Collective Decision Making Systems: An Implementation Guide

Jennifer H. Watkins
Los Alamos National Laboratory

July 1, 2008 * Sandia National Laboratory * Albuquerque, New Mexico

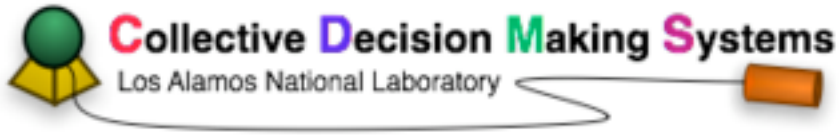

Outline

Survey of CDMS

A framework to map questions to systems

Example: prediction markets

Outline

Survey of CDMS

A framework to map questions to systems

Example: prediction markets

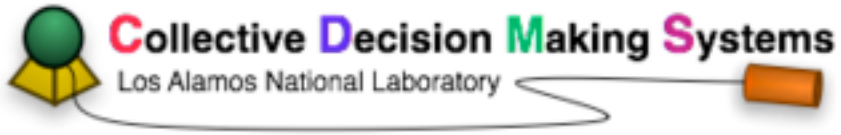

How do many minds produce knowledge?

Deliberation: a group of people meet to discuss and decide together

Vote: individuals choose an option and the most popular option is the group decision

Collaborate: a group of people meet to iteratively develop a product

Sunstein, C. R. (2006). *Infotopia: How many minds produce knowledge.* New York: OUP.

How many minds produce knowledge: Collective Decision Making System

CDMS: a system designed to enable collective intelligence

Requires a collective + an aggregator

Watkins, J.H., Rodriguez, M.A. (2008). A survey of web-based collective decision making systems. In R. Nayak and L.C. Jain (Eds.), *LNCS: Evolution of web in AI environment (pp. 245-279).* Berlin: Springer-Verlag.

Collective Intelligence

intelligence that emerges from the collaboration and competition of many individuals

- wikipedia

- As smart as the smartest person in the group
  What is the capital of Pennsylvania?
  - expert identification
  - for factual questions
  - web-based example: Innocentive

- As smart as the sum of the intelligence of the group
  Write a biography of Abraham Lincoln
  - web-based example: Wikipedia

- Smarter than the sum of the intelligence of the group
  Who will be president of the USA in 2009?
  - web-based example: Iowa Electronic Markets

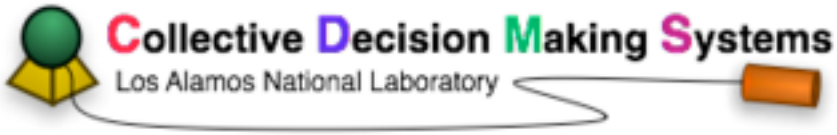

Harvesting Diversity

What is the capital of Pennsylvania?

The larger the collective, the more likely someone in the collective will know the answer

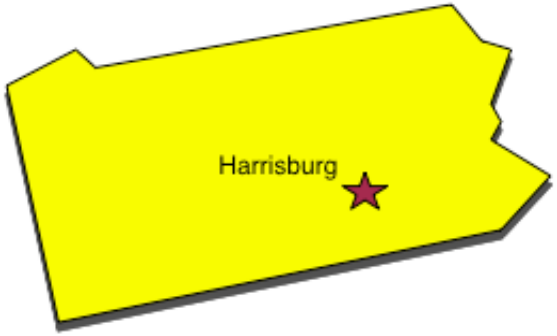

Page, S.E. (2007). *Diversity: How the power of diversity creates better groups, teams, schools, and societies. Princeton: PUP.*

Harvesting Diversity

Write a biography of Abraham Lincoln

The larger the collective, the more knowledge of Abe contained in the group and the more eyes to catch errors

Who will be president of the USA in 2009?

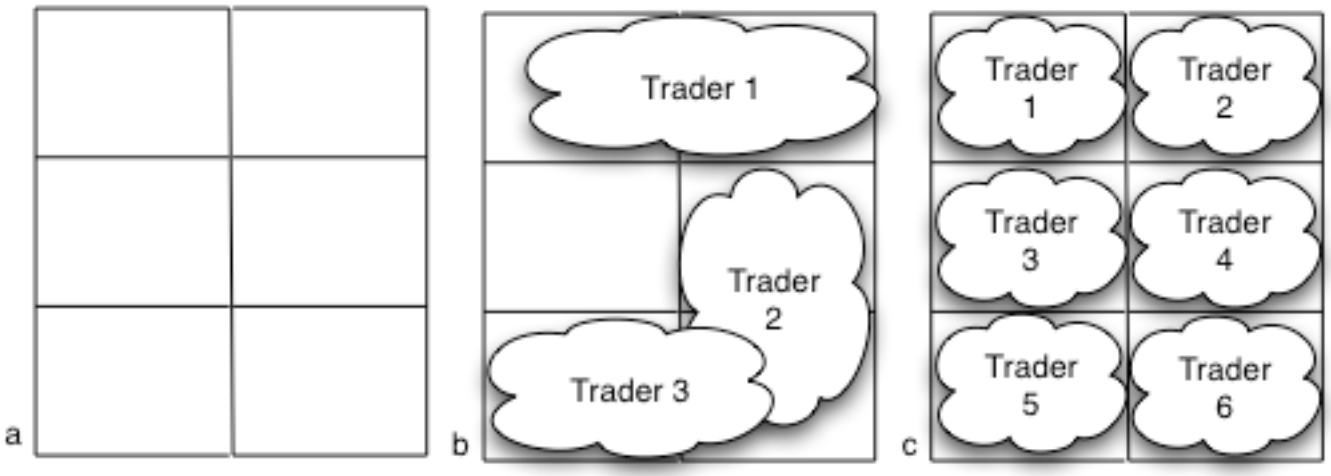

* The proper *combination* of information is key *

Collective Decision Making Systems
Los Alamos National Laboratory

Aggregators

exploit the power of a collective for a specific end

Information retrieval . governance . content creation . prediction

The web is a great way to gather a collective and to automate the aggregation process



Document Ranking

Recommender System

Wiki

Folksonomy

Vote System

Open Source

Prediction Market

Collective Decision Making Systems

Los Alamos National Laboratory

System	Aggregator	Example
document ranking	PageRank	Google search engine
folksonomy	collaborative tagging	Del.icio.us and flickr
recommender systems	collaborative filtering	Amazon.com and Netflix
vote system	plurality	Smartocracy and Cambrian house
wiki	collaborative editing	Wikipedia
open source software	collaborative development	Linux
prediction market	market scoring rule	Iowa Electronic Markets

Outline

Survey of CDMS

A framework to map questions to systems

Example: prediction markets

What technology should we invest in?

What is the most imminent threat?

What are our core capabilities?

What will the average global temperature be in 2018?

Which project is most likely to result in a breakthrough?

Is a CDMS needed?

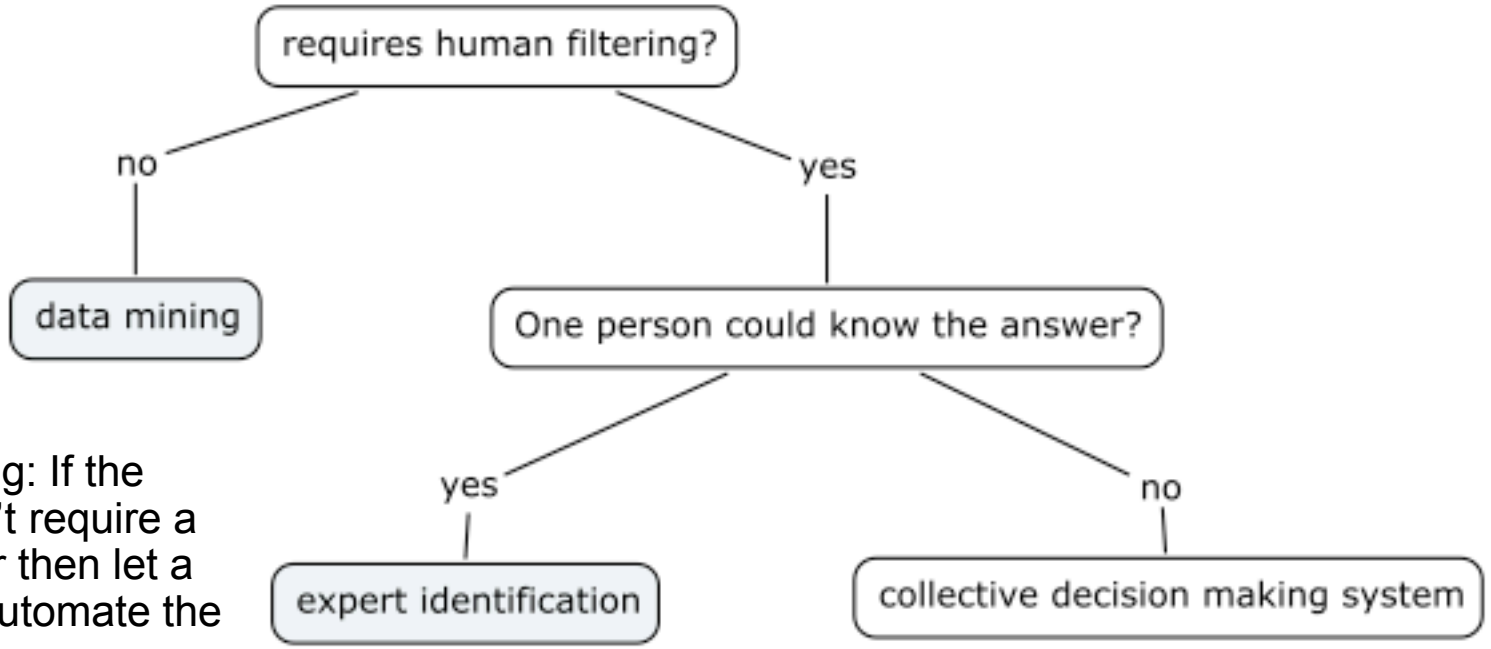

•Data mining: If the data doesn't require a human filter then let a computer automate the process

- Expert identification: There is no point in calling in a collective if only a single expert is needed. In organizations, job titles, phone books, and knowledge management tools can serve as expert identifiers

What type of CDMS is needed?

- Information retrieval: The goal is to organize a collection of resources so as to be able to retrieve specific resources from the collection later.
- Content creation: The goal is to create a product that is the result of group collaboration.

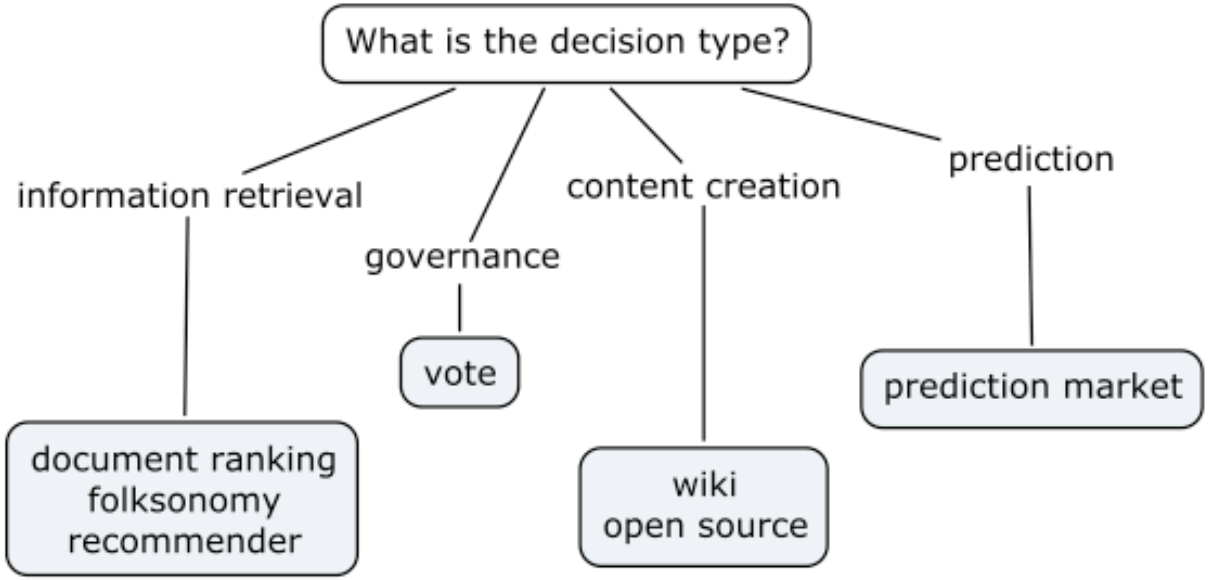

Governance: The goal is to produce a decision that fairly incorporates the values and opinions of the group.
Prediction: The goal is to predict an outcome by motivating information discovery and truthful revelation.

Problem Space

What will the average global temperature be in 2018?

Decision Type: prediction

Goal: predictive accuracy for planning and mitigation

What technology should we invest in?

Decision Type: governance

Goal: widespread satisfaction with the chosen option

Implementation

What will the average global temperature be in 2018?

Solution Space: “stocks”: less than 30°, 31°-90°, 91°-150°, …

Interface Complexity: market trading is not a widespread skill

What technology should we invest in?

Solution Space: ballot: solar power, coal, ethanol, …

Interface Complexity: voting is intended to be accessible to the entire population

Individual Features

What will the average global temperature be in 2018?

Motivation to participate: competition

Expertise in topic: necessary, garbage-in … garbage-out

Membership in group: self-selecting

What technology should we invest in?

Motivation to participate: cooperation

Expertise in topic: unnecessary

Membership in group: self-selecting

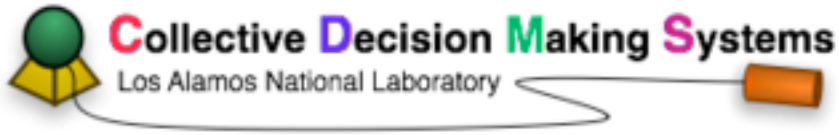

Collective Features

What will the average global temperature be in 2018?

Size of group: variable, liquidity is what matters

Diversity: coverage of relevant factors + improvement

Interaction: strategic, but truthful revelation incentivized

What technology should we invest in?

Size of group: variable

Diversity: none needed

Interaction: strategic, may result in insincere voting

Gibbard, A. (1973). Manipulation of voting schemes. *Econometrica, 41(4).*

Taxonomy of CDMS

	Document Ranking	Folksonomy	Recommender	Vote	Wiki	Open Source	Prediction Market
Problem Space							
Decision Type	information retrieval	information retrieval	information retrieval	governance	content creation	content creation	prediction
Decision Principle	centrality	frequency	similarity	frequency	consensus	consensus	trade
Goal	quality retrieval	quality retrieval	quality retrieval	satisfaction	document utility	code utility	predictive accuracy
Accuracy Metric	precision recall	precision recall	precision recall	fairness	usability	usability	forecast standard error
Implementation							
Solution Space	number of artifacts	number of artifacts	number of artifacts	ballot	creative output	creative output	disjoint + exhaustive
Interface Complexity	very restrictive	not restrictive	not restrictive	not restrictive	restrictive	very restrictive	restrictive
Skill Set	web-page design	basic skills	basic skills	basic skills	wikitext syntax	programming	market trading
Contributor/User	both	both	contributors	contributors	both	both	both
Individual Features							
Motivation	connectedness	organization	personalized advice	cooperative	critical	critical	competitive
Expertise	unnecessary	unnecessary	unnecessary	unnecessary	necessary	necessary	necessary
Membership	co-opted	self-selecting	auto/self-selecting	self-selecting	self-selecting	self-selecting	self-selecting
Collective Features							
Size	large	large	large	variable	variable	variable	variable
Diversity	coverage	coverage	coverage	none	improvement	improvement	coverage + improvement
Interaction	none	imitative	none	strategic	stigmergic	stigmergic	strategic

Watkins, J.H., Rodriguez, M.A. (2008). A survey of web-based collective decision making systems. In R. Nayak and L.C. Jain (Eds.), *LNCS: Evolution of web in AI environment (pp. 245-279).* Berlin: Springer-Verlag.

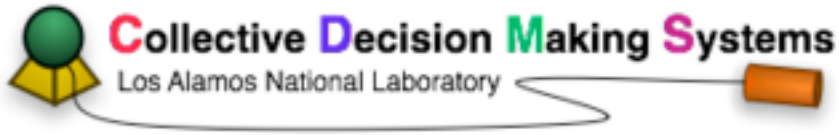

Outline

Survey of CDMS

A framework to map questions to systems

Example: prediction markets

Prediction Markets

Everybody's doing it

HP

Chrysler

GE

Eli Lilly

Google

Nokia

Best Buy

Yahoo!

Microsoft

Arcelor Mittal

Intel

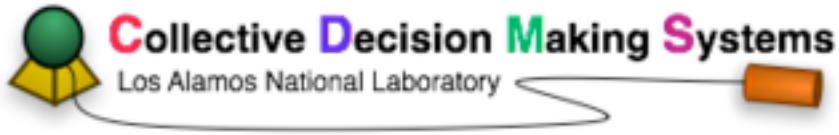

They work

- In 2004, the market odds on Intrade predicted the presidential vote of every state but Alaska. In 2006, the odds correctly indicated the outcome of every Senate race.
- Iowa Electronic Markets (IEM) in the 2004 presidential election correctly predicted the number of electoral votes by which Bush would win
- HP reports that price estimates went from a 4% error using traditional methods to a 2.5% error with BRAIN

What do companies use prediction markets for?

- HP: estimate the price of DRAM
- Google and Yahoo!: fun and research
- Microsoft: determine whether product deadlines will be reached - stop bad outcomes before they happen

Question Format

Good questions should have:

- a determinable outcome
- options that are disjoint and exhaustive
- information that is revealed through time
- a specified closing time and arbiter

- A prediction market where people only buy or sell once is a weighted vote.

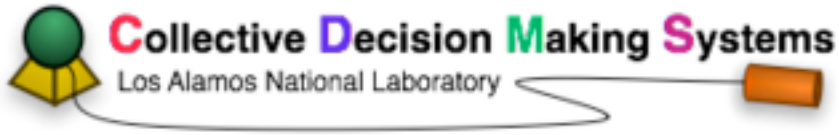

Incentive Structure

Stocks are valued between 0 and 100; therefore, prices are easily interpreted as a probability.

To earn money:

- Buy low and sell high (just like NYSE)
  - Earn the difference in price
- Hold a winning position when the market closes
  - The value of the winning position goes to 100

How do they work?

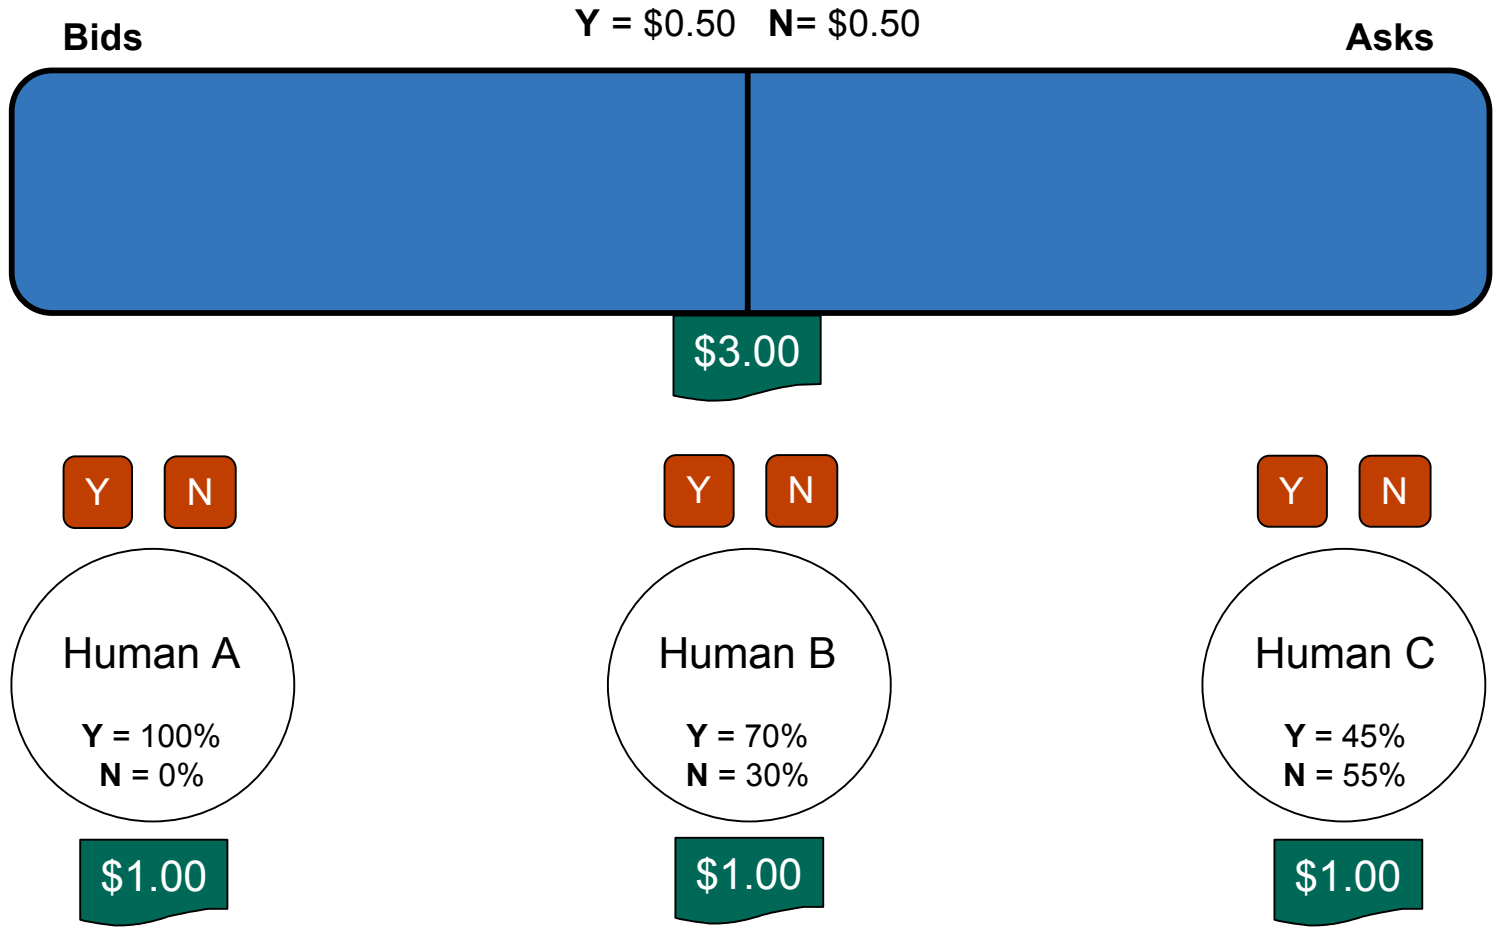

Collective Decision Making Systems
Los Alamos National Laboratory

Comparison

Prediction markets are often compared to polls

Prediction Markets	Polls
“What will happen?”	“What do you want to happen?”
Self-selecting population	Representative sample
Dynamic information	Static information
Automated weighting	One person, one vote
Incentivizes information discovery and truthful revelation	-

Accuracy

- The IEM determine accuracy primarily by comparing their results to polls
  - Be correct sooner
  - Be correct by a closer margin (measured in forecast standard error)
- Most accurately, a probability (say 80%) means that if the event were to occur 100 times, 80 of those events would result in the favored outcome, but 20% would not.

Berg, J., Nelson, F., & Rietz, T. A. (2003). Accuracy and Forecast Standard Error of Prediction Markets. University of Iowa Tech Report.

PM Aggregators

- Continuous double-auction
  - This is the standard bid-ask format familiar from traditional markets; used by IEM
- Market scoring rules (logarithmic)
  - By Robin Hanson, this market maker format encourages liquidity; used by Inkling
- Dynamic pari-mutuel
  - By David Pennock, a combination of pari-mutuel and CDA; used by Tech Buzz Game

Hanson, R. (2007). Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets, 1(1), p. 3-15.*

Pennock, D. (2004). A dynamic pari-mutuel market for hedging, wagering, and information aggregation. *ACM Conference on Electronic Commerce*. New York.

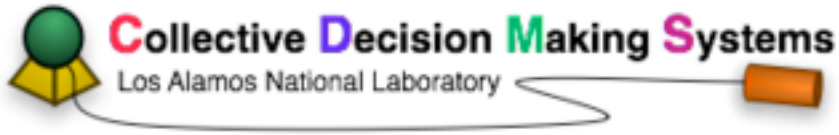

Real Money vs. Other Incentives

Play money markets perform as well as real money markets

Real money: better motivate information discovery

Play money: more efficient information aggregation, players only have wealth due to past prediction success

Other incentives: leader board, prizes

Servan-Schrieber E., Wolfers J., Pennock D., & Galebach B. (2004). Prediction markets: Does money matter? *Electronic Markets,* 14(3).

PM Providers

- Open source
  - Zocalo by Chris Hibbert in Java
  - IdeaFutures used by Foresight Exchange in Perl
- Commercial
  - Inkling
  - NewsFutures
  - ConsensusPoint

Key "Players"

- Chris Masse - Midas Oracle blog and .com
- Chris Hibbert - Zocalo writer and blogger
- Robin Hanson - mastermind of DARPA project and LMSR
- David Pennock - developer of DPM
- Justin Wolfers & Eric Zitzewitz - economists in love with PM
- Bernardo Huberman & Leslie Fine - HP BRAIN researchers

Questions?

jhw@lanl.gov

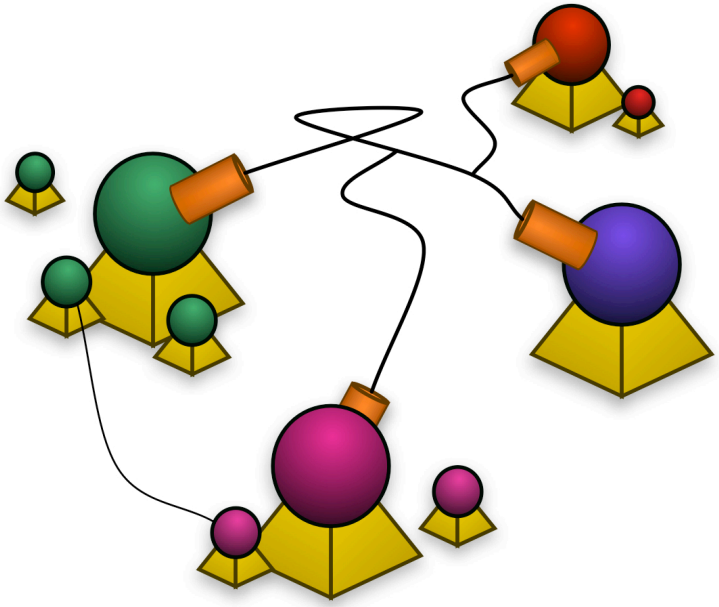

http://cdms.lanl.gov

marko@lanl.gov

Collective Decision Making Systems
Los Alamos National Laboratory