Content Creation Competition and Incentive Design

Fan Yao
Department of Computer Science
University of Virginia

Overview

- A framework for modeling strategic content creation in online recommender platforms
- Diagnostic results
  - *"How bad is top-K Recommendation"* [Yao et al., ICML23]
- Prescriptive solutions
  - *"Rethinking Incentives in Recommender Systems"* [Yao et al., Neurips23])
  - *"User Welfare Optimization with Competing Content Creators"* [In submission])

Market: booming of content creation economy

- Content creation and marketing industry exceeds 600 billion USD and is estimated to grow 16%+ annually between 2023-2028 [T23]
- Major content recommendation platforms now encourage creators to monetize directly

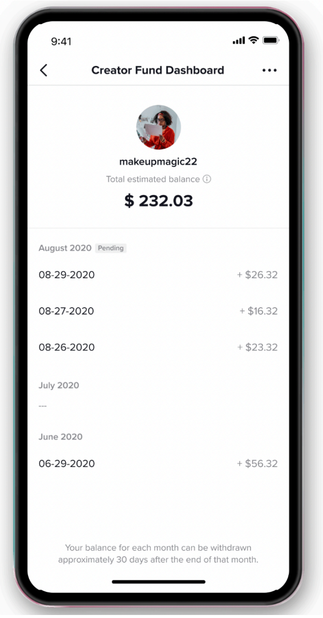

TikTok Creator Fund Program

YouTube Partner Program

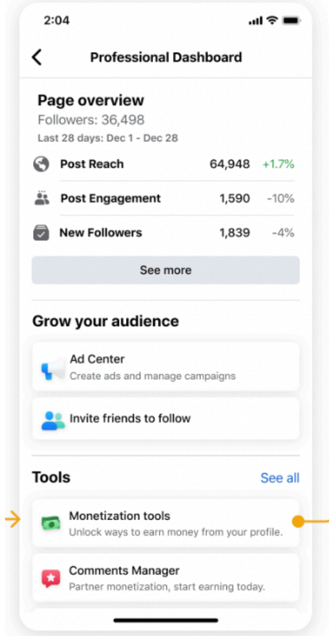

Meta Reel Ads Program

Creator: evidence of strategic creation

- When YouTube's algorithm uses view duration to evaluate the video quality, creators made longer videos [MC23]

- When Instagram poses minimum requirement for monetization program, creators create fake followers and traffics.

Challenge for recommender systems (RS)

Where does the content come from?

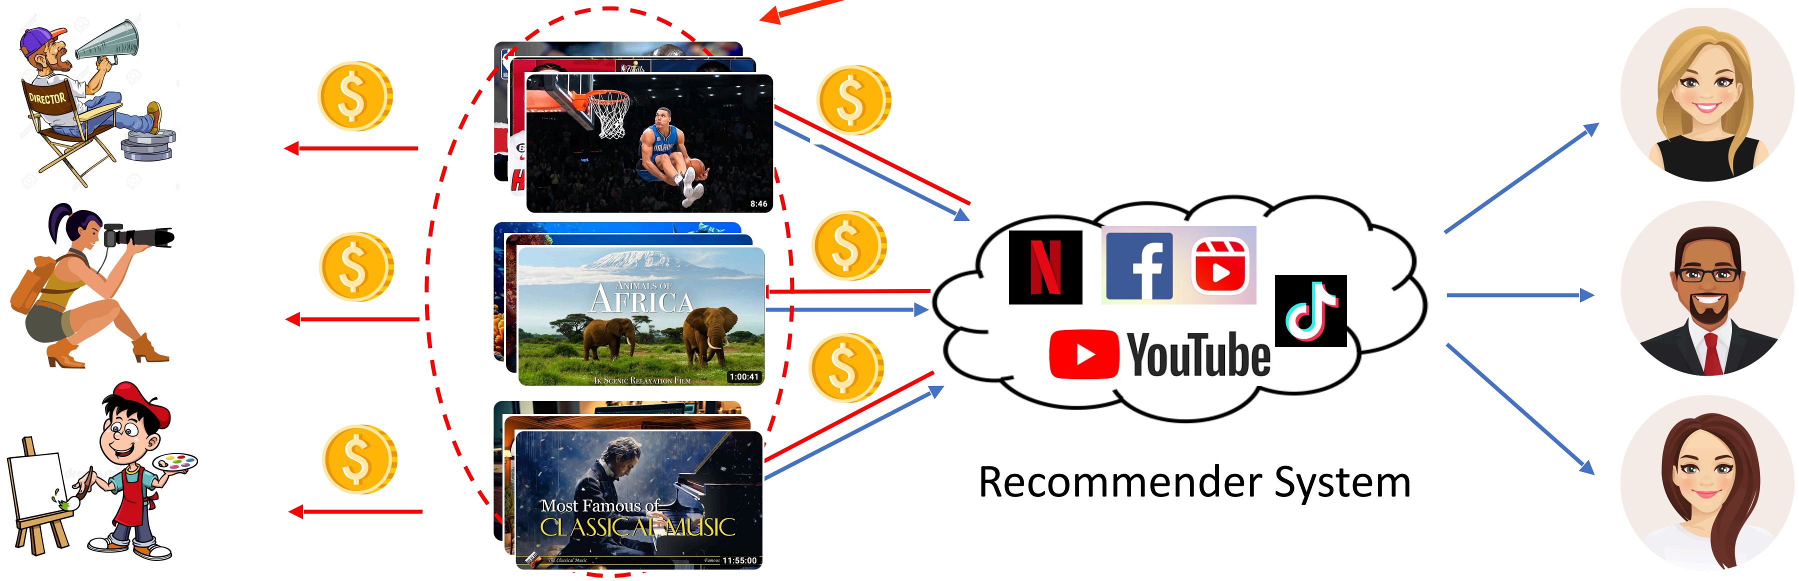

Creators
(Utility maximizing agents)

Content
Static

- **Generate personalized matching**
- **Control content supply**

Users

How can RS optimize user experience in such a strategic environment?

Model overview of Competing Content Creation (C^3)

- At each time step
  - Creators choose production strategies
  - RS observes content pool and retrieve **top-*K*** for each user
  - RS rewards creators based on user feedback
  - Creators receive their utilities and update strategies according to certain learning principle

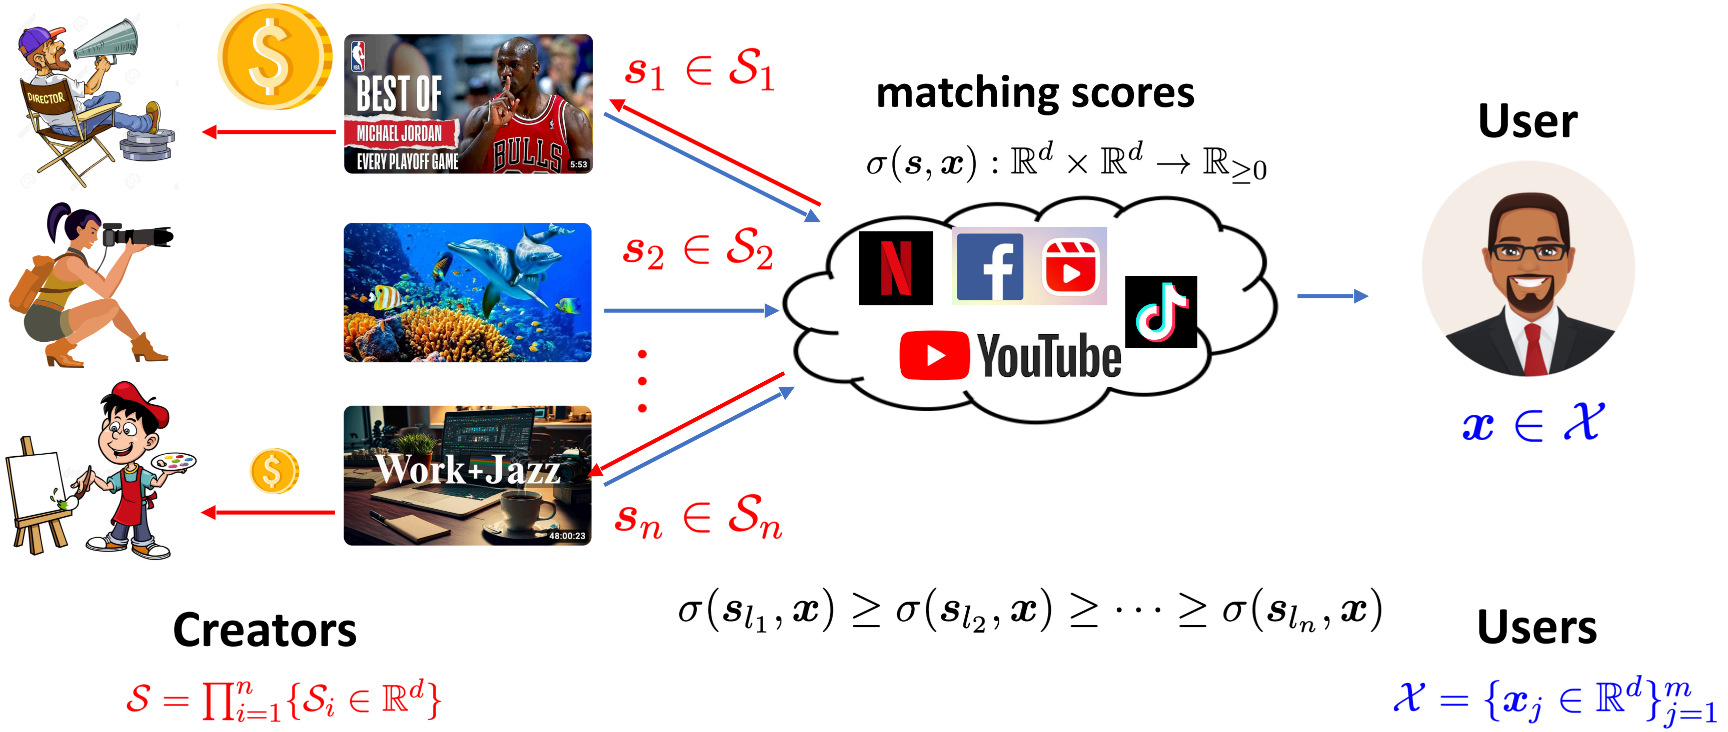

Formulation of C^3 — creator behavior

- The principle creators take to update strategies
  - No-regret learning: trial-and-error for improving his/her utility

Joint distribution of strategies at time t

$$R_i(T) = \max_{\boldsymbol{s}'_i} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{s}_{-i}\sim\boldsymbol{\alpha}^t_{-i}}[u_i(\boldsymbol{s}'_i, \boldsymbol{s}_{-i})] - \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\alpha}^t}[u_i(\boldsymbol{s})]$$

Best achievable utility for creator i   What creator i actually achieved

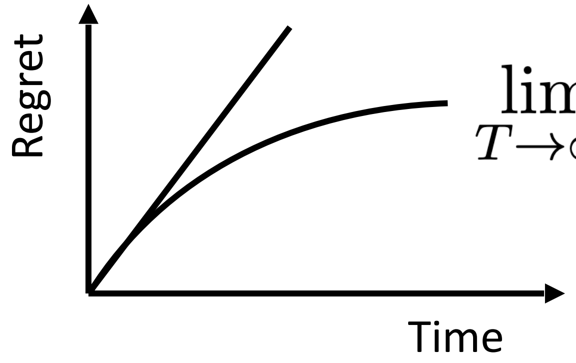

$$\lim_{T\to\infty} R_i(T)/T \to 0$$

Various efficient algorithms realize this goal in our game setting, e.g., EXP3, mirror decent, follow the regularized leader.

Formulation of C³ — user behavior

- User's choice model
  - Random utility model [Manski77]

Content utility in user's eye

$\sigma(\boldsymbol{s}_1, \boldsymbol{x}) = 0.9$

$\sigma(\boldsymbol{s}_1, \boldsymbol{x}) + \varepsilon_1$

$\sigma(\boldsymbol{s}_2, \boldsymbol{x}) = 0.6$

$\sigma(\boldsymbol{s}_3, \boldsymbol{x}) = 0.7$

$\sigma(\boldsymbol{s}_3, \boldsymbol{x}) + \varepsilon_3$

User

$\boldsymbol{x} \in \mathcal{X}$

User's choice

$$i^* = \arg\max_{\boldsymbol{s}_l \in \mathcal{T}_{\boldsymbol{x}}} \{\sigma(\boldsymbol{s}_l, \boldsymbol{x}) + \varepsilon_l\}$$

User's expected utility

$$\pi_{\boldsymbol{x}}(\boldsymbol{s}) = \mathbb{E}_{\{\varepsilon_i\}}\left[\max_{\boldsymbol{s}_l \in \mathcal{T}_{\boldsymbol{x}}} \{\sigma(\boldsymbol{s}_l, \boldsymbol{x}) + \varepsilon_l\}\right]$$

Total user welfare

$$W(\boldsymbol{s}) \triangleq \sum_{\boldsymbol{x} \in \mathcal{X}} \pi_{\boldsymbol{x}}(\boldsymbol{s})$$

Recommended top-K Slot

$$\mathcal{T}_{\boldsymbol{x}} = \{\boldsymbol{s}_1, \boldsymbol{s}_3\}$$

These are random due to the Gumbel noise $\{\varepsilon_i\}$

Formulation of C^3 — platform metrics

- Creator's utility
  - Engagement metric:

Engagement: traffic weighed by user satisfaction.

$$u_i(\boldsymbol{s}) \triangleq \mathbb{E}_{\boldsymbol{x}\in\mathcal{X},\{\varepsilon\}}\boxed{\left[(\sigma(\boldsymbol{s}_i, \boldsymbol{x}) + \varepsilon) \cdot \mathbb{I}[\boldsymbol{x} \text{ visits the creator } i]\right]}$$

- Social welfare:

$$W(\boldsymbol{s}) \triangleq \sum_{\boldsymbol{x}\in\mathcal{X}} \pi_{\boldsymbol{x}}(\boldsymbol{s}) = \sum_{i=1}^{n} u_i(\boldsymbol{s})$$

Whether the competition dynamics converge to good/bad outcome?

Main result

Theorem 1 In C^3, if creators generate content via *any* no-regret learning algorithm, we have

$$\boxed{PoA(\mathcal{G})} < 1 + \left(1 + \frac{n}{\beta \log K} \cdot \boxed{\frac{R(T)}{T}}\right) \cdot \frac{1}{c(\beta, K)}$$

goes to zero because of no-regret learning

Recommendation slots

User decision noise

where $c(\beta, K) \sim 1 + (1+\beta)\log K$.

Metric: Price of Anarchy (PoA)

$$PoA(\mathcal{G}) = \frac{\max_{\boldsymbol{s}\in\mathcal{S}} W(\boldsymbol{s})}{\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\alpha}^t}[W(\boldsymbol{s})]}$$

Maximum possible welfare

Accumulated total welfare

Main result

- *PoA* of *C*³ is upper bounded by

$$PoA(\mathcal{G}) < 1 + \frac{1}{1 + \Omega(\beta \log K)}$$

  - Less than 50% loss of efficiency for top-*K* RS!

$$PoA(\mathcal{G}) < 2, \forall \beta > 0, K \geq 1$$

  - Larger recommendation slots/explorative user behavior lead to better welfare!

$$PoA \searrow \text{ as } K \nearrow \qquad PoA \searrow \text{ as } \beta \nearrow$$

- The bound is tight

$$PoA(\mathcal{G}_0) > 1 + \frac{1}{1 + \mathcal{O}(\beta \log K)}$$

Top-K RS is not too bad, but only when…

- The welfare loss is upper bounded by *O(1/log K).*

Content creators

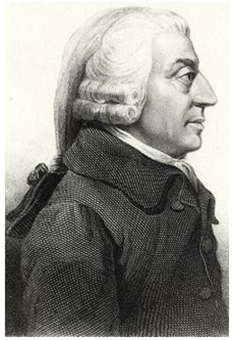

"They are led by an invisible hand to make nearly the same distribution of the necessaries of life... thus without intending it, without knowing it, advance the interest of the society, and afford means to the multiplication of the species."

-- Adam Smith, The Theory Of Moral Sentiments, 1759

- Creators follow ***no-regret learning***
- Platform provides ***right incentives***

Right incentive matters

- Our results are based on engagement-based utility functions:

$$u_i(\boldsymbol{s}) = \sum_{l=1}^{m} \underbrace{\mathbb{E}_{\{\varepsilon_i\}}\left[\sigma(\boldsymbol{s}_i, \boldsymbol{x}_l) + \varepsilon_i | \boldsymbol{x}_l \to \boldsymbol{s}_i\right]}_{\text{User's engagement/satisfaction}} \cdot \underbrace{\Pr[\boldsymbol{x}_l \to \boldsymbol{s}_i]}_{\text{Content's exposure}}$$

  - There are also systems rewarding creators only by exposure [Savy19]

$$u_i(\boldsymbol{s}) = \sum_{l=1}^{m} \Pr[\boldsymbol{x}_l \to \boldsymbol{s}_i]$$

Is it also not too bad?

NO!

Exposure metrics can be bad

Prop 1. If creators' utilities in $\mathcal{C}^3$ are set to be exposure-based, in the worse case it could happen that

$$\frac{\text{Accumulated total welfare}}{\text{Maximum possible welfare}} < \frac{1}{2}$$

Equilibrium under Exposure

Equilibrium under Engagement

$$\sigma(\boldsymbol{s}, \boldsymbol{x}) \propto 1 - \|\boldsymbol{s} - \boldsymbol{x}\|_2$$

Further improvement

- What we learned from the diagnostic result:
  - Top-*K* recommender is not too bad (O(1/log *K)* welfare loss)
  - Rewarding creators by engagement is better than by exposure.
- Can the platform design better rewarding mechanisms to make up the O(1/log *K)* fraction of loss?

New question: incentive design of RS in C^3

- Platform designs the reward (incentive) for creators

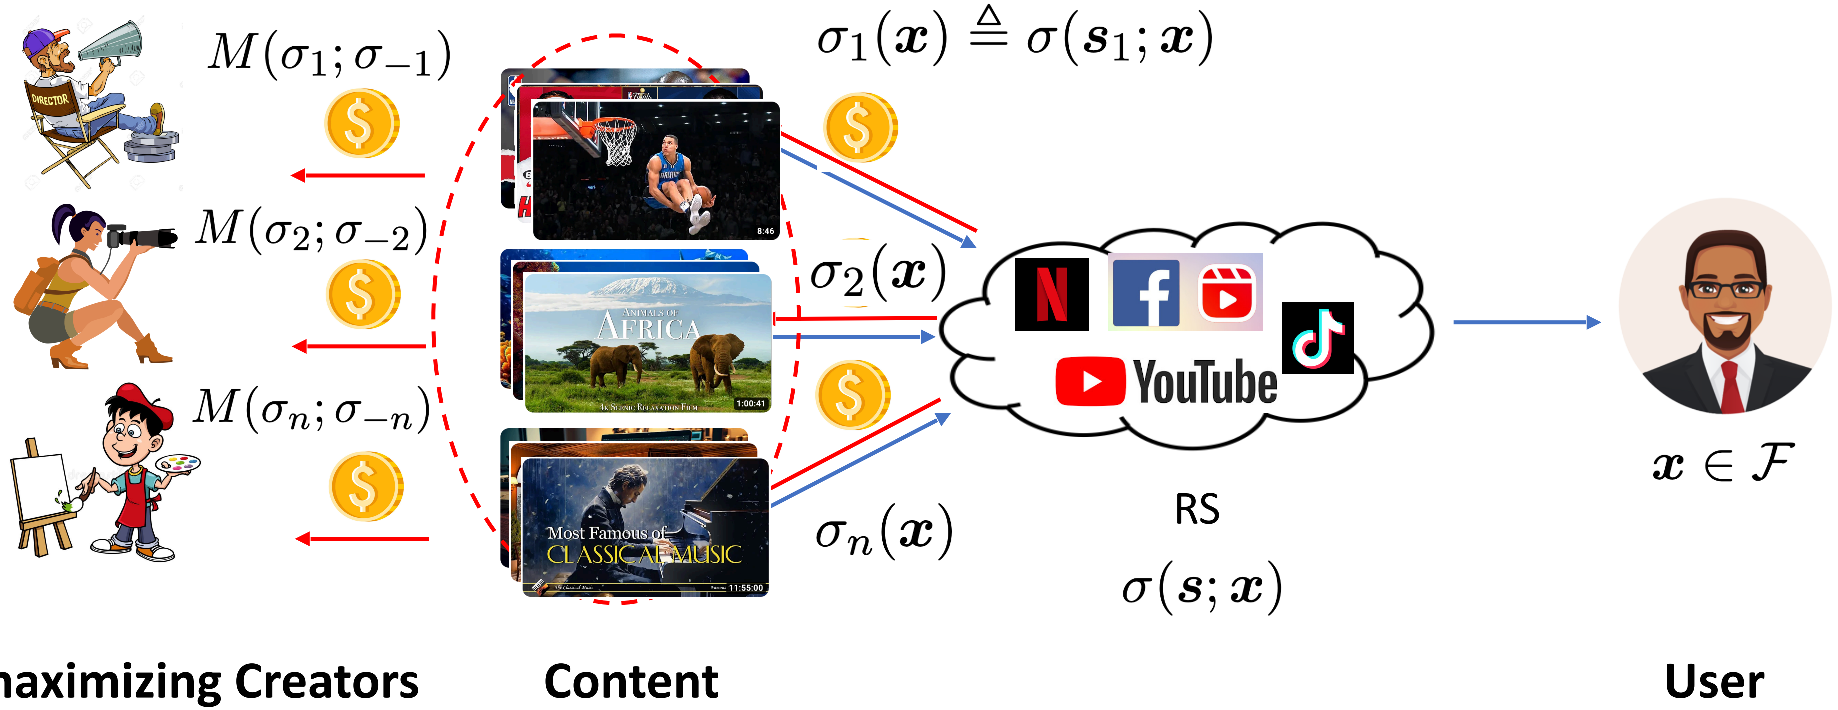

RS designs the Rewarding Mechanism *M*!

Generalized C3 with multi-choices

- Measure the *expected* user satisfaction in a session *using position bias*

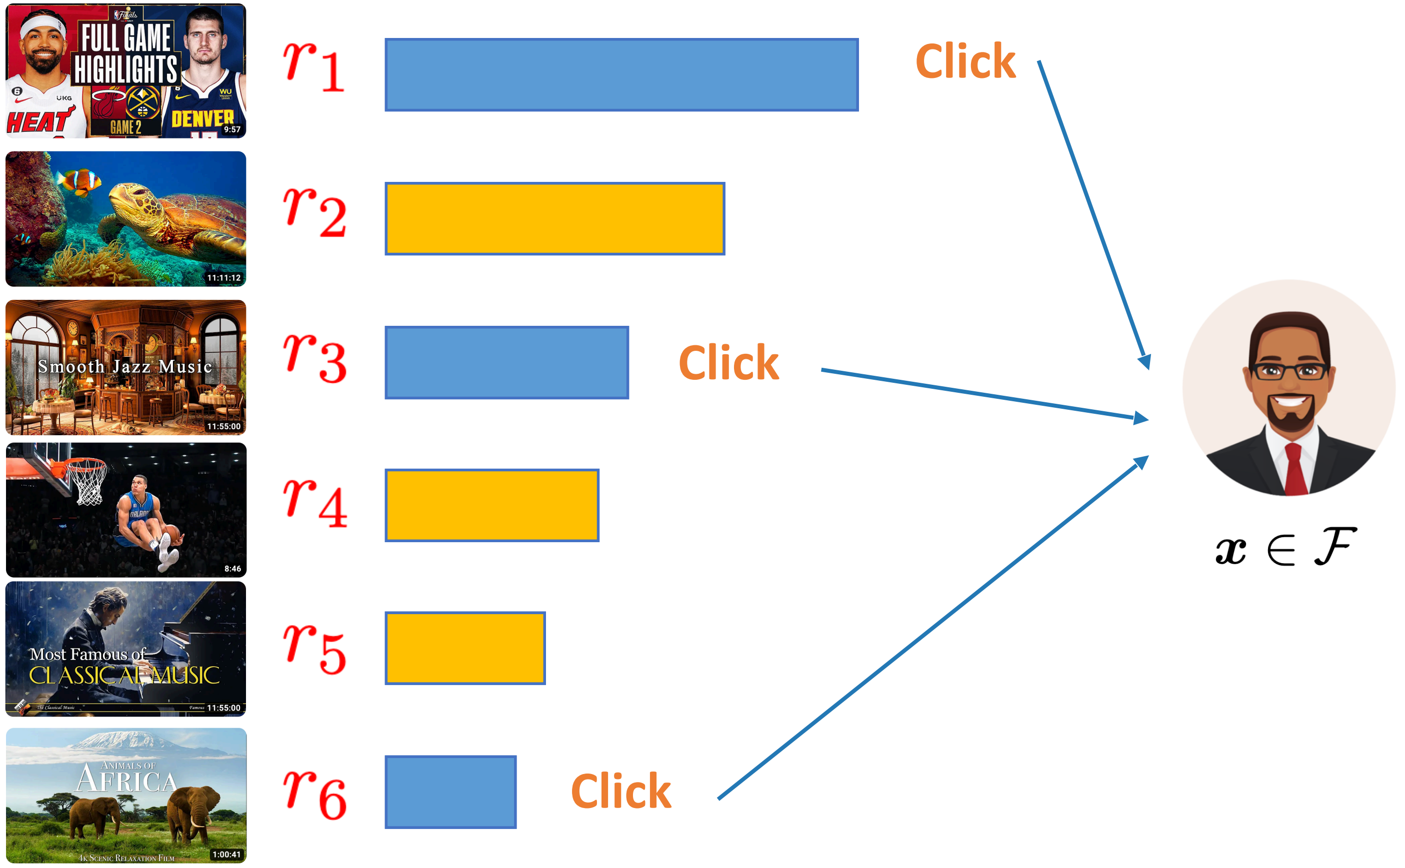

$$W(\boldsymbol{s}; \boldsymbol{x}) = \sum_{k=1}^{n} r_k \sigma_{l(k)}(\boldsymbol{x})$$

Generalized social welfare definition

- User welfare:

$$W(\boldsymbol{s};\boldsymbol{x})=\sum_{k=1}^{n} r_k\sigma_{l(k)}(\boldsymbol{x})$$

- Creator utilities:

$$u_i(\boldsymbol{s})=\mathbb{E}_{\boldsymbol{x}\in\mathcal{F}}[M(\sigma_i(\boldsymbol{x});\sigma_{-i}(\boldsymbol{x}))]-c_i(\boldsymbol{s}_i),\forall i\in[n]$$

- Social welfare:

User welfare + Creator utilities - Platform cost

$$\begin{aligned}W(\boldsymbol{s};\{r_k\}) &= \mathbb{E}_{\boldsymbol{x}\sim\mathcal{F}}[W(\boldsymbol{s};\boldsymbol{x})]+\sum_{i=1} u_i(\boldsymbol{s})-\sum_{i=1}\mathbb{E}_{\boldsymbol{x}\sim\mathcal{F}}[M(\sigma_i(\boldsymbol{x});\sigma_{-i}(\boldsymbol{x}))] \\ &= \mathbb{E}_{\boldsymbol{x}\sim\mathcal{F}}[W(\boldsymbol{s};\boldsymbol{x})]-\sum_{i=1}^{n} c_i(\boldsymbol{s}_i).\end{aligned}$$

Welfare maximization impossibility

Theorem 2 Any rewarding mechanism M that satisfies both individual-monotone and group-monotone necessarily suffers at least *1/K* fraction of welfare loss at the equilibrium in the worst case.

- The new result indicates that for a broader class of rewarding mechanisms,

$$PoA(\mathcal{G}_0) > 1 + \frac{1}{1+K}$$

- Which extended our previous result

$$PoA(\mathcal{G}_0) > 1 + \frac{1}{1+\mathcal{O}(\log K)}$$

  for the engagement-based rewarding mechanism.

Welfare maximization impossibility

Theorem 2 [Informal]. Any rewarding mechanism *M* that satisfies both individual-monotone and group-monotone necessarily suffers at least *1/K* fraction of welfare loss at the equilibrium.

- Individual-monotone
  - Fix user *x*, for each creator-*i*, higher matching score yields higher reward.
- Group-monotone
  - Fix user *x*, for all creators who have chance to be exposed to user *x*, increased matching score of any creator-*i* yields a higher collective reward.

- Many natural rewarding mechanisms satisfy both properties, e.g.,
  - (Vanilla) $M(\sigma_i;\sigma_{-i}) \propto \sigma_i$
  - Total user exposure traffic
  - Total user engagement

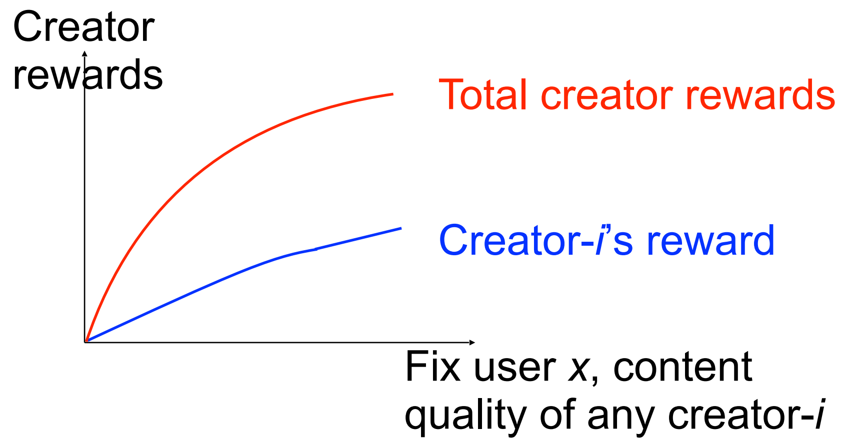

Welfare maximization impossibility

Theorem 2 [Informal]. Any rewarding mechanism *M* that satisfies both individual-monotone and group-monotone necessarily suffers at least *1/K* fraction of welfare loss at the equilibrium.

- Individual-monotone
  - Fix user *x*, for each creator-*i*, higher matching score yields higher reward.
- Group-monotone
  - Fix user *x*, for all creators who have chance to be exposed to user *x*, increased matching score of any creator-*i* yields a higher collective reward.

- We drop group-monotonicity to overcome the limitation
- Group-monotonicity is generally not satisfied in free markets, e.g., monopoly vs duopoly

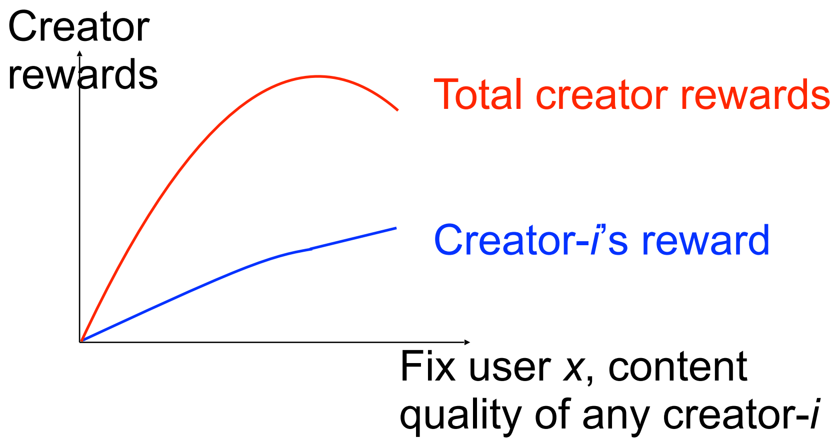

Solution: Backward Rewarding Mechanism (BRM)

We design **BRM**, a class of rewarding mechanisms that provably achieves optimal welfare.
Each instance of BRM is parameterized by a sequence of Riemann integrable functions

$$\{f_i(t) : [0,1] \to \mathbb{R}_{\geq 0}\}_{i=1}^{n}, \; s.t. \; f_1(t) \geq \cdots \geq f_n(t)$$

- Key idea: reward based on how much you are better than the next
  - Mechanism fully described by a set of functions
  - Reward = area of

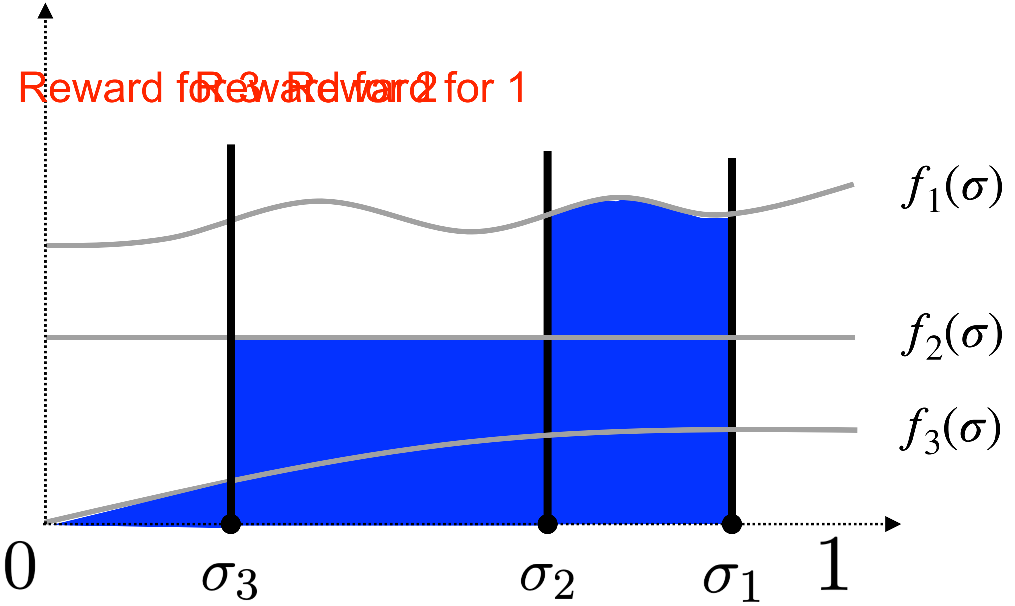

Solution: Backward Rewarding Mechanism (BRM)

We design **BRM**, a class of rewarding mechanisms that provably achieves optimal welfare.
Each instance of BRM is parameterized by a sequence of Riemann integrable functions

$$\{f_i(t) : [0,1] \to \mathbb{R}_{\geq 0}\}_{i=1}^{n}, \ s.t. \ f_1(t) \geq \cdots \geq f_n(t)$$

- Advantages
  - Compatible with the top-*K* matching principle when setting $f_{K+1} = \cdots = f_n = 0$
  - Naturally satisfies individual-monotone
  - σ_i's reward decreases when σ_{i+1}'s reward increases (competition reduces rewards)

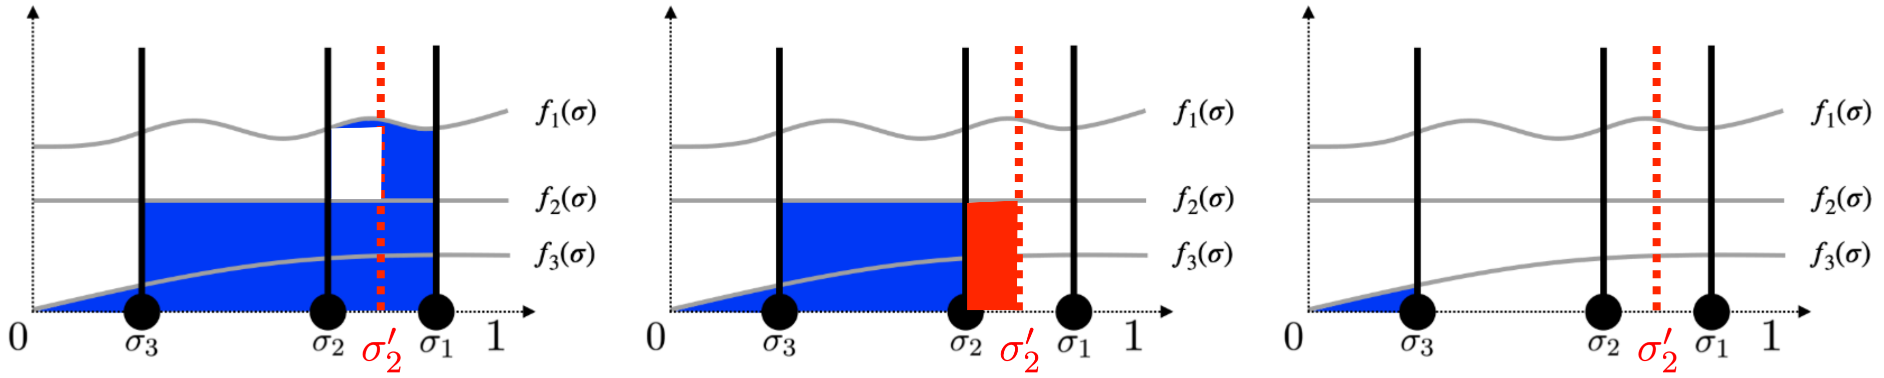

(a) Reward for σ_1 (b) Reward for σ_2 (c) Reward for σ_3

when σ_2 increases to σ_2'

Properties of BRM

Theorem 3 In C^3, BRM induces a potential game among creators, i.e., there exists a potential function P such that

$$P(\boldsymbol{s}_i', \boldsymbol{s}_{-i}) - P(\boldsymbol{s}_i, \boldsymbol{s}_{-i}) = u_i(\boldsymbol{s}_i', \boldsymbol{s}_{-i}) - u_i(\boldsymbol{s}_i, \boldsymbol{s}_{-i}).$$

- Any improvement of each creator's utility leads to an increase of a global potential function *P*!
- The global maximizer of *P* corresponds to a pure Nash equilibrium

Theorem 4 There exists a BRM instance such that the induced potential function = social welfare function, i.e., *P*=*W*.

- In fact, *P*=*W* is achieved when $[f_1, \cdots, f_n] = [r_1, \cdots, r_n]$.

Implications

- When creators selfishly optimize their utilities, their collective behavior is equivalent to maximizing some function *P*.
- If r is known, we can pick a proper BRM instance to secure a local optimum of social welfare *W*

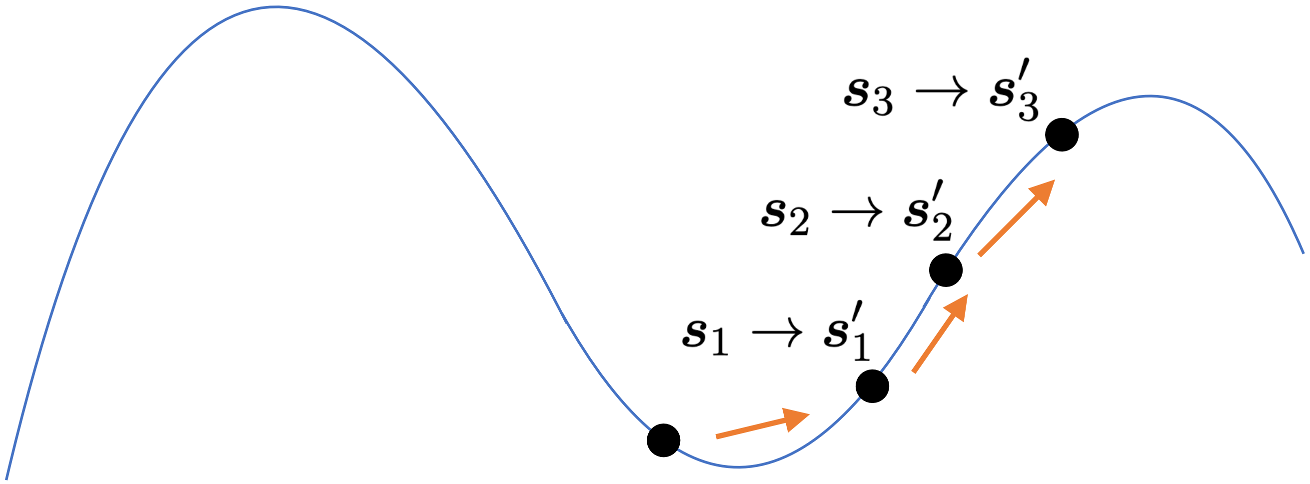

$$P(\boldsymbol{s}_1, \cdots, \boldsymbol{s}_n)$$

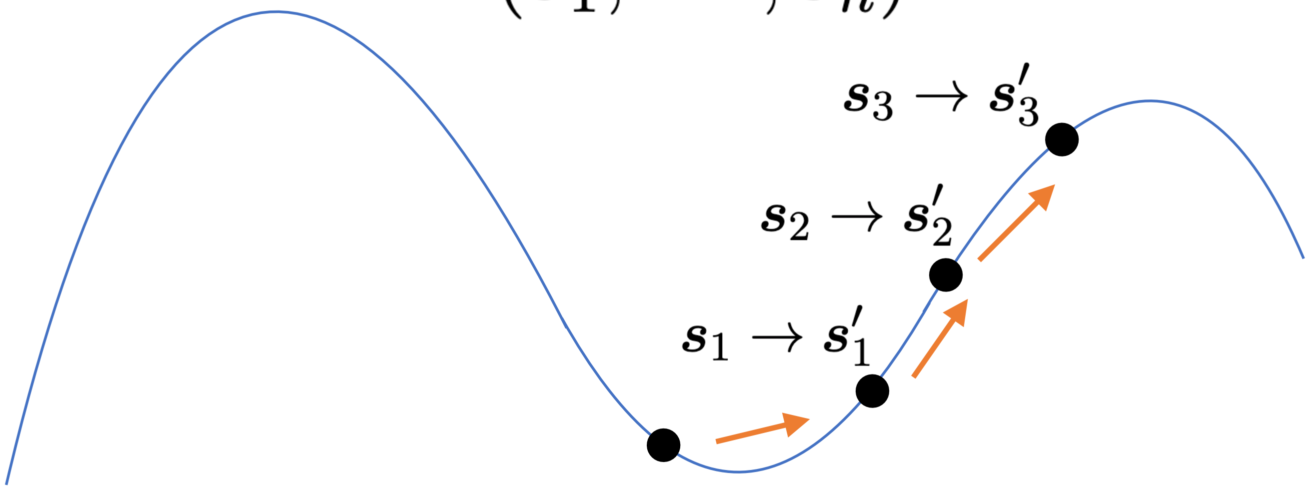

$$W = P(\boldsymbol{s}_1, \cdots, \boldsymbol{s}_n)$$

In practice

- r is usually unknown or noisy so in general *P≠W*
- We can directly optimize over the BRM space by searching for a surrogate function *P* that helps optimize *W*.

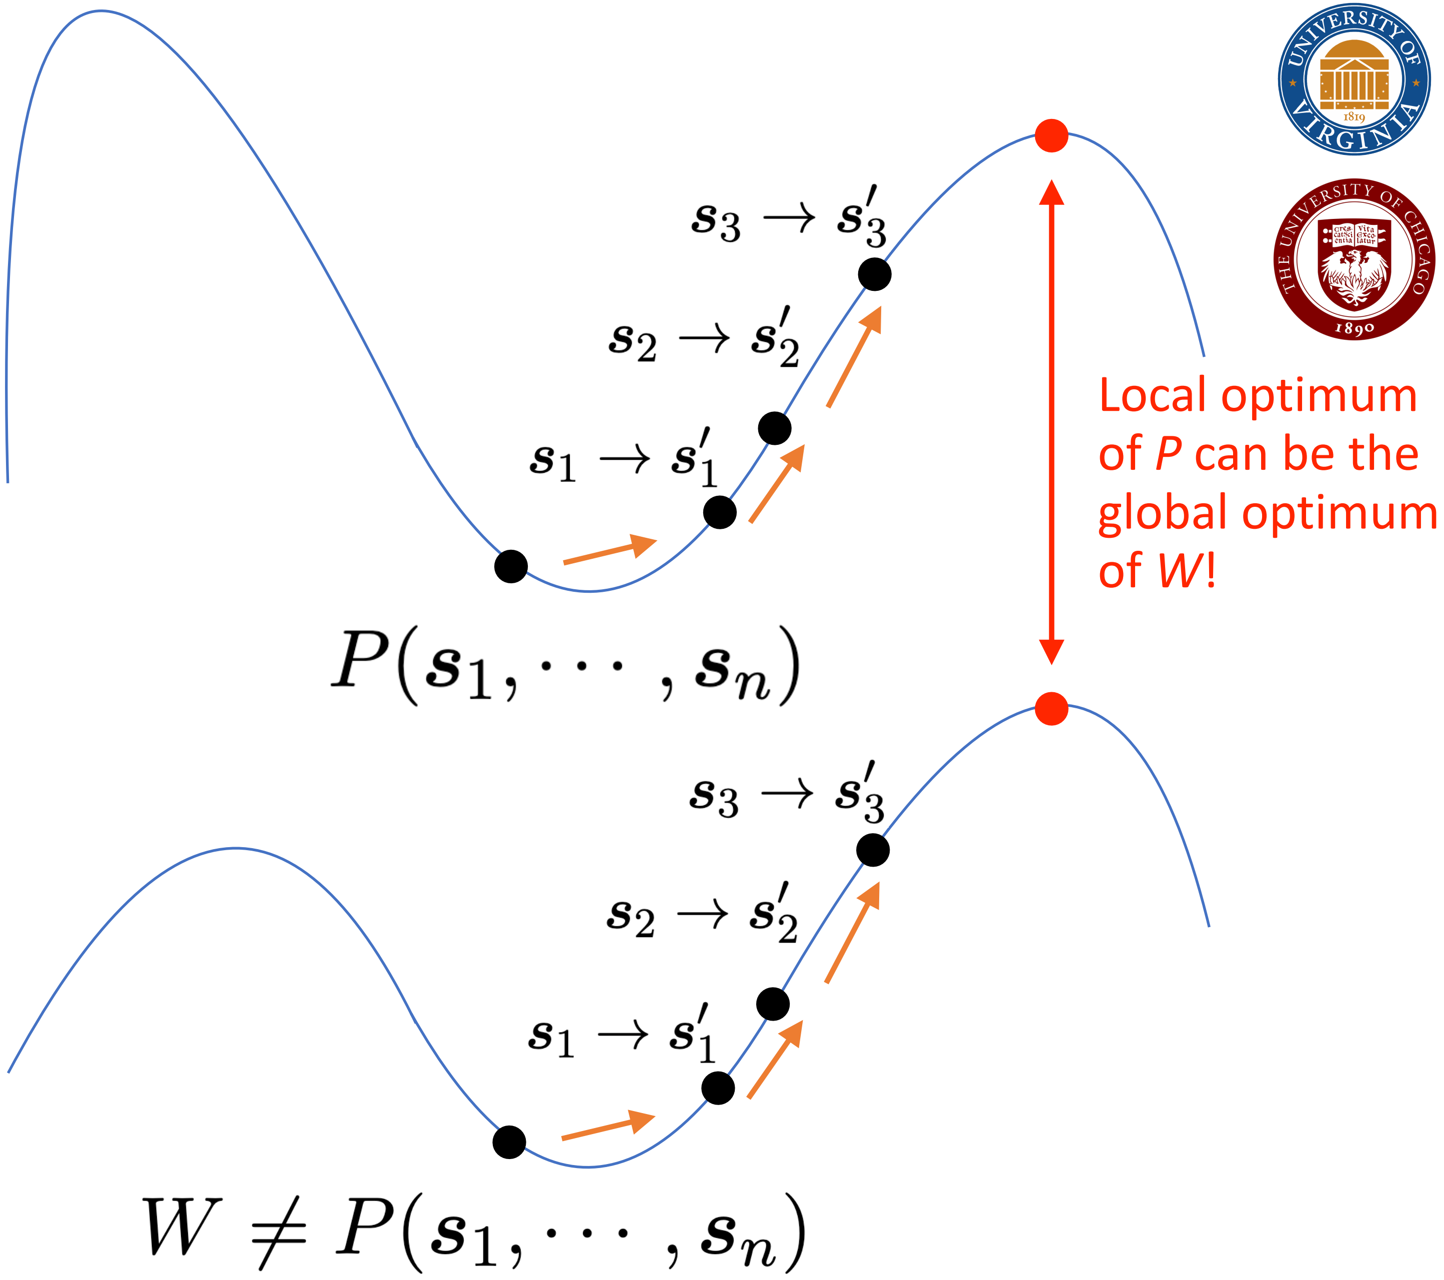

Optimize W in BRM

$$\max_{M \in BRCM} \quad W(\boldsymbol{s}^*(M))$$

$$\text{s.t.,} \quad \boldsymbol{s}^*(M) = \underset{\boldsymbol{s}}{\operatorname{argmax}}\, P(\boldsymbol{s}; M)$$

Note: we focus on BRCM, a subclass of BRM parameterized by constant functions.

- For any M⊂BRCM, it induces a potential function *P*, the resulting equilibrium among creators is *s*(M)*
- Solve it with zeroth-order optimization approaches.
    - In practice, we can directly observe *s*(M)* by letting creators evolve a period of time
    - Or we can use a proper offline simulator to estimate *s*(M)*

Empirical evaluations

- Environment constructed from MovieLens-1m data
  - Creators start at the same greedy strategy and follow projected gradient ascent
  - User population generated by Embeddings learned from the dataset

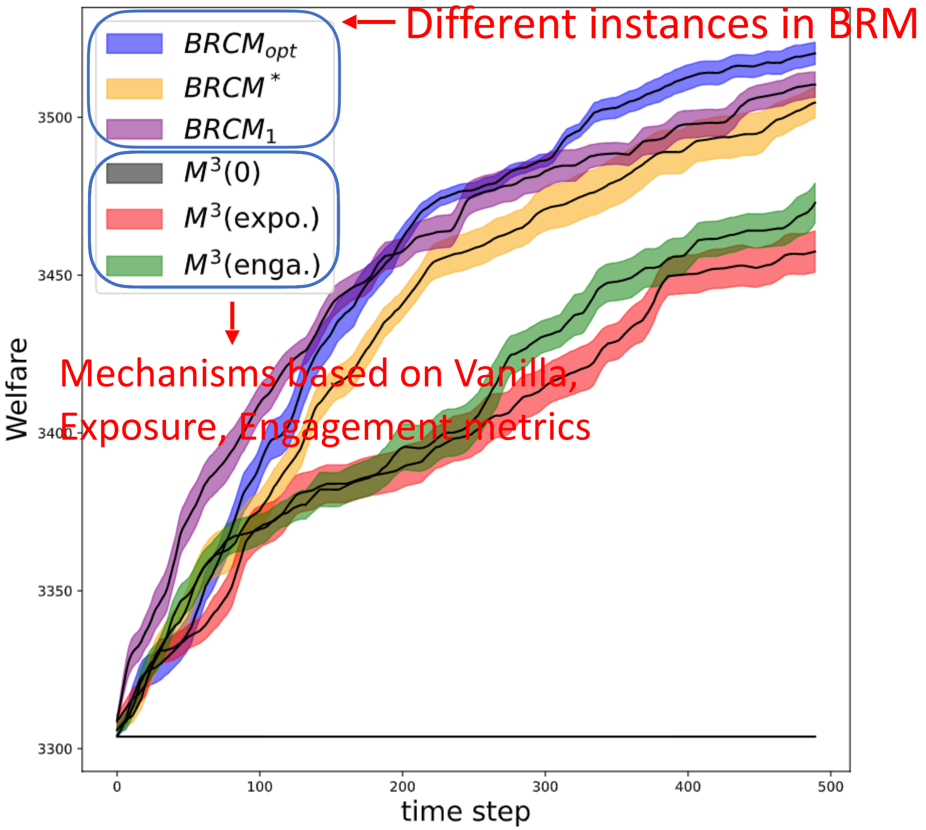

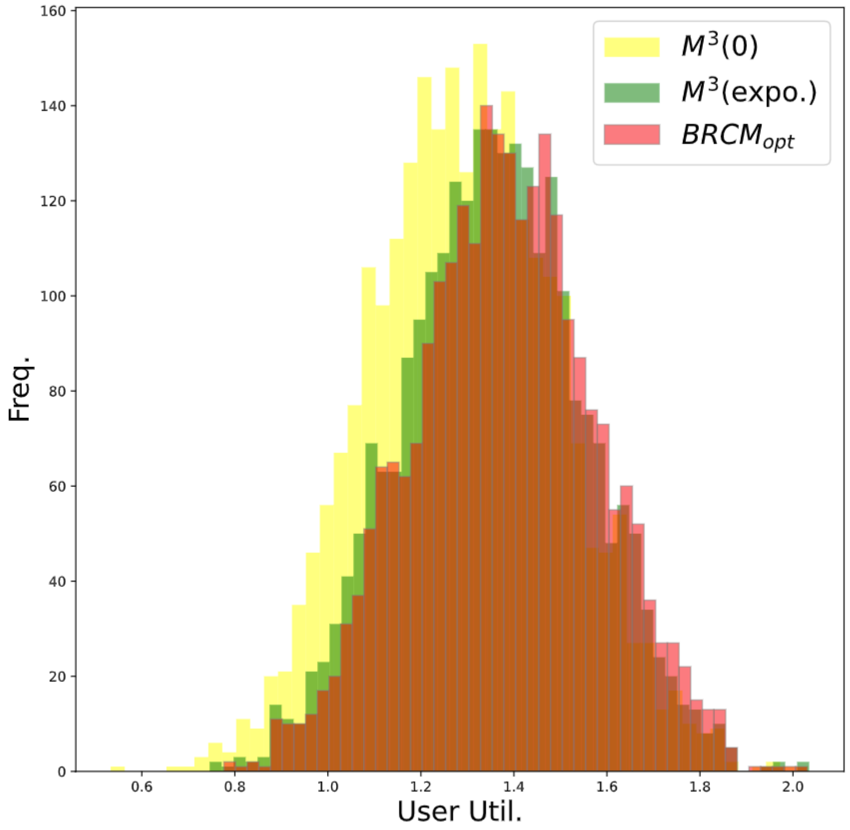

Challenges in real-world deployment

- Creators can figure out best response or achieve no-regret?

Too good to be true!

- How can we explain BRM to creators?

Platforms prefer simple merit-based rules!

Inefficiency caused by imperfect creators

- Creators do not have a holistic view of the demand distribution and can only update their strategy locally
- As a result, they might end up at a local Nash equilibrium (LNE)

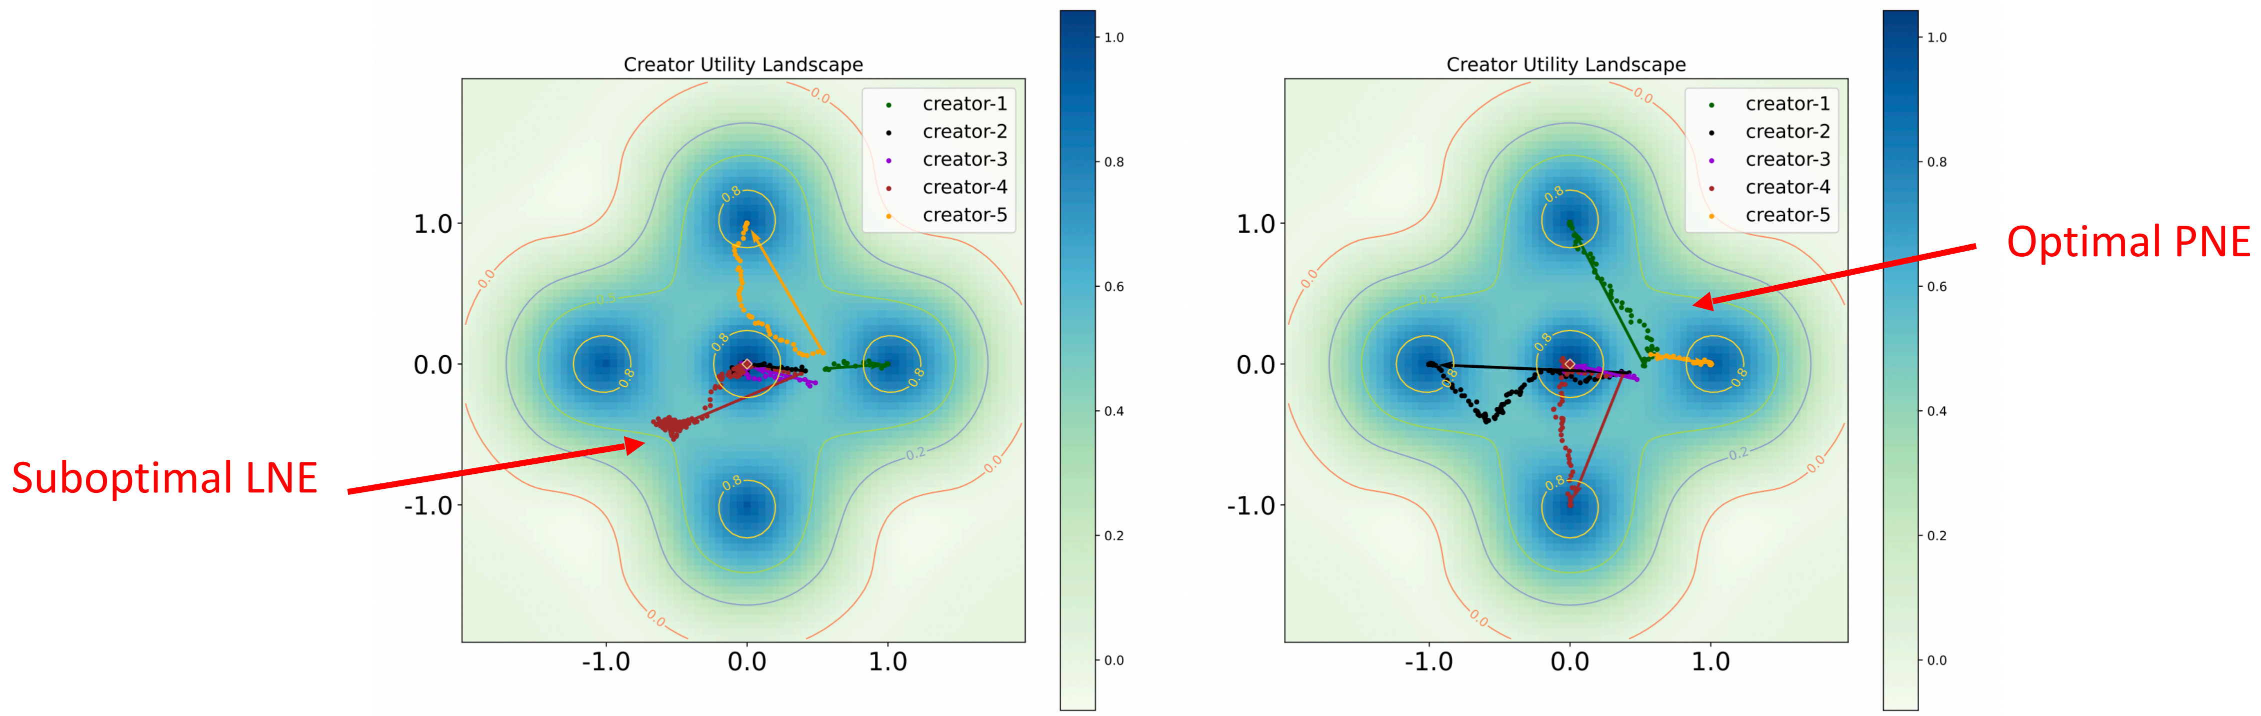

Intervention by user re-weighting

- It is the platform's responsibility to **disseminate** knowledge about users’ demand to creators
- A simple mechanism: re-weight the importance of different users

Hard to discover

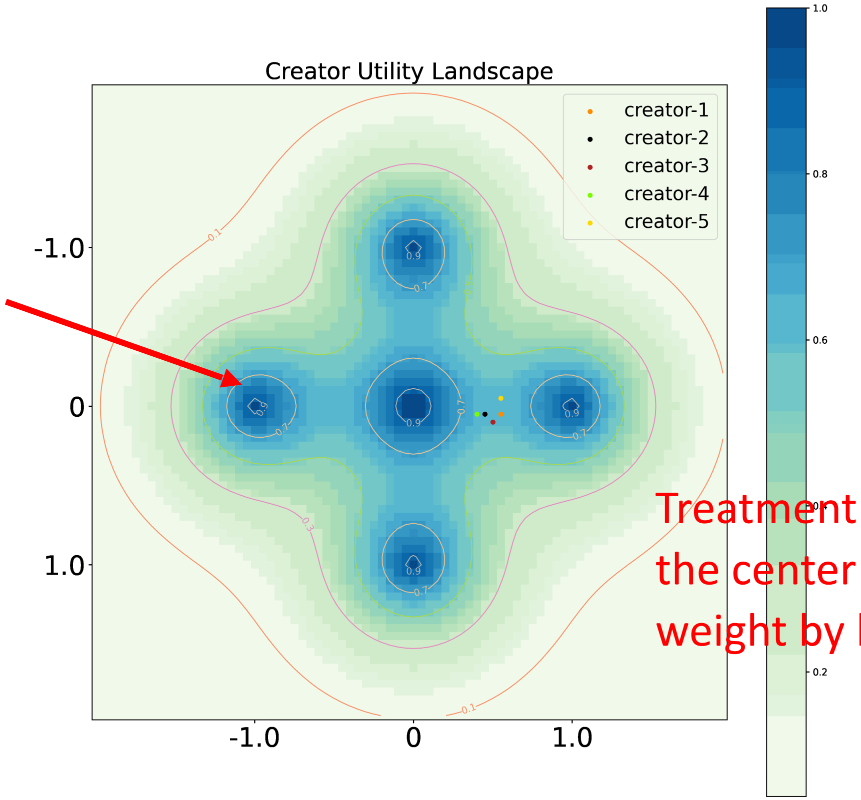

Treatment: decrease the center user’s weight by half

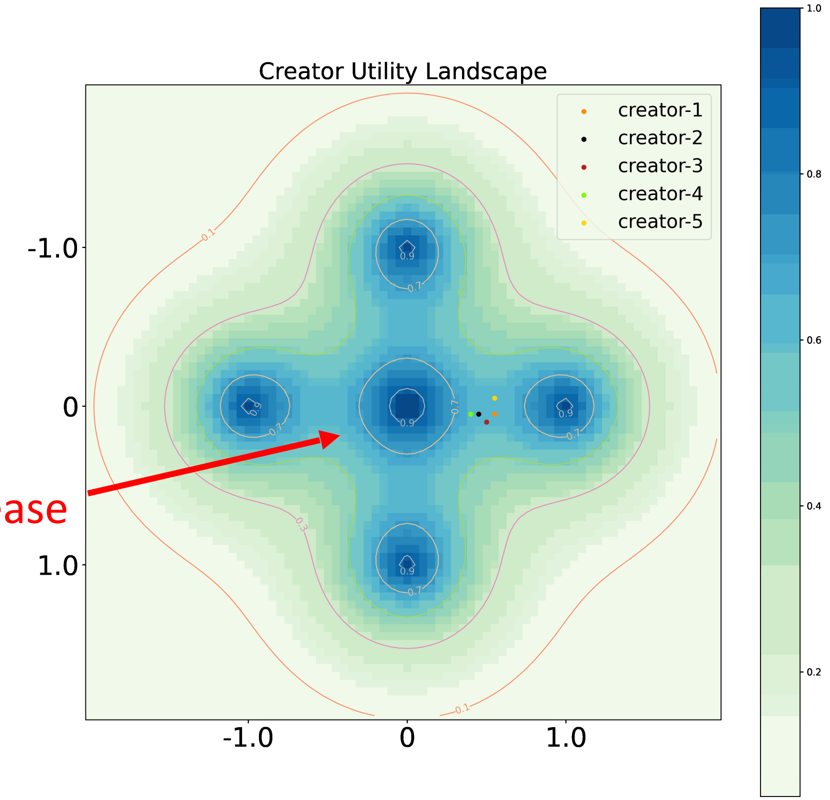

Platform's intervention mechanism

- Control what creators receive from their exposed content

$$u_i(\boldsymbol{s}) = \sum_{j=1}^{m} \mathbb{E}[\sigma(\boldsymbol{s}_i, \boldsymbol{x}_j) + \varepsilon_i | \boldsymbol{x}_j \to \boldsymbol{s}_i] \cdot \Pr[\boldsymbol{x}_j \to \boldsymbol{s}_i] \cdot w(\boldsymbol{x}_j)$$

  - User importance reweighing (UIR): when the platform possesses the flexibility to design payment incentives for creators

Finding an improving direction

Theorem 5 [Informal]. In C^3, if the number of creators is sufficiently large and the user population is perfectly separated, the following update improves the social welfare $W(s) = \mathbb{E}_{\boldsymbol{x}\sim\mathcal{X}}[\sigma(s_{\boldsymbol{x}}, \boldsymbol{x})]$ at any local Nash equilibrium s:

$$w_j' = w_j \cdot e^{-\eta\bar{\pi}(\boldsymbol{x}_j)}, \forall j \in [m],$$

where $\bar{\pi}(\boldsymbol{x}_j)$ is the expected utility of user $\boldsymbol{x}_j$ at s.

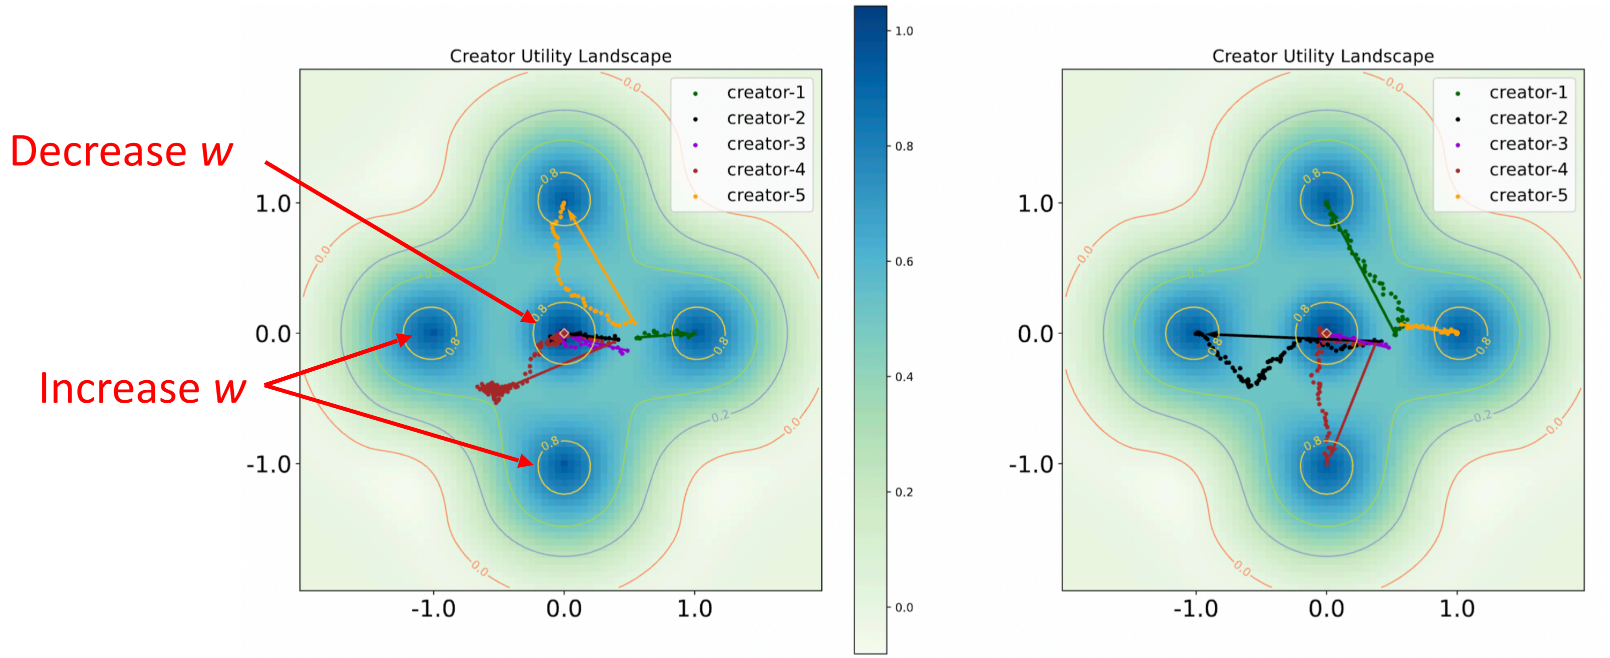

Platform's intervention mechanisms

- Welfare optimization through adaptive reweighing
  - Emphasize more on user groups who are currently under served

$$w^{(t+1)}(\boldsymbol{s}) \propto w^{(t)}(\boldsymbol{s}) \cdot \exp\left(-\alpha \bar{\pi}^{(t)}(\boldsymbol{s})\right)$$

New weights for user groups

Current utility in each user group

$$\bar{\pi}_l(\boldsymbol{s}) = \frac{1}{m|G_l|} \sum_{\boldsymbol{x} \in G_l} \sum_{i=1}^{m} \pi(\boldsymbol{x}; \boldsymbol{s}_i)$$

Evaluation on offline data

- User welfare optimization on MovieLens dataset

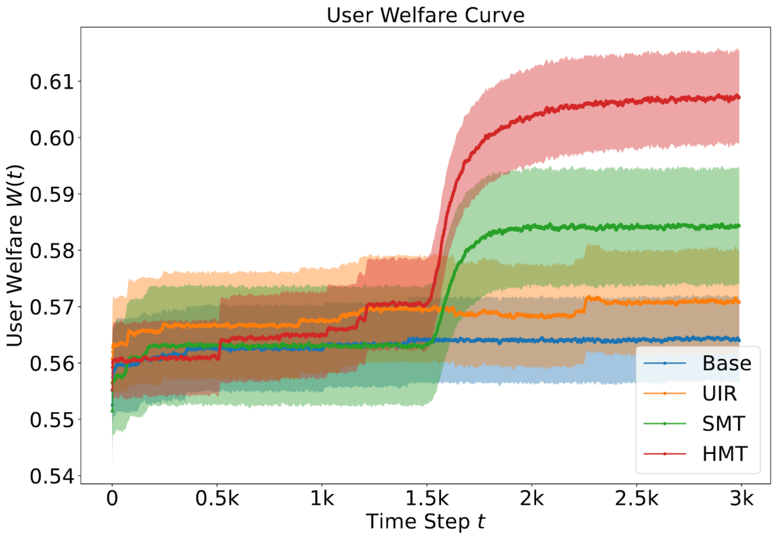

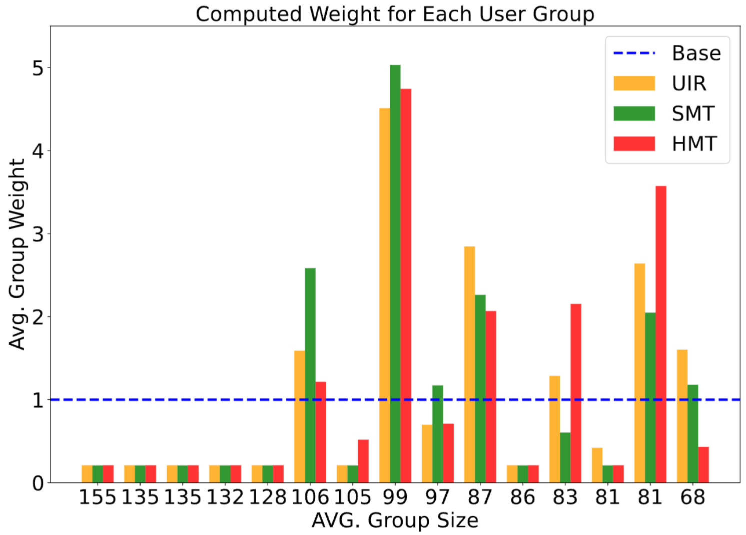

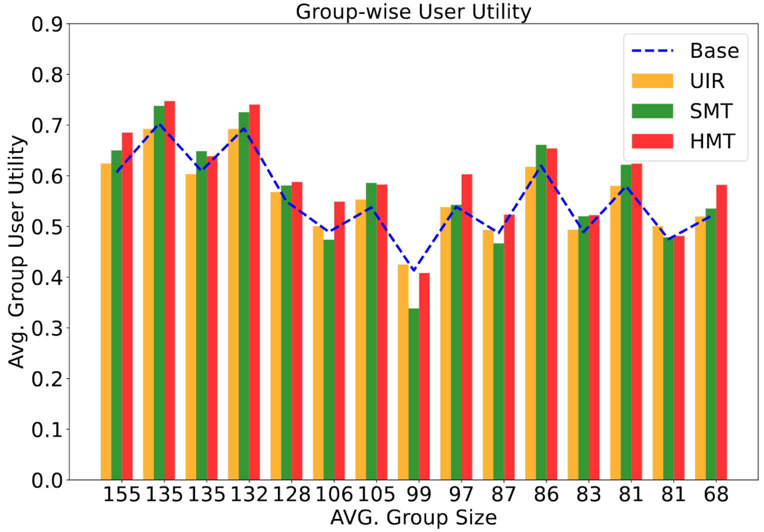

Improved total user welfare by helping creators discover users’ need

Evaluation on real traffics

- Experiment on a leading short video recommendation platform
  - welfare metric: like-through-rate
  - 3-week study
  - Symmetric A/B test: exclusively pair 3% creators with 3% users from the entire platform
  - Cluster users into 10k groups by multiple characteristics
    - Demographics: country, gender, race, occupation, etc.
    - Level of activeness: video consumption volume and watch time

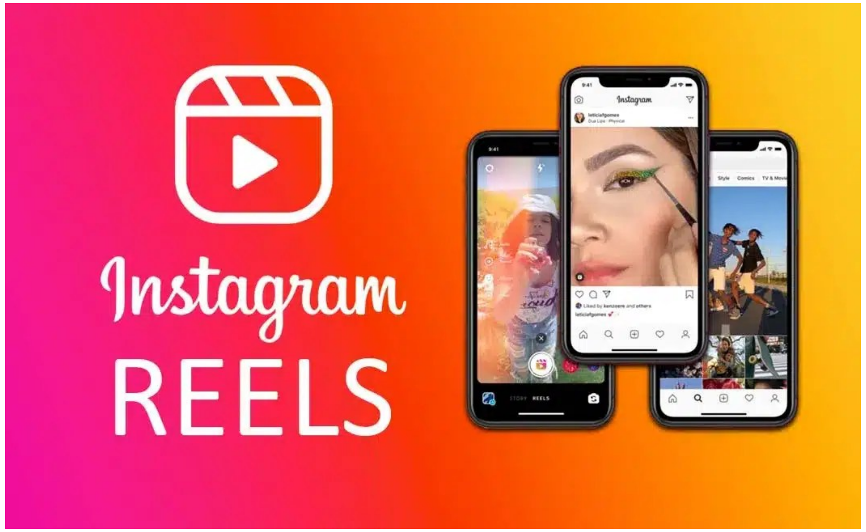

Evaluation on real traffics

- User welfare optimization targeted at Like-Through-Rate

Each cluster occupies 25% traffic

User Groups	1-5	6-20	21-74	75+	TOTAL
Like-Through-Rate	+0.43%	+1.40%	+0.75%	+1.36%	**+1.13%**
Impression	+2.64%	+0.62%	+1.42%	+0.11%	**+0.76%**

- ***3.7%*** increase in impressions on fresh content created within 2 hours
- Average number of consumed topic per user increased by ***0.71***
- An increasing trend of daily active creators
  - Head creators increased by 0.17%, others increased by 0.06%
  - 0.48% increase in the third week of experiment period

Takeaways

- What we learned
  - Welfare optimization in general is hard
  - But it is possible to design mechanisms that leads to welfare improvement
  - Human behavior modeling is a double edged sword
- Future directions:
  - Modeling the dynamics by incorporating the preference shift among users
  - How does GenAI-based creators reshape the competition dynamics

Acknowledgement

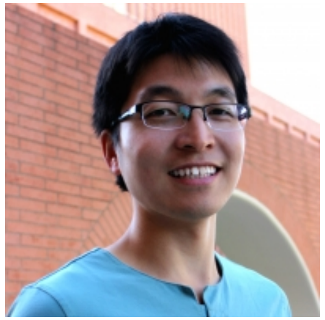


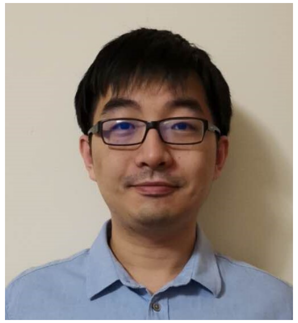

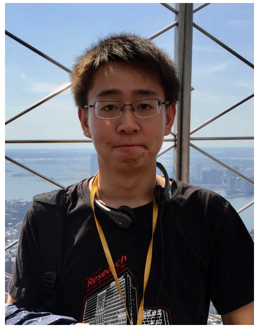

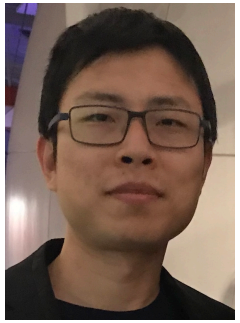


References

- Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. Computer, 42(8), 30-37.
- Hutchinson, A. (2023). Meta Expands Reels Monetization to More Creators, Updates Payout Process, SocialMediaToday
- Technavio (2023), https://www.technavio.com/report/content-marketing-market-size-industry-analysis
- Covington, P., Adams, J., & Sargin, E. (2016, September). Deep neural networks for youtube recommendations. In Proceedings of the 10th ACM conference on recommender systems (pp. 191-198).
- Kang, W. C., & McAuley, J. (2018, November). Self-attentive sequential recommendation. In 2018 IEEE international conference on data mining (ICDM) (pp. 197-206). IEEE.
- Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., & Tan, T. (2019, July). Session-based recommendation with graph neural networks. In Proceedings of the AAAI conference on artificial intelligence (Vol. 33, No. 01, pp. 346-353).
- Sharma, A., Hofman, J. M., & Watts, D. J. (2015, June). Estimating the causal impact of recommendation systems from observational data. In Proceedings of the Sixteenth ACM Conference on Economics and Computation (pp. 453-470).
- Manski, C. F. (1977). The structure of random utility models. Theory and decision, 8(3), 229.
- Fan, J., & Cheng, J. (2018). Matrix completion by deep matrix factorization. Neural Networks, 98, 34-41.
- Harper, F. M., & Konstan, J. A. (2015). The movielens datasets: History and context. Acm transactions on interactive intelligent systems (tiis), 5(4), 1-19.
- Savy. (2019) Will the new youtube algorithm impact your content?. https://savyagency.com/new-youtube-algorithm/
- McLachlan S., Cooper P. (2023) How the YouTube Algorithm Works in 2023: The Complete Guide. https://blog.hootsuite.com/how-the-youtube-algorithm-works/
- Fleder, D., & Hosanagar, K. (2009). Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. Management science, 55(5), 697-712.