Hybrid Keyword Auctions

Kamesh Munagala

Duke University

Joint work with Ashish Goel, Stanford University

Online Advertising

Pricing Models

- CPM (Cost per thousand impressions)
- CPC (Cost per click)
- CPA (Cost per acquisition)
- Conversion rates:
  - Click-through-rate (CTR), conversion from clicks to acquisitions, …

Differences between these pricing models:

- Uncertainty in conversion rates:
  - Sparse data, changing rates, …
- Stochastic fluctuations:
  - Even if the conversion rates were known exactly, the number of clicks/conversions would still vary, especially for small samples

Sponsored Search Auction

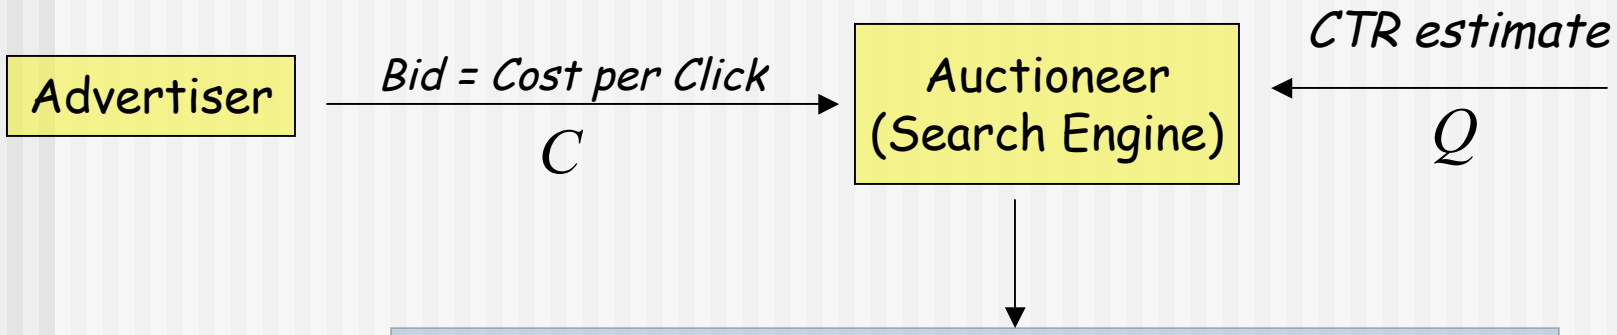

VCG Mechanism: Truthful for a single slot, assuming static CTR estimates

Can be made truthful for multiple slots [**Vickrey-Clark-Groves, Myerson81, AGM06**]

This talk will focus on single slot for proofs/examples

When Does this Work Well?

- High volume targets (keywords)
  - Good estimates of CTR
- What fraction of targets are high volume?
  - Folklore: a small fraction
  - **Motivating problem:**
    - How to better monetize the low volume keywords?

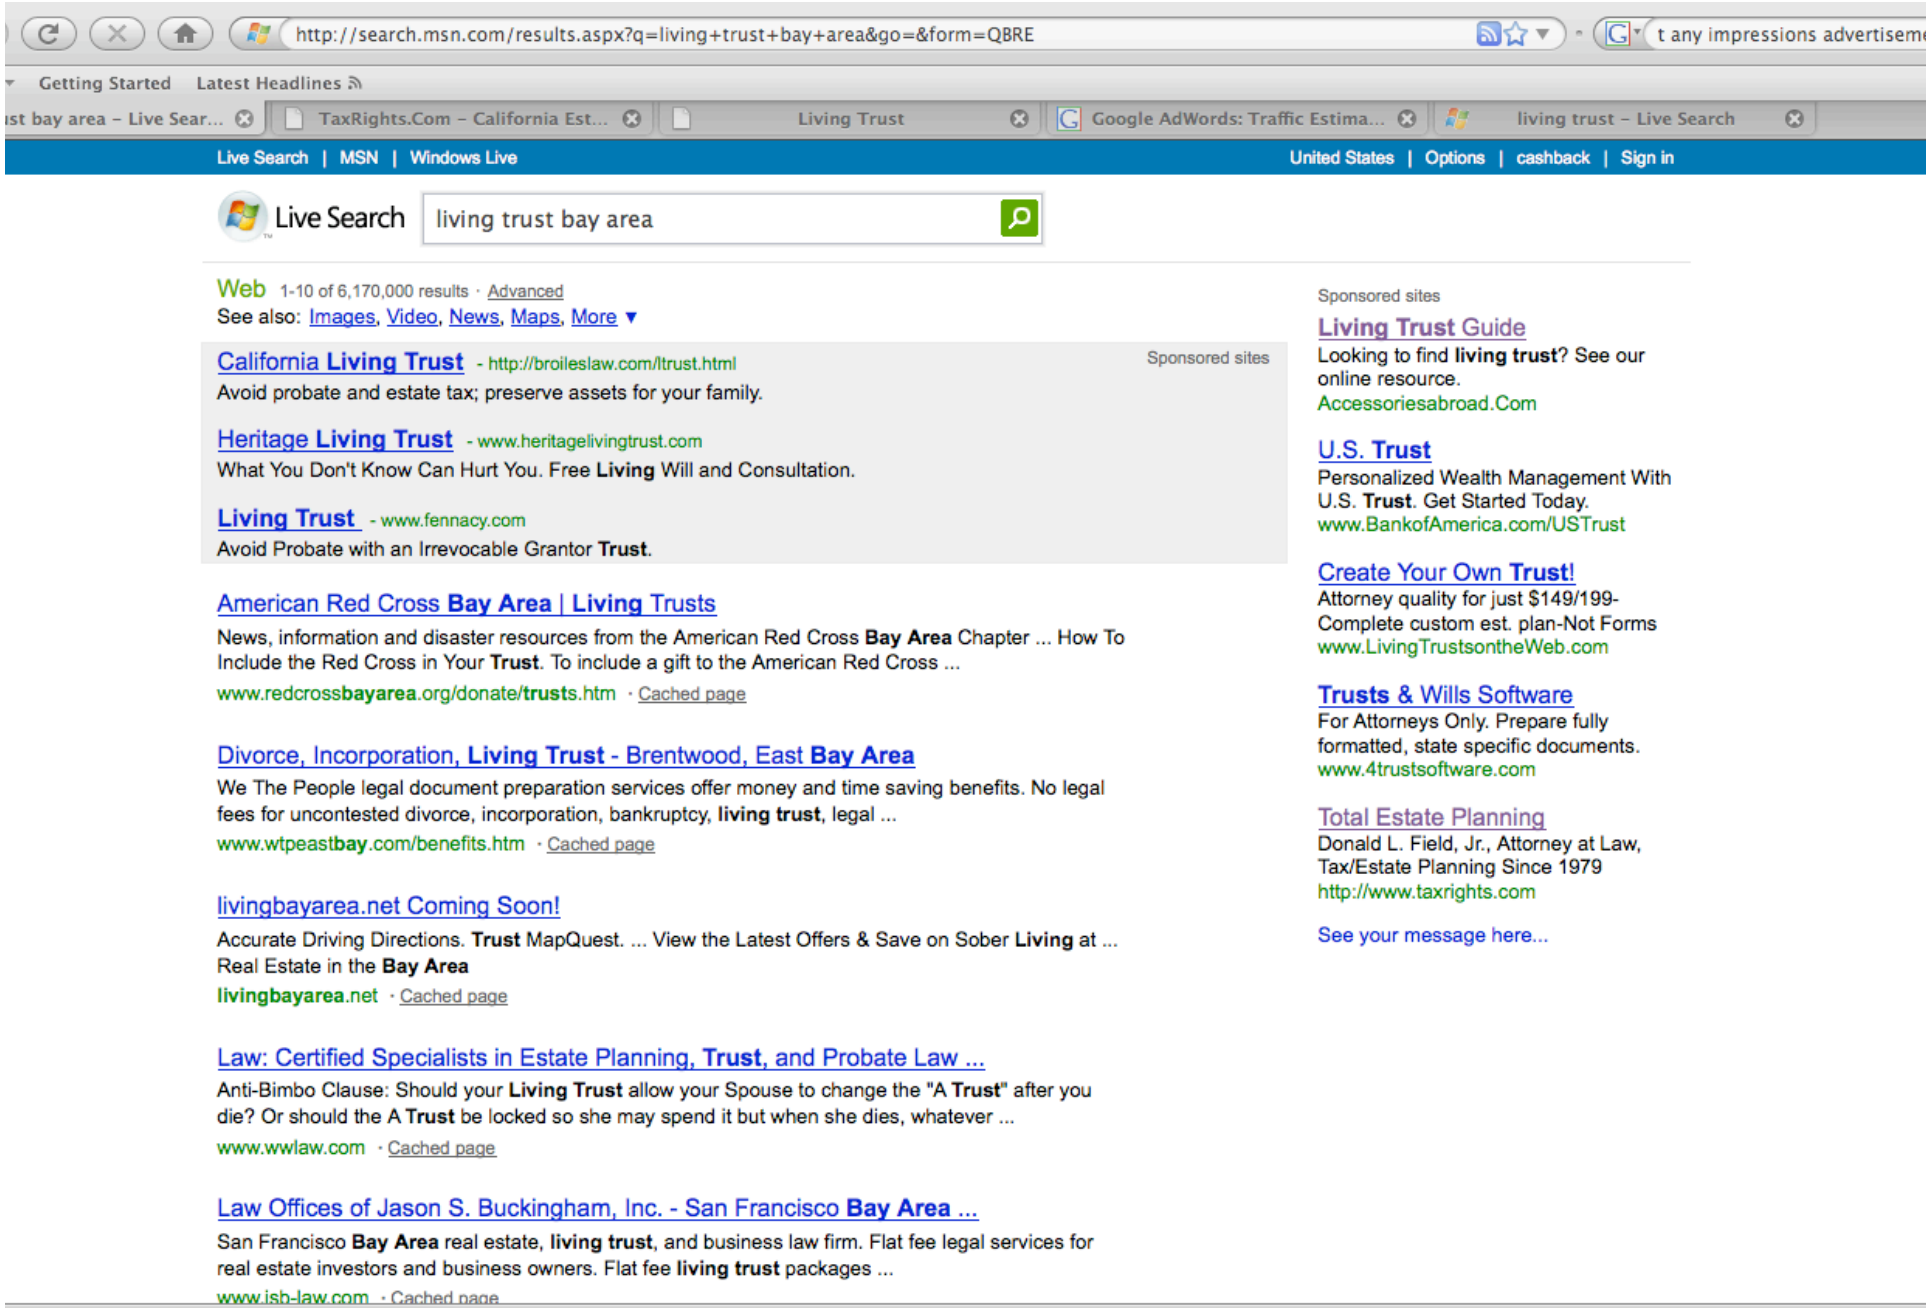
http://search.msn.com/results.aspx?q=living+trust+bay+area&go=&form=QBRE
Getting Started
Latest Headlines
TaxRights.Com – California Est...
Living Trust
Google AdWords: Traffic Estima...
living trust – Live Search
Live Search | MSN | Windows Live
United States | Options | cashback | Sign in
Live Search
living trust bay area
Web 1-10 of 6,170,000 results · Advanced
See also: Images, Video, News, Maps, More
Sponsored sites
California Living Trust - http://broileslaw.com/ltrust.html
Avoid probate and estate tax; preserve assets for your family.
Heritage Living Trust - www.heritagelivingtrust.com
What You Don't Know Can Hurt You. Free Living Will and Consultation.
Living Trust - www.fennacy.com
Avoid Probate with an Irrevocable Grantor Trust.
American Red Cross Bay Area | Living Trusts
News, information and disaster resources from the American Red Cross Bay Area Chapter ... How To Include the Red Cross in Your Trust. To include a gift to the American Red Cross ...
www.redcrossbayarea.org/donate/trusts.htm · Cached page
Divorce, Incorporation, Living Trust - Brentwood, East Bay Area
We The People legal document preparation services offer money and time saving benefits. No legal fees for uncontested divorce, incorporation, bankruptcy, living trust, legal ...
www.wtpeastbay.com/benefits.htm · Cached page
livingbayarea.net Coming Soon!
Accurate Driving Directions. Trust MapQuest. ... View the Latest Offers & Save on Sober Living at ... Real Estate in the Bay Area
livingbayarea.net · Cached page
Law: Certified Specialists in Estate Planning, Trust, and Probate Law ...
Anti-Bimbo Clause: Should your Living Trust allow your Spouse to change the "A Trust" after you die? Or should the A Trust be locked so she may spend it but when she dies, whatever ...
www.wwlaw.com · Cached page
Law Offices of Jason S. Buckingham, Inc. - San Francisco Bay Area ...
San Francisco Bay Area real estate, living trust, and business law firm. Flat fee legal services for real estate investors and business owners. Flat fee living trust packages ...
www.jsb-law.com · Cached page
Sponsored sites
Living Trust Guide
Looking to find living trust? See our online resource.
Accessoriesabroad.Com
U.S. Trust
Personalized Wealth Management With U.S. Trust. Get Started Today.
www.BankofAmerica.com/USTrust
Create Your Own Trust!
Attorney quality for just $149/199-Complete custom est. plan-Not Forms
www.LivingTrustsontheWeb.com
Trusts & Wills Software
For Attorneys Only. Prepare fully formatted, state specific documents.
www.4trustsoftware.com
Total Estate Planning
Donald L. Field, Jr., Attorney at Law, Tax/Estate Planning Since 1979
http://www.taxrights.com
See your message here...

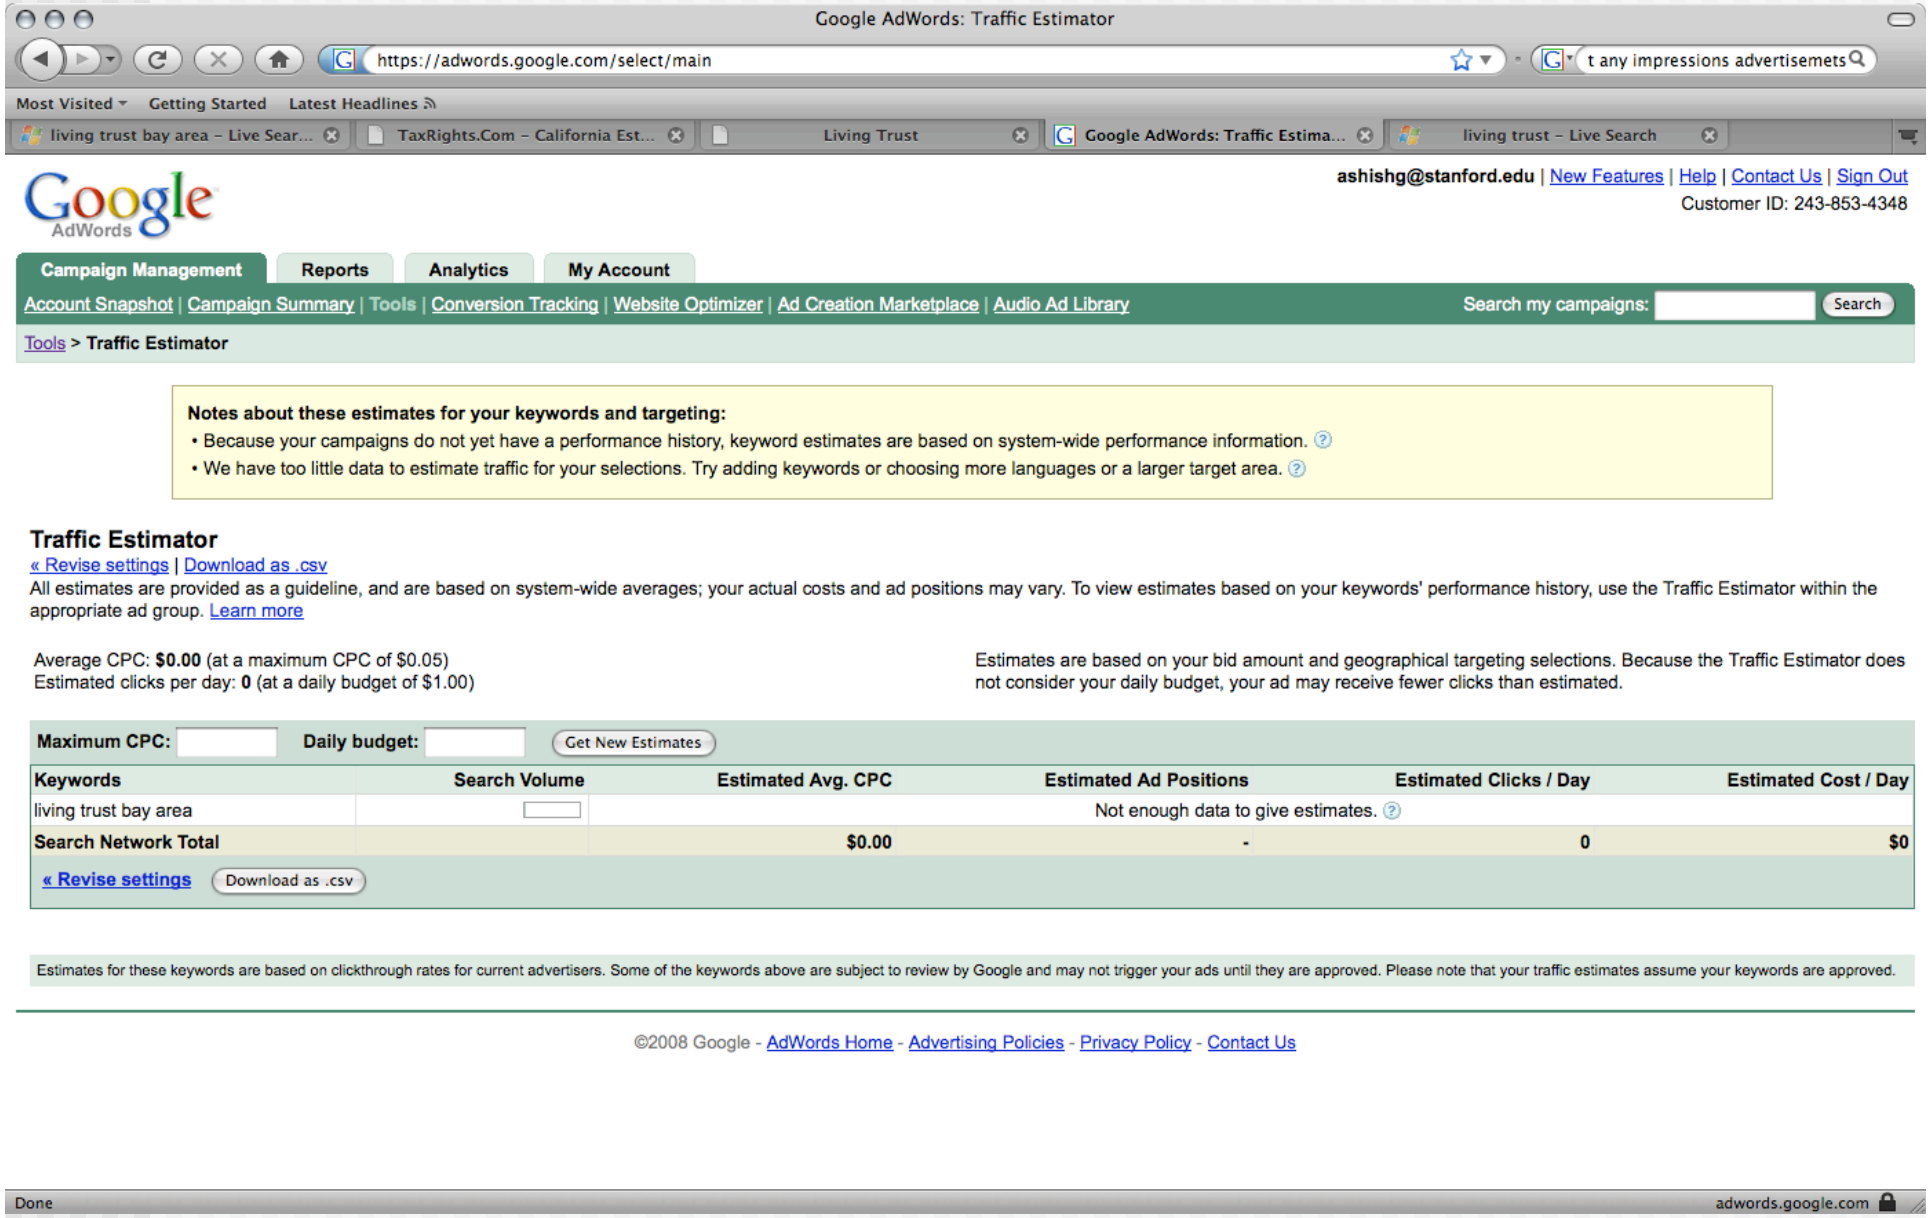
Google AdWords: Traffic Estimator
https://adwords.google.com/select/main
t any impressions advertisemets
Most Visited
Getting Started
Latest Headlines
living trust bay area – Live Sear...
TaxRights.Com – California Est...
Living Trust
Google AdWords: Traffic Estima...
living trust – Live Search
Google AdWords
ashishg@stanford.edu | New Features | Help | Contact Us | Sign Out
Customer ID: 243-853-4348
Campaign Management
Reports
Analytics
My Account
Account Snapshot | Campaign Summary | Tools | Conversion Tracking | Website Optimizer | Ad Creation Marketplace | Audio Ad Library
Search my campaigns:
Search
Tools > Traffic Estimator
Notes about these estimates for your keywords and targeting:
• Because your campaigns do not yet have a performance history, keyword estimates are based on system-wide performance information.
• We have too little data to estimate traffic for your selections. Try adding keywords or choosing more languages or a larger target area.
Traffic Estimator
« Revise settings | Download as .csv
All estimates are provided as a guideline, and are based on system-wide averages; your actual costs and ad positions may vary. To view estimates based on your keywords' performance history, use the Traffic Estimator within the appropriate ad group. Learn more
Average CPC: $0.00 (at a maximum CPC of $0.05)
Estimated clicks per day: 0 (at a daily budget of $1.00)
Estimates are based on your bid amount and geographical targeting selections. Because the Traffic Estimator does not consider your daily budget, your ad may receive fewer clicks than estimated.
Maximum CPC:
Daily budget:
Get New Estimates
Keywords
Search Volume
Estimated Avg. CPC
Estimated Ad Positions
Estimated Clicks / Day
Estimated Cost / Day
living trust bay area
Not enough data to give estimates.
Search Network Total
$0.00
-
0
$0
« Revise settings
Download as .csv
Estimates for these keywords are based on clickthrough rates for current advertisers. Some of the keywords above are subject to review by Google and may not trigger your ads until they are approved. Please note that your traffic estimates assume your keywords are approved.
©2008 Google - AdWords Home - Advertising Policies - Privacy Policy - Contact Us
Done
adwords.google.com

Possible Solutions

- Coarse ad groups to predict CTR:
  - Use performance of advertiser on possibly unrelated keywords
- Predictive models
  - Regression analysis/feature extraction
  - Taxonomies/clustering
  - Collaborative filtering

- **Our approach:** Devise richer pricing models

Hybrid Scheme

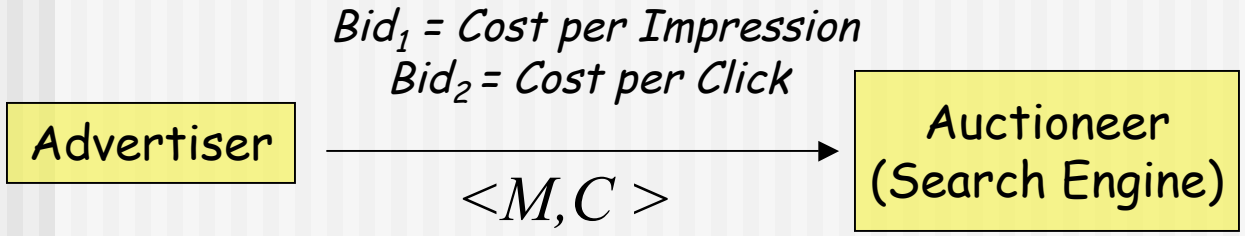

Hybrid Scheme

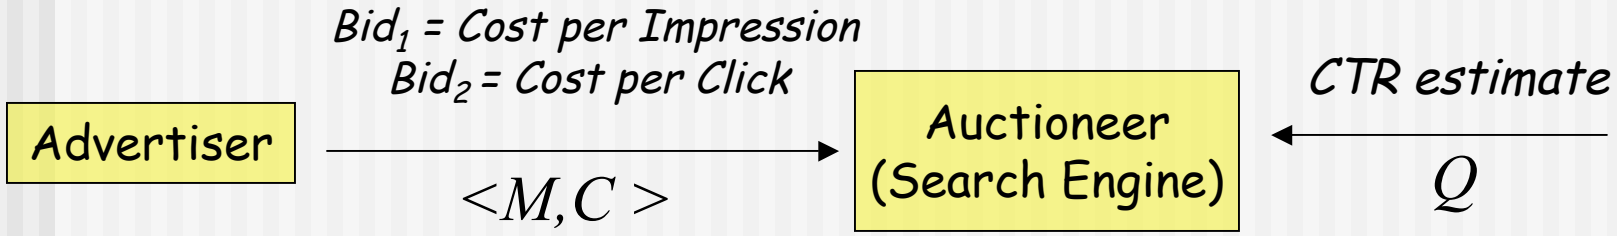

Hybrid Scheme

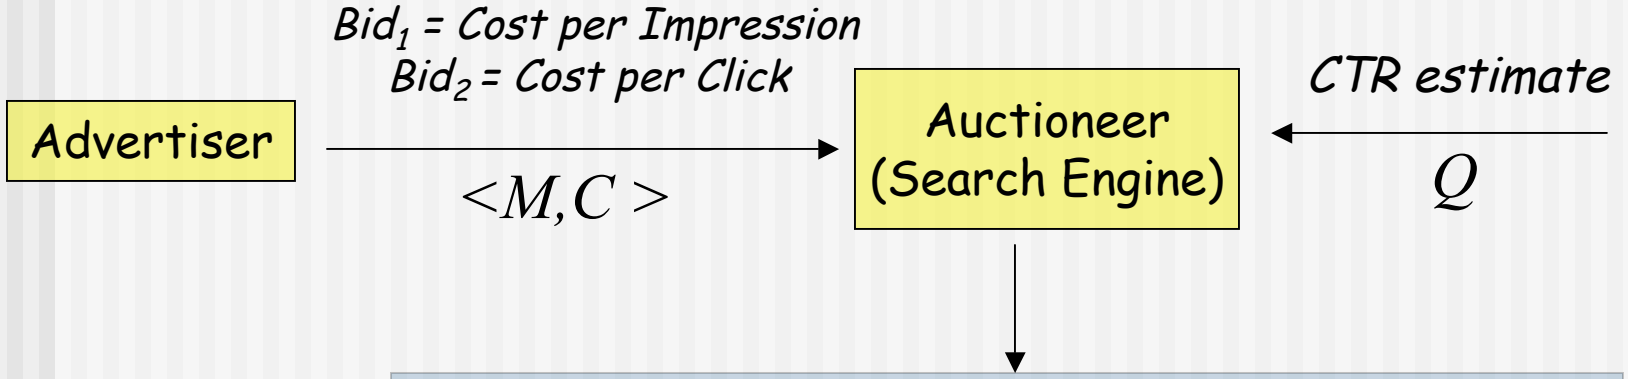

- Advertiser's score $R_i = \max \{ M_i , C_i Q_i \}$

Hybrid Scheme

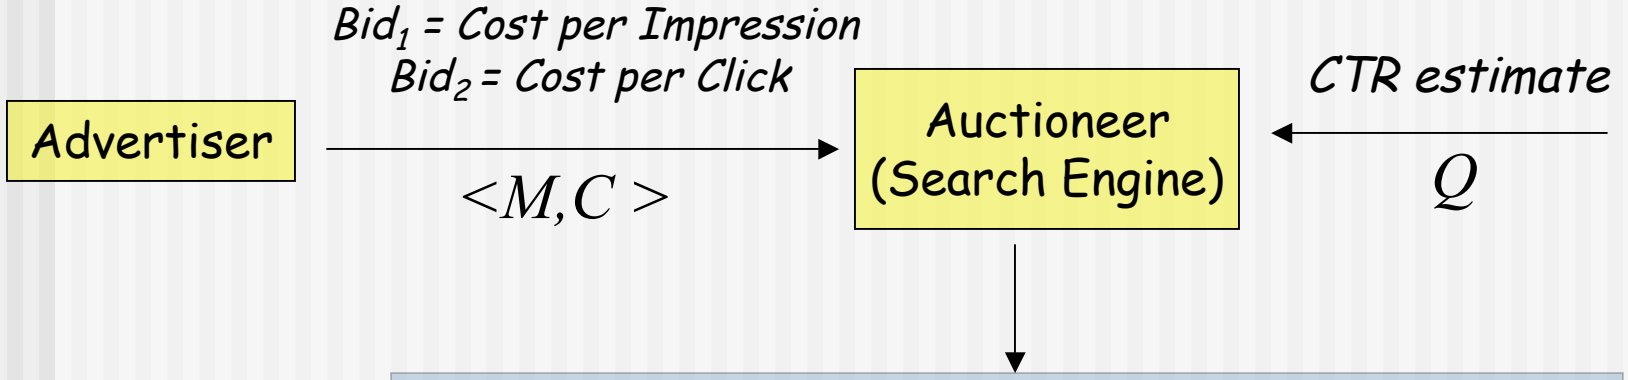

- Advertiser's score $R_i = \max \{ M_i , C_i Q_i \}$
- Order by score: $R_1 > R_2 > \ldots$

Hybrid Scheme

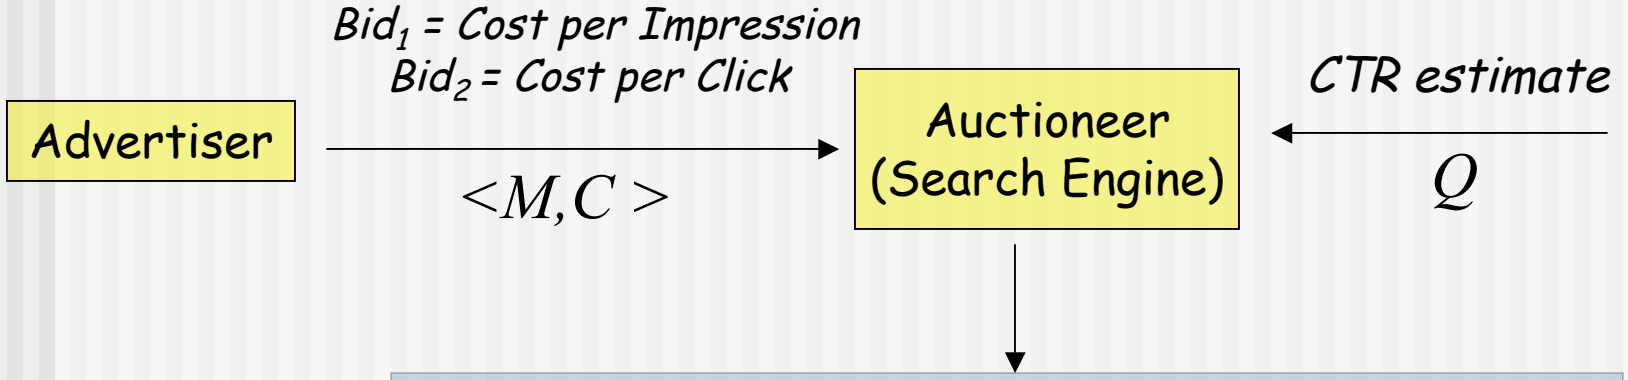

- Advertiser's score $R_i = \max \{ M_i , C_i Q_i \}$
- Order by score: $R_1 > R_2 > \ldots$
- Give impression to bidder 1:
  - If $M_1 > C_1 Q_1$ *then* charge R_2 per impression
  - If $M_1 < C_1 Q_1$ *then* charge R_2 / Q_1 per click

Why Such a Model?

- Per-impression bid:
  - Advertiser's estimate or "belief" of CTR
  - May or may not be an accurate reflection of the truth
  - Backward compatible with cost-per-click (CPC) bidding

Why Such a Model?

- Per-impression bid:
  - Advertiser's estimate or "belief" of CTR
  - May or may not be an accurate reflection of the truth
  - Backward compatible with cost-per-click (CPC) bidding
- Why would the advertiser know any better?
  - Advertiser aggregates data from various publishers
  - Has domain specific models not available to auctioneer
  - Is willing to pay a premium for internal experiments

Benefits

1. **Search engine:**
   - Better monetization of low volume keywords
2. **Advertiser:**
   - Opportunity to make the search engine converge to the correct CTR estimate ***without paying a premium***
3. **Technical:**
   a) Truthful
   b) Accounts for risk characteristics of the advertiser
   c) Allows users to implement complex strategies

Multiple Slots

- Show the top K scoring advertisers
  - Assume $R_1 > R_2 > \ldots > R_K > R_{K+1} \ldots$
- Generalized Second Price (GSP) mechanism:
  - For the i^{th} advertiser, if:
    - If $M_i > Q_i C_i$ then charge R_{i+1} per impression
    - If $M_i < Q_i C_i$ then charge R_{i+1} / Q_i per click

Multiple Slots

- Show the top K scoring advertisers
  - Assume $R_1 > R_2 > \ldots > R_K > R_{K+1} \ldots$
- Generalized Second Price (GSP) mechanism:
  - For the i^{th} advertiser, if:
    - If $M_i > Q_i C_i$ then charge R_{i+1} per impression
    - If $M_i < Q_i C_i$ then charge R_{i+1} / Q_i per click
- Can also implement VCG [Vickrey-Clark-Groves, Myerson81, AGM06]
  - Need separable CTR assumption
  - Details in the paper

Bayesian Model for CTR

True underlying CTR = p

Advertiser

Auctioneer
(Search Engine)

Bayesian Model for CTR

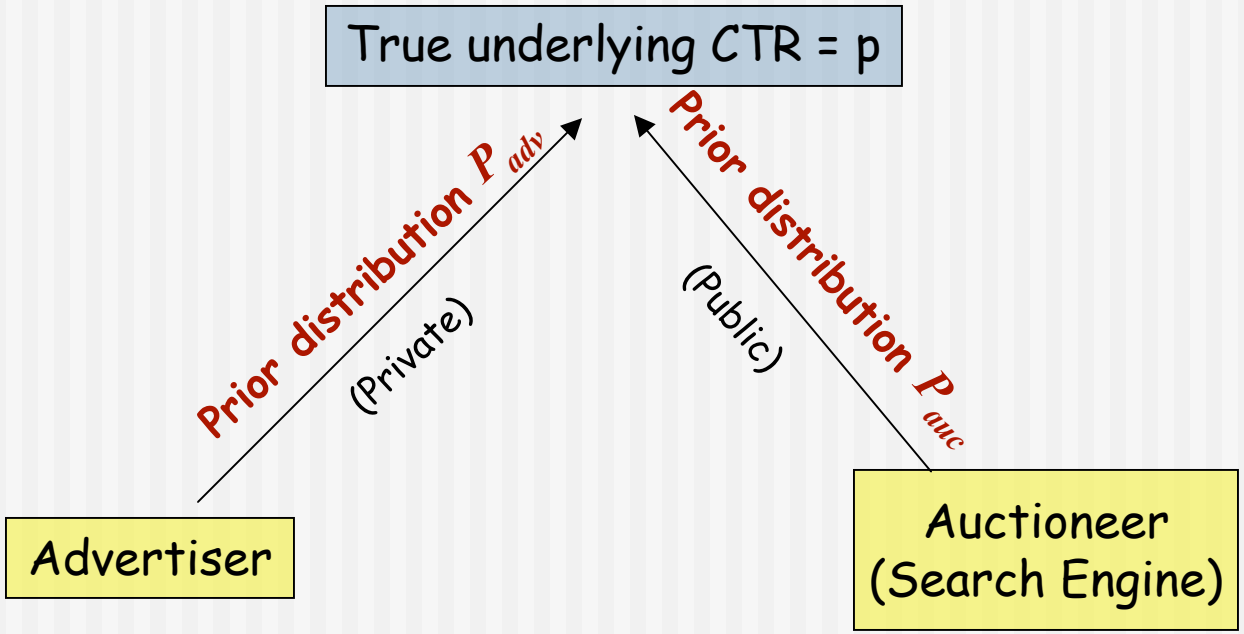

Bayesian Model for CTR

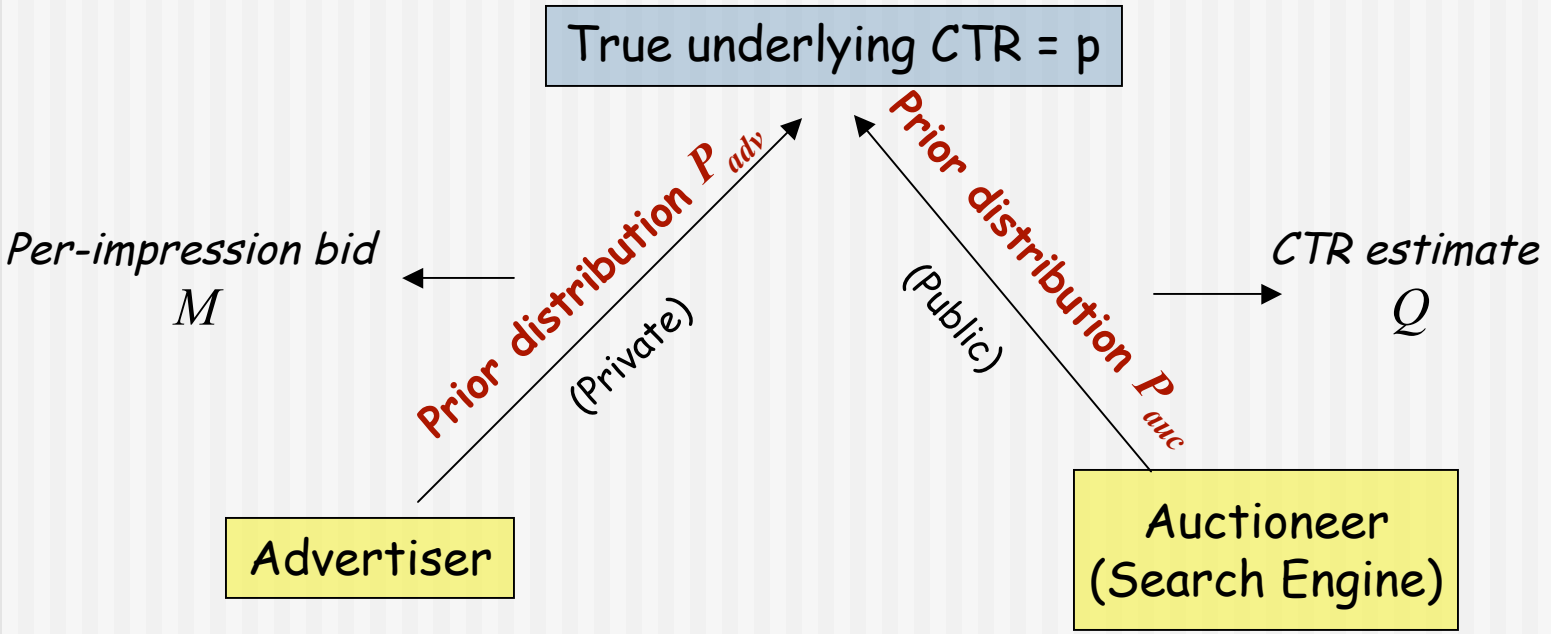

Bayesian Model for CTR

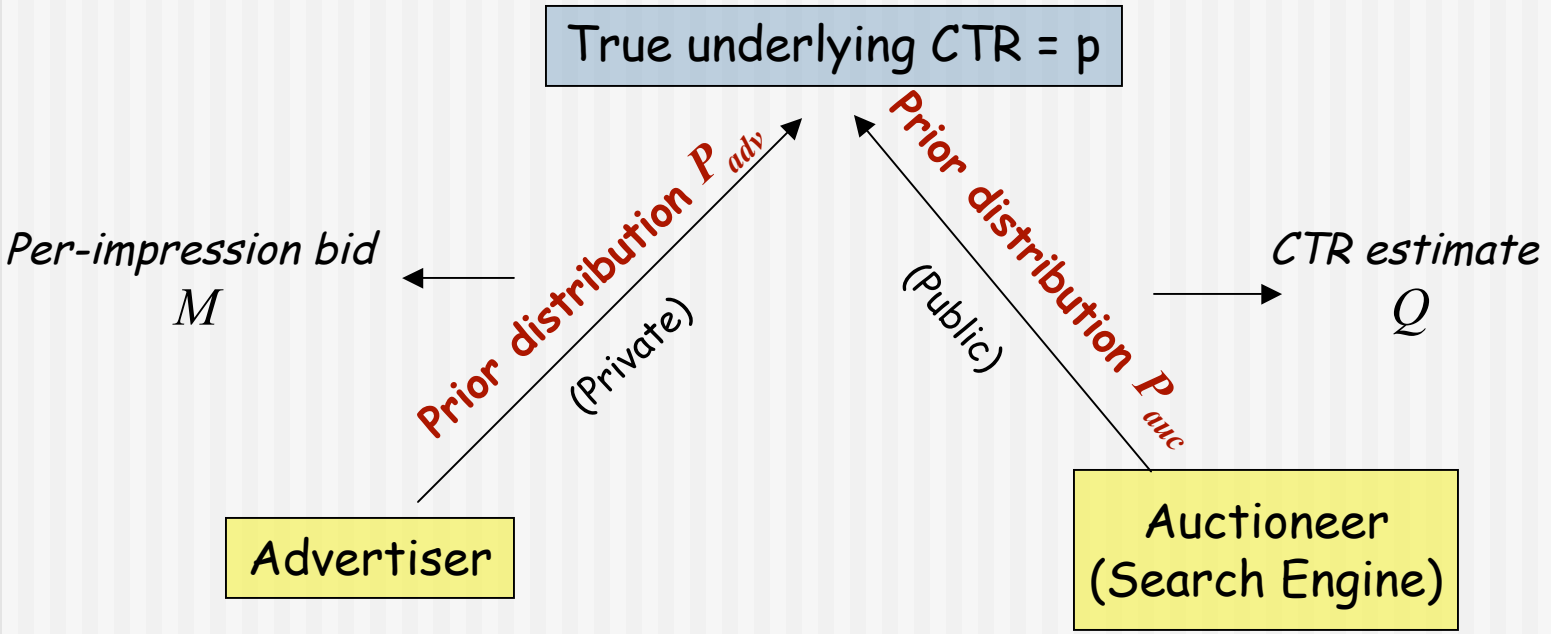

Each agent optimizes based on its current "belief" or prior:

- Beliefs updated with every impression
- Over time, become sharply concentrated around true CTR

What is a Prior?

- Simply models asymmetric information
  - Sharper prior $\Rightarrow$ More certain about true CTR p
  - **E[** Prior **]** need not be equal to p
- Main advantage of per-impression bids is when:
  - Advertiser's prior is sharper than auctioneer's
  - Limiting case: Advertiser certain about CTR p
- Priors are only for purpose of analysis
  - Mechanism is well-defined regardless of modeling assumptions

Truthfulness

- Advertiser assumes CTR follows distribution P_{adv}
- Wishes to maximize expected profit at current step
  - **E[**P_{adv}**]** = x = Expected belief about CTR
  - Utility from click = C
  - Expected profit = $C\,x$ - Expected price

Truthfulness

- Advertiser assumes CTR follows distribution P_{adv}
- Wishes to maximize expected profit at current step
  - **E[**P_{adv}**]** = x = Expected belief about CTR
  - Utility from click = C
  - Expected profit = $C\,x$ - Expected price

Let Cy = Per impression bid
R_2 = Highest other score
If $\max(Cy, C\,Q) < R_2$ then Price = 0
Else:
If $y < Q$ then: Price = $x\,R_2 / Q$
If $y > Q$ then: Price = R_2

Truthfulness

- Advertiser assumes CTR follows distribution P_{adv}
- Wishes to maximize expected profit at current step
  - **E[**P_{adv}**]** = x = Expected belief about CTR
  - Utility from click = C
  - Expected profit = $C\,x$ - Expected price

Bidding (Cx, C) is the dominant strategy

Regardless of Q used by auctioneer

Regardless of P_{adv} and true CTR p

Elicits advertiser's "expected belief" about the CTR!

Holds in many other settings (more later)

Conjugate Beta Priors

- P_{auc} for advertiser $i = Beta(\alpha, \beta)$
  - α, β are positive integers
  - Conjugate of Bernoulli distribution (CTR)
  - Expected value = $\alpha / (\alpha + \beta)$
- Bayesian prior update:
  - Probability of a click at the next step is: $\alpha / (\alpha + \beta)$
  - If click, new P_{auc} (posterior) = $Beta(\alpha+1, \beta)$
  - If no click, new P_{auc} (posterior) = $Beta(\alpha, \beta+1)$

Evolution of Beta Priors

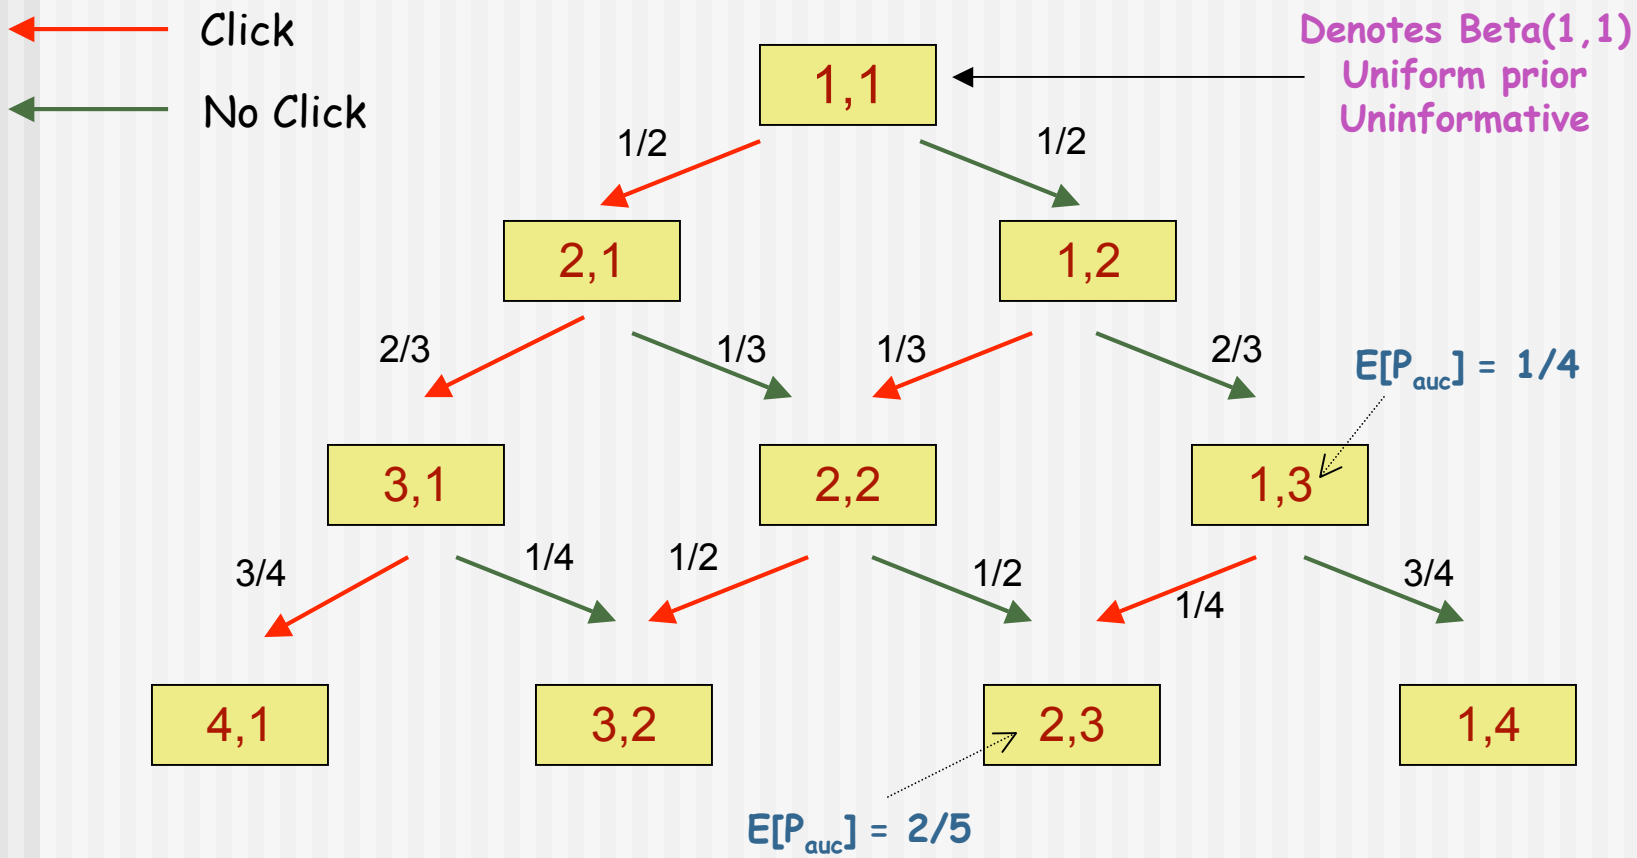

Properties

- Larger α, $\beta \Rightarrow$ Sharper concentration around p
  - Uninformative prior: $Beta(1,1) = Uniform[0,1]$

- $Q = \mathbf{E}[P_{auc}] = \alpha / (\alpha + \beta)$
  - Encodes auctioneer's "belief"
  - Could be different from true CTR p

Certain Advertiser

- Knows true CTR p and bids rationally ($M_i = p_i$)
  - $P_{adv} = p_i$ with probability 1
  - $P_{auc} = Beta(\alpha_i, \beta_i)$ and $Q_i = \mathbf{E}[P_{auc}] = \alpha_i/(\alpha_i + \beta_i)$
- Revenue properties of auctioneer:
  - Worst case: 63% of CPC scheme
  - Canonical case: $\log n$ times better than CPC scheme
- Flexibility for advertiser:
  - Can make P_{auc} converge to p without losing revenue
  - But pays huge premium for achieving this in CPC auction

Better Monetization

- Illustrative Scenario: Low volume keywords
  - n advertisers, all click-utilities $C = 1$
  - All $P_{auc} = Beta\,(1, \log n)$ so that $\mathbf{E}[P_{auc}] = Q \approx 1/\log n$
    - High variance prior
    - Some p_i close to 1 with high probability
  - Per-impression bid will elicit this high p_i
  - CPC auction allocates slot to a random advertiser

- **Theorem:** Hybrid auction can generate $\log n$ ***times more revenue*** for auctioneer than existing CPC auction

Flexibility for Advertisers

- Suppose advertiser certain about CTR = p
  - Assume $C = 1$ and $Q < p$
  - Bids truthfully and wins on per impression bids
- Hybrid scheme: Charged at most p per impression
  - Impressions shown repeatedly
  - Auctioneer's belief P_{auc} will converge to have mean p
  - Now, advertiser switches to CPC bidding
- Assume auctioneer's prior is $Beta(\alpha, \beta)$
  - $Q = \alpha/(\alpha + \beta) < p$

Flexibility for Advertisers

- If CTR converges in T impressions resulting in N clicks:
  - $(\alpha + N)/(\alpha + \beta + T) \geq p$
  - Since $Q = \alpha/(\alpha + \beta) < p$, this implies $N \geq T\,p$
  - Value gain = N; Payment for T impressions at most $T * p$
  - Hence, no loss in revenue to advertiser!
- In the existing CPC auction:
  - The advertiser would have to pay a huge premium for getting impressions and making the CTR converge

Uncertain Advertisers

- Advertiser should "pay premium" for CTR p resolving to a high value
  - What should her bidding strategy be?
  - Does it lead to a socially optimal mechanism?

- Key contribution:
  - Defining a Bayesian model for repeated auctions
  - Dominant strategy exists!

Semi-Myopic Advertiser

- Maximizes discounted utility in contiguous time horizon in which she wins the auction
  - State of other advertisers stays the same during this time
  - Once she stops getting impressions, cannot predict future
    … since future will depend on private information of other bidders!
  - Circumvents negative results in economics literature

- Private information with advertiser:
  - Discount factor γ, value C_i and prior P_{adv}
  - Discount factor models varying optimization horizons
    - Strategic vs. myopic

Dominant Hybrid Strategy

- Bidder always has a dominant hybrid strategy
  - **Bidding Index:** Computation similar to the Gittins index
  - Bidder can optimize her utility by dynamic programming

Dominant Hybrid Strategy

- Bidder always has a dominant hybrid strategy
  - **Bidding Index:** Computation similar to the Gittins index
  - Bidder can optimize her utility by dynamic programming

- Socially optimal in many reasonable scenarios:
  - Myopic advertiser: Has $\gamma_i = 0$; trusts auctioneer's prior:
    - Pure per-click bidding implements the Gittins index policy
  - If advertiser is *certain* of CTR, and Q_i is an underestimate:
    - Bidding index = Per-impression bidding, which is socially optimal
  - Implementation needs *both* per-impression and per-click bids

Summary

- Allow both per-impression and per-click bids
  - Same ideas work for CPM/CPC + CPA
- Significantly higher revenue for auctioneer
- Easy to implement
  - Hybrid advertisers can co-exist with pure per-click advertisers
  - Easy path to deployment/testing
- Many variants possible with common structure:
  - Optional hybrid bids
  - Use the “max” operator to compute score

Open Questions

- Some issues that may be exacerbated:
  - Whitewashing: Re-entering when CTR is lower than the default
  - Fake Clicks: Bid per impression initially and generate false clicks to drive up CTR estimate Q
    - Switch to per click bidding when slot is "locked in" by the high Q
- Analysis of semi-myopic model
  - Other applications of separate beliefs?
- Connections of Bayesian mechanisms to:
  - Regret bounds and learning [Nazerzadeh, Saberi, Vohra '08]
  - Best-response dynamics [Edelman, Ostrovsky, Schwarz '05]