CSE 312

Foundations of Computing II

Lecture 28: Clustering (mixture models) + glimpse of auction theory

PAUL G. ALLEN SCHOOL
OF COMPUTER SCIENCE & ENGINEERING

Anna R. Karlin

Slide Credit: Based on Stefano Tessaro's slides for 312 19au
incorporating ideas from myself ☺

Agenda

- Mixture models and clustering
- A glimpse of auction theory

Motivating application: Clustering images

Discover groups of similar images

- Ocean
- Pink flower
- Dog
- Sunset
- Clouds
- …

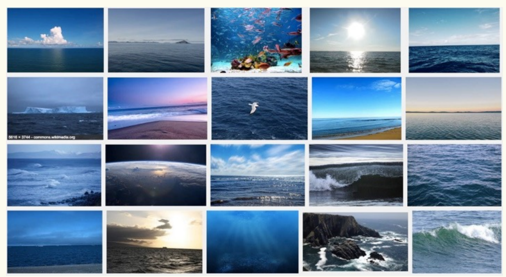



Motivates probabilistic model: Mixture model

- Take uncertainty in assignment into account
  e.g., when clustering documents, might want to say 54% chance document is world news, 45% science, 1% sports, and 0% entertainment

- Allow for cluster shapes not just centers

- Enables learning different weightings of dimensions
  - e.g., how much to weight each word in the vocabulary when computing cluster assignment

Combination of weighted Gaussians

Associate a weight π_k with each Gaussian component

$$\boldsymbol{\pi} = [\overset{\pi_1}{0.47} \;\; \overset{\pi_2}{0.26} \;\; \overset{\pi_3}{0.27}]$$

Fraction of each class in world from which we get data

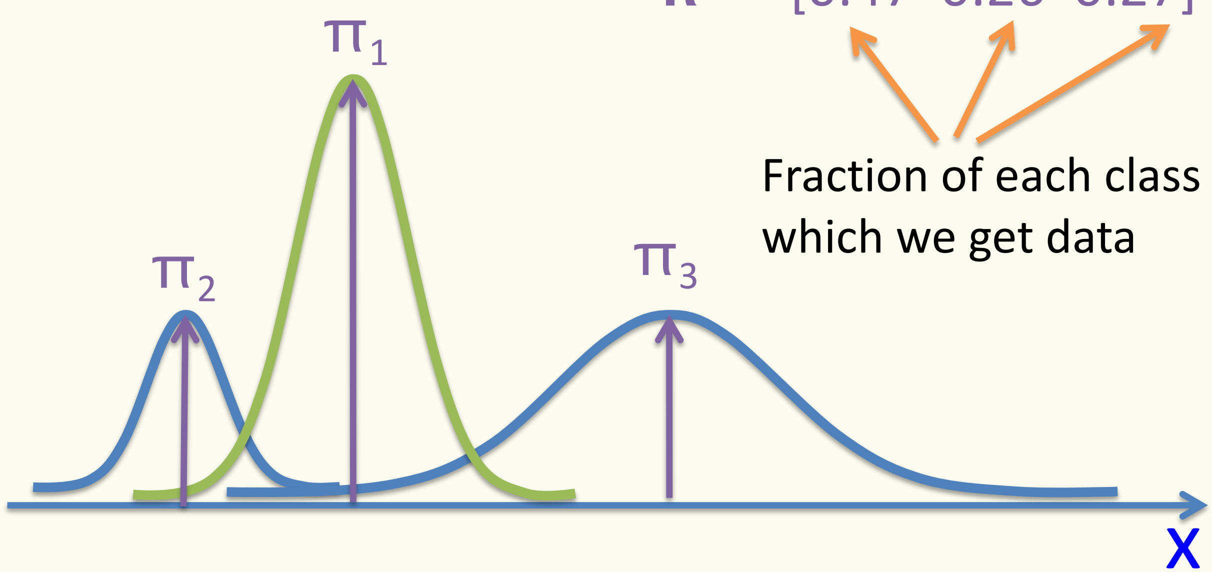

$$0 \leq \pi_k \leq 1$$

$$\sum_{k=1}^{K} \pi_k = 1$$

Mixture of Gaussians (1D)

[R = 0.05, G = 0.7, B = 0.9]

[R = 0.85, G = 0.05, B = 0.35]

[R = 0.02, G = 0.95, B = 0.4]

Each mixture component represents a unique cluster specified by: $\{\pi_k, \mu_k, \sigma_k^2\}$

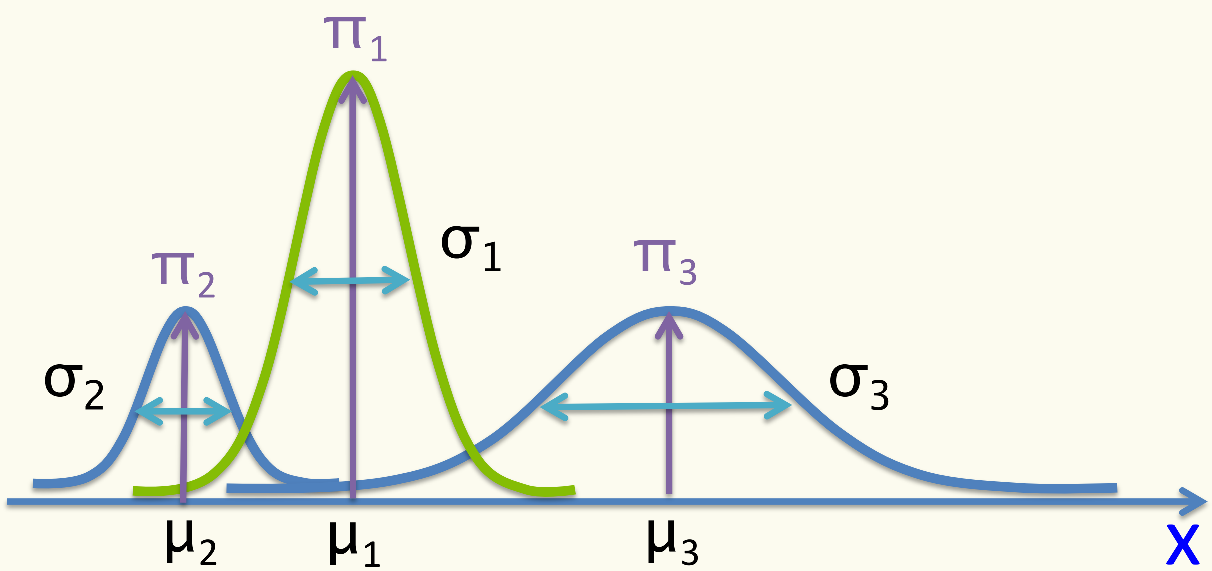

Mixture of Gaussians (general)

[R = 0.05, G = 0.7, B = 0.9]

[R = 0.85, G = 0.05, B = 0.35]

[R = 0.02, G = 0.95, B = 0.4]

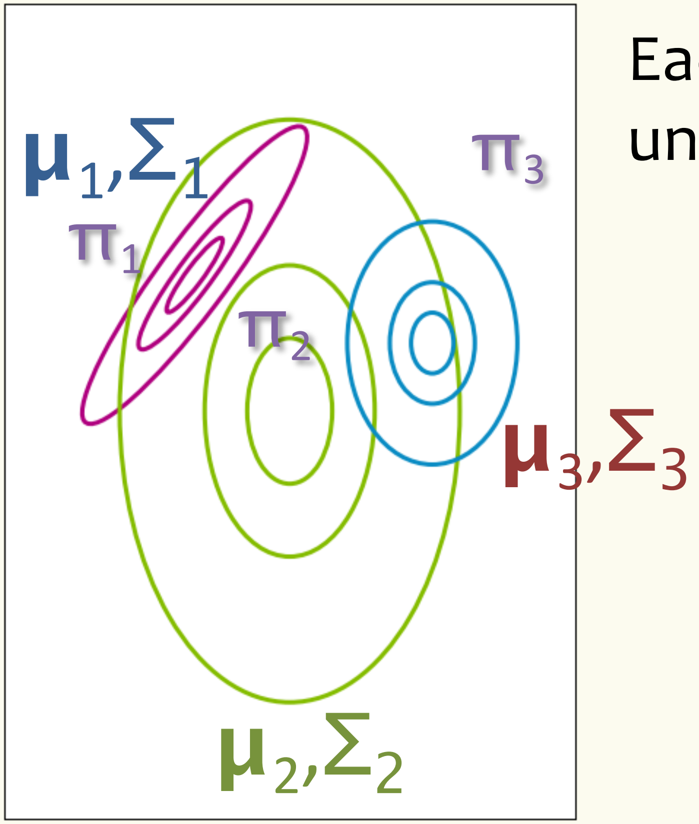

Each mixture component represents a unique cluster specified by:

$$\{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}$$

Mixture model

- K clusters, defined by the following **unknown** parameters

$$\boldsymbol{\Theta} = \{\pi_j, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j\}_{j=1}^k$$

$$\sum_{j=1}^{k} \pi_j = 1.$$

- Problem: Assume that the data comes from such a distribution, and recover the parameters of the distribution (e.g. MLE)
- Determine, for each point, the likelihood of it belonging to cluster j, for each j.

Two1-D Gaussians, with unknown mean and variance

- Easy if know the source of each data point.

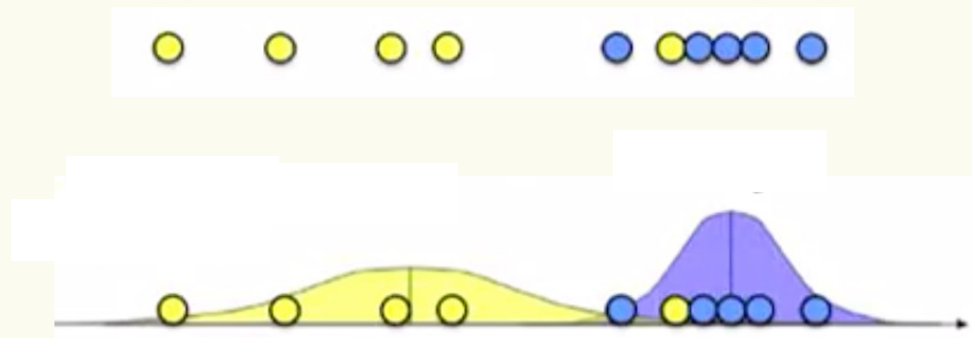

Two1-D Gaussians, with unknown mean and variance

- Easy if know the source of each data point.

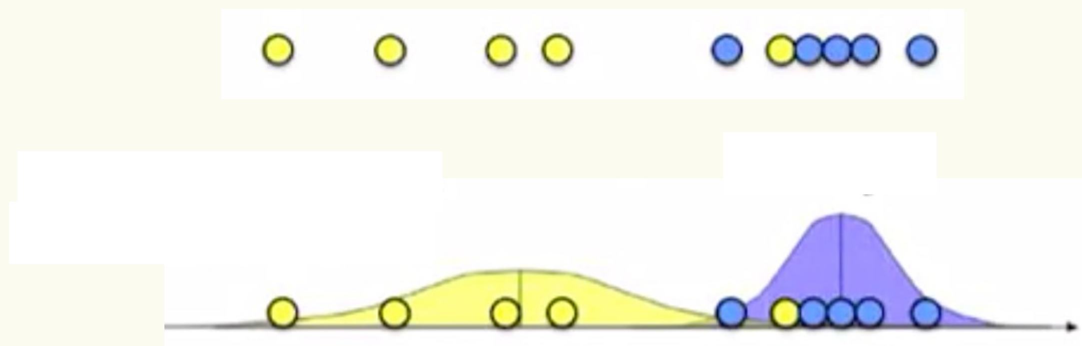

- What if we don't know the source?

Mixture model

- K clusters, defined by the following parameters

$$\Theta = \{\pi_j, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j\}_{j=1}^k$$

$$\sum_{j=1}^{k} \pi_j = 1.$$

- Problem: Assume that the data comes from such a distribution, and recover the parameters of the distribution.
- Determine, for each point, the likelihood of it belonging to cluster j, for each j.
- **PROBLEM: no closed form solution**

Expectation Maximization Algorithm

Two step approach based on following observation

- If we knew which cluster each sample was from, we could estimate all the parameters.

- If we knew all the parameters we could estimate the chance each point came from each cluster.

- EM is an iterative algorithm that alternates between these two steps.

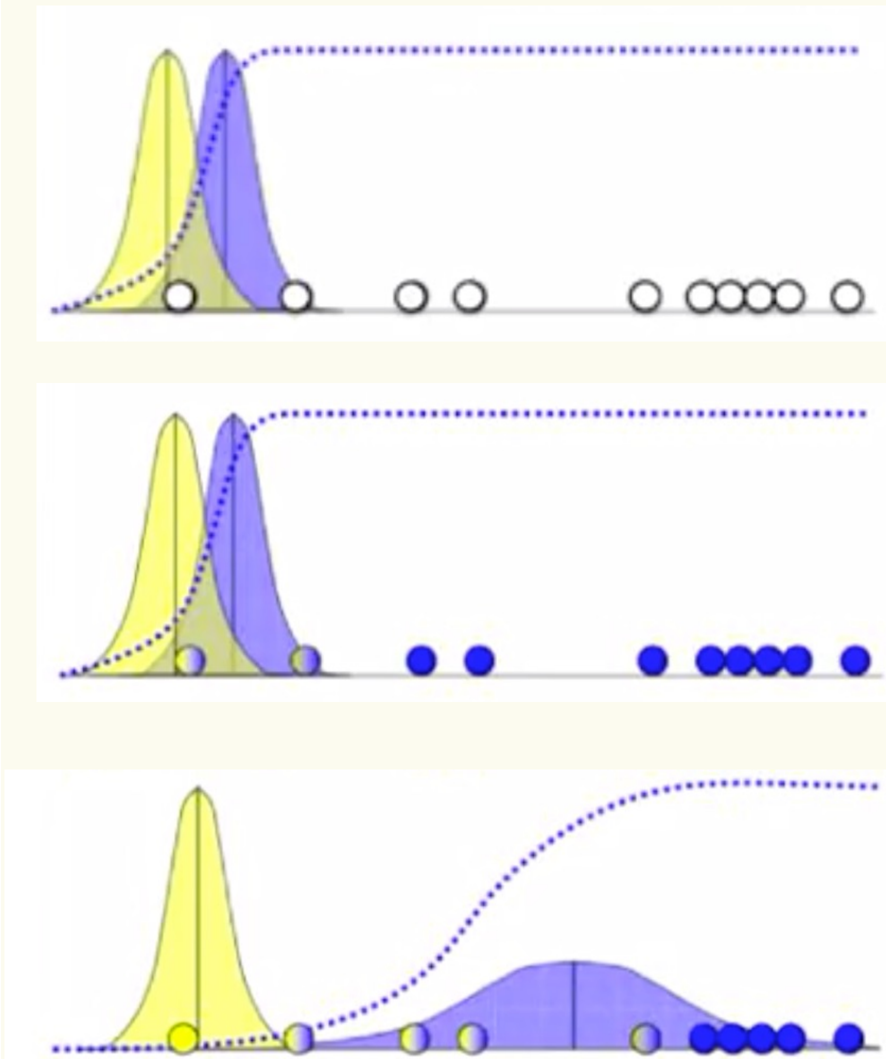

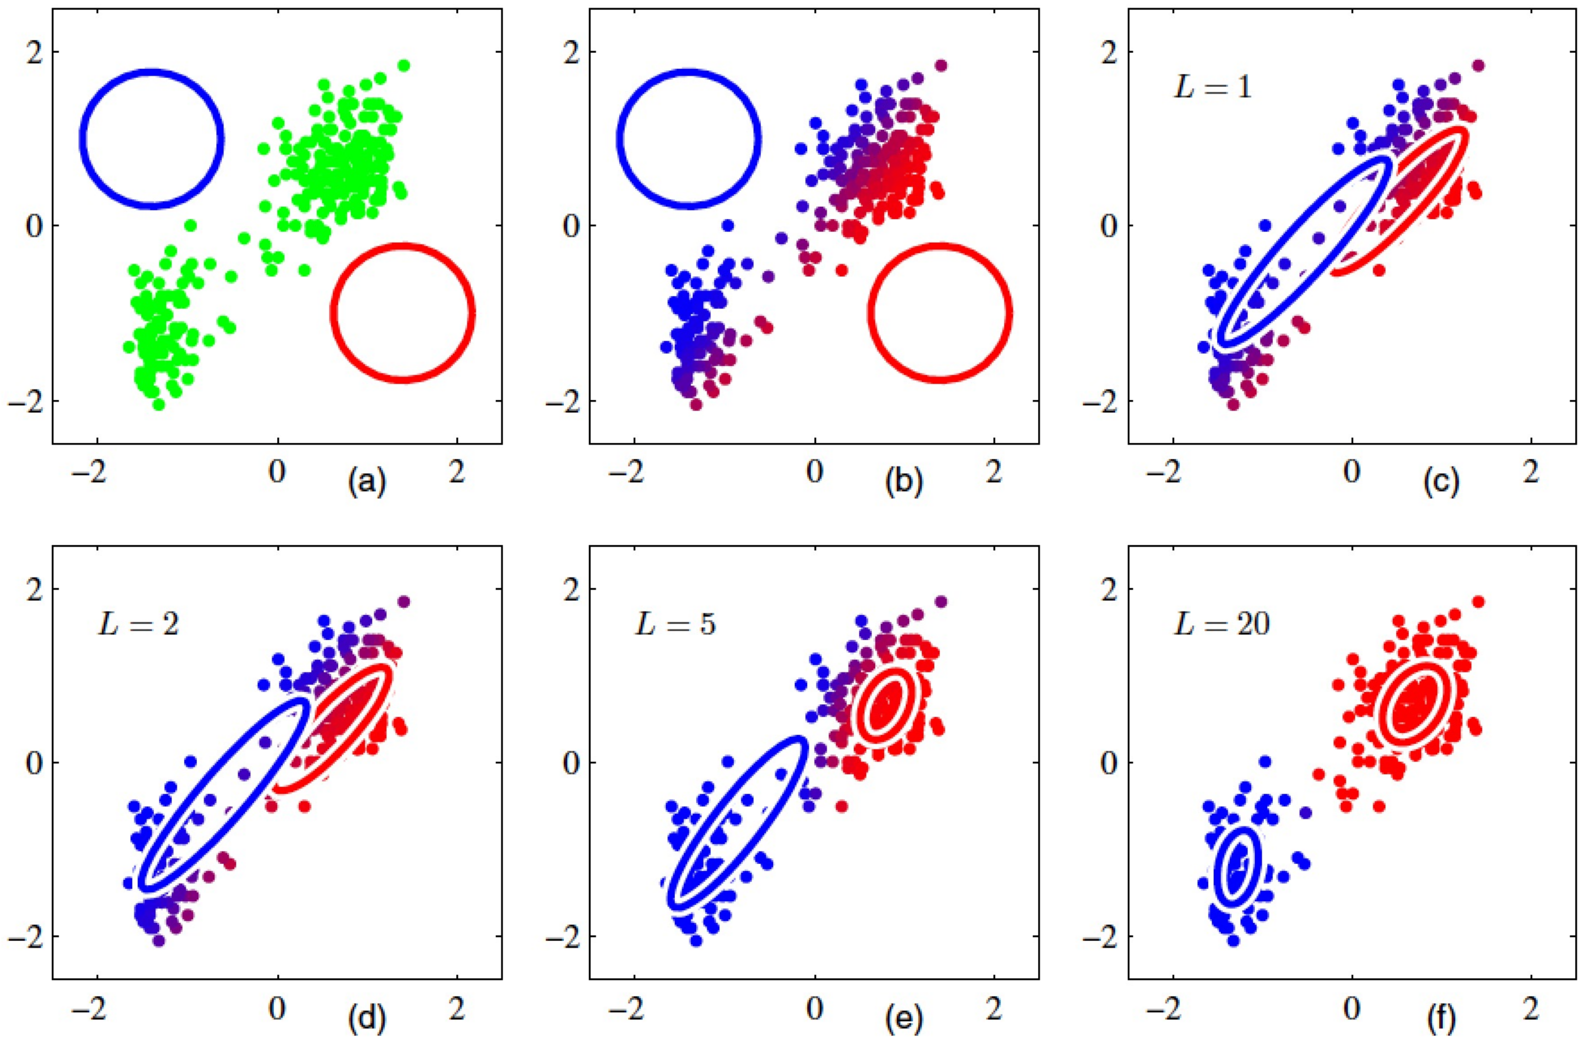
2
0
−2
−2
0
2
(a)
(b)
(c)
(d)
(e)
(f)
$L = 1$
$L = 2$
$L = 5$
$L = 20$

Agenda

- Mixture models and clustering
- 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.

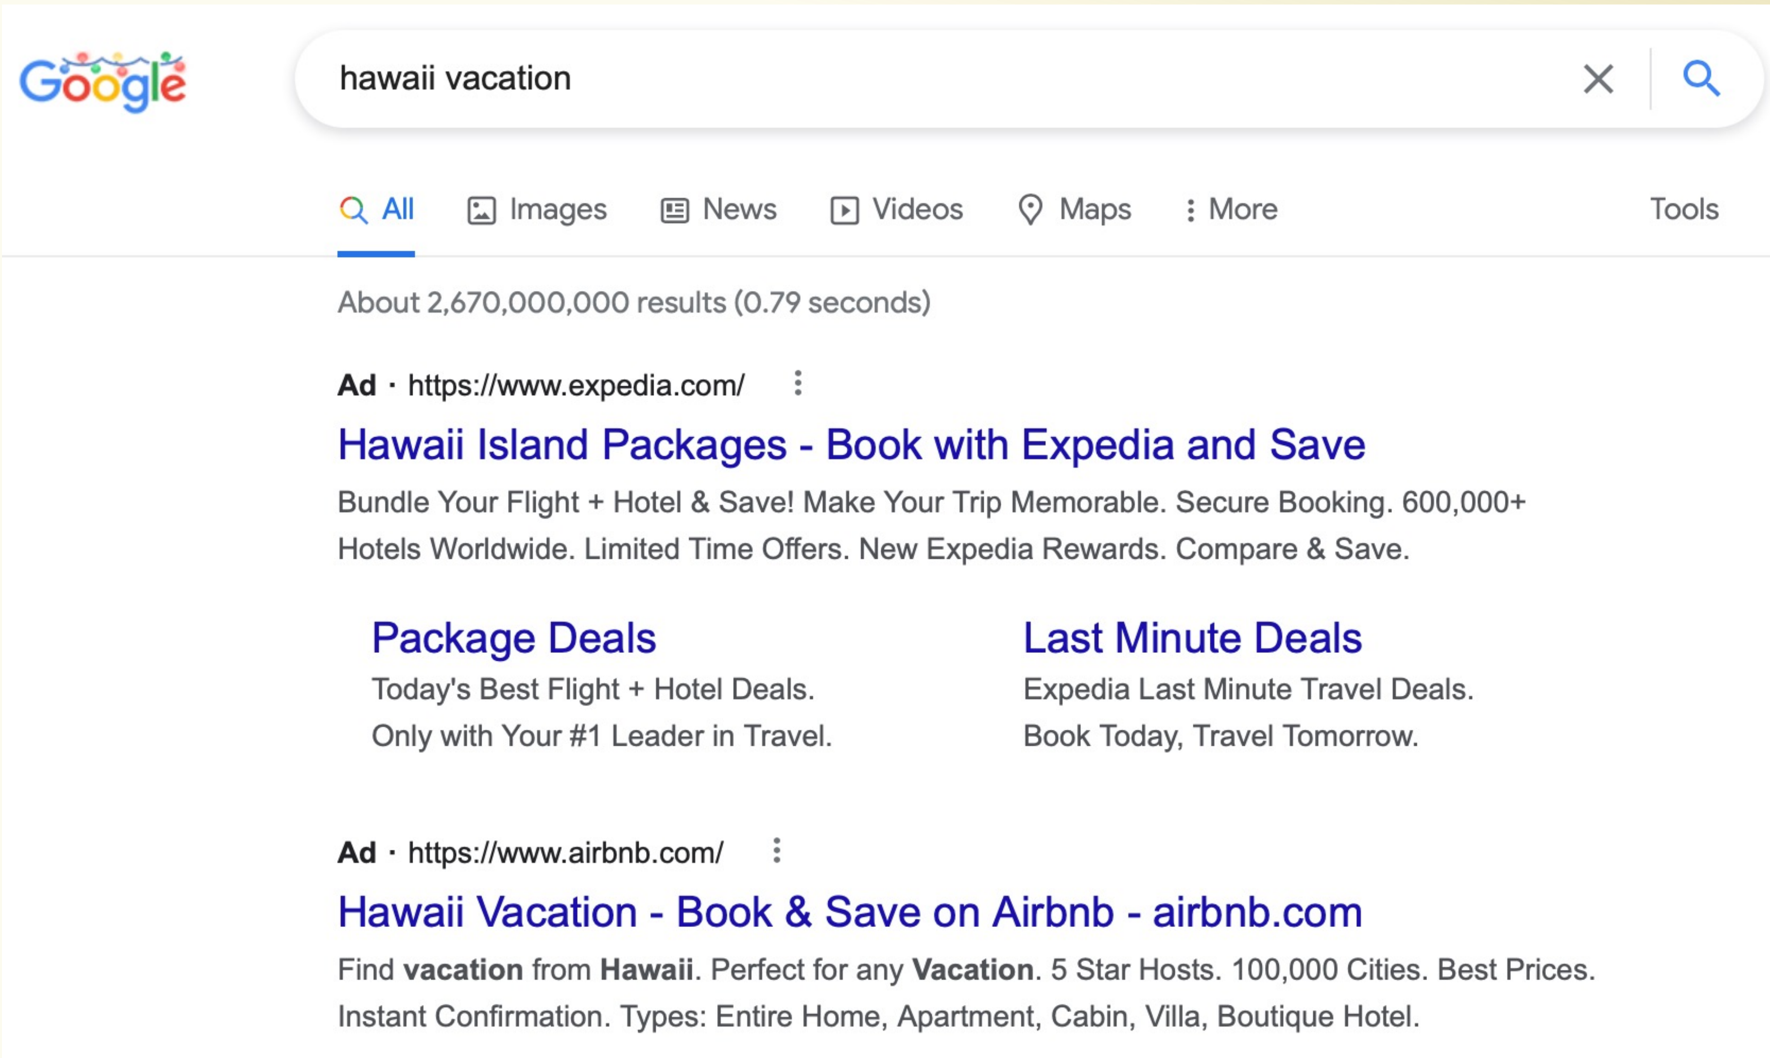
Google
hawaii vacation
All
Images
News
Videos
Maps
More
Tools
About 2,670,000,000 results (0.79 seconds)
Ad · https://www.expedia.com/
Hawaii Island Packages - Book with Expedia and Save
Bundle Your Flight + Hotel & Save! Make Your Trip Memorable. Secure Booking. 600,000+ Hotels Worldwide. Limited Time 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 model

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.

Theorem

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

Bayes-Nash equilibrium

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 $\beta_j(V_j)\ \forall j \neq i$.

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!