Fairness in Information Retrieval from an Economic Perspective

Half-day tutorial at The Web Conference (WWW) 2026

Chen Xu[1] **Clara Rus**[3] **Yuanna Liu**[3] **Marleen de Jonge**[3] **Jun Xu**[2] **Maarten de Rijke**[3]

[1]Carnegie Mellon University, [2]Renmin University of China, [3]University of Amsterdam

June 29, 2026

chenxu2@andrew.cmu.edu, c.a.rus@uva.nl, y.liu8@uva.nl, m.r.h.dejonge@uva.nl, junxu@ruc.edu.cn, m.derijke@uva.nl

Website: https://economic-fairness-ir.github.io/

Toolkit: https://github.com/XuChen0427/FairDiverse

Tutors

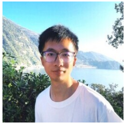

Chen Xu

Postdoc

Carnegie Mellon University

Clara Rus

Ph.D. student

University of Amsterdam

Yuanna Liu

Ph.D. student

University of Amsterdam

Marleen de Jonge

Ph.D. student

University of Amsterdam

Jun Xu

Professor

Renmin University of China

Maarten de Rijke

Professor

University of Amsterdam

Outline

1. Introduction: Fairness in IR (Yuanna, 20 min)
2. An Economic View on Fairness in IR (Chen, 30 min)
3. Economic-based Fairness Mitigation and Evaluation Strategies I (Chen, 30 min)
4. Economic-based Fairness Mitigation and Evaluation Strategies II (Clara, 30 min)
5. Economic-based Fairness Mitigation and Evaluation Strategies III (Yuanna, 30 min)
6. Open Problems, Quick Start for Learning about Fairness, and Conclusions (Clara, 20 min)

Narrative arc: **Problem** (1) → **Economic lens** (2) → mitigation along three allocation axes: **objective** (3), **scale** (4), **time** (5) → **open problems** (6).

Motivation

1. Economics provides good fairness frameworks and tools

- Economists have studied complex fairness problems for centuries. Their theory and methods can help us to structure the IR fairness problems better

Motivation

1. **Economics provides good fairness frameworks and tools**
   - Economists have studied complex fairness problems for centuries. Their theory and methods can help us to structure the IR fairness problems better
2. **Leveraging economic thinking for fairness in IR**
   - Economic theory shows that fairness is not just "the right thing to do" but often also the "**profitable choice**"

Motivation

1. **Economics provides good fairness frameworks and tools**
   - Economists have studied complex fairness problems for centuries. Their theory and methods can help us to structure the IR fairness problems better
2. **Leveraging economic thinking for fairness in IR**
   - Economic theory shows that fairness is not just "the right thing to do" but often also the "**profitable choice**"
3. **Economic perspectives point out future directions**
   - Economics highlights that we need to consider practical multi-agent scenarios and develop more rigorous, theory-driven fairness mechanisms

Overall Framework

Existing Fairness Foundations (Part I)

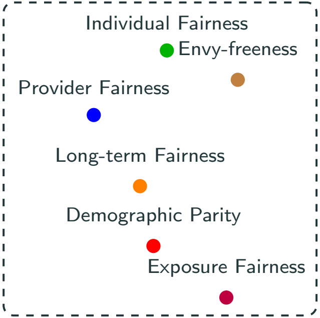

Economic Perspective on Fairness (Part II)

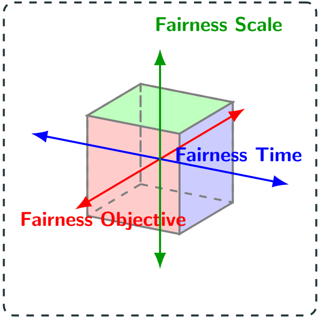

Future Fairness Directions (Part III)

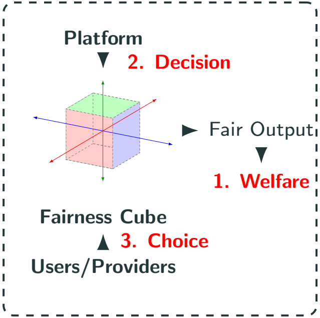

1. Introduction: Fairness in IR (Yuanna, 20 min)

Information retrieval

What is information retrieval?

- **Information retrieval (IR)** is the process of getting the right information to the right people in the right way and at the right time
- It focuses on matching user queries with documents or data items
- IR is the core technology behind **search engines** and **recommender systems**

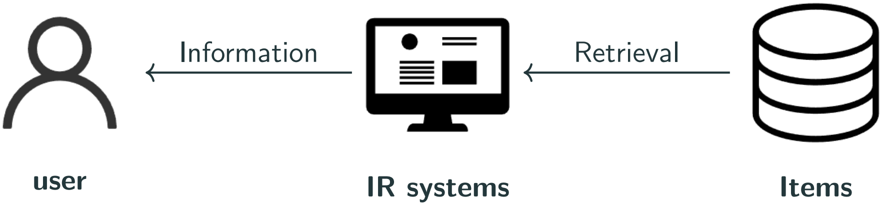

Core components

1. **Document/item collection** – Large repository of data (e.g., web pages, products)
2. **Indexing** – Efficient representation for fast search
3. **User intent understanding** – Understanding and interpreting user queries
4. **Retrieval and ranking models** – Scoring documents based on relevance
5. **Evaluation** – Measuring quality

IR is more than accuracy

- Traditional IR systems aim to maximize **ranking accuracy**

Traditional: User-centric **Now: Ecosystem-centric**

IR is more than accuracy

- Traditional IR systems aim to maximize **ranking accuracy**
- However, real-world IR systems operate in a complex **ecosystem** involving many stakeholders, such as content creators, advertisers, and even society

Traditional: User-centric **Now: Ecosystem-centric**

IR is more than accuracy

- Traditional IR systems aim to maximize **ranking accuracy**
- However, real-world IR systems operate in a complex **ecosystem** involving many stakeholders, such as content creators, advertisers, and even society
- Sustainable and responsible IR must consider all stakeholders and long-term system dynamics

Traditional: User-centric **Now: Ecosystem-centric**

Key stakeholders in IR

1. **User**
   - Seeks relevant, timely, and useful content
   - User satisfaction directly impacts system reputation

2. **Platform**
   - Operates and optimizes the IR system
   - Acts as a mediator between users and providers

3. **Provider**
   - Supplies the content or items retrieved by the system (e.g., sellers, content creators)
   - Interested in exposure, traffic, and conversions

Stakeholder interactions in IR

- **User**, **platform**, and **provider** form a dynamic ecosystem [Abdollahpouri and Burke, 2019]
- Each stakeholder has different goals and influences the system
- Balancing the goals of each stakeholder means **fairness**

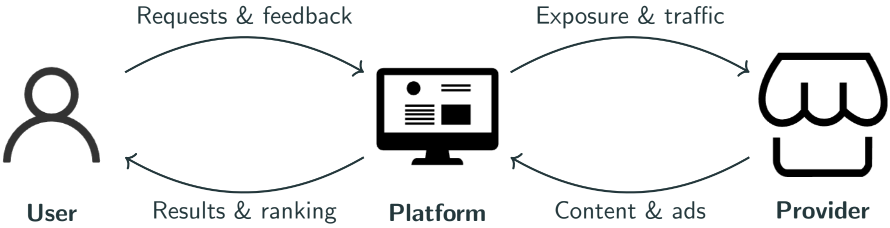

Fairness in IR

What is “beyond accuracy” in IR?

- **Definition:** Beyond-Accuracy in IR refers to a class of evaluation and modeling approaches that go beyond traditional relevance-based metrics, aiming to account for broader user and societal values

Key dimensions beyond accuracy:

- **Fairness:** Ensuring equitable or right outcomes across different groups
- **Diversity:** Promoting varied content to reduce redundancy
- **Novelty:** Encouraging discovery of unexpected but useful items
- **Transparency:** Providing users with understandable reasons behind rankings
- ...

What is fairness?

Fairness refers to the quality of treating **people** equally or in a way that is **right or reasonable** – *Cambridge Dictionary*

What is fairness?

Fairness refers to the quality of treating **people** equally or in a way that is **right or reasonable** – *Cambridge Dictionary*

Fairness has been defined in **numerous ways** across history and disciplines – from justice in social science and humanities to algorithmic fairness in IR

Taxonomy of fairness in social science

1. **Distributive justice [Lamont, 2017]**
   - Are resources (e.g., income) distributed fairly among individuals or groups?
2. **Procedural justice [Tyler and Allan Lind, 2002]**
   - Is the decision-making process transparent, consistent, and unbiased?
3. **Recognition and inclusion [Eisenstadt, 1973]**
   - Are marginalized groups fairly represented and respected?

Unfairness as harms

Unfairness often leads to **harm** by systematically disadvantaging certain individuals or groups, thereby reinforcing inequality and reducing overall welfare

Fairness in social science vs. Fairness in machine learning

Fairness in social science	Fairness in IR
Distributive justice	**Allocation harms**: How to allocate resources (e.g., computational costs, user traffic) fairly for different stakeholders? [Xu et al., 2023a]
Procedural justice	**Procedural harms**: How can we ensure models do not rely on discriminatory or harmful information when making decisions? [Lee et al., 2019]
Recognition and inclusion	**Representation harms**: How can we ensure that the model fairly represents different groups in its latent (hidden) space? [Zemel et al., 2013]

Taxonomy of fairness in IR

Allocation harms

- Individual-group fairness [Jiang et al., 2021]
- User-provider fairness [Xu et al., 2023a]
- Short-long term fairness [Xu et al., 2023b]

Procedural harms

- Controllable fairness [Lee et al., 2019]
- Explainable fairness [Ge et al., 2022]
- Transparent fairness [Lee et al., 2019]

Representational harms

- Anti-classification [Rus et al., 2023, 2024]
- Anti-subordination [Lahoti et al., 2019]

Taxonomy of fairness in IR

- **Procedural harms**

  $\Rightarrow$ Reflect constraints or flaws in the process

  $\Rightarrow$ But they matter because they are the right properties to address allocation harms

Taxonomy of fairness in IR

- **Procedural harms**
  ⇒ Reflect constraints or flaws in the process
  ⇒ But they matter because they are the right properties to address allocation harms
- **Representational harms**
  ⇒ In IR, often act as *means* to an unfair allocation
  ⇒ Not always the final objective

What do we focus on?

- In IR, we mainly focus on **allocation harms**. This is because:
  - **Allocation harm** is the central concern in IR: *Who gets ranked, recommended, or seen – and how much?*
  - Ranking slots and user traffic are scarce and impactful resources

Fairness origin and development

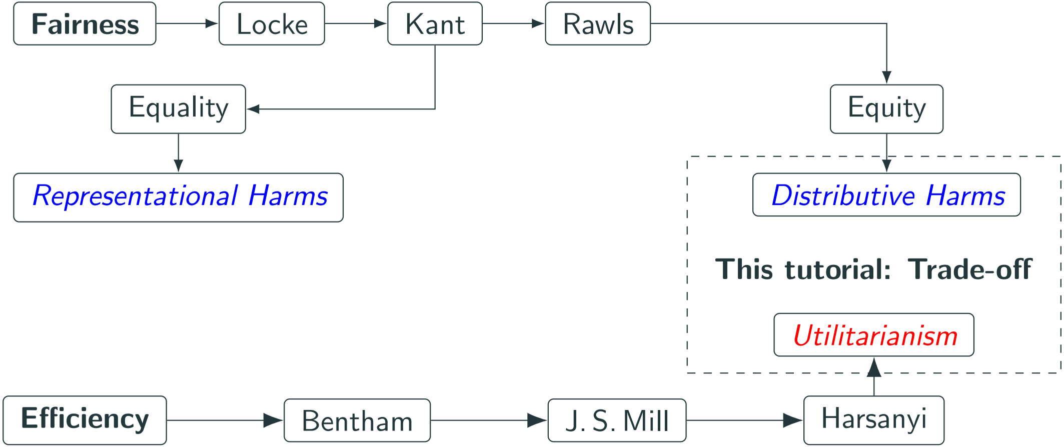

What are resources in allocation harms?

- The resource allocated in the IR could be
  - The number of item/document exposures [Xu et al., 2023a]
  - The number of item/document clicks [Xu et al., 2024, Baumann et al., 2024]
  - The utilities of user groups [Liu et al., 2024]
- The resources in IR are typically **limited** (limited ranking slots and user traffic)

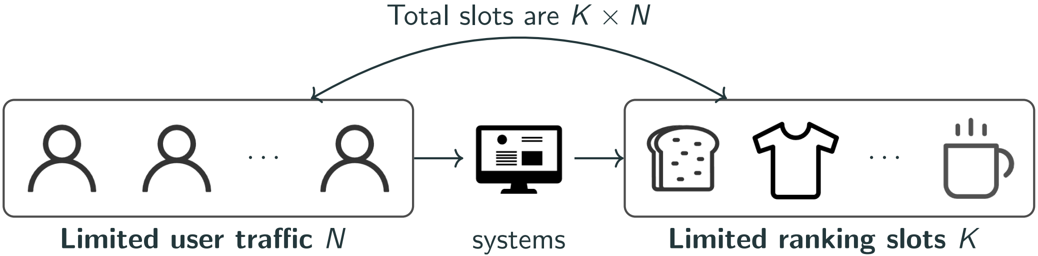

Allocation harms in IR

- Assuming N users $(u_1, u_2, \cdots, u_N)$
- Assuming M items/documents $(i_1, i_2, \cdots, i_M)$.
- IR systems can only adjust the **slots allocation matrix** X

	u_1	u_2	$\cdots$	u_N
i_1	●	●	●	
i_2	●	●	●	
i_3	●			●
$\cdots$		●		●
i_M			●	●

Ranking size $K = 3$

i_M is exposed to u_N

Allocation harms in IR

- Based on the IR resource allocation, we can define the utilities of different stakeholders, such as user groups:

user group fairness

	u_1	u_2	$\cdots$	u_N
i_1	• – 0.8	• – 0.6	• – 0.5	
i_2	• – 0.7	• – 0.7	• – 0.6	
i_3	• – 0.5			• – 0.7
$\cdots$		• – 0.8		• – 0.5
i_M			• – 0.4	• – 0.2

Utilities of user group $1 = 2.05$

$(0.8 + 0.7 + 0.5 + 0.6 + 0.7 + 0.8)/2$

Utilities of user group $n = 1.4$

$0.7 + 0.5 + 0.2$

Allocation harms in IR

- Based on the IR resource allocation, we can define the utilities of different stakeholders, such as providers:

	u_1	u_2	$\cdots$	u_N
i_1	• – 0.8	• – 0.6	• – 0.5	
i_2	• – 0.7	• – 0.7	• – 0.6	
i_3	• – 0.5			• – 0.7
$\cdots$		• – 0.8		• – 0.5
i_M			• – 0.4	• – 0.2

Provider 1
utility $= 1.95$

provider fairness

Provider m
utility $= 0.6$

Evolution of Fairness Algorithms in IR

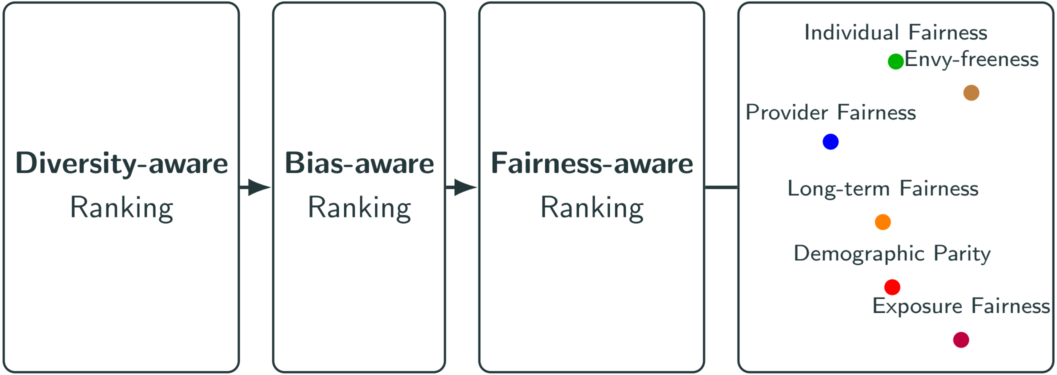

Fairness evaluation in IR

How to measure allocation harms?

- Assuming the utilities (such as exposures) of one stakeholder are

$$\boldsymbol{v} = [v_1, v_2, \cdots, v_g],$$

  where g is the stakeholder internal group number.
- A fairness evaluation function $f(\boldsymbol{v})$ is designed to measure the degree of fairness
- An example:

$$\boldsymbol{v}_1 = [1, 5, 10, 20], \quad \boldsymbol{v}_2 = [2, 4, 12, 18].$$

  How much less fair is $\boldsymbol{v}_1$ compared to $\boldsymbol{v}_2$?

Common evaluation metrics I

- Max-min fairness [Xu et al., 2023a]: ensures worst-off groups get enough utilities

$$f(\boldsymbol{v}) = \min_i (v_i).$$

- Gini Index [Do et al., 2021]: inequality by quantifying distribution disparity

$$f(\boldsymbol{v}) = \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} |v_i - v_j|}{2n \sum_{i=1}^{n} v_i}.$$

- Entropy [Jost, 2006]: captures overall diversity or uncertainty in allocation

$$f(\boldsymbol{v}) = -\sum_{i=1}^{g} v_i \log(v_i).$$

- Demographic parity [Jiang et al., 2021]: equal outcomes across groups

$$f(\boldsymbol{v}) = \sum_{i=1}^{g} |v_i - \sum_{i=1}^{g} v_i / g|.$$

Common evaluation metrics II

- Min-max ratio [Jain et al., 1984]: ratio between the best-off and worst-off groups

$$f(\boldsymbol{v}) = \min_i(v_i) / \max_i(v_i)$$

- p-norm [Bektaș and Letchford, 2020]: penalizing large deviations in utility

$$f(\boldsymbol{v}) = \left(\sum_{i=1}^{g} v_i^p\right)^{1/p}$$

- Elastic fairness [Xu et al., 2025c]: a unified fairness evaluation metric

$$f(\boldsymbol{v}) = \operatorname{sign}(1-t)\left(\sum_{i=1}^{g} \bar{\boldsymbol{v}}_i^{1-t}\right)^{(1/t)}$$

Goals of fairness-aware IR

The goal is to enforce **fairness** across stakeholders while preserving the **effectiveness and relevance** of the information retrieval process

2. An Economic View on Fairness in IR (Chen, 30 min)

Fairness concept history

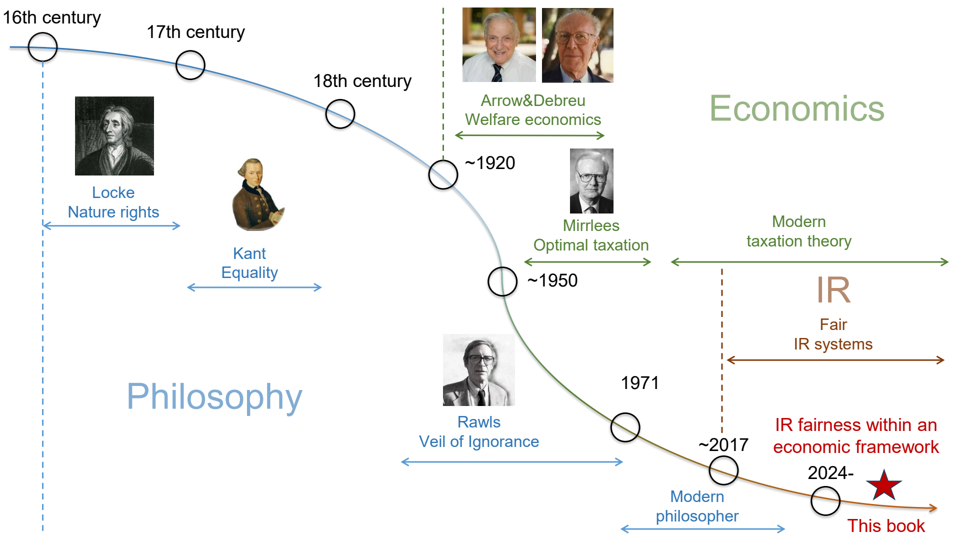

Motivation for an economic view on fairness in IR

An economic lens does not just add more complexity, more methods, and more theories: it integrates different stakeholders and justifies its relevance

- The current landscape:
  - **Fragmented objectives:** Many well-defined criteria, but no common framework for comparing them
  - **Translation gap:** Hard to quantify what fairness costs and delivers
  - **Limited generalization:** Solutions designed per setting, principles don't always transfer

Motivation for an economic view on fairness in IR

An economic lens does not just add more complexity, more methods, and more theories: it integrates different stakeholders and justifies its relevance

- The current landscape:
  - **Fragmented objectives:** Many well-defined criteria, but no common framework for comparing them
  - **Translation gap:** Hard to quantify what fairness costs and delivers
  - **Limited generalization:** Solutions designed per setting, principles don't always transfer
- Without a shared theoretical language:
  - Harder to compare approaches across the literature
  - Harder to justify sustained investment in production
  - Harder to transfer insights across settings

An economic view on information retrieval

What is a market economy?

"A market economy is an economic system in which the production and allocation of goods and services is determined by forces of **supply and demand***."*

IR systems and economic markets: A natural analogy

- IR systems do the same:
  - **Users** are the *demand side* – they express preferences through queries, clicks, engagement.
  - **Providers** are the *supply side* – they create content and compete for visibility.
  - The **platform** allocates the scarce resource – *user attention* – guided by relevance scores that function.

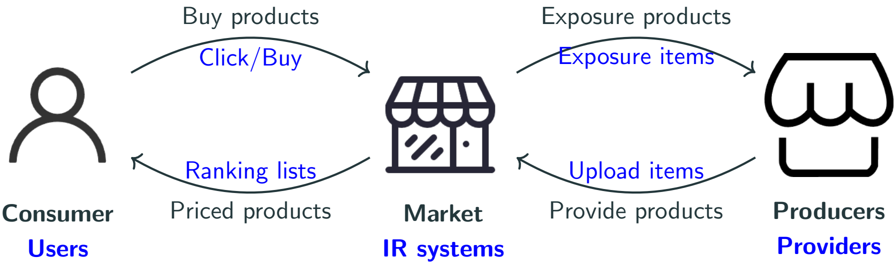

IR system as an economic market

IR system can be considered as a special **multi-sided matching economic market**

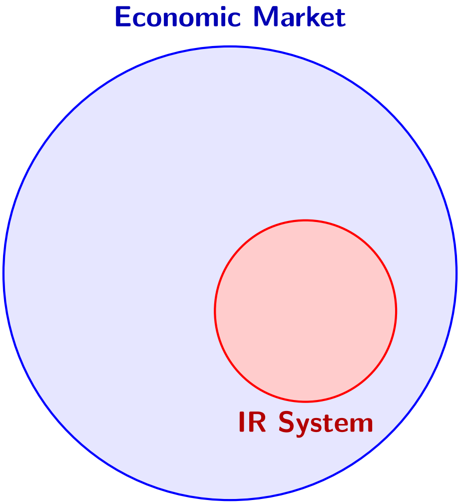

Market mechanisms vs. IR system tasks

Market mechanism	IR system analogy / task
Price mechanism	Ranking signals emerge from competition for scarce user attention [Baeza-Yates et al., 1999].
Incentive design	Stability, advertisement bidding mechanism [Yang et al., 2019], Coupons design [Yang et al., 2019].
Regulation and intervention	Platform policies enforce diversity [Dang and Croft, 2012], reduce bias [Chen et al., 2023], or increase fairness [Li et al., 2023].

Market economies include regulation by design

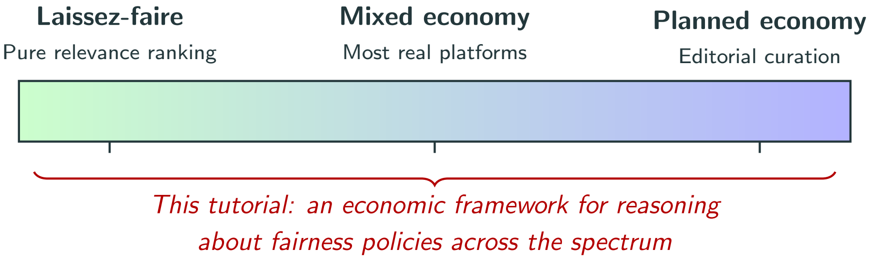

- Markets can **fail**: monopolies, externalities, extreme concentration
- When they do, governments intervene with taxes, subsidies, or rules
- IR ranking markets fail in the same way — fairness intervention is the **regulation** that corrects it

Why model IR as an economic market?

1. Import, don't invent

Centuries of theory on competing agents, scarce resources, and mediating mechanisms.

Why model IR as an economic market?

1. Import, don't invent

Centuries of theory on competing agents, scarce resources, and mediating mechanisms.

2. Operational tools, not just concepts

Taxation → redistribution | Auctions → preference elicitation | Social choice → fairness trade-offs

Why model IR as an economic market?

1. Import, don't invent

Centuries of theory on competing agents, scarce resources, and mediating mechanisms.

2. Operational tools, not just concepts

Taxation $\rightarrow$ redistribution | Auctions $\rightarrow$ preference elicitation | Social choice $\rightarrow$ fairness trade-offs

3. It works both ways

IR generates new challenges for economics: attention $\neq$ money, platforms $\neq$ governments; agents, feedback loops, and real-time bidding add new dynamics.

Recall: Fairness in IR

- In IR, we are mainly concerned with **allocation harms**
- **Unlimited** stakeholder demands vs. **limited** ranking resources

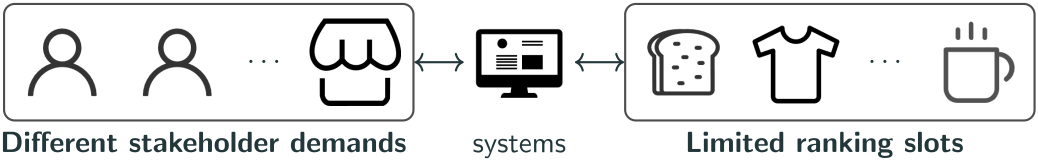

- In IR, we are mainly concerned with **allocation harms**
- **Unlimited** stakeholder demands vs. **limited** ranking resources
- Taxonomy of allocation harms [Li et al., 2021]
  - Allocation **object**: user fairness vs. provider fairness
  - Allocation **time**: short-term fairness vs. long-term fairness
  - Allocation **scale**: individual fairness vs. group fairness

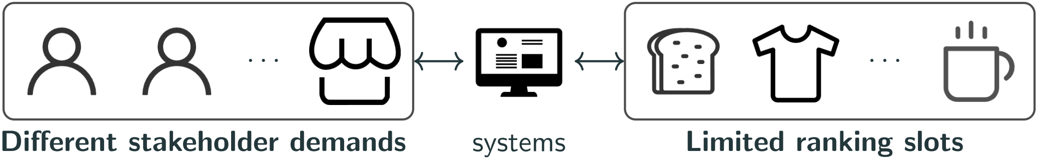

Economic concepts [Ng, 1983]

1. **Objective: supply side vs. demand side (consumer vs. producer)**
   - **Supply and demand** describe how the availability of goods and the desire to purchase them determine prices and quantities in a market.
2. **Scale: micro- vs. macro- view**
   - **Microeconomics** analyzes individual decision-making and market interactions, while **macroeconomics** focuses on economy-wide phenomena like growth, inflation, and unemployment.
3. **Time: short-term vs. long-term results**
   - **Short-term shocks** cause immediate fluctuations, while **long-term returns** reflect stable outcomes as markets adjust over time.

Taxonomy of fairness in economics

- Allocation in economics: **Allocation objective, scale and time**

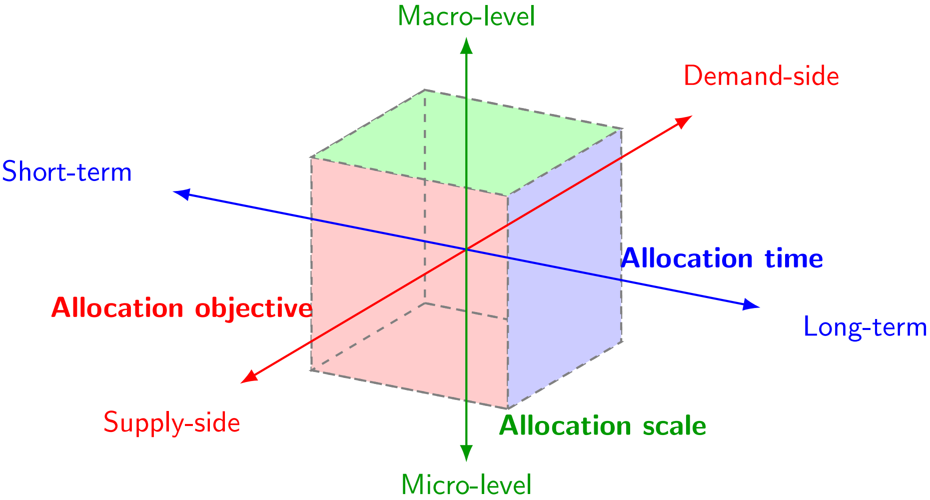

Fairness in economics

- Governments or authorities may step in to correct market failures (e.g., externalities, inequality, monopolies) through economic tools.

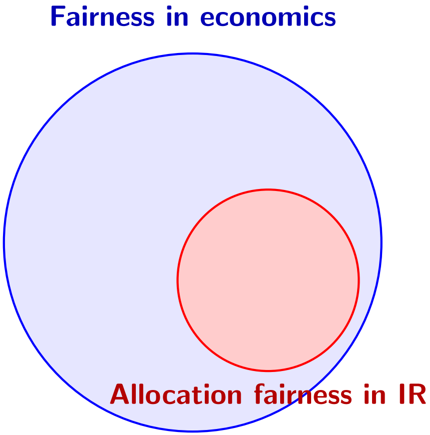

Taxonomy of fairness in IR: Alignment

- Allocation fairness in IR: **Allocation objective, scale and time**

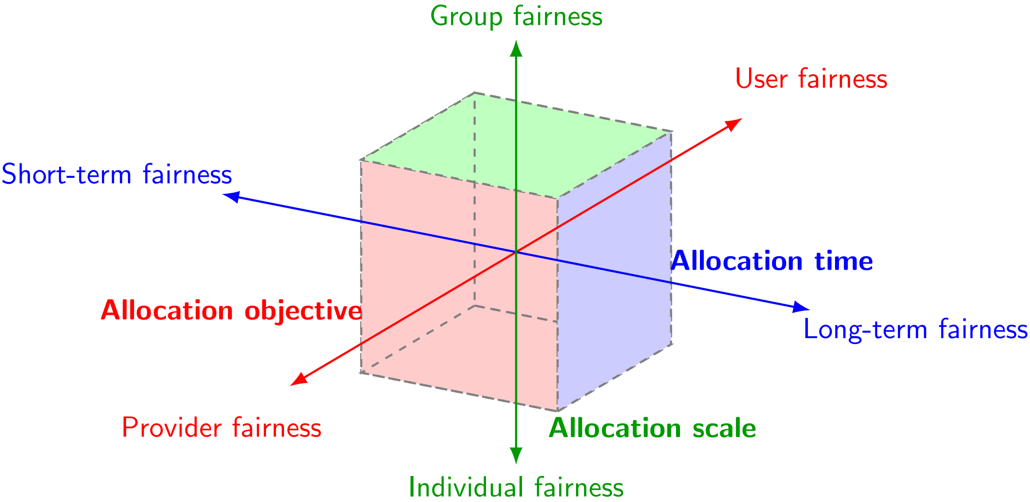

Economic Tools for Fairness

Fairness Objective — **Taxation**

Fairness Scale — **Social Choice**

Fairness Time — **Intertemporal Choice**

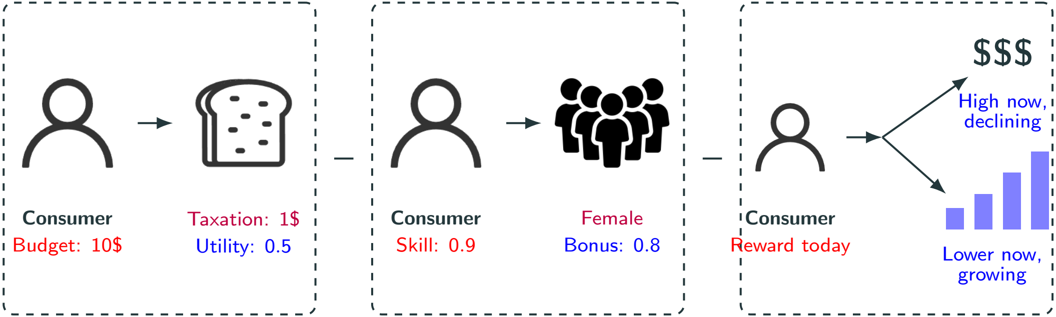

Case 1: An economic view on allocation objective

Example: Provider fairness in IR

Every user will be exposed to $k = 2$ items that have highest ranking scores:

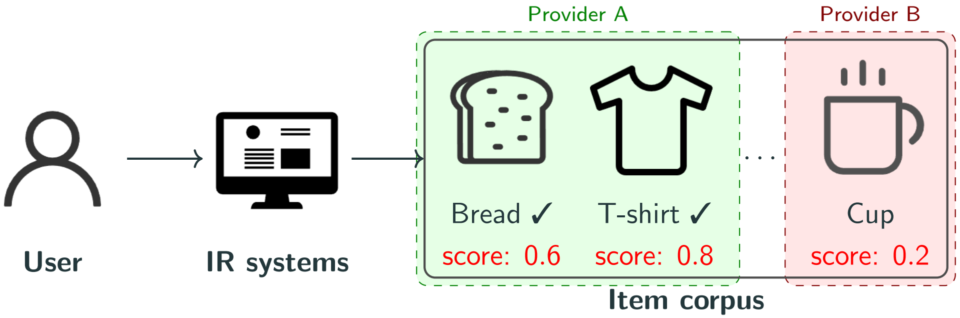

Example: Provider fairness in IR

We aim to increase the exposure of certain providers: **Through fairness score!**

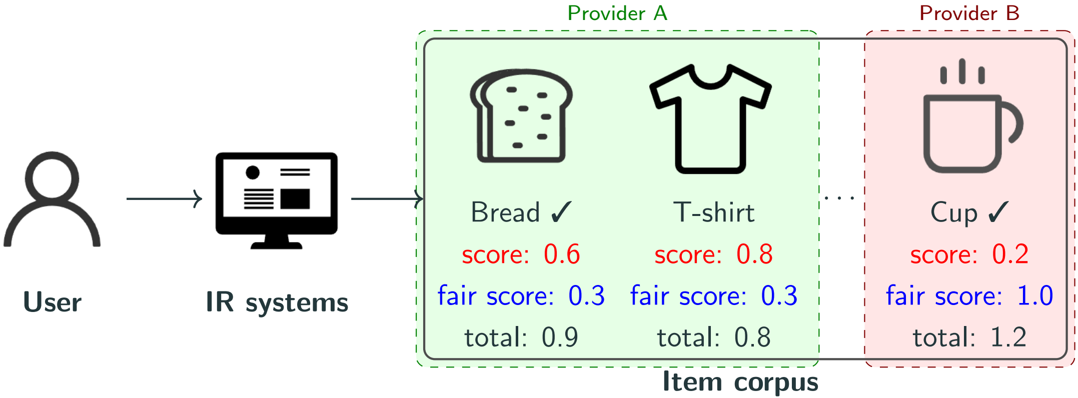

Examples: Supply-side fairness in economic market

- Users enter the market and purchase products available within it.
  - Bread: buy $10/1 = 10$ and get $0.5 \times 10 = 5$ utility
  - **T-shirt: buy** $10/5 = 2$ **and get** $3 \times 2 = 6$ **utility (win!)**
  - Cup: buy $10/2.5 = 4$ and get $1 \times 4 = 4$ utility

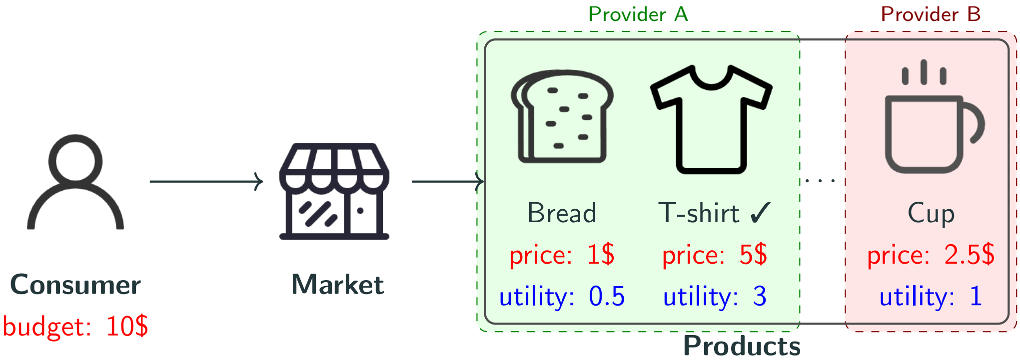

Examples: Supply-side fairness in economic market

- How can we prevent that *"the rich get richer"* **Through taxation!**
  - Bread: buy $10/2 = 5$ and get $0.5 \times 5 = 2.5$ utility
  - buy $10/10 = 1$ and get $3 \times 1 = 3$ utility
  - Cup: buy $10/2.5 = 4$ and get $1 \times 4 = 4$ utility **(win!)**

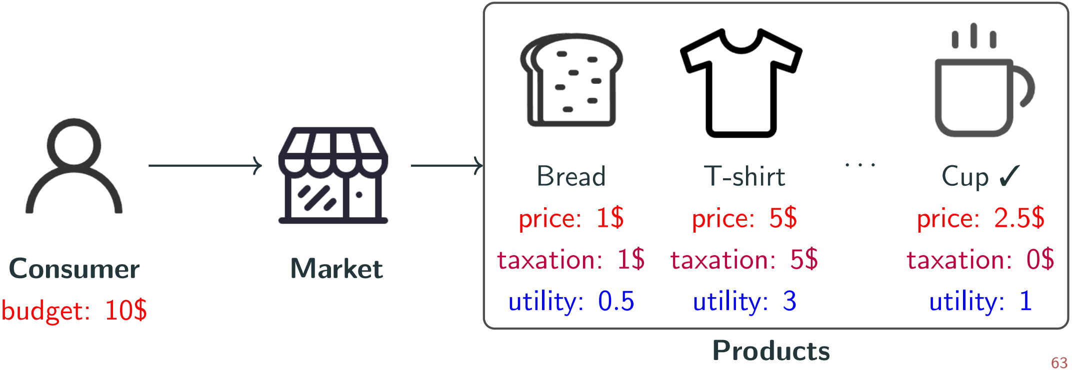

Supply-side fairness vs. provider fairness

- Supply-side fairness vs. provider fairness [Xu et al., 2024]
- **Same goal**: increasing the exposures of poor providers/demanders
- **Similar tools**: taxation mechanism as learned fairness score

Case 2: An economic perspective on allocation scale

Example: Individual fairness in employment

A firm wants to find the best hire, but cannot observe the candidates' *real* abilities

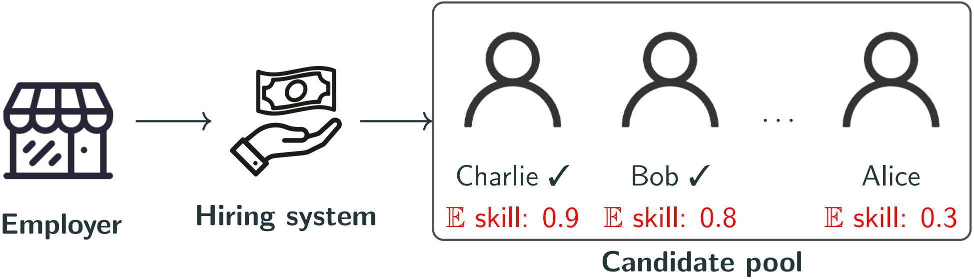

Microeconomic principle: Each person is evaluated on their expected marginal productivity; the firm applies statistical screening

Example: Group Fairness in Employment

Optimality at the firm-level $\neq$ aggregate optimality. Systemic discrimination means underemployment of the workforce!

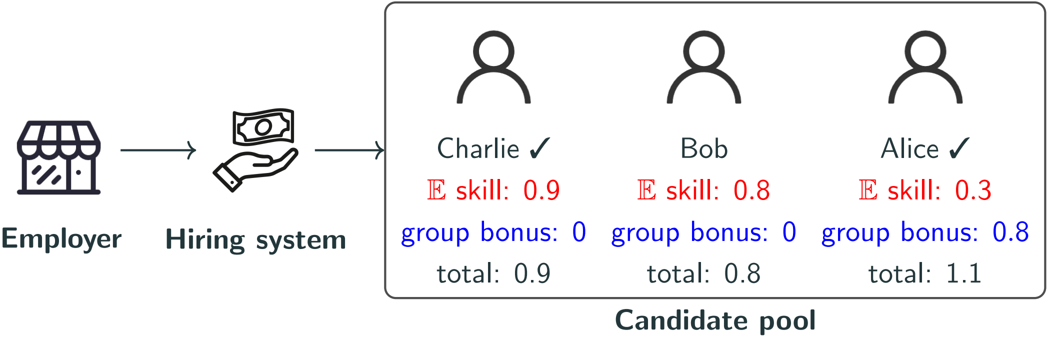

Macroeconomic principle: System-level outcomes benefit from fairness

The micro and macro dimensions are complementary

Economics addresses fairness through complementary frameworks:

 $+$  $\rightarrow$ 

Microeconomics	**Macroeconomics**	**Economic policy**
Individual merit	System outcomes	Balanced approach

Key economic frameworks that integrate these dimensions:

- **Welfare economics**: Balance efficiency and fairness in resource allocation
- **Game theory**: Achieve fair outcomes in strategic interactions
- **Social choice theory**: Aggregate individual preferences into collective decisions

Case 3: An economic view on allocation time

Multiple interactions between IR and users:

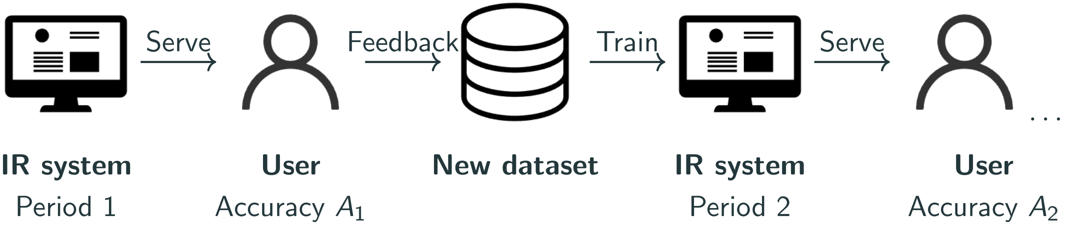

Examples: Long-term fairness in IR

- User u long-term utility reward: $R_u = A_1 + \gamma A_2 + \cdots + \gamma^n A_n$
- Using reinforcement learning (RL) to balance the long-term user reward [Ge et al., 2021]

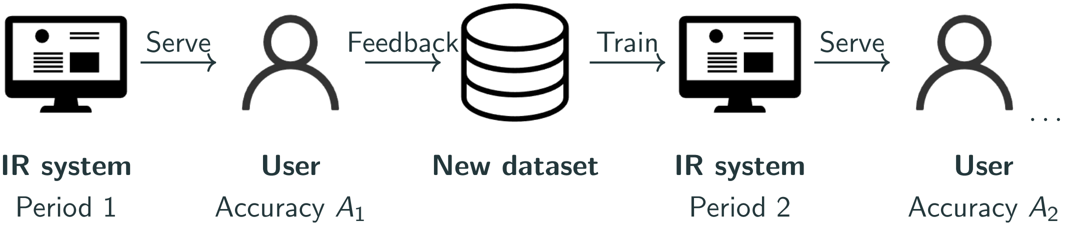

Examples: Long-term fairness in economics

Short-term focus

"Consume now"

- Maximise immediate clicks
- Providers crowd out
- Catalogue shrinks $\rightarrow$ fragility

Long-term focus

"Invest for the future"

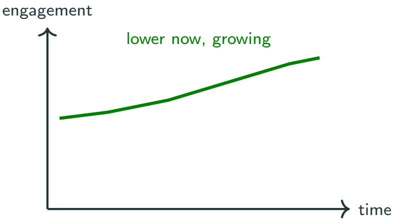

- Explore and diversify
- Providers stay viable
- Catalogue stays rich $\rightarrow$ resilience

The discount rate γ determines where the platform sits between these two extremes.

Long-term fairness in economics vs. in IR

- **Same question:** How much do we sacrifice now for a better system later?
- In economics: the **social discount rate** is a policy choice that shapes investment in infrastructure, education, climate
- In IR: the **discount factor** γ is a design choice that shapes whether the platform invests in ecosystem health or optimises for immediate engagement
- **Same insight:** This is not a technical parameter; it is a fairness decision with long-term consequences

Yuanna will develop the formal tools for this in Section 5 — including how homogeneity creates vulnerability to demand shocks, making fairness a form of risk management.

Conclusion on an eco-nomic view on fairness in IR

Fairness as allocation problem

- Fairness in IR can be viewed as **how to allocate** limited exposure or relevance to competing stakeholders (users, providers, platforms).
- The **choice of allocation** approach shapes the corresponding fairness goals and techniques.

Fairness insights from economics

1. Scarcity and trade-offs

There is no free lunch. Every fairness intervention has a cost – the question is not *whether* to trade off, but how to minimise it.

Fairness insights from economics

1. Scarcity and trade-offs

There is no free lunch. Every fairness intervention has a cost – the question is not *whether* to trade off, but how to minimise it.

2. Intertemporal thinking

A system that looks fair today can produce deeply unfair outcomes over time. Short-term decisions have long-term consequences: design for the future, not just the present.

Fairness insights from economics

1. Scarcity and trade-offs

There is no free lunch. Every fairness intervention has a cost – the question is not *whether* to trade off, but how to minimise it.

2. Intertemporal thinking

A system that looks fair today can produce deeply unfair outcomes over time. Short-term decisions have long-term consequences: design for the future, not just the present.

3. Incentive compatibility

The goal is not to enforce fairness top-down, but to design rules where doing the right thing is also the profitable choice. Fairness becomes self-sustaining.

Outlook: Token economics

What is token economics?

Token economics (*tokenomics*) studies how digital **tokens** encode incentives, access rights, and ownership in an economic system — a market-design lens born in platform and crypto-economics.

Outlook: Token economics

What is token economics?

Token economics (*tokenomics*) studies how digital **tokens** encode incentives, access rights, and ownership in an economic system — a market-design lens born in platform and crypto-economics.

- Recall (Part II): an IR system *is* an economic market, where **exposure** and **attention** are the scarce goods being allocated.
- Token economics makes this implicit "attention currency" **explicit, measurable, and transferable**.
- It reframes fairness from a constraint we *impose* to an incentive we can *price*.

Token economics meets fairness in IR

- **Exposure as a token.** Issue providers transferable exposure credits $\Rightarrow$ auditable, budgeted fairness instead of post-hoc re-ranking.
- **Incentive compatibility.** Token rewards can make fair behaviour the *profitable* choice for providers and the platform alike (recall insight 3).
- **Redistribution.** Token-based *taxation & subsidy* can support under-served stakeholders — a direct bridge to Section 3.
- **Governance.** Stakeholders who hold tokens can *vote* on the welfare function the platform optimises.

Token economics: open questions

- How to design **manipulation-resistant** token mechanisms for exposure allocation?
- Tradable tokens risk **concentration** — the "rich-get-richer" dynamic of Section 5.
- How do we measure **welfare** when fairness credits become speculative assets?
- What **transparency and regulation** keep a token market aligned with the platform's fairness goals?

Takeaway

Token economics offers a **principled, incentive-driven vocabulary** for fairness — a promising direction once we move beyond *imposing* fairness to *designing markets* for it.

Organization for next sections

- Allocation **Object**: section 3
  - Economic tool: **Taxation** for provider and user fairness
  - Application used: Next basket recommendation
  - Future and related works to explore
- Allocation **scale**: Section 4
  - Economic tool: **Micro-macro economic theory** for individual and group fairness
  - Application used: Recruitment search systems
  - Future and related works to explore
- Allocation **time**: Section 5
  - Economic tool: **Risk theory** for short-term and long-term fairness
  - Application used: Connection recommendations
  - Future and related works to explore

3. Economic-based Fairness Mitigation and Evaluation Strategies I (Chen, 30 min)

Allocation objective

- In this section, we focus on the **allocation objective**:

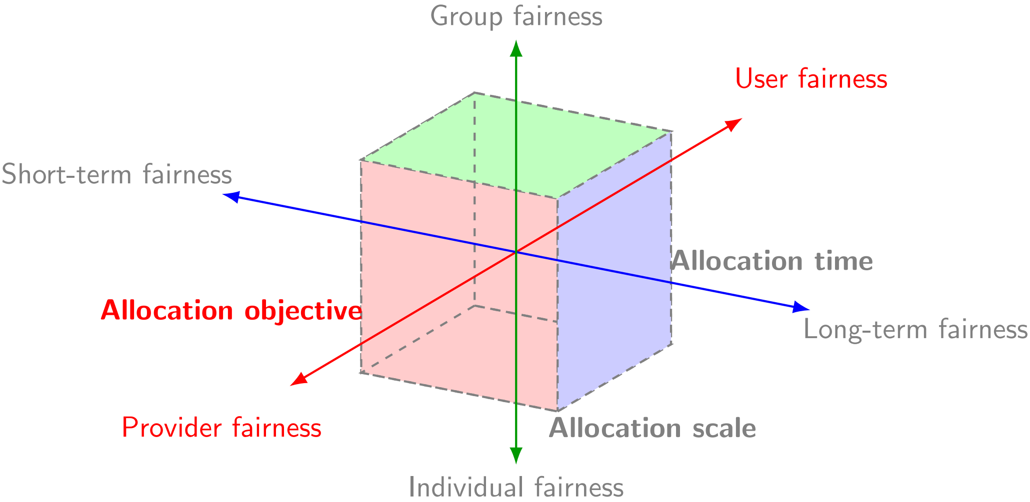

Taxation inspired user & provider fairness

Formal notation

- Assuming there are n users: $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$ arriving in IR systems

Formal notation

- Assuming there are n users: $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$ arriving in IR systems
- At each time t, the user u may input a query (search) or their profile (recommendation) u_t to the IR system.

Formal notation

- Assuming there are n users: $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$ arriving in IR systems
- At each time t, the user u may input a query (search) or their profile (recommendation) u_t to the IR system.
- Then, the IR system $f(\cdot)$ will score the item or document $i \in \mathcal{I}$ according to user's preference: $s_{u_t,i} = f(u_t, i)$

Formal notation

- Assuming there are n users: $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$ arriving in IR systems
- At each time t, the user u may input a query (search) or their profile (recommendation) u_t to the IR system.
- Then, the IR system $f(\cdot)$ will score the item or document $i \in \mathcal{I}$ according to user's preference: $s_{u_t,i} = f(u_t, i)$
- Finally, the system will generate a ranked list of size K with the highest ranking scores:

$$L_K(u_t) = \underset{S \subset \{1,2,\cdots,|\mathcal{I}|,|S|=K\}}{\arg\max} \sum_{i \in S} s_{u_t,i}$$

Recall: Fairness scoring approach

Most fairness-aware IR methods aim to utilize **fairness score** $w_{u_t,i}$ to adjust the fairness degree of users and providers: $s_{u_t,i} \rightarrow s_{u_t,i} + w_{u_t,i}$.

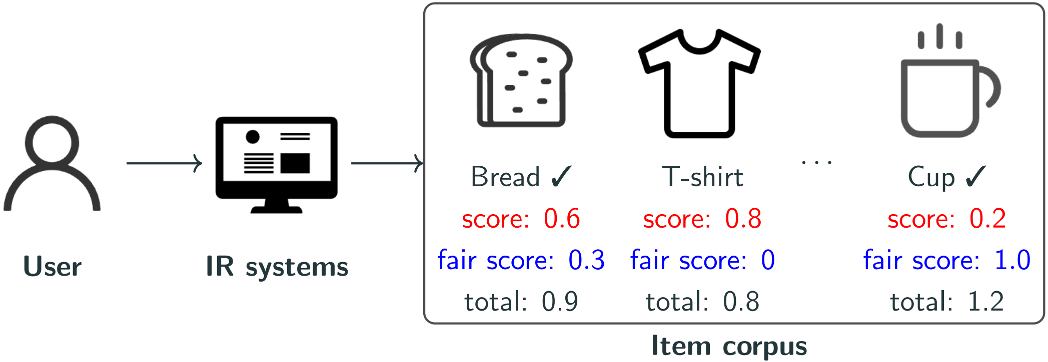

Taxation inspired fairness scoring

The fairness score $w_{u_t,i}$ can be viewed as the taxation value.

We can analyze the methods according to the taxation perspective.

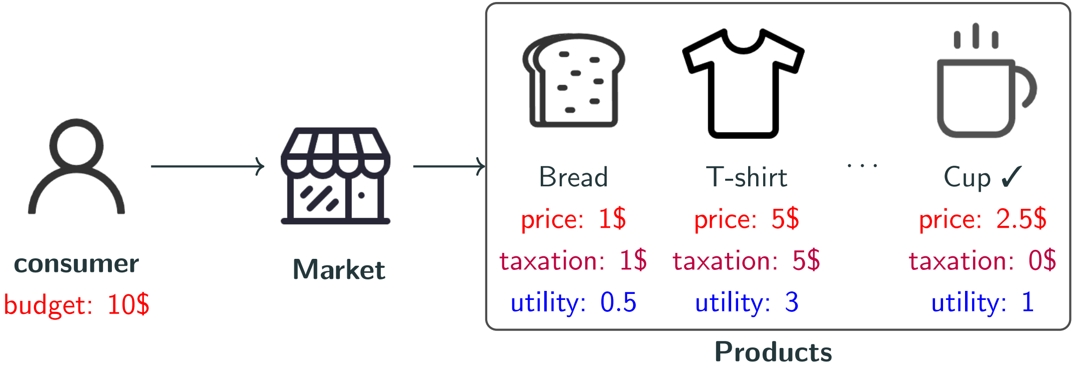

Theory: taxation and fair re-ranking are equivalent

Table 1: Analogy between the economic market and the fair re-ranking market.

Economic market	**Fair re-ranking market**
Buying products (*demand*)	Consuming attention (*users*)
Producing products (*supply*)	Producing attention (*providers*)
Binding income	Limited attention
Taxation mechanism	Fairness function
Selling price	Re-ranking score
Seeking equilibrium	Optimizing re-ranking

Theory: taxation and fair re-ranking are equivalent

Theorem (Walrasian Equilibrium in Re-ranking)

Assuming $r'(v)$ is convex, we have:

$$\max_{x_{u,i}\in\mathcal{X}} \sum_u w_u - r'(v) \iff \max_{x_{u,i}\in\mathcal{X}} \sum_u \sum_i (s_{u,i} - f_g^*) x_{u,i},$$

where $\mathcal{X} = \{x_{u,i} | \sum_i x_{u,i} \leq K,\ x_{u,i} \in [0,1]\}$
Moreover, assume that the price f_g^* is such that

$$\sum_g X_g^d(f_g^*) = \sum_g X_g^s(f_g^*) = |\mathcal{U}|K.$$

Then the allocation for demand and supply $\{X_g^d, X_g^s\}$ is Pareto optimal.

Taxation inspired fairness

Same goal: Balancing the utilities of providers and users [Xu et al., 2024].

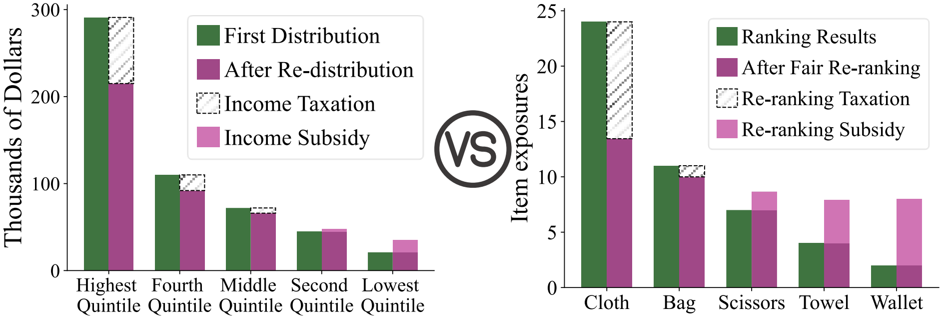

Advantages of taxation inspired fairness

1. Taxation provides a unified framework for provider and user fairness

- It helps move beyond piecemeal solutions by providing a coherent framework, making it easier to identify the strengths and limitations of existing methods.

2. Taxation inspires us to design better fairness-aware ranking models

- Taxation bridges economic fairness mechanisms with ranking systems, enabling principled, interpretable, and scalable solutions to fairness-aware IR.

Provider fairness

Insight: three taxation policies, one principle

Economics never settled on a single tax — it evolved a **family** of policies, each repairing the last. We borrow the same progression for provider fairness:

1. **Fixed taxation** (minimum wage) $\rightarrow$ **Max-min fairness**: guarantee a floor for the worst-off provider.

Insight: three taxation policies, one principle

Economics never settled on a single tax — it evolved a **family** of policies, each repairing the last. We borrow the same progression for provider fairness:

1. **Fixed taxation** (minimum wage) $\rightarrow$ **Max-min fairness**: guarantee a floor for the worst-off provider.
2. **Progressive taxation** (income brackets) $\rightarrow$ α-**fairness**: tax more those who already enjoy more exposure.

Insight: three taxation policies, one principle

Economics never settled on a single tax — it evolved a **family** of policies, each repairing the last. We borrow the same progression for provider fairness:

1. **Fixed taxation** (minimum wage) $\rightarrow$ **Max-min fairness**: guarantee a floor for the worst-off provider.
2. **Progressive taxation** (income brackets) $\rightarrow$ α-**fairness**: tax more those who already enjoy more exposure.
3. **Elasticity-aware taxation** (optimal commodity tax) $\rightarrow$ **Elastic fairness**: tax by how responsive demand is, sparing accuracy.

Insight: three taxation policies, one principle

Economics never settled on a single tax — it evolved a **family** of policies, each repairing the last. We borrow the same progression for provider fairness:

1. **Fixed taxation** (minimum wage) → **Max-min fairness**: guarantee a floor for the worst-off provider.
2. **Progressive taxation** (income brackets) → α-**fairness**: tax more those who already enjoy more exposure.
3. **Elasticity-aware taxation** (optimal commodity tax) → **Elastic fairness**: tax by how responsive demand is, sparing accuracy.

- *Read each method below as the economist's answer to the previous one's failure mode.*

Economics primer: fixed vs. progressive tax

Fixed: a minimum-wage floor

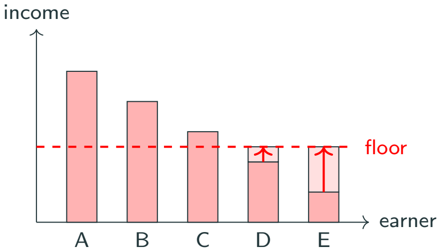

Lift only those *below* the floor; top earners untouched.

Progressive: tax brackets

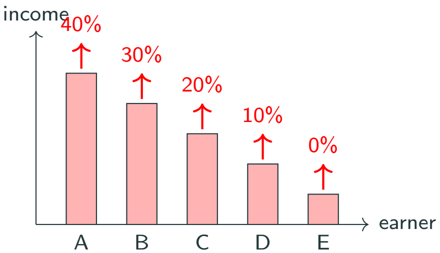

The more you earn, the higher your *rate*; revenue flows down.

1. Fixed taxation: Max-min fairness

- Aims to maximize the worst-off provider's utility (Minimum wage policy):

$$r(v_g) = \min v_g,$$

where v_g is the utility of provider g (provider g produces items $i \in I_g$).

- Accuracy-Fairness trade-off:

$$\max g(w_u) + \lambda r(v_g),$$

where $g(w_u)$ is the utility of user-side and λ is the trade-off co-efficient

Chen Xu et al. P-MMF: Provider Max-min Fairness Re-ranking in Recommender System. WWW 2023 Spotlight-Best paper candidates

1. Fixed taxation: Max-min fairness

- From an economic perspective, the taxation rate will be:

$$w_{u,i} = \begin{cases} -\lambda, & \text{if } v_g \leq v_{g'}, \forall g' \in G \\ 0, & \text{else} \end{cases} \quad (1)$$

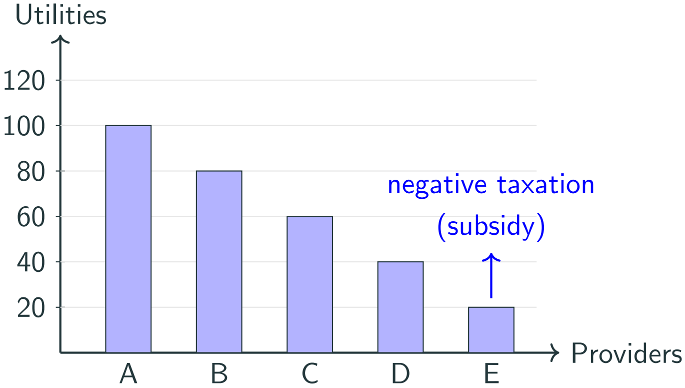

1. Fixed taxation: Max-min fairness

- To make such taxation efficient and effective in online scenarios, we design a dual-form gradient descent algorithm:

$$w_{u_t,i} = -\frac{\boldsymbol{A}^T \boldsymbol{\mu}_t}{s_{u_t,i}}, \text{ where } \boldsymbol{\mu}_t = \arg\min_{\boldsymbol{\mu} \in \mathcal{D}} \left[\langle \mathbf{g}_t, \boldsymbol{\mu} \rangle + \eta \| \boldsymbol{\mu} - \boldsymbol{\mu}_{t-1} \|_{\gamma^2}^2 \right]$$

1. Fixed taxation: Max-min fairness

- To make such taxation efficient and effective in online scenarios, we design a dual-form gradient descent algorithm:

$$w_{u_t,i} = -\frac{\boldsymbol{A}^T \boldsymbol{\mu}_t}{s_{u_t,i}}, \text{ where } \boldsymbol{\mu}_t = \arg\min_{\boldsymbol{\mu}\in\mathcal{D}} \left[\langle \mathbf{g}_t, \boldsymbol{\mu} \rangle + \eta \|\boldsymbol{\mu} - \boldsymbol{\mu}_{t-1}\|_{\gamma^2}^2 \right]$$

- Problem 1: Focus only on the worst-performing suppliers, ignoring the long tail of **middle-tier suppliers**
- Problem 2: The tax function is **complex and discontinuous**, making it difficult to choose λ

Chen Xu et al. P-MMF: Provider Max-min Fairness Re-ranking in Recommender System. WWW 2023 Spotlight-Best paper candidates

Insight: why move from fixed to progressive taxation?

Limitation of a fixed floor

- It lifts only the **worst-off** provider — the long tail of **middle-tier** providers is untouched.
- The cut-off makes the tax **discontinuous**, so tuning λ is brittle.

Economic remedy: progressivity

- Like income-tax brackets, charge a rate that **grows smoothly with exposure**.
- Every provider is adjusted by *how much they already have* — not just the single poorest.

Fixed floor $\Rightarrow$ *a smooth, exposure-dependent schedule:* this is α-fairness.

2. Progressive taxation: $\alpha-$fairness

Progressive tax: the more exposure an item receives, the higher the tax rate:

$$w_{u,i} \propto v_g^t. \tag{2}$$

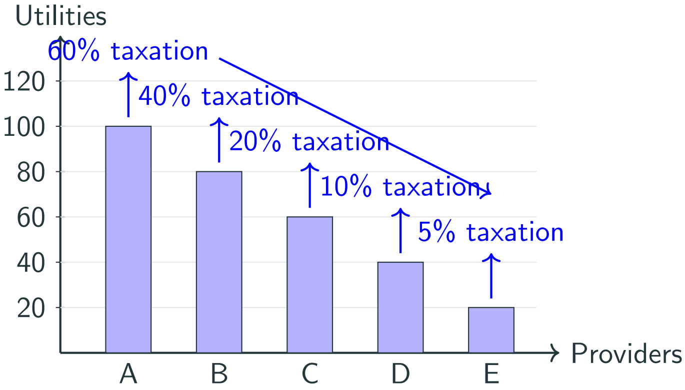

2. Progressive taxation: $\alpha-$fairness

- Optimization problem:

$$\boldsymbol{x}^*(t) = \arg\max_{\boldsymbol{x} \in \mathcal{X}_s} f(\boldsymbol{x}; t) = \begin{cases} \sum_i \gamma_i \boldsymbol{v}_i^{1-t}/(1-t) & \text{if } t \geq 0, t \neq 1 \\ \sum_i \gamma_i \log(\boldsymbol{v}_i) & \text{if } t = 1 \end{cases}$$

$$\text{s.t.} \quad \boldsymbol{v}_i = \sum_{u \in \mathcal{U}} w_{u,i} \boldsymbol{x}_{u,i}, \quad \forall i \in \mathcal{I}$$

- We utilize the optimal transportation (OT) to solve this:

$$\widetilde{\boldsymbol{x}} = \arg\min_{\boldsymbol{x} \geq 0} \langle \boldsymbol{x}, -\boldsymbol{C} \rangle + \lambda_{ot} H(\boldsymbol{x}) \quad \text{s.t.} \quad \boldsymbol{x}\mathbf{1} = K\mathbf{1}, \quad \mathbf{1}^\top \boldsymbol{x} = \boldsymbol{e}^* \quad ,$$

Chen Xu et al. A Taxation Perspective for Fair Re-ranking SIGIR 2024 Best Paper Honorable Mention

2. Progressive taxation: $\alpha-$fairness

- Problem 1: Suppliers vary in quality, and providing excessive preferential treatment to lower-quality suppliers may significantly harm user welfare.
- Problem 2: During optimization, the presence of exponentiation means that even very small errors can lead to large differences in the results (i.e., high variance).

Chen Xu et al. A Taxation Perspective for Fair Re-ranking SIGIR 2024 Best Paper Honorable Mention

Insight: why move from progressive to elasticity-aware taxation?

Limitation of progressivity

- It only sees *how much* exposure a provider has — not provider **quality**.
- Over-subsidising low-quality providers **hurts user welfare**; exponentiation adds **high variance**.

Economic remedy: elasticity

- Optimal commodity taxes (Ramsey) tax **inelastic** goods more — least distortion per unit raised.
- Tax by **price elasticity of demand**: redistribute where it costs accuracy the least.

Add a quality / responsiveness signal to the schedule: this is elastic fairness.

Economics primer: elasticity and the Ramsey rule

Inelastic good (e.g. bread)

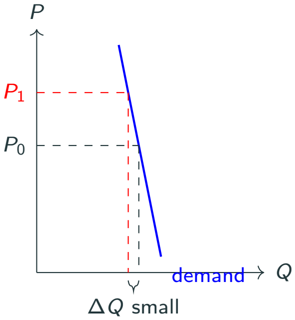

Big price rise, **tiny** quantity drop ⇒ **tax more**.

Elastic good (e.g. luxury)

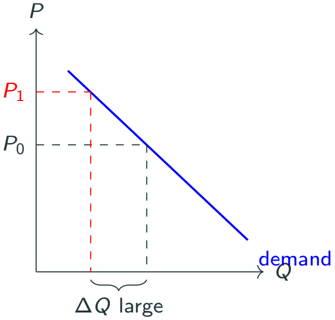

Same price rise, **large** quantity drop ⇒ **tax less**.

3. Elasticity-aware taxation: Elastic fairness

The taxation is proportional to **Price Elasticity of Demand:**

$$E = \frac{\text{\% Change in Quantity}}{\text{\% Change in Price}} = \frac{\partial Q/Q}{\partial P/P},$$

where Q is the purchase quantity of an item and P is the item price.

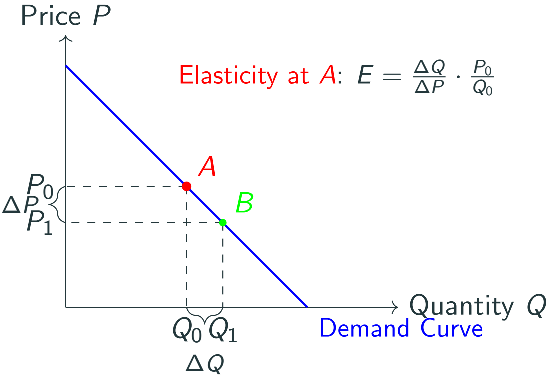

3. Elasticity-aware taxation: Elastic fairness

Elasticity tax: the higher the elasticity an item has, the higher the tax rate:

$$w_{u,i} \propto E_{i,*}. \quad (3)$$

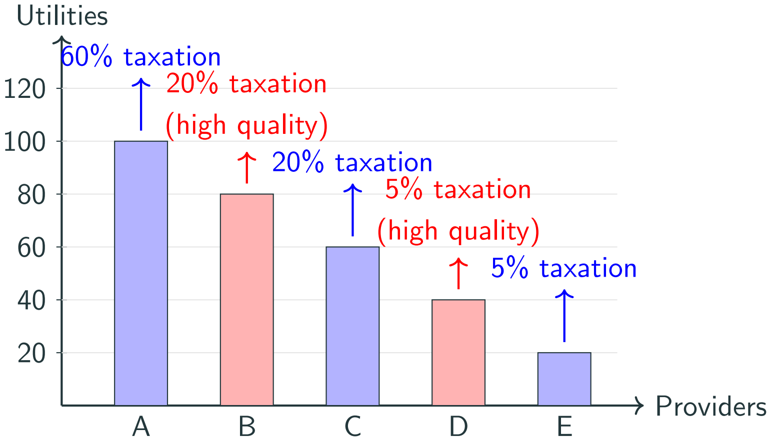

3. Elasticity-aware taxation: Elastic fairness

Theorem

The transfer ratio between fairness (commodity tax on groups) and accuracy (commodity tax transferred to users) is:

$$\gamma = \frac{\langle \nabla_{\boldsymbol{v}} L, \eta \rangle}{\langle \nabla_{\boldsymbol{v}} L, \alpha \rangle} = 1 - \frac{1}{1 + k(E_{r,p})},$$

$$k(E_{r,p}) = \frac{\sum_{p \in \mathcal{G}} \sum_{r \neq p} \boldsymbol{v}_p^{1-|t|} E_{r,p}}{\sum_{p \in \mathcal{G}} \boldsymbol{v}_p^{1-|t|}}$$

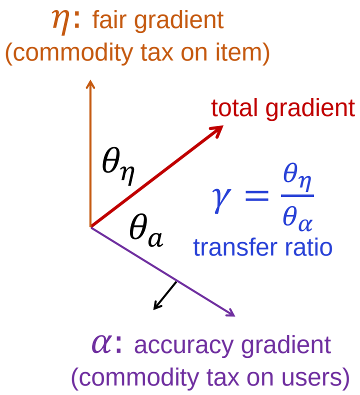

Insight: a spectrum of fairness is a spectrum of taxes

Policy	Tax rule	Helps most	Limitation
Fixed (Max-min)	flat floor on worst-off	poorest provider	ignores middle tier; discontinuous
Progressive (α)	rate $\propto$ exposure	all, by exposure	quality-blind; high variance
Elasticity-aware	rate $\propto$ elasticity	responsive demand	needs elasticity estimation

Take-away

Choosing a **fairness function** *is* choosing a **tax policy** — an explicit, auditable value judgement about how to redistribute exposure between providers and users [Xu et al., 2024, 2025b].

Application: Next basket recommendation

From theory to practice: why next-basket recommendation?

- The three policies above are **general re-ranking taxes**. Does the lens survive a messy, real task?

From theory to practice: why next-basket recommendation?

- The three policies above are **general re-ranking taxes**. Does the lens survive a messy, real task?
- **Next-basket recommendation (NBR)** is a stress test: strong **repeat bias** concentrates exposure on a few already-popular items — an extreme "rich-get-richer" market.

From theory to practice: why next-basket recommendation?

- The three policies above are **general re-ranking taxes**. Does the lens survive a messy, real task?
- **Next-basket recommendation (NBR)** is a stress test: strong **repeat bias** concentrates exposure on a few already-popular items — an extreme "rich-get-richer" market.
- We will see the *same* taxation vocabulary — tax the over-exposed, subsidise the under-exposed — applied to **two coupled objectives**: item fairness *and* repeat bias.

Next basket recommendation

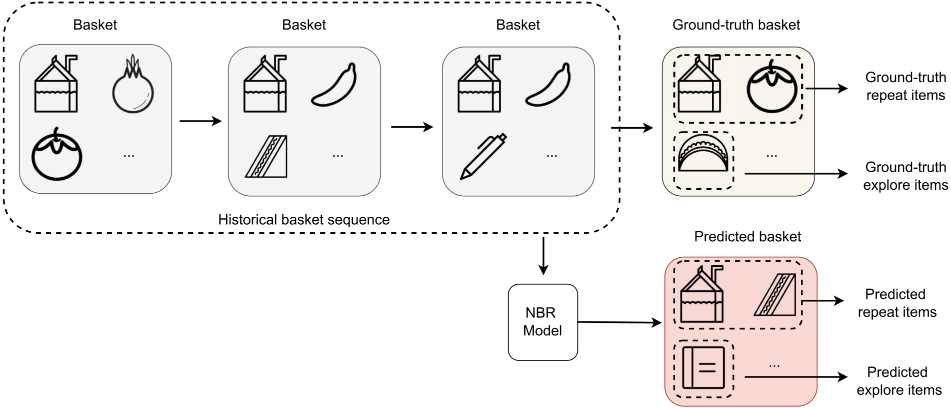

- The predicted basket contains both repeat and explore items.

Next basket recommendation

State-of-the-art NBR methods have heavy repeat bias. [Liu et al., 2025] jointly optimize item fairness and repeat bias via mixed-integer linear programming.

- Repeat-bias-aware item fairness optimization (RAIF):

$$\max \quad f(x) + \alpha g(x) - \lambda RepRatio(x)$$

A taxation perspective on RAIF

- Higher taxation rate α on the unprotected group
- Another taxation rate λ on the repeated items

$$\max \quad f(x) + \alpha g(x) - \lambda RepRatio(x)$$

Future and related work

Carefully choose fairness function

- Different fairness objectives **taxes on different types of users/providers** [Xu et al., 2025b]!

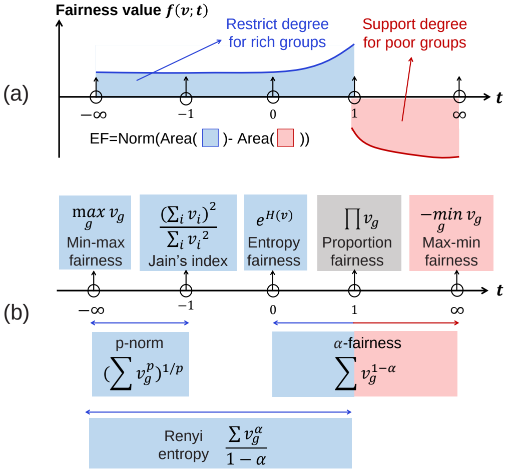

Carefully choose fairness function

- Different fairness objectives **taxes on different types of users/providers** [Xu et al., 2025b]!
- Different fairness objectives **have different taxation properties** [Xu et al., 2024].

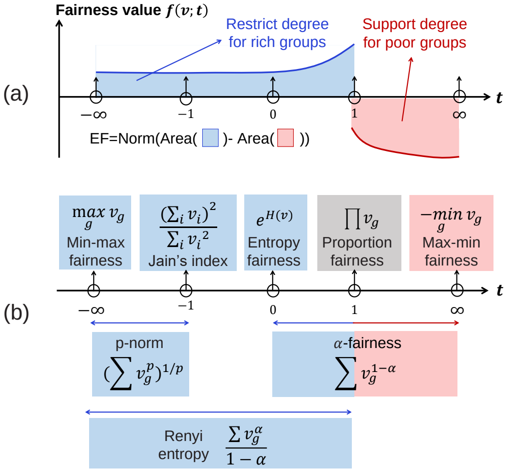

Evaluation

1. Evaluation metrics

- To measure algorithm convergence performance, we need to make sure the taxation policy (fairness objective) be same.
- To assess an algorithm's fairness, we should analyze the shifts in utility experienced by every user or provider, rather than only relying on a single overall metric.

Evaluation

1. Evaluation metrics

- To measure algorithm convergence performance, we need to make sure the taxation policy (fairness objective) be same.
- To assess an algorithm's fairness, we should analyze the shifts in utility experienced by every user or provider, rather than only relying on a single overall metric.

2. Evaluation properties

- Economic principles tell us that, beyond just looking at a single fairness metric, we also need to consider the inherent properties of fairness algorithms, such as continuity.

Better tools

- Taxation can be regarded as a tool to theoretically analyze the **accuracy-fairness trade-offs** in IR [Xu et al., 2025b].

Better tools

- Taxation can be regarded as a tool to theoretically analyze the **accuracy-fairness trade-offs** in IR [Xu et al., 2025b].
- Taxation theory can inform real-world systems, suggesting the need for **mixed taxation policies** tailored to different applications.

Better tools

- Taxation can be regarded as a tool to theoretically analyze the **accuracy-fairness trade-offs** in IR [Xu et al., 2025b].
- Taxation theory can inform real-world systems, suggesting the need for **mixed taxation policies** tailored to different applications.
- Inspired by taxation mechanisms, IR systems can **adopt diverse taxation strategies**—for instance, taxing user traffic to fund essential infrastructure and other foundational services.

Future: from SEO to GEO

- As retrieval gives way to **generative engines** (RAG, AI search, chat answers), the exposure surface shifts: ranked lists $\rightarrow$ **generated answers and citations**.

Future: from SEO to GEO

- As retrieval gives way to **generative engines** (RAG, AI search, chat answers), the exposure surface shifts: ranked lists $\rightarrow$ **generated answers and citations**.
- Providers adapt accordingly — from **SEO** (Search Engine Optimization) to **GEO** (Generative Engine Optimization) [Aggarwal et al., 2024]: shaping content to be *retrieved, cited, and paraphrased* by the model. Recent work even *automates* this loop [Wu et al., 2025].

Future: from SEO to GEO

- As retrieval gives way to **generative engines** (RAG, AI search, chat answers), the exposure surface shifts: ranked lists $\rightarrow$ **generated answers and citations**.
- Providers adapt accordingly — from **SEO** (Search Engine Optimization) to **GEO** (Generative Engine Optimization) [Aggarwal et al., 2024]: shaping content to be *retrieved, cited, and paraphrased* by the model. Recent work even *automates* this loop [Wu et al., 2025].
- New fairness questions arise:
  - Whose content gets **cited / represented** in the generated answer?
  - Does GEO reward genuine quality, or merely **model-pleasing** content?
  - Exposure becomes **implicit and unranked** — how do we even measure it?

A token-economics view of GEO fairness

Recall the **token-economics** outlook (Part II): make exposure an explicit, transferable currency.

- **Citation as a token.** Treat each citation / attribution in a generated answer as a unit of exposure that can be *budgeted, audited, and redistributed*.

Outlook

GEO turns provider behaviour into a market; **token economics** + **taxation** give us a principled language for fairness in that market.

A token-economics view of GEO fairness

Recall the **token-economics** outlook (Part II): make exposure an explicit, transferable currency.

- **Citation as a token.** Treat each citation / attribution in a generated answer as a unit of exposure that can be *budgeted, audited, and redistributed*.
- **Taxation still applies.** Over-cited sources can be *taxed* and under-represented providers *subsidised* — the same fixed / progressive / elasticity policies, now levied over citations.

Outlook

GEO turns provider behaviour into a market; **token economics** + **taxation** give us a principled language for fairness in that market.

A token-economics view of GEO fairness

Recall the **token-economics** outlook (Part II): make exposure an explicit, transferable currency.

- **Citation as a token.** Treat each citation / attribution in a generated answer as a unit of exposure that can be *budgeted, audited, and redistributed*.
- **Taxation still applies.** Over-cited sources can be *taxed* and under-represented providers *subsidised* — the same fixed / progressive / elasticity policies, now levied over citations.
- **Incentive compatibility.** Token rewards can steer GEO toward genuine quality, so that being fair becomes the *profitable* GEO strategy.

Outlook

GEO turns provider behaviour into a market; **token economics** + **taxation** give us a principled language for fairness in that market.

Fairness in IR on allocation objective: Related work

Provider fairness:

- FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms
- FairSync: Ensuring Amortized Group Exposure in Distributed Recommendation Retrieval

User fairness:

- User Fairness in Recommender Systems

Two-sided fairness:

- CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems
- Intersectional Two-sided Fairness in Recommendation

Q&A

Website

Toolkit

Contact information: chenxu0427ruc@gmail.com

Break (Coming Section 4-6)

1. Introduction: Fairness in IR (Yuanna, 20 min)
2. An Economic View on Fairness in IR (Chen, 30 min)
3. Economic-based Fairness Mitigation and Evaluation Strategies I (Chen, 30 min)
4. Economic-based Fairness Mitigation and Evaluation Strategies II (Clara, 30 min)
5. Economic-based Fairness Mitigation and Evaluation Strategies III (Yuanna, 30 min)
6. Open Problems, Quick Start for Learning about Fairness, and Conclusions (Clara, 20 min)

4. Economic-based Fairness Mitigation and Evaluation Strategies II (Clara, 30 min)

Allocation scale

- In this section, we focus on **Allocation scale**

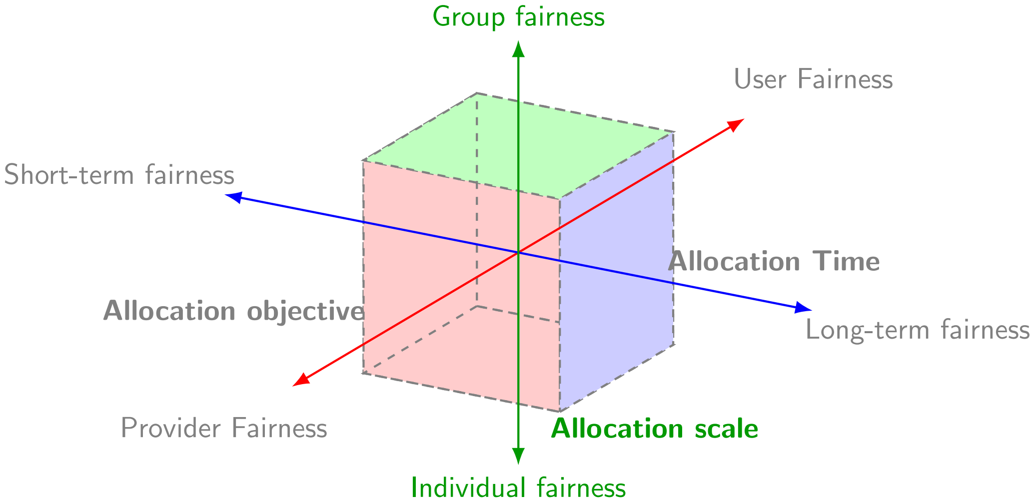

Micro-macro economic inspired individual & group fairness

Individual and group fairness

Individual fairness: Individuals who are similar with respect to a particular task should receive similar outcomes [Dwork et al., 2012].

Group fairness: Members of different protected groups should be treated the same.

Economic Lens: Individual vs. group fairness

Economic parallel: Microeconomics vs. macroeconomics

Economists have studied a similar dichotomy between local level optimization and aggregate level outcomes using micro- and macroeconomics.

Micro vs. macro objectives

- Microeconomics focuses on individual behavior and incentives
  - Individuals, firms, local optimization
  - Key idea: merit-based allocation (e.g. productivity $\rightarrow$ reward)
- Macroeconomics focuses on system-level outcomes
  - Aggregates, growth, stability, equity
  - Key Idea: optimize welfare, diversity

Microeconomic approach

- **Individual Fairness:** Each person receives treatment based on their specific circumstances
- **Pareto Efficiency:** No individual can be made better off without making another worse off
- **Personalized Allocation:** Resources distributed based on individual merit/need

Macroeconomic approach

- **Group Fairness:** Focus on aggregate outcomes of the system and across demographic groups
- **Distributional Justice:** Ensuring equal group-level statistical parity
- **Market Equilibrium:** Balancing overall system fairness

What we gain from this economic lense

Often group and individual fairness are viewed as competing and independent goals.

Economic view: Individual decisions and behaviors (micro level) collectively shape system-wide outcomes (macro level), while macro-level conditions (such as inequality, growth, or systemic biases) in turn influence individual opportunities and choices.

Can help understand the relationship between group and individual fairness.

How does this map to IR?

In IR, we have multiple **stakeholders**:

- **Users** - individuals with an information need (e.g. candidates, consumers).
- **Items** - entities being ranked/recommended (e.g. documents, products, **people**).
- **Providers** - parties offering or supplying items (e.g. companies, publishers).

Individual fairness: Similar users/items/providers should receive similar outcomes.

Group fairness: Groups of users/items/providers should receive proportional or equal outcomes.

How does this map to IR?

Individual fairness: Similar users/items/providers should receive similar outcomes.

Group fairness: Groups of users/items/providers should receive proportional or equal outcomes.

- How to define similar outcomes in IR?
- How to define similar individuals? How to divide the groups?
- How to achieve group/individual fairness in IR and how does the economic view help?

Individual fairness

Individuals who are similar with respect to a particular task should receive similar outcomes [Dwork et al., 2012].

Individual fairness in IR: Similar users/items/providers should receive similar outcomes.

Input similarity

How to define similarity among individuals?

Input Similarity

How to define **similarity** among individuals?

Input similarity is measured as the **distance** between individuals in the feature space.

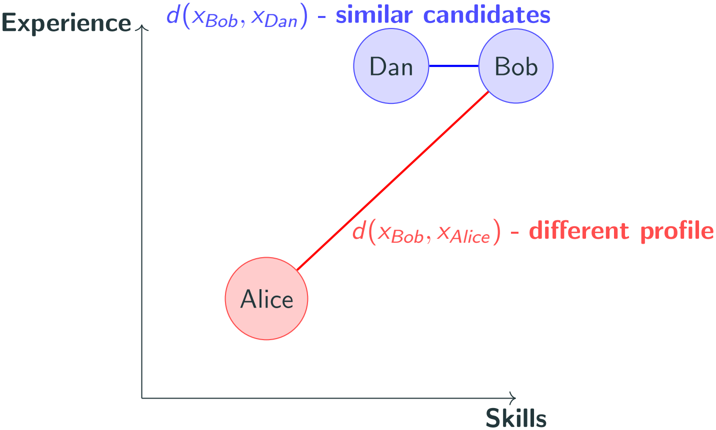

Output similarity

How to define similarity in the outcomes for individuals?

Output similarity

Output similarity is defined relative to each stakeholder's need:

- **Items**: similar items should get similar levels of exposure over time [Biega et al., 2018, Lahoti et al., 2019, Rus et al., 2024].
- **Users**: similar users should receive similar recommendations [Chawla and Jagadeesan, 2022].

Individual Fairness and Output Similarity: User

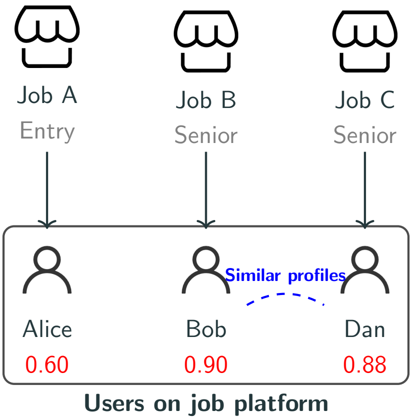

Individual fairness: Bob and Dan, with similar skill levels, should receive similarly senior-level job recommendations, unlike Alice.

Output similarity: Items

Items should receive **similar levels of exposure** across time.

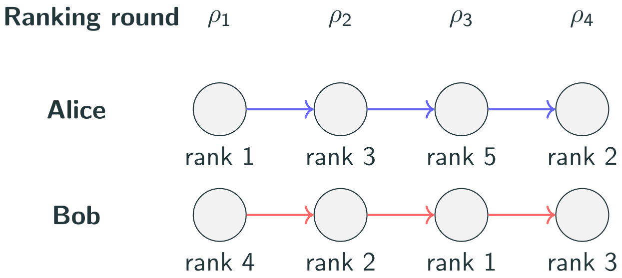

Cumulative exposure

$\sum_{t=1}^{T} \frac{1}{\log_2(\text{rank}_t + 1)} \approx$ equal for Alice and Bob

Individual fairness: Items

An individually fair ranking system should give similar candidates similar exposure over time [Dwork et al., 2012, Lahoti et al., 2019, Rus et al., 2024].

$$|\text{Cumulative Exposure}(x_i) - \text{Cumulative Exposure}(x_j)| \leq L \cdot d_X(x_i, x_j)$$

- Cumulative Exposure(x): the attention or visibility individual x receives across time
- $d_X(x_i, x_j)$: similarity metric between individuals (e.g., feature distance)
- L: Lipschitz constant - controls how much exposure difference is allowed for a given dissimilarity

Example: Individual fairness

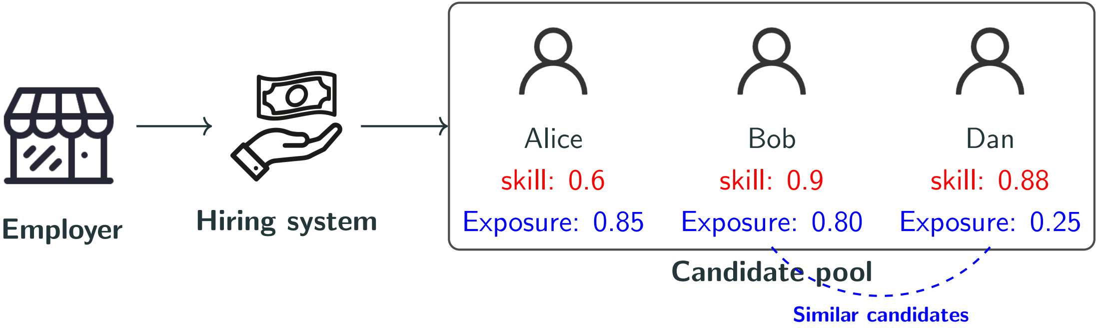

Fairness violation: Bob and Dan have nearly identical skill levels, but Bob receives exposure similar to Alice.

Achieving individual fairness: Lipschitz fairness constraint

- Define an **input similarity metric** d_X between individuals.
- Define an **output similarity metric** d_Y between individuals.
- Optimize the ranking function $f(x)$ under fairness constraints $g(x)$.

$$\begin{aligned} \max \quad & f(x) \\ s.t. \quad & g(x) \leq c, \end{aligned}$$

where $g(x)$ is defined as

Lipschitz fairness constraint

$$d_Y\left(x_i, x_j\right) \leq L \cdot d_X(x_i, x_j) \quad \forall\left(x_i, x_j\right)$$

Individual fairness: Challenges

Defining an input similarity function

- Requires a task-specific, ethically-grounded distance metric between individuals.
- In practice, it's difficult to know which features are truly "non-sensitive".
- **Proxy problem:** Non-sensitive features may still encode sensitive information.
  - Example: years of experience could be a proxy to age or gender

Consequence: This definition of **individual fairness** requires strong assumptions and domain knowledge to avoid fairness-washing.

A different view on individual fairness

Goal: Ensure that each individual receives attention proportional to their relevance over time [Biega et al., 2018, Singh and Joachims, 2018, 2019, Heuss et al., 2022].

Equity of attention [Biega et al., 2018]

For each subject i, over a sequence of rankings $\rho_1, \ldots, \rho_m$:

$$\frac{\sum_{\ell=1}^{m} a_i^{\ell}}{\sum_{\ell=1}^{m} r_i^{\ell}} = \text{constant}, \quad \forall i$$

- a_i^{ℓ}: attention (exposure) in ranking ρ_ℓ
- r_i^{ℓ}: relevance score in that round

Achieving individual fairness: Equity of attention

Use integer linear programming (ILP) to generate a new ranking $\rho_{\ell}*$ that:

$$\begin{aligned} \min \quad & \sum_i |A_i - R_i| \\ \text{s.t.} \quad & \text{NDCG@k}(\rho^j, \rho^{j*}) \geq c, \quad \forall j = 1, \dots, m \end{aligned}$$

where A_i and R_i are cumulative attention and relevance over m rankings (ρ) for an individual

$$\text{NDCG@k}(\rho, \rho^*) = \frac{\text{DCG@k}(\rho^*)}{\text{DCG@k}(\rho)}$$

Challenges

It is crucial to ensure that the utility or relevance function is **objective** and does not reinforce existing biases.

Group fairness

Members of different **protected groups** should be treated the same.

Group fairness: Groups of users/items/providers should receive proportional or equal outcomes.

How to define the groups?

- **Protected attributes**: gender, race, age ...
- **Task-specific attributes**: seniority levels, job types, user tiers ...
- **Popularity**: popular vs niche items
- **Behavioral groups**: active vs. passive users, frequent vs. infrequent buyers ..

Beyond binary groups

- **Multiple** attributes
- **Non-binary** groups
- **Intersectional** groups: Intersectional discrimination arises when an individual faces discrimination due to their **membership in multiple groups** at the same time [Wang et al., 2022, Gohar and Cheng, 2023].

Intersectionality

Original

τ
d
e
a
c
b
f
g
h
i

$\tau 1'$
d
e
c
h
i

Male White
Female White
Male Black
Female black
Female Asian

τ1' - at least *two* candidates from each gender and *one* of each race group group → *returns a candidate from the female, male, white and black group.*

Not all intersectional groups are present in the top → *no representant from the Male Black group.*

Group fairness

Members of different protected groups should be treated the same.

- Demographic parity:

$$P(\hat{Y} = 1 \mid A = a) = P(\hat{Y} = 1 \mid A = b)$$

- Equal opportunity:

$$P(\hat{Y} = 1 \mid Y = 1, A = a) = P(\hat{Y} = 1 \mid Y = 1, A = b)$$

- Equalized odds:

$$P(\hat{Y} = 1 \mid Y = y, A = a) = P(\hat{Y} = 1 \mid Y = y, A = b) \quad \text{for all } y \in \{0, 1\}$$

Group fair outcomes

- **Items**: Groups of items should receive proportional/equal **exposure**.
- **Users**: Groups of users should receive **equal quality** of recommendations, ensuring no group is systematically disadvantaged.

Group fair outcomes

- **Items**: Groups of items should receive proportional/equal **exposure**.
- **Users**: Groups of users should receive **equal quality** of recommendations, ensuring no group is systematically disadvantaged.

In this part we focus on the item side! Check out Economic-based Fairness Mitigation and Evaluation Strategies I (User Fairness)

Example: Group fairness

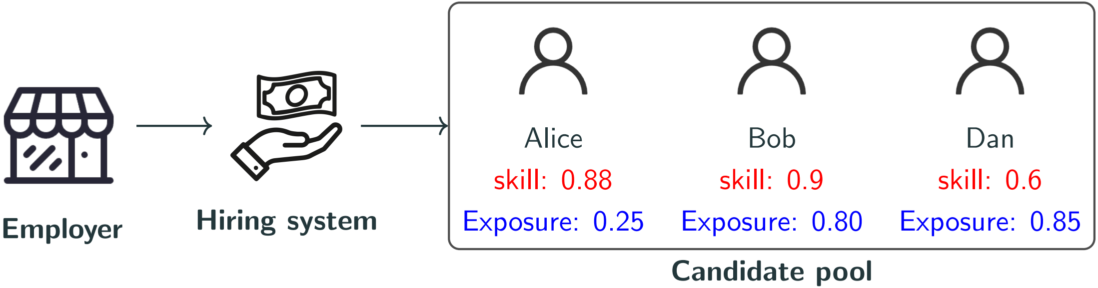

Even though Alice is more skilled than Dan, she receives lower exposure - ranking favoring one group in the top of the ranking.

Group fairness in rankings

Small a difference in relevance can lead to a large difference in exposure (an opportunity) for the group of females [Singh and Joachims, 2018].

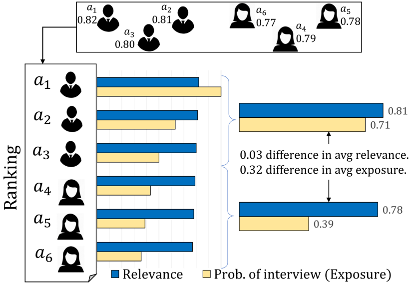

Group fairness: Members of different protected groups should receive similar/proportional exposure.

Achieving group fairness

Goal: Generate a ranked list which balances utility and group fairness.

$$\begin{aligned} \max \quad & f(x) \\ s.t. \quad & g(x) \leq c, \end{aligned}$$

where $g(x)$ is the fairness constraint and $f(x)$ is the ranking function.

Achieving group fairness: FA*IR [Zehlike et al., 2017]

Fairness constraint: At each position i in the top-k list, the number of protected candidates should be at least as high as the expected number in a fair distribution.

Approach:

- Create a ranked list for each protected and non-protected group.
- At each position i, if the current ranking has fewer protected candidates than the lower bound $\Rightarrow$ select next most relevant protected candidate.
- Otherwise, select next most relevant candidate (protected or not).

Example: Group fairness vs. individuals

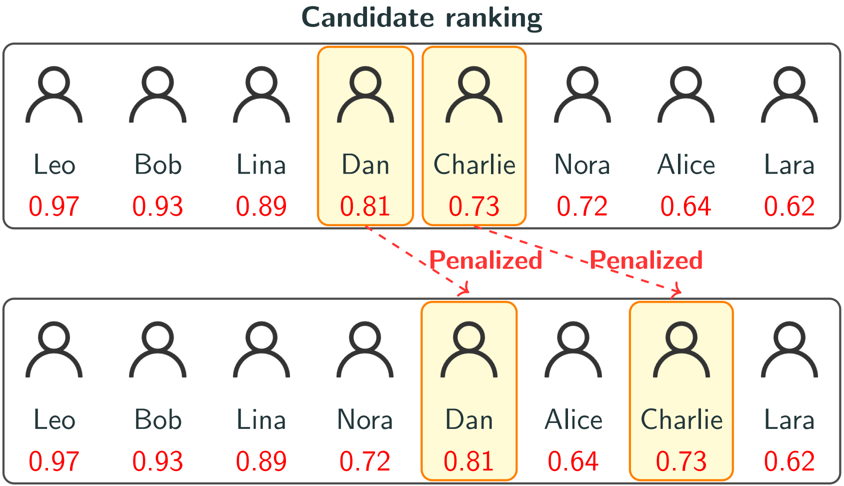

Group fair constraint: have at least $k/2$ individuals of each gender in top-k ($k \geq 3$)

Individual Fairness under Group-Fairness Constraints

Challenge: Enforcing group-fairness often hurts high-scoring individuals.

Goal: Minimize the amount of individual unfairness when enforcing group fairness [García-Soriano and Bonchi, 2021].

Approach: Rawls's theory of justice [John et al., 1971] – arranging social and financial inequalities to the benefit of the worst-of.

Individual fairness under group-fairness constraints

$$\begin{aligned} \max_{P} \quad & \min_{u \in \mathcal{U}} \mathbb{E}_{r \sim P}\left[V(r, u)\right] \\ \text{s.t.} \quad & \mathbb{E}_{r \sim P}\left[g(r)\right] \leq c \end{aligned}$$

where P is a probability distribution over rankings.

$V(r, u)$ is the received utility of individual u in ranking r,

and $g(r)$ is the fairness constraint applied to ranking r.

Individual fairness under group-fairness constraint

Deterministic group fairness ranking:

$r' = \langle$ Leo, Bob, Lina, Nora, Dan, Alice, Charlie, Lara $\rangle$

Worst-off utility: V(r,Charlie) = -2

Probability Distribution over Fair Rankings (P):

$r_1 = \langle$ Leo, Dan, Lina, Lara, Bob, Nora, Charlie, Alice $\rangle$ $\quad \mathbb{P}(r_1) = \frac{1}{4}$

$r_2 = \langle$ Bob, Leo, Lina, Nora, Dan, Alice, Lara, Charlie $\rangle$ $\quad \mathbb{P}(r_2) = \frac{1}{2}$

$r_3 = \langle$ Bob, Leo, Lina, Lara, Charlie, Nora, Dan, Alice $\rangle$ $\quad \mathbb{P}(r_3) = \frac{1}{16}$

$r_4 = \langle$ Charlie, Leo, Lina, Lara, Bob, Nora, Dan, Alice $\rangle$ $\quad \mathbb{P}(r_4) = \frac{3}{16}$

Worst-off expected utility: all users have $\mathbb{E}[V(r, u)] \geq -0.75$

Economic perspective in IR

Individual fairness (micro view)

- Focus on **pairwise treatment** of individuals.
- Ensures **similar individuals** receive **similar outcomes**.

$$|\text{Exposure}(i) - \text{Exposure}(j)| \cdot \frac{1}{d_X(i,j)} \leq c$$

Economic view: Like microeconomics, focusing on individual outcomes.

Group fairness (macro view)

- Focus on **aggregated outcomes** across groups.
- Ignores within-group differences.

$$\left| \frac{1}{|G_a|} \sum_{i \in G_a} \text{Exposure}(i) - \frac{1}{|G_b|} \sum_{i \in G_b} \text{Exposure}(i) \right|$$

Economic view: Like macroeconomics, focusing on group-level outcomes.

How is this useful?

The economic perspective offers new approaches to fairness by drawing connections between individual and group-level concerns.

By adopting this economic view, we can better understand the **trade-offs** between group and individual fairness and design fairness-aware systems that account for both levels simultaneously.

Social choice for fairness in recommendation

For example, recent work [Aird et al., 2023, 2024a,b, Sonboli et al., 2020] uses **social choice theory**, which formalizes how to aggregate individual preferences into collective decisions.

Approach: Fairness concerns and individual preferences are represented as agents which interact through social choice.

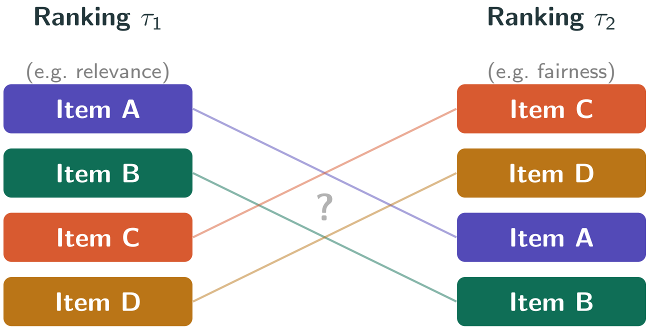

Social choice for recommendation fairness: SCRUF-D

- The recommendation system is modeled as a **multi-agent system** with two types of agents:
  - **User agents:** Represent individual user preferences.
  - **Fairness agents:** Represent different fairness principles (e.g., exposure parity, diversity) and can evaluate or re-rank recommendations for fairness.
- **Stage 1: Allocation of fairness agent** When a user arrives, a suitable fairness agent (or multiple) is assigned to the user.
- **Stage 2: Aggregation** Lists from user agents and fairness agents are aggregated via a **social choice rule** (e.g., Borda Count).

Application: Recruitment search systems

Recruitment system

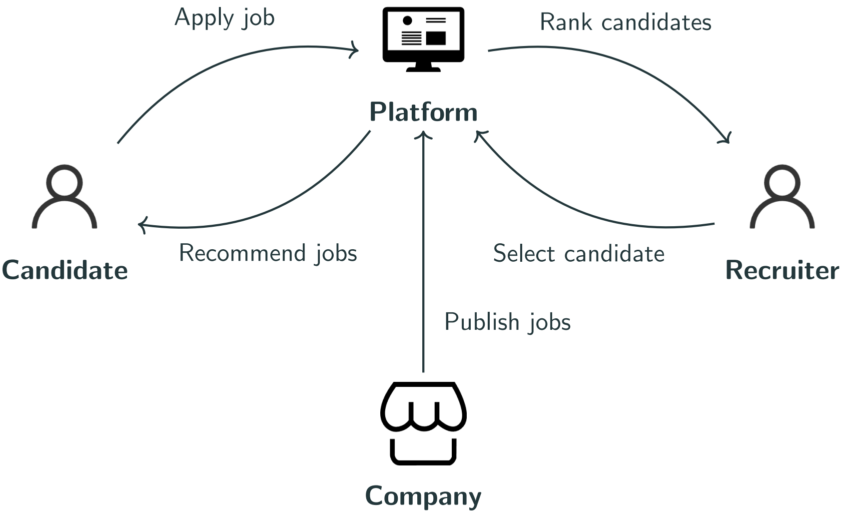

Fairness concerns

- Groups of candidates defined by protected attributes are often subject to discrimination in the interaction with the:
  - **platform**: not being exposed to well-payed jobs [Rus et al., 2022]
  - **recruiter**: not being in the top-k of the list, thus not being selected for an interview

Fairness concerns

Most existing approaches focus on **group fairness**, often ignoring individual qualifications and needs. This can unintentionally amplify existing stereotypes and biases.

Economic tools: Leverage **social choice theory** to incorporate individual qualifications while achieving group-fair outcomes.

Future and related work

Future work

- **Individual fairness** remains under-explored compared to group fairness.
- **Group fairness** approaches typically focus on a single binary protected attribute.
- The relationship and trade-offs between **group fairness** and **individual fairness** need further investigation.
- Adopting an **economic perspective** (e.g., micro- and macroeconomics, social choice theory) can provide new insights and solutions.

Related work

Individual fairness

- Evaluation Measures of Individual Item Fairness for Recommender Systems: A Critical Study
- Fair Ranking as Fair Division: Impact-Based Individual Fairness in Ranking
- Operationalizing Individual Fairness with Pairwise Fair Representations

Group fairness:

- Fair Top-k Ranking with multiple protected groups
- Balanced Ranking with Diversity Constraints
- Causal Intersectionality for Fair Ranking

5. Economic-based Fairness Mitigation and Evaluation Strategies III (Yuanna, 30 min)

Allocation time

- In this section, we focus on **allocation time**

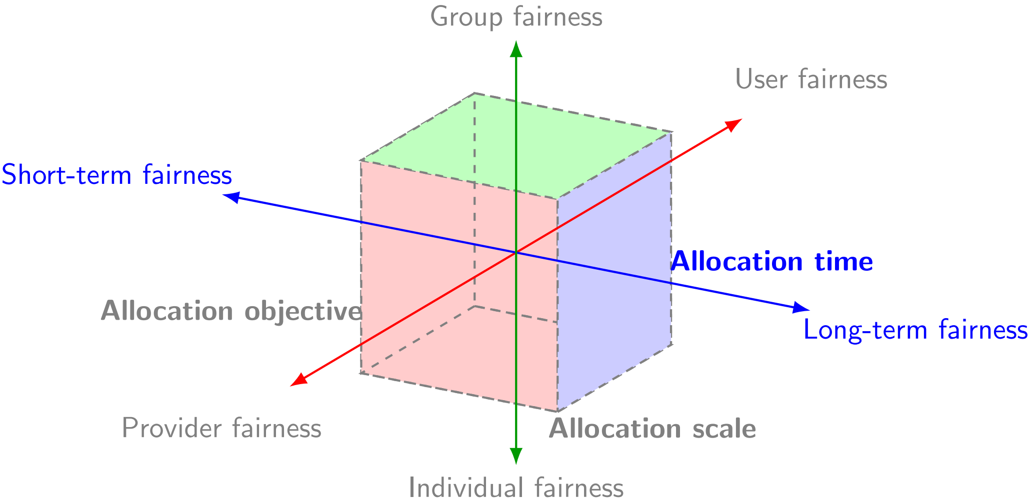

Dynamic allocation inspired short & long-term fairness

Dynamic interactions among stakeholders in IR

- **User**, **platform**, **items** and **provider** form a dynamic ecosystem [Abdollahpouri and Burke, 2019].
- Maintaining fairness for each of the changing stakeholders.

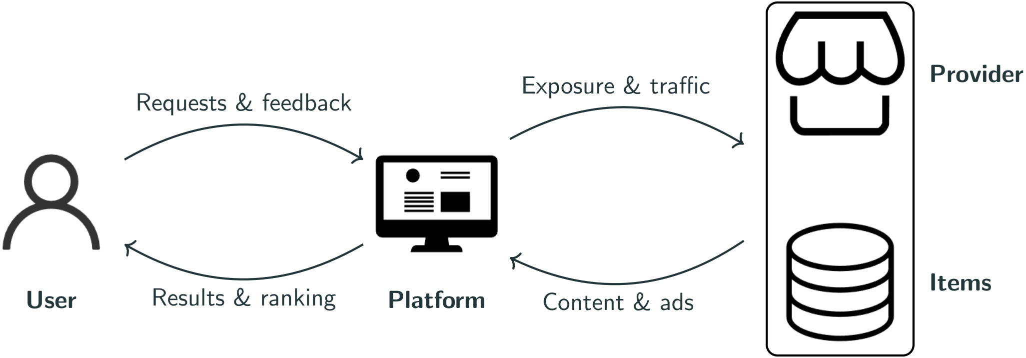

Short & long-term fairness in IR

- Short-term fairness (static fairness): most of work are situated in a **static** or one-shot setting, and the model provides a **one-time fairness solution** based on fairness-constrained optimization.

Short & long-term fairness in IR

- Short-term fairness (static fairness): most of work are situated in a **static** or one-shot setting, and the model provides a **one-time fairness solution** based on fairness-constrained optimization.
- Long-term fairness (dynamic fairness): due to the dynamic nature of IR systems, **attributes of each stakeholder** will change over time.
  - **Users** & user preference shift
  - **Ranking model** in the feedback loop
  - **Item** popularity, rating, content information, stock availability
  - **Provider** behavior

Formulation of long-term fairness in IR

Optimize ranking model and maintain the fairness constraint during time period $t = 1, 2, \ldots, T$.

$$\begin{aligned} \max \quad & \sum_t \gamma_r^t \, f(x) \rightarrow \text{accumulated reward w/ time discount} \\ s.t. \quad & \sum_t \gamma_c^t \, g(x) \leq c \rightarrow \text{accumulated fairness-related variable w/ time discount} \end{aligned}$$

or

$$\max \quad \sum_t (\gamma_r^t \, f(x) + \lambda \, (\gamma_c^t \, g(x))),$$

where $f(x)$ is the ranking utility function and $g(x)$ is the fairness-related function; $\gamma_r^t, \gamma_c^t \in [0, 1]$ are time discount rate.

Economic intuition of IR platform

Economic Intuition

Platforms must balance immediate utility vs. long-term fairness

Short-term Focus:

- Maximize current engagement
- Show popular/relevant items
- High immediate utility

Long-term Focus:

- Maintain fair exposure
- Include diverse/niche items
- Sustainable ecosystem

Ranking optimization through economic time discounting

An economist would see this as a **dynamic optimization problem**:

The platform chooses ranking r_t at each time t so that it is maximizing expected utility of the platform's engagement E over time:

$$\max_{r_t} \mathbb{E}\left[\int_0^T \underbrace{e^{-\rho t}}_{\text{Discount factor}} u(E_t)\, dt\right]$$

A higher discount rate ρ reflects a stronger preference for immediate engagement and utility over long-term outcomes.

Platform-specific calibration: Tunable Discount Rates

The discount rate ρ in our optimization framework can be **adjusted based on platform priorities**:

$$\max_{r_t} \mathbb{E}\left[\int_0^T e^{-\rho t} u(E_t) dt\right]$$

- **High** ρ: Short-term focused platforms (startups, growth phase)
  - Prioritize immediate engagement and user acquisition
  - Accept higher long-term fairness risks
- **Low** ρ: Long-term focused platforms (established, regulated)
  - Emphasize sustainable ecosystem health
  - Invest more in fairness and diversity

Engagement can be modeled as an uncertain time process

The platform's **engagement** E_t can be modeled **as a dynamic process** dependent on the platform's rankings and fairness.

$$\Delta E_t = f(r_t)\Delta t - \beta g(\boldsymbol{r})\xi_t \Delta t$$

- $f(r_t)$ is the immediate engagement outcome of ranking r_t
- $g(\boldsymbol{r})$ is the platform's fairness policy
- ξ_t is a random demand shock that can be positive or negative.

An unfair platform becomes more homogeneous and is therefore more vulnerable to shocks in consumer demand. This threatens *long-term* engagement of the platform.

What we gain from this economic lens:

- The discount rate ρ reflects the 'impatience' of the platform. A higher ρ prioritizes immediate utility, while a lower ρ promotes long-term fairness and sustainability.
- Future engagement depends on both current rankings and long-term fairness, due to vulnerability to demand changes.
- By summing over (discounted) future rewards, resilience of the platform is naturally taken into account.

Long-term fairness in IR

Long-term fairness methods that specifically model dynamic attributes of each stakeholder:

- **Item** popularity
- Performance disparity between **user** groups
- **Ranking model** in the feedback loop
- **Provider** behavior

Long-term fairness in IR: Item popularity

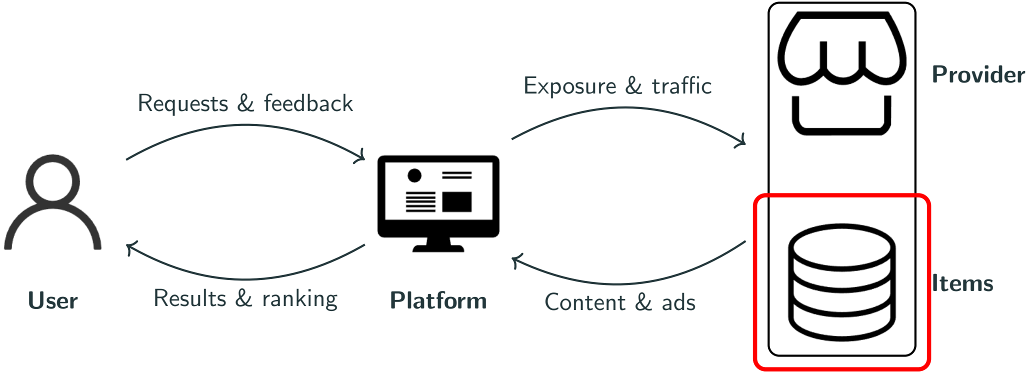

Long-term fairness in IR: Item popularity

In the dynamic recommender systems, **item popularity** may change over time due to the recommendation policy and user engagement [Ge et al., 2021].

Target: maintain long-term fairness of item exposure with changing group labels.

- Problem formulation: Constrained Markov decision process
  - State $\mathcal{S}$: user features (e.g., user's recent click history)
  - Action $\mathcal{A}$: recommendation list
  - Reward $\mathcal{R}$: user feedback, i.e., click, purchase
  - Cost $\mathcal{C}$: the number of recommended items that come from popular group
  - Discount rate of reward γ_r; discount rate of cost γ_c.

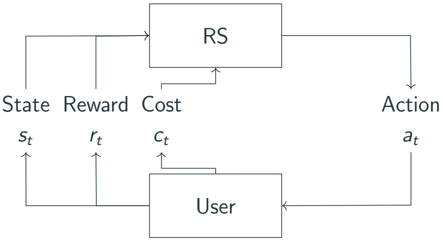

Long-term fairness in IR: Item popularity

- Fairness constrained policy optimization (FCPO)

$$\begin{aligned} \max_{\pi} \quad & J_R(\pi) \\ \text{subject to} \quad & J_C(\pi) \leq d \end{aligned}$$

  - Cumulative reward $J_R(\pi)$
  - Cumulative cost $J_C(\pi)$
  - Limit d: the limit is computed by fairness constraints $\frac{\text{Exposure}_t(G_0)}{\text{Exposure}_t(G_1)} \leq \alpha$
  - Aim to learn a policy π that maximizes reward while satisfying the fairness constraint.

Long-term fairness in IR: User fairness

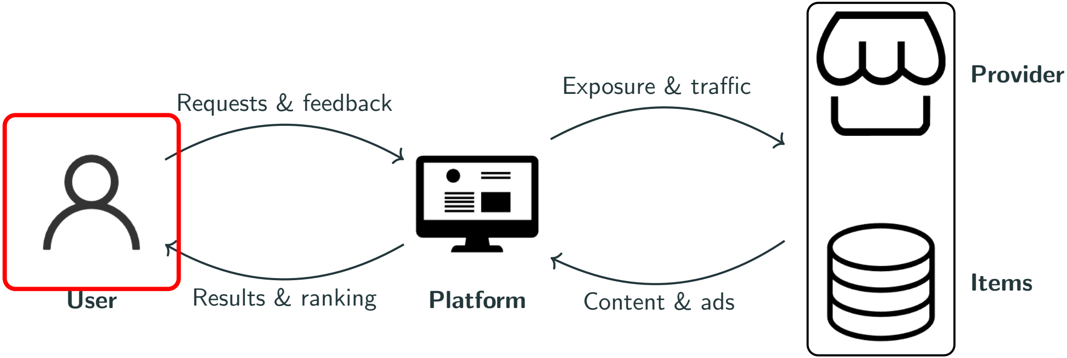

Long-term fairness in IR: User fairness

Neglecting user fairness during dynamic adaptation leads to **performance disparity** between user groups persisting or even expanding over time [Yoo et al., 2024].

$\mathcal{D}_1$	$\mathcal{D}_2$	$\mathcal{D}_n$
Female: 0.45	Female: 0.42	Female: 0.44
Male: 0.5	Male: 0.52	Male: 0.53
PD_1: 0.05	PD_2: 0.1	PD_n: 0.09

- Performance disparity: $\mathrm{PD}_t = \mathrm{Perf}(\mathcal{D}_t^{\text{test}}|\text{male}) - \mathrm{Perf}(\mathcal{D}_t^{\text{test}}|\text{female})$

Long-term fairness in IR: User fairness

- Problem formulation: incremental fine-tuning
- FAir Dynamic rEcommender (FADE) fine-tunes the model parameters incrementally over time only with the new data $\mathcal{D}_t$.
- Loss: $\mathcal{L}^{\mathcal{D}_t} = \mathcal{L}^{\mathcal{D}_t}_{\text{rec}} + \lambda \mathcal{L}^{\mathcal{D}_t}_{\text{fair}}$
  - $\mathcal{L}^{\mathcal{D}_t}_{\text{rec}}$ uses BPR loss
  - $\mathcal{L}^{\mathcal{D}_t}_{\text{fair}}$ is computed based on differentiable Hit (DH).
  - Model update: $\mathcal{W}_t := \mathcal{W}_t - \eta \nabla_{\mathcal{W}_t} (\mathcal{L}^{\mathcal{D}_t}_{\text{rec}} + \lambda \mathcal{L}^{\mathcal{D}_t}_{\text{fair}})$

Long-term fairness in IR: RS model in feedback loop

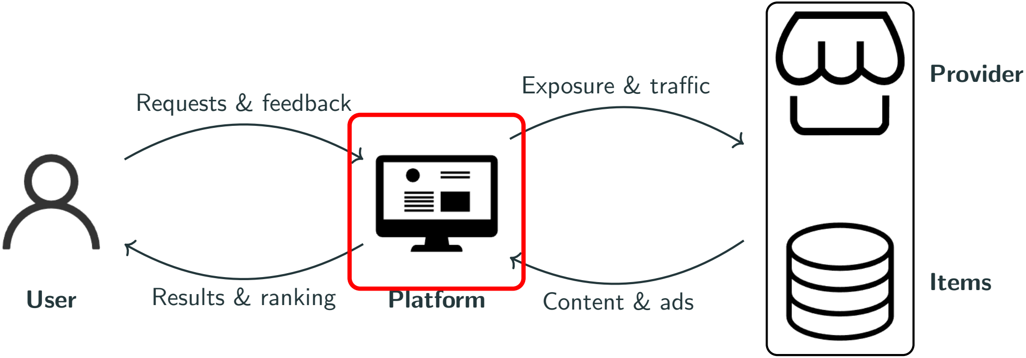

Recommendation feedback loops (RFL) will influence the provider Max-Min Fairness in the long term since RS can only receive feedback on exposed items, while **unexposed items** are considered as negative samples [Xu et al., 2023b].

- Problem formulation: Repeated resource allocation problem under batched bandit setting

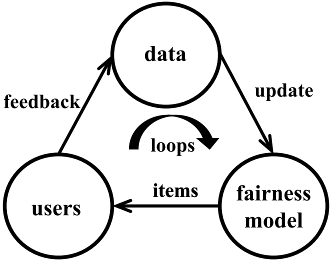

Long-term fairness in IR: RS model in feedback loop

- LTP-MMF: for a batch of users, accuracy-fairness-exploration score: $R = f(x) + \lambda g(x) + e(u, i)$. Then, collect users' feedback to update accuracy module.
- UCB module: explores the feedback of unexposed items.

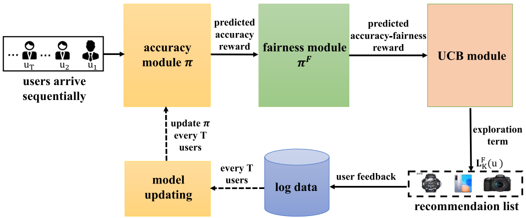

Long-term fairness in IR: Provider behavior

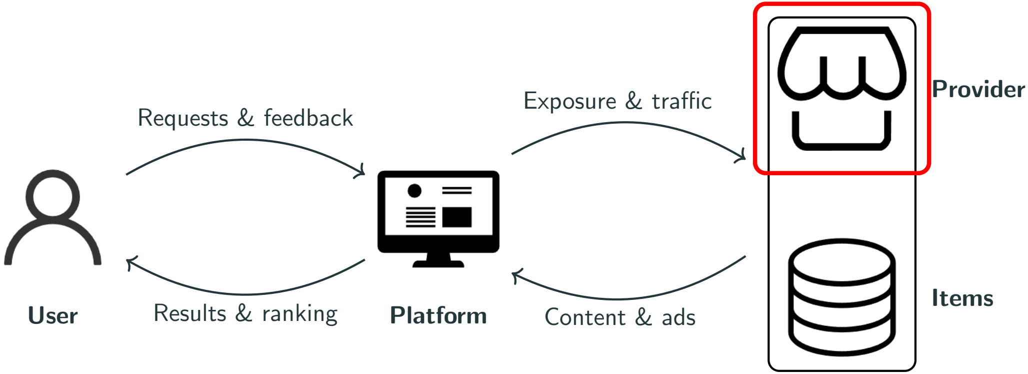

Long-term fairness in IR: Provider behavior

Content providers cannot **remain viable** unless they receive a certain level of user engagement. Myopic policies often drive the dynamical system to a poor equilibrium, with low user social welfare and poor provider diversity [Mladenov et al., 2020].

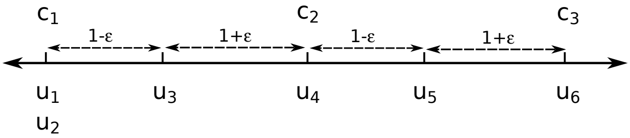

$\epsilon < 0.5, v_c = 2$

Myopic policy

- $c_1 : u_1, u_2, u_3$
  $c_2 : u_4, u_5$
  $c_3 : u_6 \Rightarrow c_3$ quit
- future reward: $8 + 2\epsilon$

Long-term policy

- $c_1 : u_1, u_2$
  $c_2 : u_3, u_4$
  $c_3 : u_5, u_6$
- future reward: 10 - 2ϵ

Long-term fairness in IR: Provider behavior

- Problem formulation: epoch-based optimal constrained matching problem

$$\begin{aligned} \max_{\pi} \quad & \sum_{u \in \mathcal{U}} f(u|\pi) \\ \text{s.t.} \quad & g(c) \geq v_c, \forall c \end{aligned}$$

  - objective: maximize social welfare (user utility) over the epoch
  - constraint: ensure that any matched provider remains viable

Application: Connection recommendation

People you may know

- Real-world connection recommendation systems are dynamic systems with feedback loops. Recommendations alter users' social graphs and influence future recommendation strategies [Akpinar et al., 2022].
- **Questions:**
  - What consequences and side effects can result from ignoring fairness interventions?
  - In the long run, are existing fairness interventions truly effective?

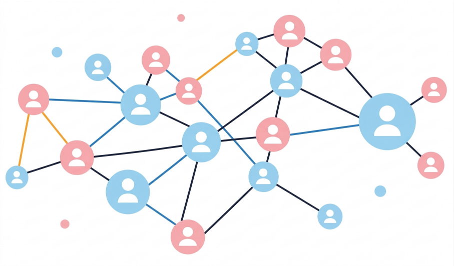

Rich get richer

- 65% male vs. 35% female
- No intervention
- Destination-side interventions:
  - Demographic parity of exposure (DP)
  - Dynamic parity of utility (Dyn)

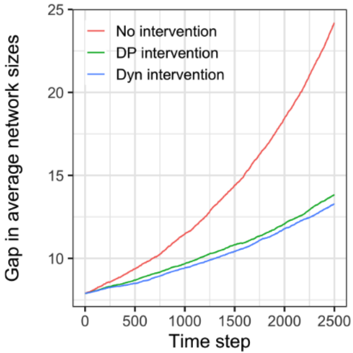

Lessons

- Only applying fairness interventions on the destination side (i.e., list of recommended people) fails to prevent the long-term amplification of inequality.

Lessons

- Only applying fairness interventions on the destination side (i.e., list of recommended people) fails to prevent the long-term amplification of inequality.
- **Source-side bias**:
  - Initial differences in the distributions of users' similarity and their number of common connections
  - Users with larger networks are served connection recommendations more frequently, and are more likely to form connections

Future and related work

Future work

All these long-term fairness works update RS model and consider the dynamic change of a certain stakeholder.

- Long-term fairness requires additional algorithm designs to maintain the sustainability of the system.
- Long-term fairness algorithms can draw on tools such as **dynamic optimization** in economics.
- How to model/simulate the changes of **multi-stakeholders**?
- How to use **LLM-powered agent** to simulate the long-term behavior of each stakeholder?

Long-term fairness in IR: Related work

RS model in feedback loop:

- Controlling Fairness and Bias in Dynamic Learning-to-Rank
- Maximizing Marginal Fairness for Dynamic Learning to Rank

Provider behavior:

- CreAgent: Towards Long-Term Evaluation of Recommender System under Platform-Creator Information Asymmetry

Survey:

- Long-Term Fairness Inquiries and Pursuits in Machine Learning: A Survey of Notions, Methods, and Challenges

6. Open Problems, Quick Start for Learning about Fairness, and Conclusions (Clara, 20 min)

Open problems

Future directions of fairness from an economic perspective

- Economics suggests future directions for fairness-aware IR
- Three levels of fairness [Rosenfeld and Xu, 2025]:
  - **Level-1**: Designing fair welfare functions (most papers)
  - **Level-2**: Incorporating platform Policys (few papers)
  - **Level-3**: Considering user/provider choices (few papers)

Current fairness-aware IR style

Adjust IR systems to meet fairness requirements!

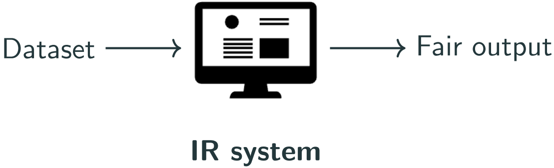

Level 1: Designing a fair welfare function

- **Level-1**: How to design a better welfare evaluation function?

Dataset ⟶  ⟶ Fair output ⟶ **1. Welfare functions**

Fair IR system

Level 1: Designing a fair welfare function

Objective: Can we design a unified fair welfare function for stakeholders?

- For **single** stakeholder (user, provider) [Xu et al., 2025b]
- For **multi-sided** stakeholders

Level 1: Designing a fair welfare function

Objective: Can we design a unified fair welfare function for stakeholders?

- For **single** stakeholder (user, provider) [Xu et al., 2025b]
- For **multi-sided** stakeholders

Scale: Can we design a unified fair aggregation function?

- **Single** layer aggregation (time, category)
- **Hierarchical** aggregation

Level 1: Designing a fair welfare function

Objective: Can we design a unified fair welfare function for stakeholders?

- For **single** stakeholder (user, provider) [Xu et al., 2025b]
- For **multi-sided** stakeholders

Scale: Can we design a unified fair aggregation function?

- **Single** layer aggregation (time, category)
- **Hierarchical** aggregation

Time: Can we design a unified long-term fair function?

- **Accumulated** fairness constraint

Level 2: Incorporating platform Policys

- **Level 2**: Incorporating platform policies: from **predictions** to **actions**

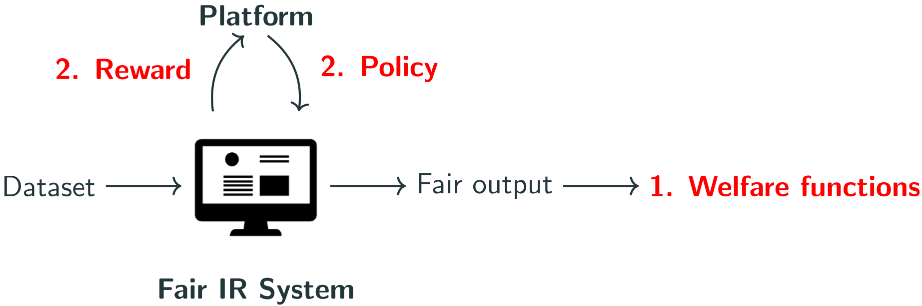

Level 2: Incorporating platform policies

Objective: Platform needs adapt different policy for stakeholders

- Incorporating platform and user/provider objectives

Level 2: Incorporating platform policies

Objective: Platform needs adapt different policy for stakeholders

- Incorporating platform and user/provider objectives

Scale: Platform policy influences different scales of stakeholders

- **Simulating and modeling** different scale of stakeholders

Level 2: Incorporating platform policies

Objective: Platform needs adapt different policy for stakeholders

- Incorporating platform and user/provider objectives

Scale: Platform policy influences different scales of stakeholders

- **Simulating and modeling** different scale of stakeholders

Time: Platform policy will influence both short and long-term fairness

- **Simulating and modeling** the change of platform policy

Level 3: Considering user/provider choices

- **Level 3**: User and provider are rational: change action according to utilities

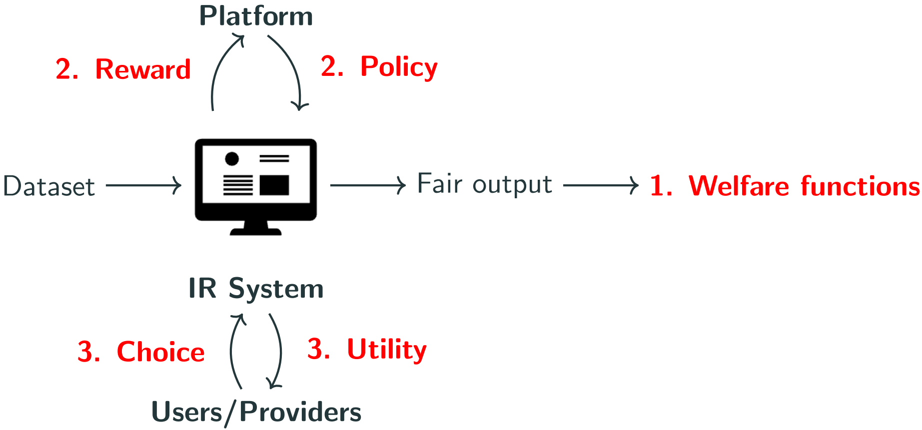

Level 3: Considering user/provider choices

Objective: Objective needs to consider user/provider's choice

- Game theory-inspired fairness objective for users/providers

Level 3: Considering user/provider choices

Objective: Objective needs to consider user/provider's choice

- Game theory-inspired fairness objective for users/providers

Scale: Different scale stakeholders make different choice

- **Micro**-individual behavior patterns
- **Macro**-group behavior patterns

Level 3: Considering user/provider choices

Objective: Objective needs to consider user/provider's choice

- Game theory-inspired fairness objective for users/providers

Scale: Different scale stakeholders make different choice

- **Micro**-individual behavior patterns
- **Macro**-group behavior patterns

Time: Choices of users and providers evolve over time

- Fairness **equilibrium** remains stable and aligned with the predefined objectives

Quick start for learning about fairness in IR

Toolkits: FairDiverse

- We have developed an easily-usable toolkit *FairDiverse* [Xu et al., 2025a] for learning about fairness in IR
- GitHub: https://github.com/XuChen0427/FairDiverse

Toolkits: FairDiverse

- We have developed an easily-usable toolkit *FairDiverse* [Xu et al., 2025a] for learning about fairness in IR
- GitHub: https://github.com/XuChen0427/FairDiverse
- Advantages
  - Contains **29** fairness algorithms across **16** base models for two fundamental IR tasks – search and recommendation
  - Contains **tens of fairness datasets** for fairness tasks
  - Offers **multiple APIs** (such as evaluation metrics) to enable IR researchers to quickly develop their own fairness IR models

Existing toolkits

Comparison of **FairDiverse** with existing toolkits:

Features	RecBole	FFB	Fairlearn	AIF360	Aequitas	**FairDiverse**
Recommendation	✓	✗	✗	✗	✗	✓
Search	✗	✗	✗	✗	✗	✓
Pre-processing	✗	✗	✓	✓	✓	✓
In-processing	✓	✓	✓	✓	✓	✓
Post-processing	✗	✗	✓	✓	✓	✓
Number of models	4	6	6	15	10	**29**

Toolkits: FairDiverse

- End-to-end coverage: From **data collection**, **data processing**, and **model training** to **result evaluation**
- Helps researchers understand & apply fairness in **structured, reproducible** way
- Helps researchers develop their **own fairness-aware IR models**

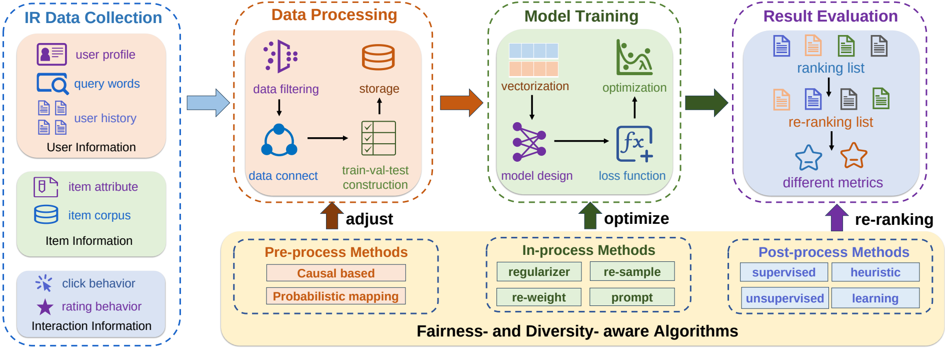

Conclusions

Economics provides a helpful framework for understanding fairness in IR

- Allocation objective, scale, and time

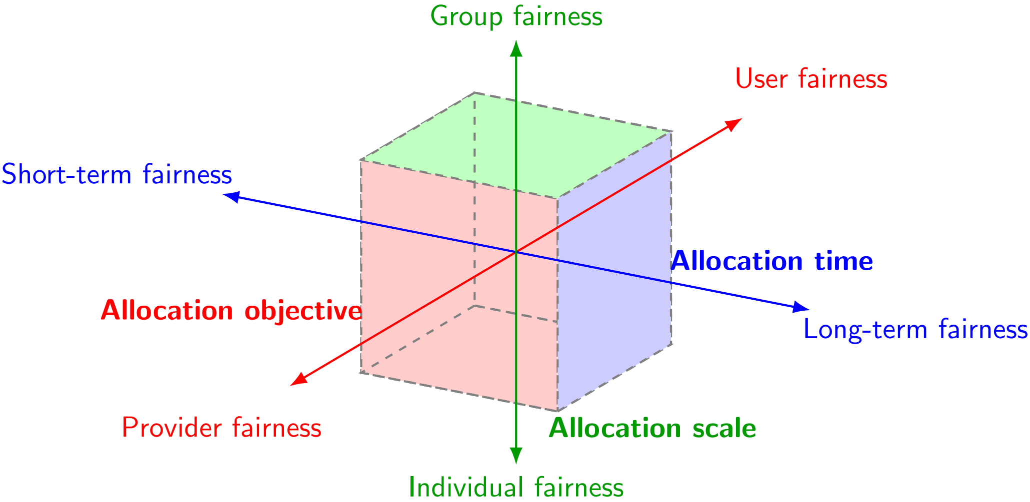

Economics perspective provides new tools

- Taxation, risk-return, game theory, social choice

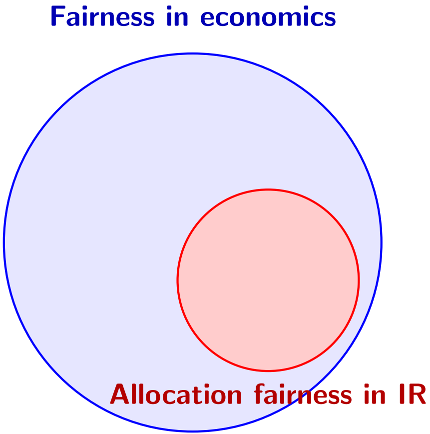

Applications

Fairness Objective
Next-basket
Recommendation

Fairness Scale
Recruitment
Search

Fairness Time
Connection
Recommendation

Using economic thinking for fairness in IR

- Fairness is not just "the right thing" but often also the "**profitable choice**"
- Fairness can be seen as a form of **anticipatory consumption**: it discounts future value to be accounted for in the present

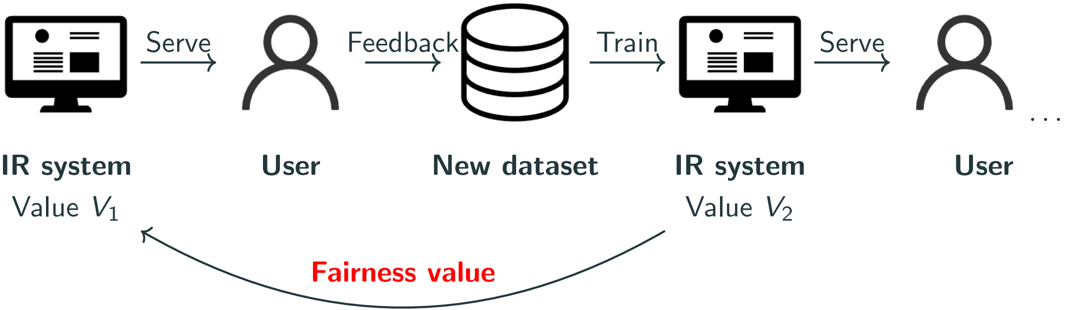

Economic perspective hints at future firections

- Three levels of fairness problem

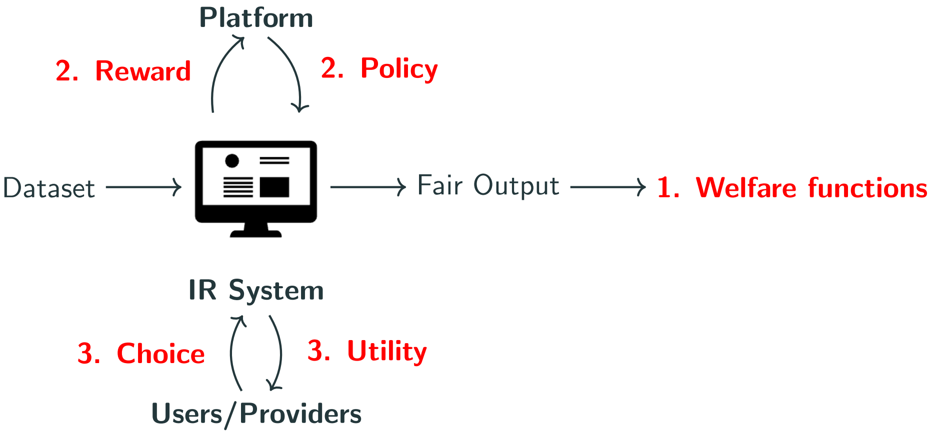

Our Tutorial Containing Messages

Summary (Part I)
Existing Fairness Concept
Existing Fair IR Algorithms

New Perspective (Part II-III)
Economics Market Cube
Economic Tools for Fair IR

Future Work (Part IV)
Three-levels of Fairness
Welfare, Policy, Choice

Future Plan

- This tutorial is only a first step toward a broader research agenda on economic perspectives on fairness in IR.
- We are currently planning to develop these materials into a book, where we will present more detail.
- The book and related resources will be gradually made available on our website in the future.

More updates will be shared at: https://economic-fairness-ir.github.io/

Related materials for exploring fairness in IR

Surveys:

- A Survey on the Fairness of Recommender Systems
- Fairness in Recommendation: Foundations, Methods and Applications
- Fairness in Ranking: A Survey
- Bias and Unfairness in Information Retrieval Systems: New Challenges in the LLM Era

Open toolkit:

- FairDiverse, RecBole2.0

Q&A

Website **FairDiverse toolkit**

Contact information: chenxu0427ruc@gmail.com

References i

H. Abdollahpouri and R. Burke. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158*, 2019.

P. Aggarwal, V. Murahari, T. Rajpurohit, A. Kalyan, K. Narasimhan, and A. Deshpande. GEO: Generative engine optimization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, 2024.

A. Aird, C. All, P. Farastu, E. Stefancova, J. Sun, N. Mattei, and R. Burke. Exploring social choice mechanisms for recommendation fairness in scruf. *arXiv preprint arXiv:2309.08621*, 2023.

A. Aird, P. Farastu, J. Sun, E. Stefancová, C. All, A. Voida, N. Mattei, and R. Burke. Dynamic fairness-aware recommendation through multi-agent social choice. *ACM Transactions on Recommender Systems*, 3(2):1–35, 2024a.

A. Aird, E. Štefancová, C. All, A. Voida, M. Homola, N. Mattei, and R. Burke. Social choice for heterogeneous fairness in recommendation. In *Proceedings of the 18th ACM Conference on Recommender Systems*, pages 1096–1101, 2024b.

N.-J. Akpinar, C. DiCiccio, P. Nandy, and K. Basu. Long-term dynamics of fairness intervention in connection recommender systems. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 22–35, 2022.

R. Baeza-Yates, B. Ribeiro-Neto, et al. *Modern information retrieval*, volume 463. ACM press New York, 1999.

J. Baumann, P. Sapiezynski, C. Heitz, and A. Hannak. Fairness in online ad delivery. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, page 1418–1432, New York, NY, USA, 2024. Association for Computing Machinery.

T. Bektaș and A. N. Letchford. Using ℓp-norms for fairness in combinatorial optimisation. *Computers & Operations Research*, 120:104975, 2020.

A. J. Biega, K. P. Gummadi, and G. Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international ACM Sigir Conference on Research & Development in Information Retrieval*, pages 405–414, 2018.

S. Chawla and M. Jagadeesan. Individual fairness in advertising auctions through inverse proportionality. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, pages 42–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.

J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, and X. He. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, 41(3):1–39, 2023.

V. Dang and W. B. Croft. Diversity by proportionality: an election-based approach to search result diversification. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 65–74, 2012.

V. Do, S. Corbett-Davies, J. Atif, and N. Usunier. Two-sided fairness in rankings via lorenz dominance. *Advances in Neural Information Processing Systems*, 34:8596–8608, 2021.

C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226, 2012.

S. N. Eisenstadt. Post-traditional societies and the continuity and reconstruction of tradition. *Daedalus*, pages 1–27, 1973.

D. García-Soriano and F. Bonchi. Maxmin-fair ranking: individual fairness under group-fairness constraints. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 436–446, 2021.

Y. Ge, S. Liu, R. Gao, Y. Xian, Y. Li, X. Zhao, C. Pei, F. Sun, J. Ge, W. Ou, et al. Towards long-term fairness in recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*, pages 445–453, 2021.

Y. Ge, J. Tan, Y. Zhu, Y. Xia, J. Luo, S. Liu, Z. Fu, S. Geng, Z. Li, and Y. Zhang. Explainable fairness in recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 681–691, 2022.

U. Gohar and L. Cheng. A survey on intersectional fairness in machine learning: Notions, mitigation, and challenges. *arXiv preprint arXiv:2305.06969*, 2023.

M. Heuss, F. Sarvi, and M. de Rijke. Fairness of exposure in light of incomplete exposure estimation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 759–769, 2022.

References v

R. K. Jain, D.-M. W. Chiu, W. R. Hawe, et al. A quantitative measure of fairness and discrimination. *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA*, 21(1):2022–2023, 1984.

Z. Jiang, X. Han, C. Fan, F. Yang, A. Mostafavi, and X. Hu. Generalized demographic parity for group fairness. In *International Conference on Learning Representations*, 2021.

R. John et al. A theory of justice. 1971.

L. Jost. Entropy and diversity. *Oikos*, 113(2):363–375, 2006.

P. Lahoti, K. P. Gummadi, and G. Weikum. ifair: Learning individually fair data representations for algorithmic decision making. In *2019 ieee 35th international conference on data engineering (icde)*, pages 1334–1345. IEEE, 2019.

J. Lamont. *Distributive Justice*. Routledge, 2017.

M. K. Lee, A. Jain, H. J. Cha, S. Ojha, and D. Kusbit. Procedural justice in algorithmic fairness: Leveraging transparency and outcome control for fair algorithmic mediation. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–26, 2019.

Y. Li, Y. Ge, and Y. Zhang. Tutorial on fairness of machine learning in recommender systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, page 2654–2657, 2021.

Y. Li, H. Chen, S. Xu, Y. Ge, J. Tan, S. Liu, and Y. Zhang. Fairness in recommendation: Foundations, methods and applications. *ACM Transactions on Intelligent Systems and Technology*, 14(5):Article No.: 95, 2023.

Q. Liu, X. Feng, T. Gu, and X. Liu. Faircrs: Towards user-oriented fairness in conversational recommendation systems. In *Proceedings of the 18th ACM Conference on Recommender Systems*, RecSys '24, page 126–136, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400705052.

Y. Liu, M. Li, M. Aliannejadi, and M. de Rijke. Repeat-bias-aware optimization of beyond-accuracy metrics for next basket recommendation. *arXiv e-prints*, pages arXiv–2501, 2025.

M. Mladenov, E. Creager, O. Ben-Porat, K. Swersky, R. Zemel, and C. Boutilier. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *International Conference on Machine Learning*, pages 6987–6998. PMLR, 2020.

Y.-K. Ng. *Welfare economics*. Springer, 1983.

N. Rosenfeld and H. Xu. Machine learning should maximize welfare, but not by (only) maximizing accuracy, 2025. URL https://arxiv.org/abs/2502.11981.

C. Rus, J. Luppes, H. Oosterhuis, and G. H. Schoenmacker. Closing the gender wage gap: Adversarial fairness in job recommendation. *arXiv preprint arXiv:2209.09592*, 2022.

C. Rus, M. de Rijke, and A. Yates. Counterfactual representations for intersectional fair ranking in recruitment. In *HR@ RecSys*, 2023.

C. Rus, A. Yates, and M. de Rijke. A study of pre-processing fairness intervention methods for ranking people. In *European Conference on Information Retrieval*, pages 336–350. Springer, 2024.

A. Singh and T. Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM international conference on knowledge discovery & data mining*, pages 2219–2228, 2018.

A. Singh and T. Joachims. Policy learning for fairness in ranking. *Advances in Neural Information Processing Systems*, 32, 2019.

N. Sonboli, R. Burke, N. Mattei, F. Eskandanian, and T. Gao. " and the winner is...": Dynamic lotteries for multi-group fairness-aware recommendation. *arXiv preprint arXiv:2009.02590*, 2020.

T. R. Tyler and E. Allan Lind. Procedural justice. In *Handbook of justice research in law*, pages 65–92. Springer, 2002.

A. Wang, V. V. Ramaswamy, and O. Russakovsky. Towards intersectionality in machine learning: Including more identities, handling underrepresentation, and performing evaluation. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, pages 336–349, 2022.

Y. Wu, S. Zhong, Y. Kim, and C. Xiong. What generative search engines like and how to optimize web content cooperatively. *arXiv preprint arXiv:2510.11438*, 2025.

C. Xu, S. Chen, J. Xu, W. Shen, X. Zhang, G. Wang, and Z. Dong. P-mmf: Provider max-min fairness re-ranking in recommender system. In *Proceedings of the ACM Web Conference 2023*, pages 3701–3711, 2023a.

C. Xu, X. Ye, J. Xu, X. Zhang, W. Shen, and J.-R. Wen. Ltp-mmf: Towards long-term provider max-min fairness under recommendation feedback loops. *arXiv preprint arXiv:2308.05902*, 2023b.

C. Xu, X. Ye, W. Wang, L. Pang, J. Xu, and T.-S. Chua. A taxation perspective for fair re-ranking. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, page 1494–1503, 2024.

C. Xu, Z. Deng, C. Rus, X. Ye, Y. Liu, J. Xu, Z. Dou, J.-R. Wen, and M. de Rijke. Fairdiverse: A comprehensive toolkit for fair and diverse information retrieval algorithms, 2025a.

C. Xu, J. Zhao, W. Wang, L. Pang, J. Xu, T.-S. Chua, and M. de Rijke. Understanding accuracy-fairness trade-offs in re-ranking through elasticity in economics, 2025b.

C. Xu, J. Zhao, W. Wang, L. Pang, J. Xu, T.-S. Chua, and M. de Rijke. Understanding accuracy-fairness trade-offs in re-ranking through elasticity in economics, 2025c.

X. Yang, Y. Li, H. Wang, D. Wu, Q. Tan, J. Xu, and K. Gai. Bid optimization by multivariable control in display advertising. In *Proceedings of the 25th ACM international conference on knowledge discovery & data mining*, pages 1966–1974, 2019.

H. Yoo, Z. Zeng, J. Kang, R. Qiu, D. Zhou, Z. Liu, F. Wang, C. Xu, E. Chan, and H. Tong. Ensuring user-side fairness in dynamic recommender systems. In *Proceedings of the ACM Web Conference 2024*, pages 3667–3678, 2024.

M. Zehlike, F. Bonchi, C. Castillo, S. Hajian, M. Megahed, and R. Baeza-Yates. Fa* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1569–1578, 2017.

References x

R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *International conference on machine learning*, pages 325–333. PMLR, 2013.