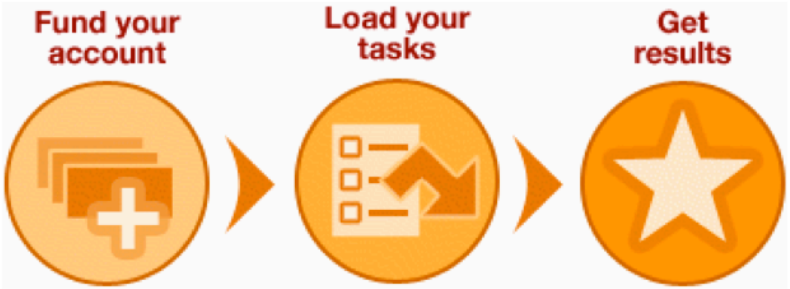

Get Started

UCG Questions

- How can we incentivize users to produce the highest quality content?
- How can content/answers from multiple unreliable users be aggregated most accurately?
- How should tasks in online labor markets be priced to maximize revenue or social welfare?

UCG Questions

- How can we incentivize users to produce the highest quality content?

- How can content/answers from multiple unreliable users be aggregated most accurately?

- How should tasks in online labor markets be priced to maximize revenue or social welfare?

total benefit or "utility" to all users

Reputation Mechanisms

Reputation Mechanisms

- Can be used for sanctioning or signaling

Questions about Reputation

- How can we design reputation systems that lead to the highest revenue or social welfare?
- How can “whitewashing” be prevented?
- How can manipulation be discouraged?
- Can social network structure be incorporated into reputation mechanisms?

Other Social Computing Systems

- How can we incentivize high quality contributions to question and answer forums?

- How can we incentivize high quality contributions in games with a purpose?

A Brief Introduction to Game Theory (Part 1)