

Playing Games by Thinking Ahead

Adrian Vetta


McGill University

(With Vahab Mirrokni and Nithum Thain)

We are **not** interested in prescribing how games *should* be played.

We **are** interested in analysing how games *really are* played.

We **will** analyse how *some* games really are played.



Judea Pearl:

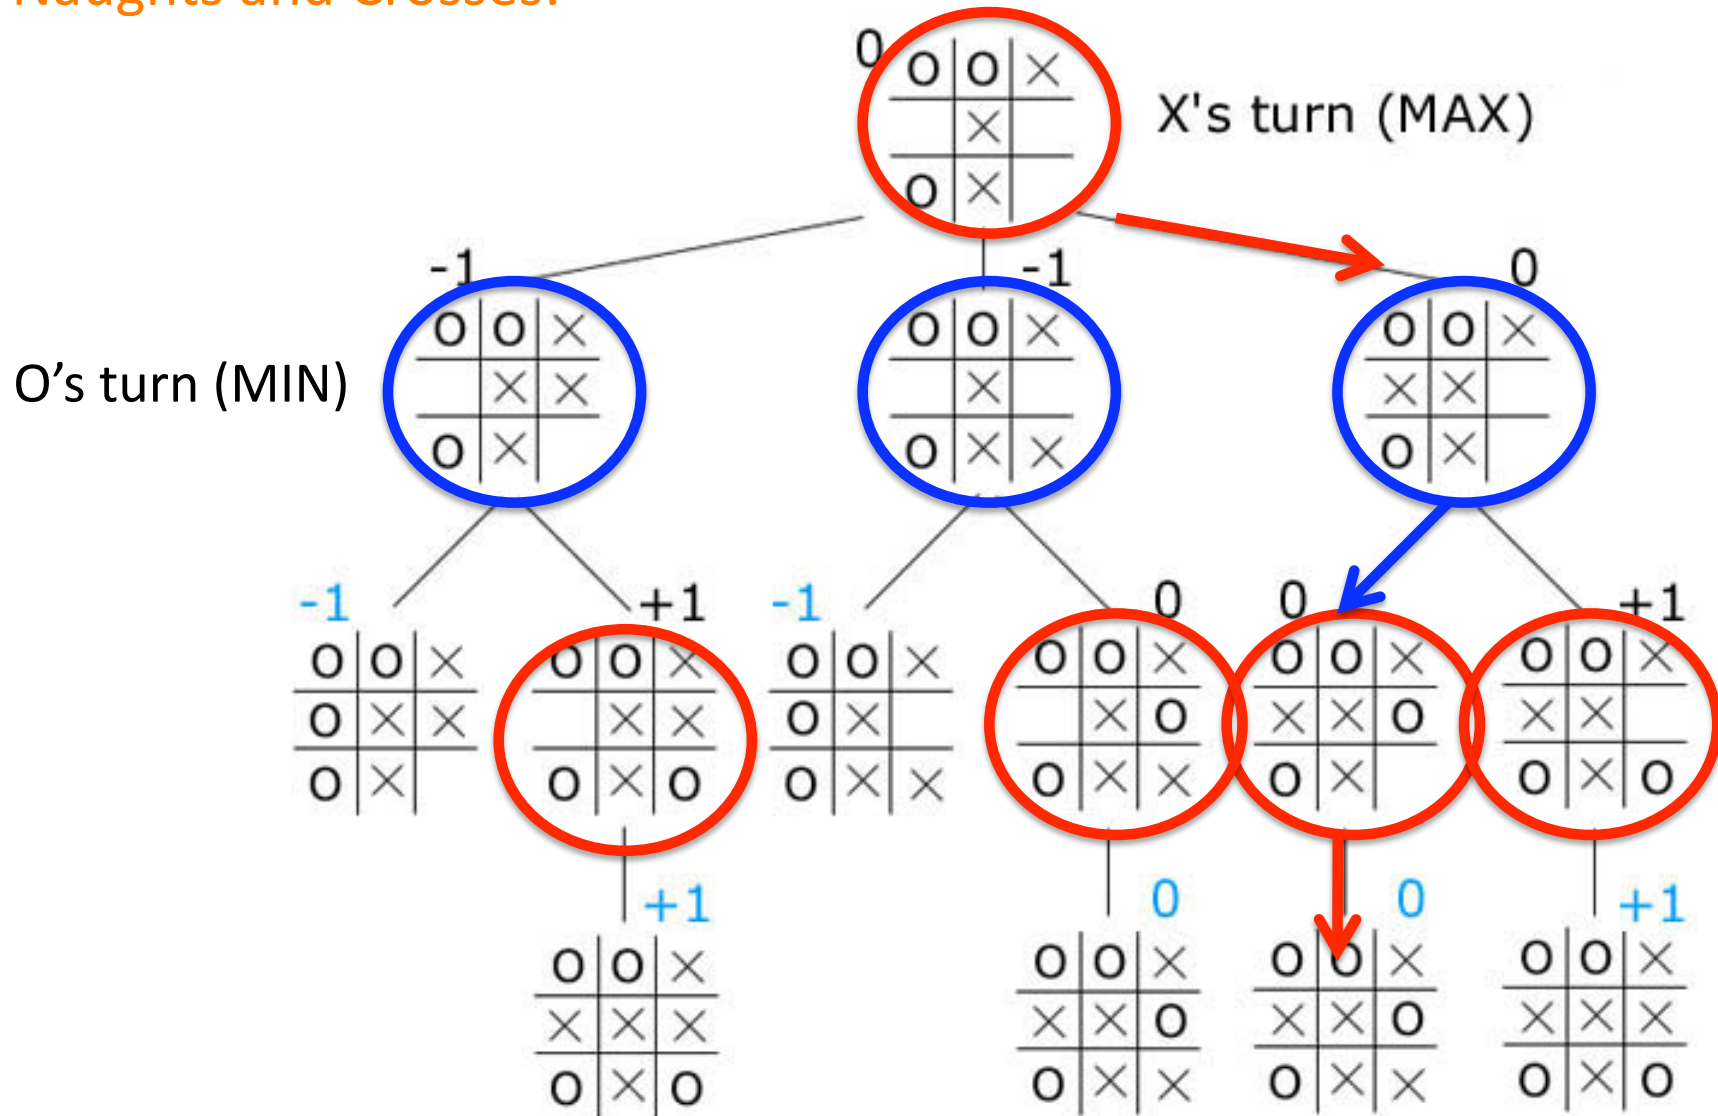
“Almost all game-playing programs use variants of the **lookahead** (minimax) heuristic.”

Overview of Talk


1. The Lookahead Method.
2. A Bit of a Digression.
3. Some Results.

Backwards Induction

Naughts and Crosses:

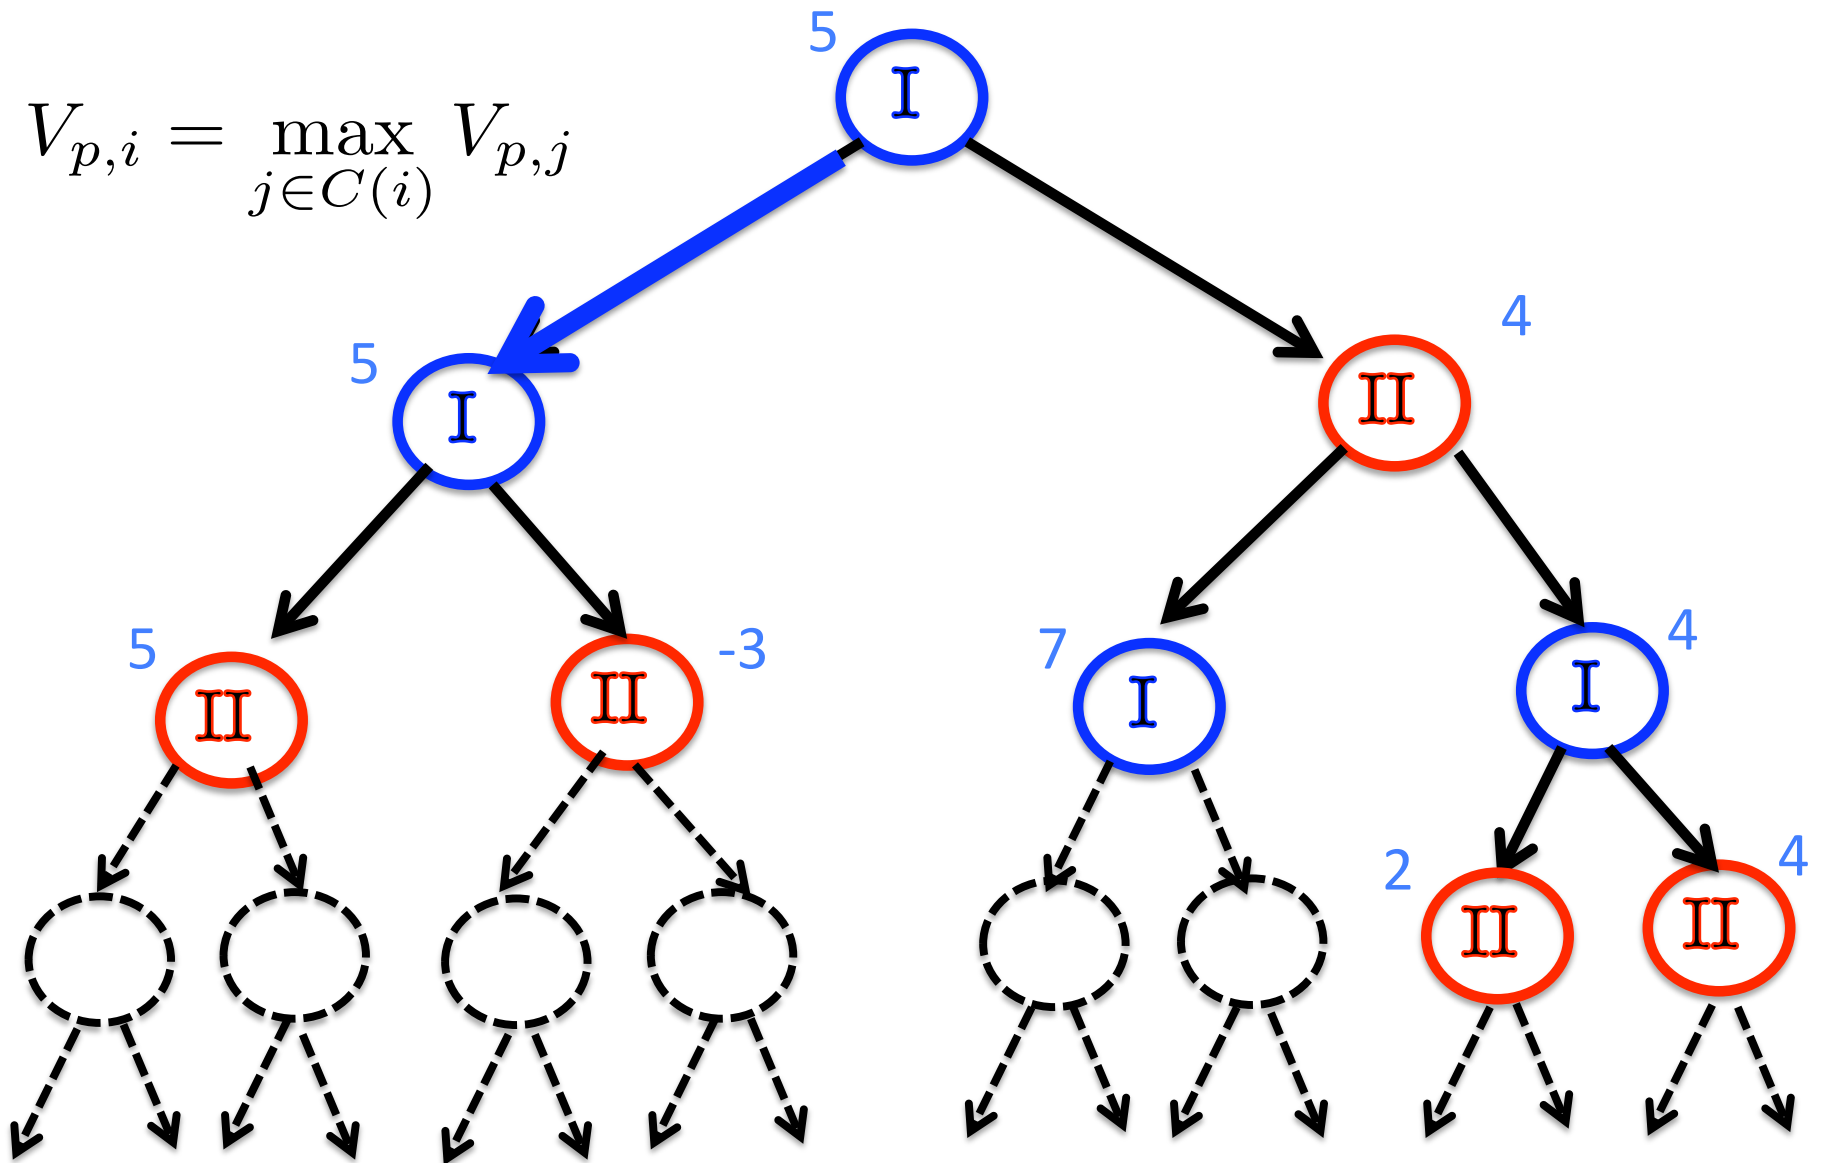


What if you can't think as far ahead as the leaves?



Estimate values for leaves of search tree and work backwards.

$$V_{p,i} = \max_{j \in C(i)} V_{p,j}$$



Special Cases

- *Backwards Induction*
 - Zermelo's Method
- *Best Response Dynamics*
 - 1-Lookahead Search (*Nash Equilibria*)
- *Leader-Follower Behaviours*
 - Asymmetric Computational Power


Adaptability

The actual implementation of the method will vary with the game and with the players:

- **Search Trees:** Vary with experience, computational abilities, etc. They are also dynamic.*
- **Node Evaluations Functions:** Are payoffs accumulated; does only the final outcome matter? (Leaf Model vs Path Model.)
- **Order of Moves:** Fixed, Random, Worst-Case?
- **Utilities or Not?**

* Here we will assume the search trees are BFS trees of depth k .

Unpredictability



Lookahead Search


- The **lookahead method** was formally first proposed by Claude Shannon in 1950.
- Shannon considered it a *practical way* for machines to tackle complex problems that require:

“general principles, something of the nature of judgement, and considerable trial and error, rather than a strict, unalterable computing process”



Chess


Shannon described in detail how the lookahead method could be applied by a computer to play chess.




C. Shannon, "Programming a computer for playing chess", *Philosophical Magazine*, Series 7, **41(314)**, pp256-275, 1950.



Humans & Chess



In a 1946 psychology thesis, Adriaan de Groot studied the thought processes of **human chess players**.



He found that they **all** used the lookahead search heuristic!*

Indeed, De Groot's findings had a large influence on Shannon's subsequent work.

*Experts were better at evaluation positions and deciding how to grow the search tree.


Analysis

- **Objective:** We wish to analyse the consequences when agents use the lookahead method in an assortment of games.
 - *Adword Auctions, Traffic Routing, Bandwidth Sharing, Industrial Organisation, etc.*
- **Quality of Solutions:** To evaluate outcomes, we will examine the **quality of equilibria** when lookahead search is used.
- **Dynamics:** These methods can be extended to measure the expected quality of short-run dynamic solutions.
 - To do this, you need to analyse polynomial-length random walks* on the state graph of the game.

* Random depending upon how the lookahead method is implemented.

Rational Choice Theory

- A rational agent (economic man) makes decisions via utility **optimization**.
- Economic men may not exist but this does not matter provided agents act **as if** they are rational.



Milton Friedman


Example: To save time optimising, I decide to allocate 30% of my budget to housing, 10% to food, 5% to beer, etc.

Conclusion: I am a rational consumer with a Cobb-Douglas utility function.

Bounded Rationality

- Herb Simon, due to considerations of computational power and predictive ability, argued in the 1950s that:

“The task is to replace the global rationality of economic man with a kind of rational behaviour that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist.”



Bounded Rationality: Heuristics

- Simon believed that
 - Agents **do not** optimise in decision-making.
- Instead, he thought that
 - Agents use heuristics in decision-making.


Satisficing

- One heuristic Simon presented was **satisficing**.
 - Agents search for feasible solutions.
 - The search stops when a desired aspiration level is achieved.*
 - The found satisficing solution is chosen.
- Note, for agents of bounded rationality, the form of the search will heavily influence the final decision.
- In contrast, the search is irrelevant for rational agents, as they will make the optimal decision regardless.

* The aspiration level may change over time and depending upon how the search is going.

Human Problem Solving

- Interestingly, the seminal work of Newell and Simon on human cognition was also heavily influenced by De Groot's work.*



* In fact, Herb Simon sent his student George Baylor to help translate De Groot's work into English.

Bounded Rationality & the Lookahead Method

Lookahead Search clearly fits within Simon's framework:

- **Search:** By local search tree.
- **Stopping Rule:** Dependent on experience, computational power, etc.
- **Decision Rule:** By Backwards Induction.

1. Optimisation under Constraints

- One approach is to optimise subject to constraints imposed by time, computation, money etc.
- This can be in the form of an optimisation program or an optimisation via search.
 - e.g. Stop searching when the future costs exceed the future benefits.
- But this approach can be even more complicated than the original optimisation problem!
 - i.e. It doesn't fit with Simon's original ideas.

2. Heuristics and Biases

- The Heuristics & Biases Program examines human irrationality.



Daniel Kahneman



Amos Tversky

- Human use *heuristics* that typically do not satisfy simple laws of logic and probability.
- How and why do such errors occur?
- Can we use these insights to model human behaviour?
e.g. Prospect Theory

Anchoring

- In human decision-making there is a bias to rely (*anchor*) on one specific piece of information.

- Estimates given for 10! vary widely with ordering.

e.g.

$$10 \times 9 \times 8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$$

or

$$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8 \times 9 \times 10$$

- After writing down the first few digits of their Social Security numbers, people with larger numbers bid higher in an auction!

The Law of Small Numbers

- People assume that small random samples will have similar characteristics to the whole population.
- **Gambler's Fallacy:** After a run of losses a win is *more* likely.
- **Pattern Spotting:** Overconfidence in early trends.
- **Clustering:** Clusters are *unlikely* in random data.
- **Medical Trials:** Significant results can be validated using additional *small* trials.

Representatives

- People analyse events according to how **representative** they are of parent populations.

- Steve is very shy and withdrawn, invariably helpful, but with little interest in people.

- Therefore Steve is a librarian not a farmer.

- Bill is intelligent, but unimaginative, compulsive and generally lifeless. In school he was strong in mathematics but weak in social studies and humanities.

- Therefore Bill is likely to be an accountant.
- He is unlikely to play jazz for a hobby.
- He is quite likely to be an accountant **and** play jazz for a hobby.

3. Fast and Frugal Heuristics

- Yes, humans do use decision-making heuristics...
- ...but, don't judge heuristics by their coherence with the laws of logic or probability.
- The purpose of a heuristic is not to be consistent but to perform well at its task.
- So judge a heuristic by its **performance!**

Fast and Frugal School

- Humans often use simple heuristics that are *Fast* (Time) and *Frugal* (Information).
- These heuristics are often very effective.
- Moreover, they are extremely adaptable to new environments, information, or problems.



Gerd Gigerenzer.

Catching a Ball

Which approach is more effective?

- **Optimisation:** Calculate trajectory based upon style of throw, velocity, spin, wind resistance, quality of the ball, etc. Then move to the best spot to catch it.

or

- **Heuristic:** Move towards ball such that your angle of gaze remains constant.



Modern Portfolio Theory

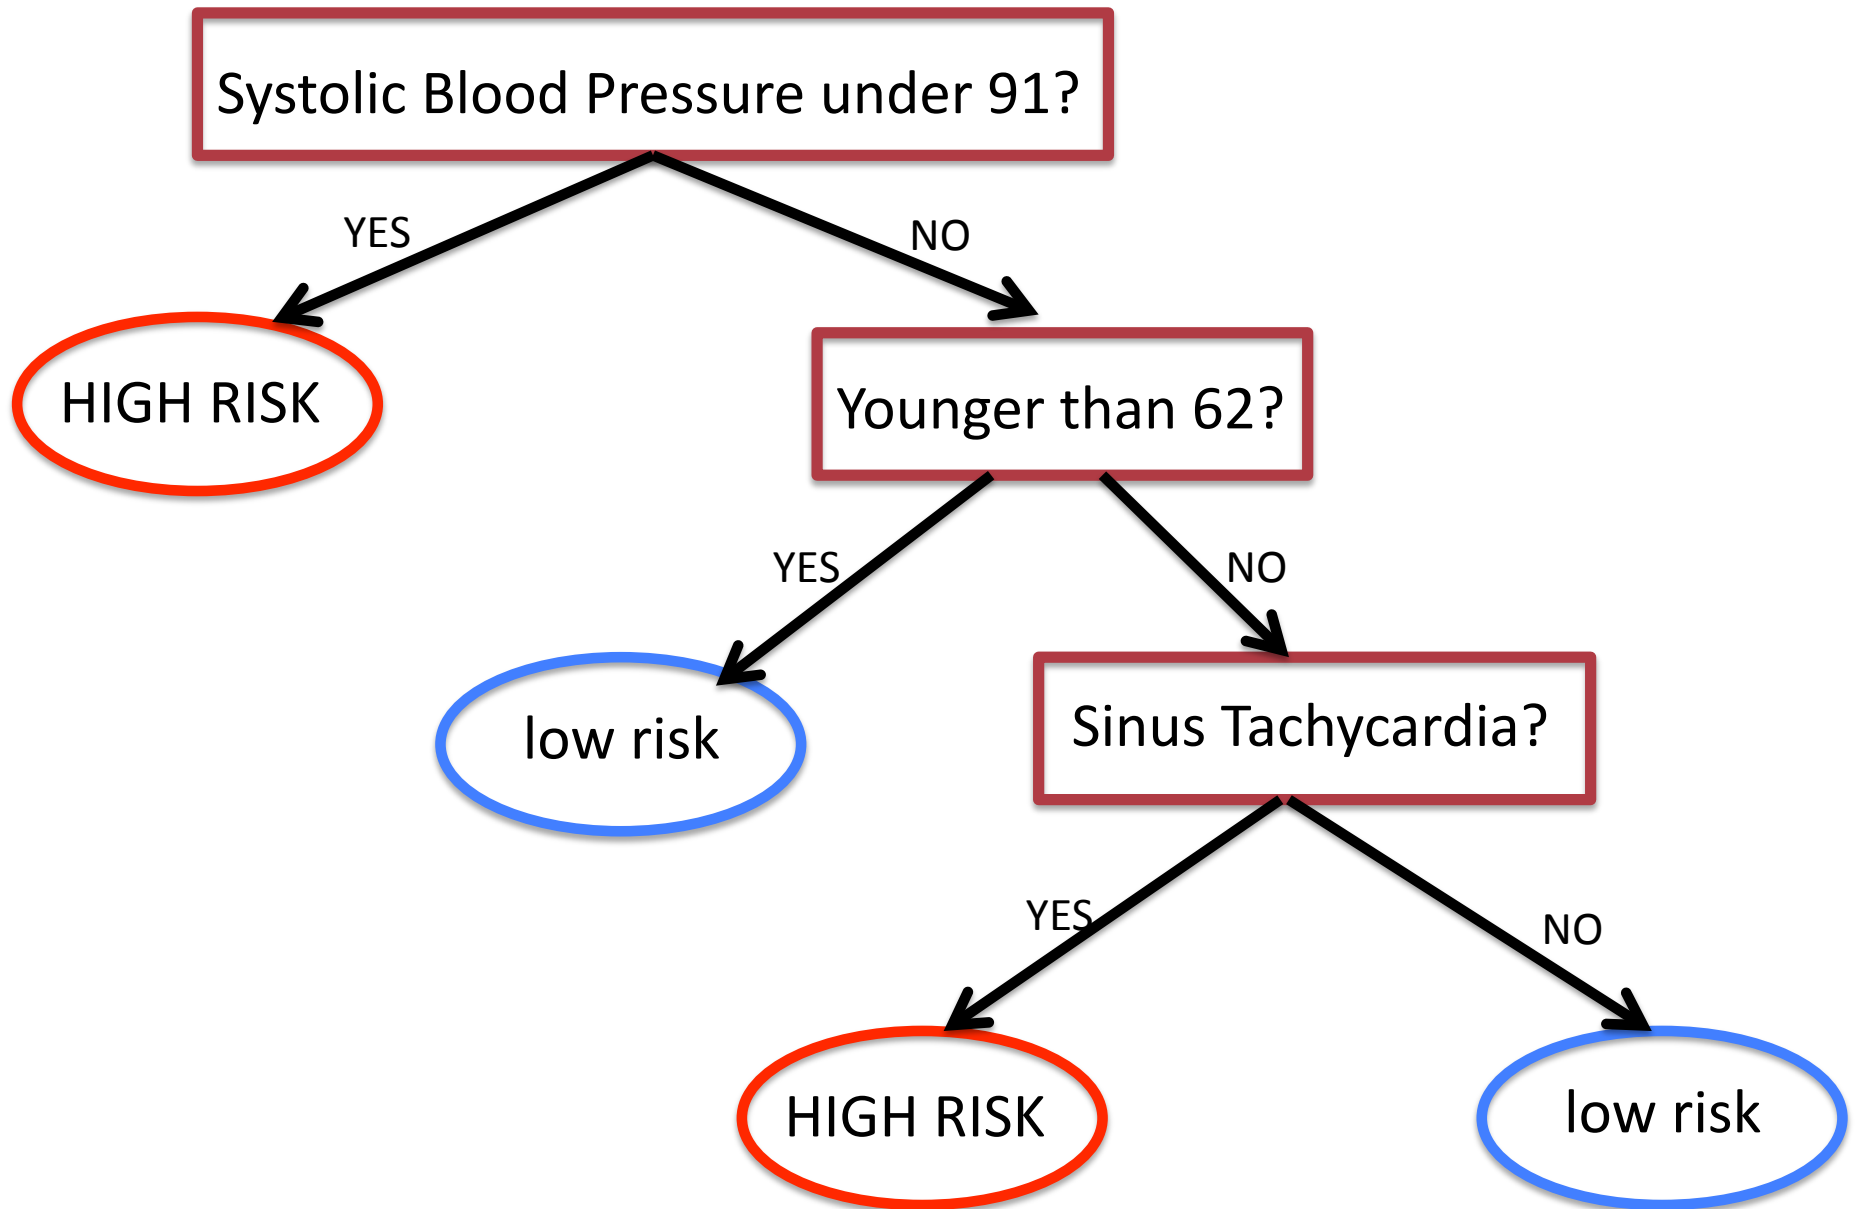
- Harry Markowitz pioneered *Modern Portfolio Theory* in the 1950s.
- He showed how design portfolios to maximise returns and minimise risk.
- How well did this method do for his own retirement plan?
 - He didn't use it!
 - He used the **1/N heuristic**: *split your money equally amongst each of the N assets.*



Take the Best!

- Given a set of cues that may be relevant for your task.
- Rank the cues in terms of importance.
- Choose the option that does best against the top cue.
 - Recurse if ties.
- In tests, this heuristic typically outperforms multiple regression, especially on new data.
 - Multiple Regression overfits to test data.

Heart Attacks



Our Work

- We wish to analyse the consequences when agents use the lookahead method in an assortment of games.

e.g. Adword Auctions, Traffic Routing, Bandwidth Sharing, Industrial Organisation, etc.


- Our focus is on quantitative performance guarantees.
- And the consequences are?

Sometimes good, sometimes bad, sometimes indifferent!

The Cournot Model of Oligopoly

Strategies: The players choose quantities q_1 and q_2 .

Price Function:



Cost Functions: The players have marginal costs c .

Equilibrium: Player i produces $q_i = \frac{1}{3}(a - c)$

Lookahead Equilibrium

- For the path model, what happens when the players use k -lookahead search?
- As k increases, output increases and quickly converges to:

$$q_i = 0.370(a - c)$$

- This 11% increase in output gives a 12% increase in the social surplus.

The Stackelberg Model

Strategies: As in the Cournot model, the players choose quantities q_1 and q_2 .

Commitment: Player 1 is the **leader** and picks a quantity first. Player 2 is the **follower**.

Equilibrium: Player 1 produces $q_1 = \frac{1}{2}(a - 2c_1 + c_2)$
Player 2 produces $q_2 = \frac{1}{4}(a - 3c_2 + 2c_1)$

Stackelberg Behaviour

- For the leaf model, if Player I uses 2-lookahead, but Player II only uses 1-lookahead then the outcome is a Stackelberg equilibrium.
- Thus Leader-Follower behaviours can be induced by asymmetric computational abilities!

Adword Auctions



Web [Images](#) [Video](#) [News](#) [Maps](#) [more »](#)

Current account

Search

[Advanced Search](#)
[Preferences](#)

Web

Results 1 - 10 of about 662,000,000 for **Current account** [\[definition\]](#). (0.16 seconds)

Try your search on [Yahoo](#), [Ask](#), [AllTheWeb](#), [MSN](#), [Lycos](#), [Technorati](#), [Feedster](#), [Wikipedia](#), [Bloglines](#), [Altavista](#)

[Barclays current account](#)

Sponsored Links

[www.barclays.co.uk](#) with up to £330 worth of benefits Special offer now on - apply online

[Top 5 Current Account](#)

[Accounts.Compare-And-Save.co.uk](#) Compare The Top 5 High Interest **Current** Accounts & Apply Online

[current account](#)

[www.firstdirect.com](#) Transfer money and pay bills online £500 overdraft £100 when you switch

Sponsored Links

[Compare Current Accounts](#)

Find the top UK **current** accounts. Compare services and apply online.
[www.find.co.uk/current_accounts](#)

[Nationwide Savings](#)

Earn up to 6.25% gross p.a./AER. Save up to £250 each month.
[www.nationwide.co.uk/savings](#)

[Top 10 Bank Accounts UK](#)

Compare & Apply 300 UK **Bank Account** - Savings - **Current** Accounts - ISAs
[www.which-bank-account-4u.co.uk](#)

[Earn More With Halifax](#)

Switch Today For 60x More With Our High Interest **Current Account!**
[Halifax.co.uk/Bank-Accounts](#)

[Citibank UK](#)

Get a free £500 overdraft and earn

1. [Current account](#) - Wikipedia, the free encyclopedia

This article is about the macroeconomic **current account**. ... Action to reduce a substantial **current account** deficit usually involves increasing exports or ...

[en.wikipedia.org/wiki/Current_account](#) - 23k - [Cached](#) - [Similar pages](#) - [Note this](#) - [Filter](#)

2. [Current Accounts | compare UK bank accounts - moneysupermarket UK](#)

Compare **Current** Accounts: Compare your **Current Account** overdraft with other ... Our **Current Account** guide will help you compare bank accounts and you choose ...

[www.moneysupermarket.com/currentaccounts/](#) - 41k -

[Cached](#) - [Similar pages](#) - [Note this](#) - [Filter](#)

3. [current account](#) Definition

current account - definition of **current account** - The net flow of **current** transactions, including goods, services, and interest payments, between countries.

[www.investopedia.com/5403/current-account.html](#) - 20k -

Generalized Second-Price Auctions

- Sponsored Slots are ranked by their **click-through rates**:

$$c_1 > c_2 > c_3 > \dots > c_T$$

- Each agent has a valuation v^i but bids b^i .
- The t^{th} highest bidder wins slot t , but only pays the $t+1^{\text{st}}$ highest bid.

Generalized Second-Price Auctions

- Despite the name (and the advertising), these are **not** truthful auctions.
- Moreover, there are Nash equilibria whose social values are arbitrarily bad compared to optimal allocations.
- We will analyse 2-lookahead equilibria in the leaf model.

Safely-Aggressive Bidding

- Suppose agent i bid suffices for slot t .
- A bid is **safely-aggressive (balanced)** if it is as high as possible s.t. no agent in a higher slot can hurt i by undercutting her.
- Balanced bidding is apparently a commonly used strategy:
 - Bidding high increases chances of a better slot.
 - Pushes up prices for competitors.
 - But bidding too high is risky, and this alleviates a lot of risk.

Balanced Bidding

- A safely-aggressive (**balanced**) bid satisfies:

$$b^i = \left(1 - \frac{c_t}{c_{t-1}}\right)v^i + \frac{c_t}{c_{t-1}}b_{t+1}$$

- Note that a losing bidder bids $b^i = v^i$, as does the highest bidder.
- Thus winning bids are all higher than losing valuations.

Output Truthful Allocations

- An allocation is **output truthful** if the agent with the i^{th} highest valuation wins the i^{th} slot.

Lemma 1. An allocation is socially optimal
if and only if
it is output truthful.

Proof. Assume agent i is in slot i . If there are agents with $v^t < v^{t+1}$ then switching their slots increases social welfare as

$$c_t \cdot v^{t+1} + c_{t+1} \cdot v^t > c_t \cdot v^t + c_{t+1} \cdot v^{t+1}$$



1-Lookahead

Lemma 2. Assume agent i is in slot i . If there is an agent with $v^t < v^{t+1}$ then agent t **myopically** prefers slot $t+1$.

Proof. If not, then

$$\begin{aligned}c_{t+1}(v^t - b_{t+2}) &\leq c_t(v^t - b_{t+1}) \\ &\leq c_t \left(v^t - \left(1 - \frac{c_{t+1}}{c_t}\right)v^{t+1} - \frac{c_{t+1}}{c_t}b_{t+2} \right) \\ &< c_t \left(\frac{c_{t+1}}{c_t}v^t - \frac{c_{t+1}}{c_t}b_{t+2} \right) \\ &= c_{t+1}(v^t - b_{t+2})\end{aligned}$$



Worst-Case 2-Lookahead

Theorem. Any 2-lookahead equilibrium is socially optimal (in the worst case, leaf model).

Proof.

- The T agents with the highest valuations win the T slots.
 - As with balanced bidding the losing agents bid their values.
- The lowest valuation winner i myopically improves by moving to slot T . - Iteratively apply Lemma 2.
- This move has the same 2-lookahead value.
 - Other winning bidders cannot hurt her as she made a balanced bid.
 - The losing bidder have lower valuations than her winning bid.
- Staying in slot i has lower 2-lookahead value.
 - Its myopic value is worse than slot T .
 - Its (worst case) 2-lookahead value can only be worse.




Average-Case 2-Lookahead


- Some bad news...

Theorem. There are 2-lookahead equilibria that are not socially optimal (in the average case, leaf model). 

- But, in contrast to Nash equilibria, we are always guaranteed a *good* solution...

Theorem. Any 2-lookahead equilibrium has a social value within a factor 2 of optimal (average case, leaf model). 


Cost Sharing Games



- Agents choose source-destination paths.
- The cost of a link is shared equally between the agents using it.
- Nash Equilibria can be a factor N from optimal.

Cooperative Behaviours

- But using lookahead we can get uncoordinated “cooperative” behaviour.

Theorem. With k -lookahead, the worst case guarantee in cost sharing games is $O(N/k)$. 

Other Results

- Congestion Games.
 - Selfish Routing.
- Valid Utility Games.
 - Facility Location Games.
 - Market Sharing Games.
 - Combinatorial Auctions.
 - Distributed Caching.
 - Traffic Routing.

Thank You!