

CSE 312


Foundations of Computing II

Lecture 24: Finish Markov Chains, applications

www.slido.com/2167587

Office hours
today 2pm
CSE1203

Agenda

- Markov Chains 
- Two applications:
 - Markov Chain Monte Carlo
 - Pagerank
- A glimpse of auction theory

A typical day in my life

How do we interpret this diagram?

At each time step t

– I can be in one of 3 **states**


- Work, Surf, Email

– If I am in some state s at time t

- the **labels of out-edges** of s **give the probabilities** of my moving to each of the states at time $t + 1$ (as well as staying the same)


– so **labels on out-edges sum to 1**

e.g. If I am in **Email**, there is a 50-50 chance I will be in each of **Work** or **Email** at the next time step, but I will never be in state **Surf** in the next step.



This kind of random process is called a **Markov Chain**

An organized way to understand the distribution of $X^{(t)}$




Write as a tuple $(q_W^{(t)}, q_S^{(t)}, q_E^{(t)})$ a.k.a. a row vector:

$$\vec{q}^{(t)} = [q_W^{(t)}, q_S^{(t)}, q_E^{(t)}]$$

t	0	1	2
$q_W^{(t)} = P(X^{(t)} = \text{Work})$	1	0.4	$= 0.4 \cdot 0.4 + 0.6 \cdot 0.1 = 0.16 + 0.06 = 0.22$
$q_S^{(t)} = P(X^{(t)} = \text{Surf})$	0	0.6	$= 0.4 \cdot 0.6 + 0.6 \cdot 0.6 = 0.24 + 0.36 = 0.60$
$q_E^{(t)} = P(X^{(t)} = \text{Email})$	0	0	$= 0.4 \cdot 0 + 0.6 \cdot 0.3 = 0 + 0.18 = 0.18$

$X^{(t)}$ state at time t .

An organized way to understand the distribution of $X^{(t)}$



$[q_W^{(t)}, q_S^{(t)}, q_E^{(t)}]$


$$\begin{matrix}
 & \begin{matrix} W & S & E \end{matrix} \\
 \begin{matrix} W \\ S \\ E \end{matrix} & \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.1 & 0.6 & 0.3 \\ 0.5 & 0 & 0.5 \end{bmatrix}
 \end{matrix}$$

Write as a “transition probability matrix” M

- one row/col per state. Row=now, Col=next
- each row sums to 1

t	0	1	2
$q_W^{(t)} = P(X^{(t)} = \text{Work})$	1	0.4	$= 0.4 \cdot 0.4 + 0.6 \cdot 0.1 = 0.16 + 0.06 = 0.22$
$q_S^{(t)} = P(X^{(t)} = \text{Surf})$	0	0.6	$= 0.4 \cdot 0.6 + 0.6 \cdot 0.6 = 0.24 + 0.36 = 0.60$
$q_E^{(t)} = P(X^{(t)} = \text{Email})$	0	0	$= 0.4 \cdot 0 + 0.6 \cdot 0.3 = 0 + 0.18 = 0.18$

An organized way to understand the distribution of $X^{(t)}$




$$[q_W^{(t)}, q_S^{(t)}, q_E^{(t)}] \begin{matrix} M \\ \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.1 & 0.6 & 0.3 \\ 0.5 & 0 & 0.5 \end{bmatrix} \end{matrix} = [q_W^{(t+1)}, q_S^{(t+1)}, q_E^{(t+1)}]$$

Vector-matrix multiplication

$$[0.4, 0.6, 0] \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.1 & 0.6 & 0.3 \\ 0.5 & 0 & 0.5 \end{bmatrix} = [0.22, 0.60, 0.18]$$

$q_W^{(1)} = 0.4$	$q_W^{(2)} = 0.4 \cdot 0.4 + 0.6 \cdot 0.1 = 0.16 + 0.06 = 0.22$
$q_S^{(1)} = 0.6$	$q_S^{(2)} = 0.4 \cdot 0.6 + 0.6 \cdot 0.6 = 0.24 + 0.36 = 0.60$
$q_E^{(1)} = 0$	$q_E^{(2)} = 0.4 \cdot 0 + 0.6 \cdot 0.3 = 0 + 0.18 = 0.18$

An organized way to understand the distribution of $X^{(t)}$



Vector-matrix
multiplication

$$[q_W^{(t)}, q_S^{(t)}, q_E^{(t)}] \begin{matrix} M \\ \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.1 & 0.6 & 0.3 \\ 0.5 & 0 & 0.5 \end{bmatrix} \end{matrix} = [q_W^{(t+1)}, q_S^{(t+1)}, q_E^{(t+1)}]$$

$$q_W^{(t)} \cdot 0.4 + q_S^{(t)} \cdot 0.1 + q_E^{(t)} \cdot 0.5 = q_W^{(t+1)}$$


$$q_W^{(t)} \cdot 0.6 + q_S^{(t)} \cdot 0.6 + q_E^{(t)} \cdot 0 = q_S^{(t+1)}$$

$$q_W^{(t)} \cdot 0 + q_S^{(t)} \cdot 0.3 + q_E^{(t)} \cdot 0.5 = q_E^{(t+1)}$$


Write $\mathbf{q}^{(t)} = [q_W^{(t)}, q_S^{(t)}, q_E^{(t)}]$

Then for all $t \geq 0$,

$$\mathbf{q}^{(t)} M = \mathbf{q}^{(t+1)}$$




By induction ... we can derive



$$M = \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.1 & 0.6 & 0.3 \\ 0.5 & 0 & 0.5 \end{bmatrix}$$

$$q^{(t)} = q^{(0)} M^t \text{ for all } t \geq 0$$

Many interesting questions about Markov Chains



1. What is the probability that I am in state s at time 1?
2. What is the probability that I am in state s at time 2?
3. What is the probability that I am in state s at some time t far in the future?


Given: In state **Work** at time $t = 0$

$$q^{(t)} = q^{(0)} M^t \text{ for all } t \geq 0$$

What does M^t look like for really big t ?

$$q^{(t)} = q^{(0)} M^t \text{ for all } t \geq 0$$

M^t as t grows



M

	W	S	E
W	0.4	0.6	0
S	0.1	0.6	0.3
E	0.5	0	0.5

M^2

	W	S	E
W	.22	.6	.18
S	.25	.42	.33
E	.45	.3	.25

M^3

	W	S	E
W	.238	.492	.270
S	.307	.402	.291
E	.335	.450	.215

M^{10}

	W	S	E
W	.2940	.4413	.2648
S	.2942	.4411	.2648
E	.2942	.4413	.2648

M^{30}

	W	S	E
W	.29411764705	.44117647059	.26470588235
S	.29411764706	.44117647058	.26470588235
E	.29411764706	.44117647059	.26470588235


M^{60}

	W	S	E
W	.294117647058823	.441176470588235	.264705882352941
S	.294117647068823	.441176470588235	.264705882352941
E	.294117647068823	.441176470588235	.264705882352941

What does this say about $q^{(t)}$?



M^t as t grows



$$q^{(60)} = q^{(0)} M^{60}$$

$[q_W^{(0)}, q_S^{(0)}, q_E^{(0)}] \cdot \begin{pmatrix} \text{W} & \text{S} & \text{E} \\ \text{.294117647058823} & \text{.441176470588235} & \text{.264705882352941} \\ \text{.294117647068823} & \text{.441176470588235} & \text{.264705882352941} \\ \text{.294117647068823} & \text{.441176470588235} & \text{.264705882352941} \end{pmatrix} = [q_W^{(60)}, q_S^{(60)}, q_E^{(60)}]$

$(a+b+c) \cdot 0.294\dots$


0.29411

Observation

If $q^{(t+1)}$ = $q^{(t)}$ then it will never change again!

Since for all $t \geq 0$, $q^{(t+1)} = \underline{q^{(t)}} M$


$$\textcircled{(a, b, c)} \stackrel{\uparrow}{=} (\underline{a}, b, c) M$$



Observation

If $\mathbf{q}^{(t+1)} = \mathbf{q}^{(t)}$ then it will never change again!

Since for all $t \geq 0$, $\mathbf{q}^{(t+1)} = \mathbf{q}^{(t)} \mathbf{M}$



Called a **stationary distribution** and has a special name

$$\boldsymbol{\pi} = (\pi_W, \pi_S, \pi_E)$$

Solution to $\boldsymbol{\pi} = \boldsymbol{\pi} \mathbf{M}$

$$\pi = \pi M$$

$$I\pi = I\pi M$$


Solving for Stationary Distribution

$$M = \begin{pmatrix} .4 & .6 & 0 \\ .1 & .6 & .3 \\ .5 & 0 & .5 \end{pmatrix}$$

Stationary Distribution satisfies

- $\pi = \pi M$, where $\pi = (\pi_W, \pi_S, \pi_E)$
- $\pi_W + \pi_S + \pi_E = 1$

$$\rightarrow \pi_W = \frac{10}{34}, \pi_S = \frac{15}{34}, \pi_E = \frac{9}{34}$$



As $t \rightarrow \infty$, $q^{(t)} \rightarrow \pi$ no matter what distribution $q^{(0)}$ is !!

Markov Chains in general

Finite Markov chain

- A set of n **states** $\{1, 2, 3, \dots, n\}$
- The state at time t is denoted by $X^{(t)}$
- A **transition matrix** M , dimension $n \times n$

$$M_{ij} = P(X^{(t+1)} = j \mid X^{(t)} = i)$$

- $q^{(t)} = (q_1^{(t)}, q_2^{(t)}, \dots, q_n^{(t)})$ where $q_i^{(t)} = P(X^{(t)} = i)$
- Transition: LTP $\Rightarrow q^{(t+1)} = q^{(t)} M$ so $q^{(t)} = q^{(0)} M^t$

Stationary Distribution of a Markov Chain

Definition. The **stationary distribution of a Markov Chain** with n states is the n -dimensional row vector π (which must be a probability distribution; that is, it must be nonnegative and sum to 1) such that

$$\pi M = \pi$$

Intuition: Distribution over states at next step is the same as the distribution over states at the current step

Fundamental Theorem of Markov Chains

Recall $\mathbf{q}^{(t)}$ is the distribution of being at each state at time t computed by $\mathbf{q}^{(t)} = \mathbf{q}^{(0)} \mathbf{M}^t$. As t gets large $\mathbf{q}^{(t)} \approx \mathbf{q}^{(t+1)}$.

Fundamental Theorem of Markov Chains: For a Markov Chain that is aperiodic* and irreducible*, with transition probabilities \mathbf{M} and for any starting distribution $\mathbf{q}^{(0)}$ over the states

$$\lim_{t \rightarrow \infty} \mathbf{q}^{(0)} \mathbf{M}^t = \boldsymbol{\pi}$$


where $\boldsymbol{\pi}$ is the stationary distribution of \mathbf{M} (i.e., $\boldsymbol{\pi} \mathbf{M} = \boldsymbol{\pi}$)

**These concepts are beyond us but they turn out to cover a very large class of Markov chains of practical importance.*

Another Example: Random Walks


Suppose we start at node 1, and at each step transition to a neighboring node with equal probability.

This is called a “random walk” on this graph.



Example: Random Walks on an Undirected Graph

Start by defining transition probs.




	To 1	To 2	To 3	To 4	To 5
From 1					
From 2					
From 3					
From 4					
From 5					

$$M_{ij} = P(X^{(t+1)} = j \mid X^{(t)} = i)$$

Agenda

- Markov Chains
- Two applications: 
 - Markov Chain Monte Carlo
 - Pagerank
- A glimpse of auction theory



Markov Chain Monte Carlo

- Technique for sampling from high dimensional distributions
- Computational method for studying large, very complex sets.
- In some cases, a technique for getting good approximate solutions to complex optimization problems
- Used in every field of science and engineering
- “To someone working in my part of the world, asking about applications of MCMC is like asking about applications of the quadratic formula. The results are really used in every aspect of scientific inquiry” --- Persi Diaconis, Stanford


2^n

$O(n)$



100

n^3



MCMC

- Idea: simulate a random walk that moves among possible configurations of a system and will converge to a useful distribution over these configurations.
- Example: Knapsack Problem (NP-complete)
 - Input: collection of n items, for each item
 - Value v_i $i=1, \dots, n$
 - Weight w_i $i=1, \dots, n$
 - Goal: output subset S of items of maximum total value, that has total weight $\leq W$.

MCMC for knapsack

- Define a Markov chain with states being possible solutions and transition probabilities that have higher probabilities on “good solutions”
- Simulate the Markov chain for many iterations until reach a “good” state.

MCMC for Knapsack Problem

Algorithm 1 MCMC for 0-1 Knapsack Problem

```
1: subset  $\leftarrow$  vector of  $n$  zeros (indexed by 0 to  $n - 1$ ), where subset is always a binary vector in  $\{0, 1\}^n$  that
   represents whether or not we have each item. (This means that we initially start with an empty knapsack).
2: best_subset  $\leftarrow$  subset
3: for  $t = 1, \dots, \text{NUM\_ITER}$  do
4:    $k \leftarrow$  a uniformly random integer in  $\{0, 1, \dots, n - 1\}$ .
5:   new_subset  $\leftarrow$  subset but with subset[ $k$ ] flipped ( $0 \rightarrow 1$  or  $1 \rightarrow 0$ ).
6:    $\Delta \leftarrow$  value(new_subset)  $-$  value(subset)
7:   if new_subset satisfies weight constraint (total weight  $\leq W$ ) then
8:     if  $\Delta > 0$  OR ( $T > 0$  AND  $\text{Unif}(0, 1) < e^{\Delta/T}$ ) then
9:       subset  $\leftarrow$  new_subset
10:  if value(subset)  $>$  value(best_subset) then
11:    best_subset  $\leftarrow$  subset
```



Brain Break

PageRank: Some History

The year was 1997

- Bill Clinton in the White House
- Deep Blue beat world chess champion (Kasparov)

The Internet was not like it was today. Finding stuff was hard!

- In Nov 1997, only one of the top 4 search engines actually found itself when you searched for it

The Problem

Search engines worked by matching words in your queries to documents.

Not bad in theory, but in practice there are lots of documents that match a query.

- Search for 'Bill Clinton', top result is 'Bill Clinton Joke of the Day'
- Susceptible to spammers and advertisers

The Fix: Ranking Results

- Start by doing filtering to relevant documents (with decent textual match).
- Then **rank** the results based on some measure of ‘quality’ or ‘authority’.

Key question: How to define ‘quality’ or ‘authority’?

Enter two groups:

- Jon Kleinberg (professor at Cornell)
- Larry Page and Sergey Brin (Ph.D. students at Stanford)

Both groups had the same brilliant idea

Larry Page and Sergey Brin (Ph.D. students at Stanford)

- Took the idea and founded Google, making billions



Jon Kleinberg (professor at Cornell)

- MacArthur genius prize, Nevanlinna Prize and many other academic honors




PageRank - Idea

Take into account the directed graph structure of the web.

Use **hyperlink analysis** to compute what pages are high quality or have high authority.


Trust the Internet itself define what is useful via its links.



PageRank - Idea

Idea 1: Think of each link as a citation
“vote of quality”

Rank pages by in-degree?




PageRank - Idea

Idea 1: Think of each link as a citation
“vote of quality”

Rank pages by in-degree?


Problems:

- Spamming
- Some linkers are not discriminating
- Not all links created equal



PageRank - Idea


Idea 2 : Perhaps we should weight the links somehow and then use the weights of the in-links to rank pages




Inching towards PageRank

1. Web page has high quality if it's linked to by lots of high quality pages
2. A page is high quality if it links to lots of high quality pages

That's a recursive definition!



Inching towards PageRank

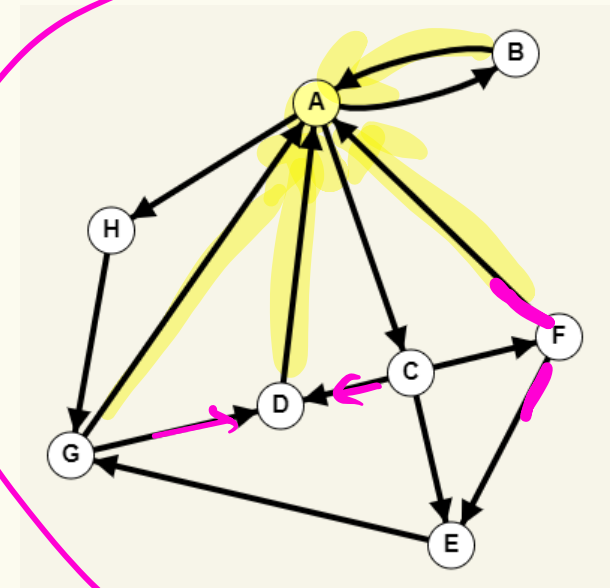


- If web page x has d outgoing links, one of which goes to y , this contributes $1/d$ to the importance of y

- But $1/d$ of what?

We want to take into account the importance of x too...

... so it actually contributes $1/d$ of the importance of x



q_x : quality/authority of page x

$$q_A = q_B \cdot 1 + q_F \frac{1}{2} + q_D \cdot 1 + q_G \frac{1}{2}$$

$$q_D = q_C \cdot \frac{1}{3} + q_G \frac{1}{2}$$

$i \rightarrow j$

we docs i j web \wedge

$$\begin{pmatrix} \frac{1}{d} & \frac{1}{d} & \frac{1}{d} \end{pmatrix}$$

This gives the following equations

Idea: Use the transition matrix M defined by a random walk on the web to compute quality of webpages.

Namely: Find q such that $qM = q$ **Seem familiar?**



This gives the following equations

Idea: Use the transition matrix M defined by a *random walk* on the web to compute quality of webpages.

Namely: Find q such that $qM = q$ **Seem familiar?**




This is the stationary distribution for the Markov chain defined by a random web surfer

- Starts at some node (webpage) and randomly follows a link to another.
- Use stationary distribution of her surfing patterns after a long time as notion of quality



Issues with PageRank

- How to handle dangling nodes (dead ends that don't link to anything)?
- How to handle Rank sinks – group of pages that only link to each other?



Both solutions can be solved by “teleportation”

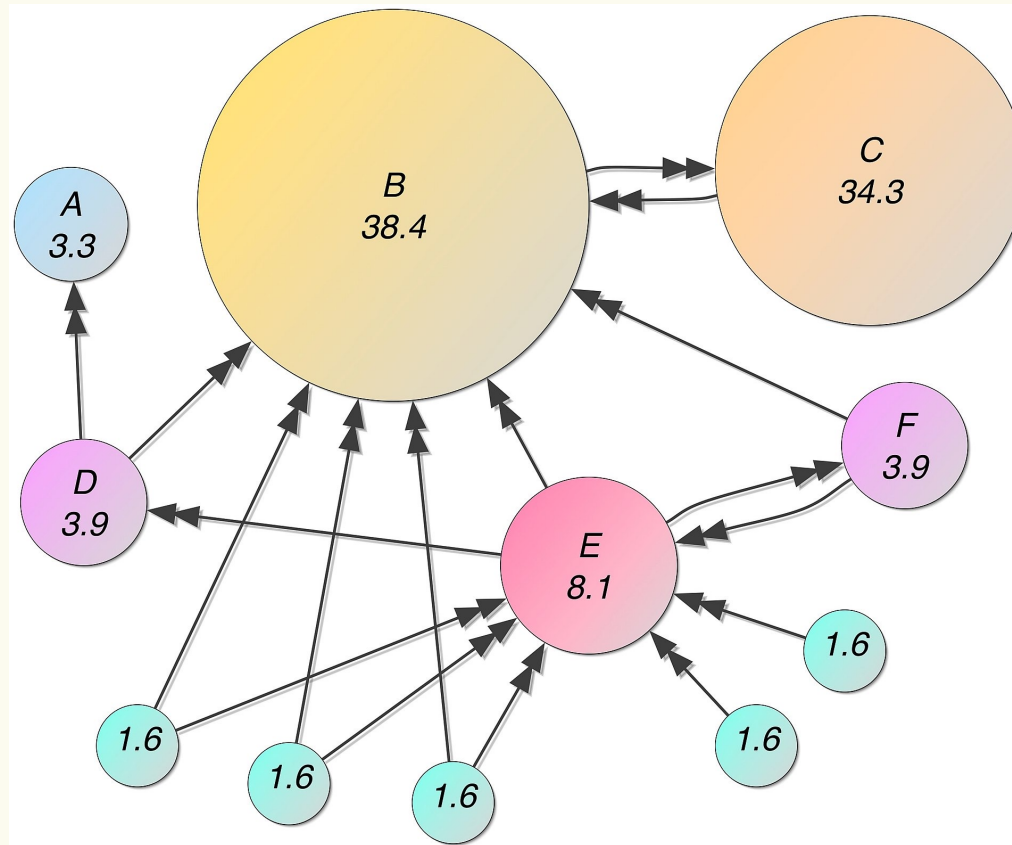
Final PageRank Algorithm



1. Make a Markov Chain with one state for each webpage on the Internet with the transition probabilities $M_{ij} = \frac{1}{outdeg(i)}$.
2. Use a modified random walk. At each point in time if the surfer is at some webpage i :
 - If i has outlinks:
 - With probability p , take a step to one of the neighbors of i (equally likely)
 - With probability $1 - p$, “teleport” to a uniformly random page in the whole Internet.
 - Otherwise, always “teleport”
3. Compute stationary distribution π of this perturbed Markov chain.
4. Define the PageRank of a webpage i as the stationary probability π_i .
5. Find all pages with decent textual match to search and then order those pages by PageRank!

M'

PageRank - Example



It Gets More Complicated

While this basic algorithm was the defining idea that launched Google on their path to success, this is far from the end to optimizing search

Nowadays, Google and other web search engines have a LOT more secret sauce to rank pages, most of which they don't reveal 1) for competitive advantage and 2) to avoid gaming of their algorithms.

Agenda

- Markov Chains
- Two applications:
 - Markov Chain Monte Carlo
 - Pagerank
- A glimpse of auction theory 

Auctions

- Some goods on eBay and amazon are sold via auction.
- Companies like Google and Facebook make most of their money by selling ads.
- The ads are sold via auction.
 - Advertisers submit bids for certain “keywords”

Facebook Ads bidding... 🤔 Is this an auction?

Yes! That's the first thing you need to understand to master bidding management of Facebook Ads. When you're creating a new campaign, you're joining a huge, worldwide auction.

You'll be competing with hundreds of thousands of advertisers to buy what Facebook is selling: Real estate on the News Feed, Messenger, Audience Network, and mobile apps to display your ads to the users.





hawaii vacation



All

Images

News

Videos

Maps

More

Tools

About 2,670,000,000 results (0.79 seconds)

Ad · <https://www.expedia.com/>

Hawaii Island Packages - Book with Expedia and Save

Bundle Your Flight + Hotel & Save! Make Your Trip Memorable. Secure Booking. 600,000+ Hotels Worldwide. Limited Time Offers. New Expedia Rewards. Compare & Save.

Package Deals

Today's Best Flight + Hotel Deals.
Only with Your #1 Leader in Travel.

Last Minute Deals

Expedia Last Minute Travel Deals.
Book Today, Travel Tomorrow.

Ad · <https://www.airbnb.com/>

Hawaii Vacation - Book & Save on Airbnb - airbnb.com

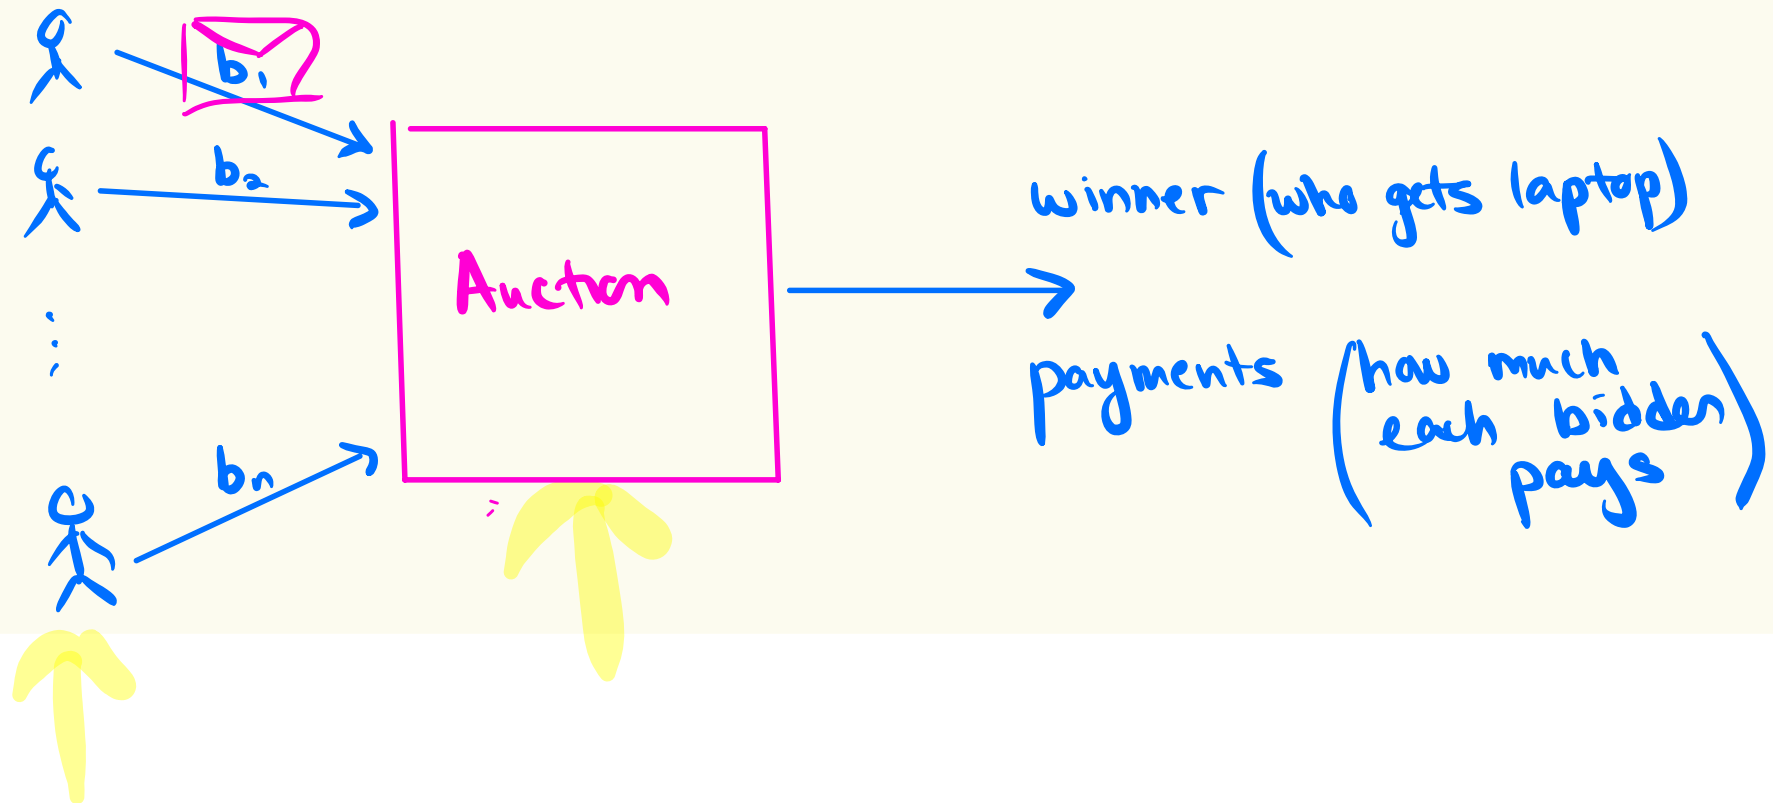
Find **vacation** from **Hawaii**. Perfect for any **Vacation**. 5 Star Hosts. 100,000 Cities. Best Prices. Instant Confirmation. Types: Entire Home, Apartment, Cabin, Villa, Boutique Hotel.

An auction is a ...

- Game
 - Players: advertisers
 - Strategy choices for each player: possible bids
 - Rules of the game – made up by Google/Facebook/whoever is running the auction
- What do we expect to happen? How do we analyze mathematically?

Special case: Sealed bid single item auction

- Say I decide to run an auction to sell my laptop and I let you be the bidders.
- If I want to make as much money as possible – what should I choose as the rules of the auction?



Special case: Sealed bid single item auction

- Say I decide to run an auction to sell my laptop and I let you be the bidders.
- If I want to make as much money as possible – what should the rules of the auction be?

Some possibilities:

- **First price auction:** highest bidder wins; pays what they bid.
- **Second price auction:** highest bidder wins; pays second highest bid.
- **All pay auction:** highest bidder wins; all bidders pay what they bid.

Which of these will make me the most money?

Special case: Sealed Bid single item auction

Some possibilities:

- **First price auction:** highest bidder wins; pays what they bid.
- **Second price auction:** highest bidder wins; pays second highest bid.
- **All pay auction:** highest bidder wins: all bidders pay what they bid.

Bidder	1	2	3	4
Bids	100	81	35	24
Payments	1st price	0	0	0
	2nd price	81	0	0
	All pay	100	81	35

Bidder model

\$100

Each bidder has a value, say v_i for bidder i .

Bidder is trying to maximize their "utility" – the value of the item they get – price they pay.

100	-	98	= \$2
100	-	101	= -\$1
0	*	56	= -\$56



Theorem

A second price auction is **truthful**. In other words, it is always in each bidder's best interest to bid their true value.

$$V \sim \textcircled{F}$$

↑
CDF

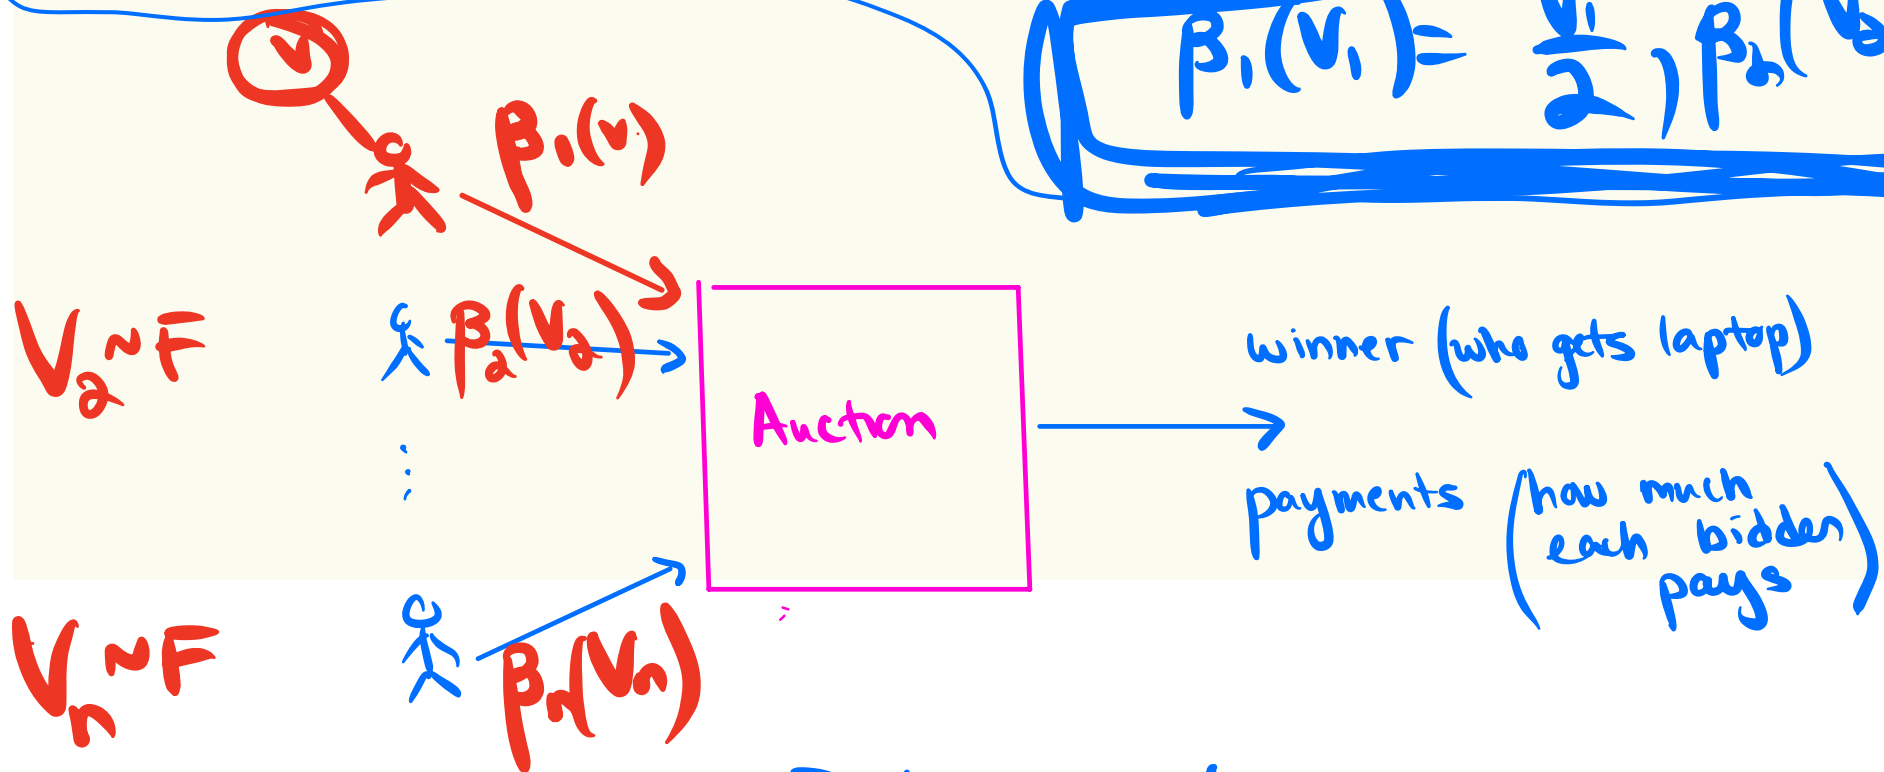
Bayes-Nash equilibrium

bidding strategy value \rightarrow bids.

Suppose that $V_1 \sim F_1, V_2 \sim F_2, \dots, V_n \sim F_n$.

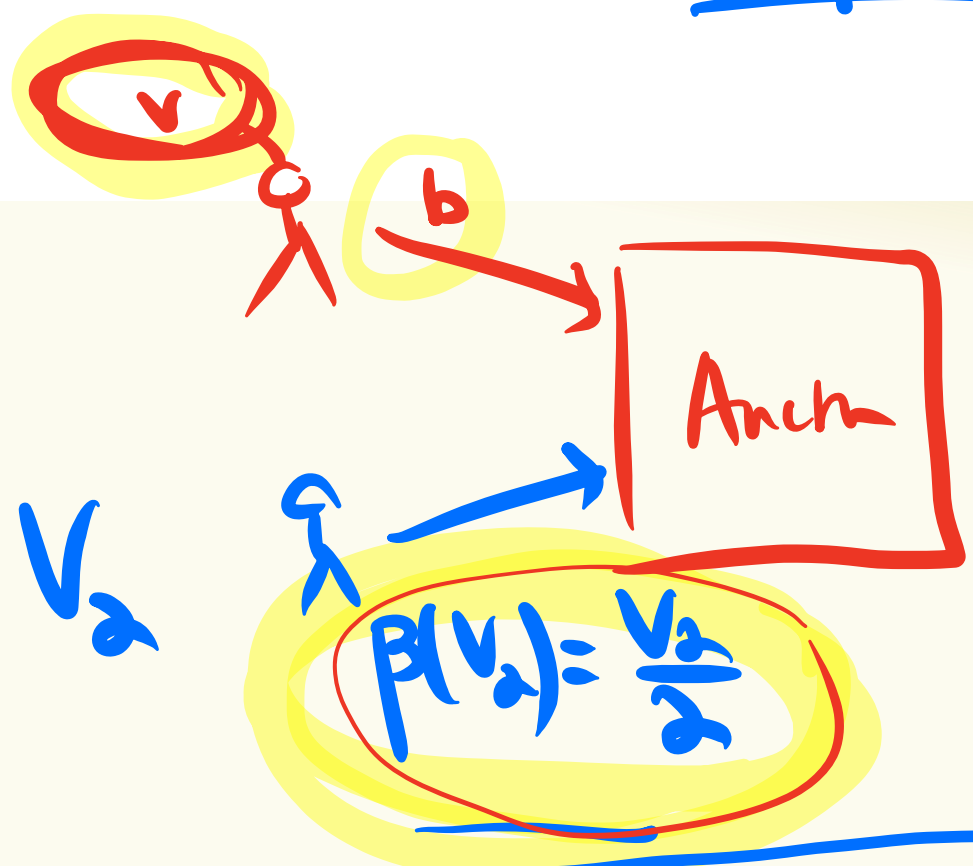
A bidding strategy $\beta_i(\cdot)$ is a **Bayes-Nash equilibrium** if $\beta_i(v_i)$ is a **best response in expectation** to $\beta_j(V_j) \forall j \neq i$.

$$\beta_1(v_1) = \frac{v_1}{2}, \beta_2(v_2) = \frac{v_2}{2}$$



$$v_1, v_2 \sim U[0, 100]$$

$$F(x) = \frac{x}{100}$$



$$E(\text{utility if I bid } b) = (v - b) \Pr(\text{I win the auction})$$

$$= (v - b) \Pr(b > \frac{v_2}{2})$$

$$\Pr(v_2 < 2b)$$
$$\frac{2b}{100}$$

$$(v - b) \cdot \frac{b}{50}$$

$$\frac{d}{db} = (-1) \frac{b}{50} + (v - b) \left(\frac{1}{50} \right)$$

$$= 0$$


$$v - 2b = 0$$

$$v = 2b$$

$$b = \frac{v}{2}$$

2 bidders
 V_1, V_2

	2nd price	1st price	all-pay auction.
v	$\beta(v) = v$	$\beta(v) = \frac{v}{2}$	$\beta(v) = \frac{v^2}{200}$
Exp auctioneer revenue.	$E[\min(V_1, V_2)]$	$E[\max(\frac{V_1}{2}, \frac{V_2}{2})]$	$E[\frac{V_1^2}{200} + \frac{V_2^2}{200}]$
V_1, V_2	$\Downarrow \frac{100}{3}$	$\Downarrow \frac{100}{3}$	$\Downarrow \frac{100}{3}$



$$E(\max(v_1, v_2)) = \frac{2}{3} 100$$

Revenue Equivalence Theorem

In equilibrium, no matter what distribution the bids are drawn from, the expected auctioneer revenue is the same in all three auctions!