# Frontiers of Behavioral Auction Theory

E. Glen Weyl

Department of Economics
Princeton University

Guest lecture
COS444 Electronic Auctions
Professor Ken Steiglitz
April 24, 2008

# Office hours

I will hold office hours related to this guest lecture from 4pm-6pm in the Bendheim Center for Finance today. I also will be available by appointment (my email is eweyl@princeton.edu).

# Introduction

- What this lecture is not
  - A presentation of classic auction theory
  - An application of theory to real world problems
  - An empirical analysis
- What this lecture is:
  1. A casual overview of many different topics
  2. A bit of psychology, a bit of economics
  3. A bit of very recent research and work in progress
  4. Random, confused(?) but hopefully not confusing ideas
  5. Extremely biased towards what I find interesting
  6. Purpose: inspire ideas for independent work

# What is behavioral auction theory?

Not much yet! But with the help of folks like you...

- Classical econ/game theory assumes strong rationality
  1. Coherent aims and goals (internal unity)
  2. Selfishness
  3. Time consistency (no temptation)
  4. Goal-oriented (expected-utility maximizing)
  5. "Objectively" rational information processing
  6. Equilibrium (common knowledge of this rationality)
- But people aren't like this! Ergo "behavioral economics"
  - 1970's: psychologists' experiments falsify assumptions
  - 1990's: economists alter models for psychological realism
  - Yet not much in auction theory!
  - Such high game theory hasn't yet been challenged

# Agenda

In hopes to right this...

- Three ways of relaxing strong rationality (useful elsewhere)
  1. Prospect theory and risk preferences
  2. Information biases
  3. Disequilibrium
- Given you a brief introduction to each
- Discuss how these *might* be and (a few cases) have been used to enrich auction theory

# What's wrong with expected utility?

- Most auction theory uses expected utility
- But two paradoxes show that people don't act like this

# Rabin's Paradox (Rabin 2000)

- Here's a proposition:
  - Flip a coin: if heads, I take $100 from you; if tails I give you $110
  - Who will take this?
  - Would you continue to feel this way if you were rich?
- Another proposition:
  - I flip a coin: if heads, I take $1000 from you; if tails I give you $1,000,000,000,000,000,000
  - Would you take this?
- Then you aren't an EU maximizer!

**Lesson**

*People care about change in wealth not just final wealth.*

# The Allais Paradox (Allais 1953)

- You have two choices
  1. $1 million for sure
  2. $1 million dollars with 89% probability, $ 5 million with 10% and nothing otherwise
- Here's another set of two choices
  1. 11% chance of $ 1 million, otherwise nothing
  2. 10% chance of $ 5 million, otherwise nothing
- But both of these are basically the same!
  - All that changes is the 89% you can't affect
- This is inconsistent with expected utility

**Lesson**

*Difference between certainty and 99% chance $>$ than between 11% and 10%.*

# The Doctor's Paradox

- You are a doctor. Two dilemmas:
  1. Save 300 people for sure or 50-50 chance of saving 600 people or saving no one.
  2. Death of 300 people for sure or 50-50 chance of no one dying and 600 people dying.
  3. Action for 50-50 chance of saving 400 people or killing 300?

**Lesson**

*Gains treated differently than losses (reference point matters).*

# Kahneman and Tversky's "Prospect Theory"

- Kahneman and Tversky (1979) tried to systematize these insights
  1. Gains and losses, not final wealth
  2. Non-linear probabilities
  3. Kink at the origin...
  4. Concave for gains, convex for losses
- But what is the reference point?
  - Current wealth?
  - Social comparison?
  - Pure framing?

# Köszegi-Rabin (2006) model of reference point

- A day in my life in Washington
  - Who knows what food I want, just something good
  - Hear about good sushi
  - But the place is closed!
  - Go to bad sushi, rather than another good place
- Reference point determined by your (rational) expectations of your own actions
- Losses/gains narrowly framed (sushi v. money v. food)
- Expecting to receive something and expecting to pay a lot for it both make it is worth more

# How is this relevant to auctions?

- Auctions involve much risk and EU affects analysis
- Some counter-intuitive (or falsified) predictions of auction theory
  1. All-pay: bidders with little chance bid positive amount
  2. Dutch auction has same revenue as first-price (shown false by Lucking-Reiley 1999) even with risk-aversion
- Can Prospect Theory help explain why counter-intuitive?

# All-pay

- Why don't we think people with little chance of winning would bid?
  1. Little chance= negligible chance
  2. Do not expect to win, or if do pay little, so worth little
  3. Loss much more likely than gain, weighted more
- All of these are Prospect Theory ideas

# Dutch vs. First-Price

- Why does Dutch make more than First-Price?
- Could just be excitement, but here's another story
  - As price starts to fall, chances rise for the highest valuation people of winning
    $\Longrightarrow$ In Köszegi-Rabin valuation rises
  - Thus they bid higher
- This makes an additional prediction: it is in the middle range that Dutch does better than First-Price
  - Testable with current data (Lucking-Reiley 1999)
- Consistent with fact that in field Dutch better, in experiments First-Price better, as depends on having real object, not fungible money

# Other potential applications/predictions

- Risk aversion makes first-price more attractive
  - Also under PT, but different reason:
    - Risk in price you pay matters (reference point)
- English vs. 2nd price with private values
  - What do dynamics do to expectations?
  - Which direction does it go in?
  - What other predictions can we generate?
- Disclosure, reserve prices, participation costs
  - Effects of all of these depend on risk attitudes
- Multi-unit auctions
  - Expectations of future prices crucial (classical)
  - What does Köszegi-Rabin add?
- Optimal mechanism (technical issues)

# The real winner's curse

Much of this comes from a paper of mine "Biasing Auctions"

- You've talked about the winner's curse
- But classical auction theory assumes people adjust for it
- Famous example
  1. Company has "value" $v$ uniform on $[0, 100]$
  2. Whatever its value, it is worth $\frac{3v}{2}$ to you
  3. You make me an offer $b \in [0, 100]$
  4. I accept if $b > v$
  5. What should you offer?
  6. 0!
- Most people miss this

# Empirical evidence on the winner's curse

- In the lab, people overbid in common values auctions
  1. First shown by Bazerman and Samuelson (1983)
  2. Kagel and Levin (1986)
  3. Book by Kagel and Levin (2002) book surveys
- In the field some evidence as well
  1. Notion of winner's curse first motivated by evidence from Capen, Clapp and Campbell (1971)
  2. Hendricks, Porter and Boudreau (1987) find mixed evidence
  3. Others fail to replicate
  4. Survey by Thaler (1988): mixed
- Why?
- Two broad classes of explanations
  1. Overconfidence: people think they know more than they do
  2. Disregard: people think others know less (or think less about what others know)

# Various theories of disregard

- Most theories of winner's curse are disregard-type
- I don't think about the other person's information
- Three examples
  1. Cursed Equilibrium (Eyster and Rabin 2006): I don't think about the informational content of others' actions
  2. Coarse Thinking (Mullainathan, Schwartzstein and Shleifer 2008): I act like I am not sure others have any information
  3. Pure disregard: I act like others have less information
- Now a bit of motivation

# My favorite examples (Matt Rabin's so clever)

1. In the 2000 campaign, George W. Bush said he had not used cocaine for the last 25 years
   - But when asked whether in the last 26 years he said "I won't talk about the ancient past"
   - Yet after hearing this, most people said less that 50-50 chance that he used cocaine between 25-26 years ago!
2. Suppose that a mutual fund company's advertisement says "We value you, the consumer."
   - What do you infer?
   - They must have performed terribly last year!

**Lesson**

*People do not infer full information in others' actions, particularly when not salient.*

# Cursed equilibrium

- Eyster and Rabin (2006) formalize this idea
  1. Correct belief about (marginal) distribution of actions
  2. Underestimate correlation between actions and information
  3. Mistakenly believe that with probability $\lambda$ others random
  4. Equilibrium simple and when $\lambda = 0$, people rational

# Two other forms of disregard

1. "Coarse thinking" (Mullainathan, Schwartzstein and Shleifer 2008)
   - Sometimes others know things, sometimes not
   - If we fail to see the difference between these settings then when the person knows something we may think they know nothing with some probability
2. Pure disregard
   - We may think others are fools
   - More noise in their signals than there really is

# Overconfidence

1. Disregard is one kind of bias
2. Another is overconfidence
   - Give me a 95% confidence interval for closing price of crude oil yesterday
   - People wrong much more than 1 out of 20 times
   - People are "fooled by randomness
3. Another: the "curse of knowledge"
   - What fraction of Princeton students go to Wall Street?
   - Asked a European: what do you think they would guess?
   - Most say about 10%
   - It's hard to separate how others think from how you think

**Lesson**

*People have trouble thinking beyond their world (to randomness they don't know or to thoughts of those unlike them).*

# Competing explanations (Weyl 2006)

- Both of these can explain the winner's curse:
  1. Disregard: don't pay attention to the other's information
  2. Overconfidence: already know, so don't care what they say
- Questions
  1. Does this mean anything for design?
  2. In a real auction setting, how to distinguish?
  3. What does distinction mean for auction design?

# Effect on disclosure principle

- Classic prescription: if an auction designer can commit to disclosing information, she should
- This depends crucially on "no speculation"
  - Common prior $\Longrightarrow$ people don't bet
- Either of these can undermine this
- With non-common prior, uncertainty induces speculation

$\Longrightarrow$ Auctioneer takes advantage

- Open question: how to take most advantage?
  - Depends on when people willing to bet and how

# How to tell them apart

- Despite similarities, very different implications for *who* suffers winner's curse
- Under disregard, who gets burned?
  - People neglect content of others' actions...
  - When do others' actions have most content?
  - When others' bids are surprisingly low!
  - So bidders in mid range suffer winner's curse
  - High enough bidders actually bid lower

  $\Longrightarrow$ Compresses spread
- Under overconfidence?
  - People think they know what it is worth
  - So no regression to the mean
  - Most exaggerated with *highest* bidders

  $\Longrightarrow$ Increases spread!
- Also effects of more bidders

# For design, why does cause matter?

Very different implications for auction design!

1. Disregard
   - People believe others know little
     $\Longrightarrow$ Act randomly
   - Those with high valuation expect less competition than there actually is
     $\Longrightarrow$ Reinforces revenue ranking
2. Overconfidence
   - People think they know value perfectly
     $\Longrightarrow$ Others more correlated to them (curse of knowledge)
   - High valuation expect more competition
     $\Longrightarrow$ Reverses revenue ranking!

Depends crucially on link between overconfidence vs. disregard and belief in low vs. high correlation

# But this is just the beginning...

- Many types of biases...
- And many types of auctions
- Computer scientists: information theory important!

# What is Nash equilibrium assuming?

- Nash equilibrium involves very strong assumptions
  1. Rationality and expected utility (see part 1)
  2. Common priors (see part 2)
  3. Common knowledge of rationality
     - I am rational
     - I know you are rational
     - I know you know I am rational
     - And so on...
- Does this last matter?

# Muddy faces

Forgive me if you know this story

- N people standing in a room, all have mud on face
- If I know I have mud, leave from embarrassment
  - $\Longrightarrow$ To be nice, no one tells anyone else
    - Door opens once per minute
- Someone walks in, says "there is someone with mud"
- What happens?
  - After *N* rounds everyone leaves
  - Inductive argument
- But everyone know what the person said...
- Why did it matter?
  - I knew... didn't know you knew he knew

**Lesson**

*Public knowledge* $\ll$ *common knowledge.*

# Guess the average

- Everyone, write down a number between 1 and 100
  - Real number, not necessarily integer
- Whoever is closest to 2/3 of average wins $ 5
  - If equal, lottery
- Nash eq: everyone chooses 0
  - But no one does this...anyone who does is “rational” fool!
- People only do so many stages of reasoning:
  1. Level 0 (L0): I choose something random
  2. L1: Others choose random, I choose 44
  3. L2: Others choose 44, I choose 30
  4. And so on...
- Assume some distribution over these types, gives us model

# Crawford and Iriberri (2007)

- This too can help explain winner's curse (like disregard)
- Predicts that some under bid...
- Fits data better (but more degrees of freedom)
- Nice framework (many ways of playing with it)
- Only applied so far to a few contexts
- And only so far allow L1's and L2's

# Learning bid distributions

- Traditional auction theory assumes bidders know and agree on distribution of values, and that the auction designer knows this as well
- There is no way this is true
- Once we leave this world, many interesting questions
- Here are a few recent papers on this (much more to do!)
  - For example, very simple approach is asymmetric information (no one has done this!)

# Machine learning for bidders

- For those who don't know, machine learning is statistics
  - Try to use past data to predict future outcomes/distributions
- You are bidder, want to use this to learn about how to bid
- Schapire, Stone, McAllester, Littman and Csirik (2002)
  - Use data from past auctions to forecast distribution of bids
  - Choose optimal bid given this distribution
- Much more sophisticated things could be done with better statistics, econometrics, etc.
- Also, what if bidders behave as if learning like this?
- This approach neglects strategic considerations...somewhere in between might be interesting

# Machine learning for auctioneers

- What if an auctioneer wants to learn?
- Can he set up rules that incorporate learning?
- Two settings
  1. Online auctions (Blum, Kumar, Rudra and Wu 2004)
     - One item being sold
     - But sequence of bidders arrive (like EBay)
     - Design incentive compatible rules involving learning
  2. Repeated sales (Blum and Hartline 2005)
     - Many people come to you wanting to buy same thing
     - No competition at any time, but learning makes like auction
     - Learn about what price to charge
- With more sophisticated incentives, could do better
- Other realistic settings involving learning?

# Learning about yourself

- Sometimes people don't even know their own value!
- This can give a reason for sniping (Rasmusen 2006)
- Explains data from EBay auctions (Nekipelov 2007)

# The way to a computer scientist's heart...

- In game/auction theory rationality means infinite computation
- But even rational people have limits (bounds)
- Some auctions are *very* complex
- Should auctions try to stay simple?
- How should this be traded off against efficiency?
- How can we make it simple to calculate allocations?
- What are limits on communication?
- A few directions...

# Some reasons for simplicity

- Complex mechanisms often not very robust (Wilson 1987)
- Guide bidders to (a desirable) equilibrium (Milgrom 2007)
- Equity concerns (Pathak and Sönmez 2008)
- Participation
- Costs of finding optimal bid
- Costs of computing allocation
- But how to quantify complexity
- Two approaches:
  1. Communication
  2. Computing outcomes
- But human mind hard to capture
  - Simple input vs. transparent allocation rule
- A bit on the two approaches

# Computational mechanism design

- We want to assign $M$ items among many $N$ bidders
- This gives something like $\begin{pmatrix} 2^M \\ N \end{pmatrix}$ possibilities
  - BAD!!!

**Lesson**

*Even with non-strategic bidders, computational problems arise.*

- Trade-off computability v. efficiency (and incentives)
  - Saving grace: incentive compatibility gets easier as auction gets large (Pathak and Kojima 2007)
- Those who know both CS and econ are in big demand!
  - Best work combines clever knowledge of both
  - Also: automatically designing auctions (Conitzer and Sandholm 2002)

# Communication and auctions

- One way to think about simplicity constraints is communication capacity
  - What if bidders can only transmit limited information?
    - With big auctions, things go to hell (Nisan and Segal 2004)
    - With small auctions, things are fine (Rosenblum, Nisan and Segal 2005)
- Yeah but....
- Still big open question: how to quantify simplicity reasonably?

# Wrapping up...

- Many fronts on which auction theory falls short in understanding human behavior
  1. Risk preferences
  2. Information processing
  3. Rationality and equilibrium
- Also several not mentioned
  1. Social preferences
  2. “Fun” and social influence
  3. What can you think of?
- This is where you come in!
- Many, many senior theses to be written (published!)