

# Recommender Ecosystems: A Mechanism Design Perspective on Holistic Modeling and Optimization

Craig Boutilier, Martin Mladenov, Guy Tennenholtz

Google Research

cboutilier@google.com, mmladenov@google.com, guytenn@google.com

## Abstract

Modern recommender systems (RSs) lie at the heart of complex *recommender ecosystems* (RESs) that couple the behavior of users, content providers, vendors, advertisers, and other actors. Despite this, the focus of much RS research and deployment is on the *local, myopic* optimization of the recommendations made to individual users. This comes at a significant cost to the *long-term utility* that RSs generate for their users. We argue that modeling the incentives and behaviors of these actors, and the interactions among them induced by the RS, is needed to maximize value and improve overall ecosystem health. Moreover, we propose the use of economic *mechanism design* (MD), largely overlooked in RS research, as a framework for developing such models. That said, one cannot apply “vanilla” MD to RES optimization out of the box—the use of MD raises a number of subtle and interesting research challenges. We outline a number of these here, emphasizing the need to develop nonstandard approaches to MD that intersect with numerous areas of research, including preference modeling, reinforcement learning and exploration, behavioral economics, and generative AI, among others.

## Why Model Recommender Ecosystems?

Recommender systems (RSs) play an ever-increasing role in our daily lives, mediating the search for information, the consumption of content, the purchase of goods and services, and even communication between individuals. This ubiquity amplifies the importance of research into effective models and algorithms that ensure RSs act in the best interests of their users and society at large. The majority of RS research and practice focuses on the *local, myopic* optimization of recommendations to *individual users*, in which a recommendation to one user does not consider the impact it might have on other users or interested actors. This overlooks the fact that most RSs lie at the heart of complex *recommender ecosystems* (RESs), in which the RS mediates and induces *interactions* between and among users, content creators, vendors, etc. These interactions, in turn, can impact—both positively and negatively—the ability of the RS to make high-quality recommendations in the future, a fact we illustrate below. (See Abdollahpouri et al. (2020) a comprehensive overview of the multiagent perspective on RSs.)

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This leads to a set of research questions pertaining to such RESs: (i) How should we analyze and model agent interactions in an RES? (ii) How do we optimize RS policies in the face of such dynamic interactions? (iii) What criteria and objectives should be used to guide this optimization?

In this research challenges paper, we outline a research agenda intended to develop a deeper understanding of RESs, and encourage research into methods, models and algorithms for the design of RS policies that maximize long-term user utility and overall social welfare across the ecosystem. An important part of this agenda is the incorporation of concepts, methods and perspectives from economic *mechanism design* (MD) (Hurwicz 1960). MD has played at best a minor role in the design of RSs,<sup>1</sup> in large part due to the local, single-user focus typical of RS research. When moving to RESs, modeling and optimization in the presence of interacting agents must account for: (i) their incentives and behaviors; (ii) potential information asymmetry; (iii) the potential for strategic behavior; and (iv) tradeoffs in the value generated for different agents. The design of mechanisms (here, RS policies) in such settings is precisely the province of MD. However, the complexity of RESs poses unique challenges for traditional MD, some of which we outline.

We begin by outlining a very stylized RS model to ground our discussion, illustrating how *local* RS policies—those that ignore RES interactions—degrade long-term user utility. Next we describe key concepts of classical MD, and how these map to elements of an RES. We then argue that decidedly “nonstandard” MD models are needed, describing four classes of research challenges: (1) we first outline the subtleties of agent incentives, preference modeling and elicitation; (2) we then discuss private information and its role in

<sup>1</sup>With notable exceptions, a few of which we list. Of special note is work using MD, game-theoretic or equilibrium analysis (Ben-Porat and Tennenholtz 2018; Mladenov et al. 2020; Hron et al. 2022; Jagadeesan, Garg, and Steinhardt 2022; Kurland and Tennenholtz 2022; Ben-Porat et al. 2022; Cen and Ilyas 2022; Liu, Mania, and Jordan 2020). Also of relevance is work on fairness in RSs (Akpınar et al. 2022; Asudeh et al. 2019; Basu et al. 2020; Biega, Gummadi, and Weikum 2018; Heuss, Sarvi, and de Rijke 2022; Wu et al. 2022; Mehrotra et al. 2018); and work on phenomena such as filter bubbles, polarization, and popularity bias (Pariser 2011; Abebe et al. 2018; Amelkin, Bullo, and Singh 2017; Ribeiro et al. 2020; Abdollahpouri et al. 2019).




Figure 1: (a) An initial myopic/local RS matching of users to content/creators; (b) the equilibrium induced by myopic matching; (c) a non-myopic, ecosystem-aware matching (in equilibrium).

agent decision making; (3) next we address the potential for strategic behavior by agents; and (4) finally we examine the challenge of designing social choice functions. We refer to a greatly expanded version of this paper (Boutilier, Mladenov, and Tennenholtz 2023) for a more in-depth discussion of the challenges discussed here, related problems, and additional research directions.

### Some Stage Setting

**The Stylized Recommender.** We assume a stylized RS, which embeds its users  $\mathcal{U}$  and items  $\mathcal{I}$  in some latent embedding space—we equate a user  $u$  and item  $i$  with their respective embeddings. The RS uses cosine similarity, dot products, or (inverse) distances to measure user  $u$ ’s *affinity* for item  $i$ . For now, we treat affinity as user utility.<sup>2</sup>

For expository purposes, we assume items embody some form of *content* (e.g., news, music, video) and that each item  $i \in \mathcal{I}$  is generated by some content provider or creator  $c(i) \in \mathcal{C}$  (Konstan et al. 1997; Jacobson et al. 2016), though everything that follows applies to other settings (e.g., product recommendation).<sup>3</sup> For ease of exposition, we assume each  $c \in \mathcal{C}$  creates a single item (which can be altered, see below) and equate creators with their items. Similarly, user affinity for items varies with their request, context, etc., and embeddings are dependent on these factors, but for simplicity, treat  $u$  as a latent point rather than a set or distribution.

**Ecosystem Interactions: An Illustration.** We first illustrate how even simple dynamics can induce complex multi-agent interactions that require non-trivial optimization by the RS. We adapt the stylized scenario of Mladenov et al. (2020) to show how ignoring ecosystem interactions prevents the RS from maximizing long-run user utility. Fig. 1(a) uses (inverse) distance to reflect affinity between users (circles) and the content of specific creators (rectangles): we have two creators (P1 (red), P2 (blue)), and eight users (five (red) of whom are closer to P1, and three (blue) closer to P2).

As users request recommendations, an *omniscient*, but otherwise typical, RS matches each request to the creator with which the user has greatest affinity (see Fig. 1(a)). This matching is *myopic*—it maximizes immediate affinity for the user—and *local*—it does not consider impact the recommendation might have on other users. If each user issues a single request during a given time period, P1 engages with

<sup>2</sup>This captures, many forms of collaborative filtering, e.g., matrix factorization (Salakhutdinov and Mnih 2007) or dual encoders/two-tower models (Yi et al. 2019; Yang et al. 2020).

<sup>3</sup>For instance, if the items are products for sale, one should read “vendor” for “creator,” “sales” for “engagement,” etc.

five users, and P2 three. Now suppose that a creator who does not attract four users per period abandons the RS (e.g., for monetary or social reasons). Here P2 abandons the system after the first period, leaving the RS no choice but to recommend P1 to the three blue users at all subsequent periods (Fig. 1(b)). At this point, the system is in *equilibrium*.

Unfortunately, this equilibrium leaves blue users worse off than had P2 remained, showing that ignoring *creator behavior* induced by the RS policy can negatively impact *users*. By contrast, a *non-local, non-myopic* matching (see Fig. 1(c) where red user  $u$  is matched to P2) anticipates the distribution of user requests, and incentive-induced creator behavior, to optimize overall user welfare in the *long run*. Since  $u$  is nearly indifferent to P1 and P2, its small sacrifice in utility enables a large increase in the long-run utility of the blue users. This equilibrium maximizes *user (utilitarian) social welfare*, and is arguably more “fair”. However, it imposes a small cost on  $u$  and P1. Indeed, *optimality* of this equilibrium depends on the choice of objective, or in MD terms, the *social choice function*.

This example suggests treating an RS policy more holistically as an (online or offline) matching problem (Mehta 2013; Su, Bayoumi, and Joachims 2022). Of course, it is unrealistic in the many assumptions required to induce equilibrium. Among them: extremely simple creator preferences and dynamics; stationary, simplistic user preferences for content; no user dynamics; full information on the part of the RS about user preferences and creator incentives/behavior; lack of strategic behavior by creators or users; no outside options (e.g., other RS platforms); and a simplistic objective function. Developing RS policies that work well when we relax these assumptions poses numerous research challenges, and is critical if RSs are to act in the best interests of users, creators, and society as a whole.

### Why Mechanism Design?

The example above illustrates that the design of an RS in the ecosystem context requires some form of optimization that accounts for: (i) the preferences and incentives of *all* actors (e.g., users and creators) that engage with the RS; (ii) how these preferences—and the RS policy—influence their behavior; and (iii) suitable tradeoffs over agent preferences. Moreover, when we relax the strong informational assumptions embodied in the stylized model above, the RS must: (i) handle incomplete and noisy information about the preferences and other private information held by these actors (e.g., their beliefs, abilities, behavioral and decision making processes); (ii) take steps to extract or refine this information when possible; and (iii) incentivize the actors to behave in a way that advances the overall objectives of the RS.

This form of optimization is precisely the province of *market and mechanism design* (Hurwicz and Reiter 2006) which involves policies for interacting with—and making decisions that impact—a set of self-interested agents—each of whom holds some private information—to optimize some objective that depends on the preferences of those agents. We first describe classical MD, then outline how the MD perspective can help in the design of complex RESs.

**Classical Mechanism Design.** Informally, classical MD encompasses several key ingredients. We have an *action space*  $A$ , where each  $a \in A$  induces a distribution over some *outcome space*  $O$ . Each *agent*  $g \in G$  has preferences over outcomes, or a *utility function*  $v_g$ , which is unknown to the *mechanism*  $M$ .<sup>4</sup>  $M$  also has a *social choice (or welfare) function (SCF)*  $C$  which, given any vector  $\mathbf{v} = \{v_g\}_{g \in G}$  of agent utility functions, determines the “social utility”  $C(\mathbf{v}, o)$  of any  $o \in O$ . The goal of  $M$  is to *implement* the SCF by taking the action  $a_{\mathbf{v}}^* = \arg \max_{a \in A} \mathbb{E}_O C(\mathbf{v}, o)$ . To do so,  $M$  must take steps to observe, elicit, estimate or otherwise determine  $\mathbf{v}$  by interacting with the agent by (a) making available a set of *strategies/agent actions*; and (b) specifying a mapping from the strategies selected by each agent to a choice of outcome  $o \in O$ . An agent  $g$  may act strategically to manipulate the selected  $o$  to maximize its own utility  $v_g(o)$ , at the expense of  $C$  (and other agents). As such, one attempts to design  $M$  so that, if *agent strategies are selected to be in (some form of) equilibrium*, it will (via  $M$ ’s outcome mapping) induce the outcome dictated by the target SCF. This is known as *implementation of the SCF*.

**Mapping RESs into MD.** MD invokes critical concepts that can be applied to the design of complex RESs, including: (i) estimation of agent incentives; (ii) tradeoffs across agent preferences in the SCF; (iii) information asymmetry/sharing; and (iv) agent strategic behavior.

The mapping of an RES into a “standard” MD formalization is *prima facie* straightforward. The *agents* (in the MD sense) are those impacted by the actions of the RS. Here we focus on users and content creators, but other actors of interest include content distributors, product vendors, advertisers, external organizations, regulatory bodies, other RS platforms, etc. MD actions are *RS actions*. In our simple setting above, these are recommendations to users; but actions should be taken to be the *joint* set of recommendations made to all users. More broadly, actions are the set of *RS policies*, a term we use to emphasize the multi-stage nature of the RS’s choices (e.g., across users or across time). These are distinct from the *agent actions* available to users, creators, etc. (i.e., *strategies* in MD). The outcome space reflects the range of effects RS actions can have on users and creators (in the MD sense, the *joint* outcome over all actors). In simple settings this might include the content consumed (user) or user engagement (creator). Agent utilities reflect their preferences over such outcomes (usually the parts of the outcome “relevant” to them). Finally, the SCF encodes the RS’s objectives.

The full complexity of realistic RESs, of course, presents a rather different picture. When designing RS policies to maximize the long-term benefit to users, creators and other participants, the space of outcomes must be enriched substantially. Models of user and creator preferences are much more complex as a result, which renders preference assessment more difficult; this, in turn requires RS policies to actively explore/ elicit and act under incomplete information. Finally, committing to an explicit SCF is challenging given

<sup>4</sup>More generally, each agent has private information encoded in its *type*, which reflects any information not known to the designer or other agents (e.g., a content creator’s “skill”).

this preference complexity, as well as the need to account for social objectives. We now turn to some of the research challenges that need to be addressed to manage this complexity.

## Agent Preferences & Incentives

We first outline some of the challenges associated with managing realistic agent preferences in complex RESs

**Outcome Spaces.** Maximizing utility for users, creators and other agents requires capturing the true complexity of outcomes induced by the RS, beyond one-shot engagement metrics. For instance, the utility-bearing outcomes for a user often comprises an extended *consumption stream*. Utility for the sequence may not be decomposable into scores of isolated items (e.g., a user who values topic diversity). Thus, the RS action space is complex, as it must adapt its recommendations suitably; it comprises the set of *policies* that map the current user interaction history (in general, that of *all* users) to the choice of recommendation for a specific user.

In some cases, the (utility-bearing) outcomes may not be directly observable to the RS. For example, if a user accepts a product recommendation by purchasing the product (e.g., a camera), their ultimate utility will depend on the degree to which the product, its features and its functionality serves their needs over time (i.e., means vs. ends objectives (Keeney 1992)). *Latent factors* also play a role in user outcomes, e.g., a user’s true satisfaction with a music playlist is generally not observable. Some latent factors may reflect user utility, while others may be directly outcome-related, e.g., a user’s state of knowledge, influenced by consumption of news content on some issue.

For creators, outcomes may extend beyond cumulative user engagement, and capture temporal or other properties of that engagement over time. A creator may value smoothness in the user traffic, or user diversity across their audience. This requires an enriched outcome space beyond that typically considered in either MD or RS research.

**Passive Preference Assessment.** Unlike classic MD preference revelation, RSs often estimate user preferences from the past behavior (e.g., clicks, views, consumption) of both the user in question and *other users* (e.g., collaborative filtering (Koren, Rendle, and Bell 2011)). Of course, outcome-space complexity makes such estimation challenging.

The problem of incentivizing agents to reveal private information truthfully when it is to be *aggregated* in a predictive model is studied under *incentive-compatible machine learning* (Dekel, Fischer, and Procaccia 2010; Vorobeychik 2023). Adapting such models to the complexity of RESs would be fruitful, though strategic reporting on the part of users may be rare in RESs. That said, generalization error should be accounted for in any assessment of user preferences. Moreover, the assumption that user behavior (e.g., item choice) indicates their preferences ignores potential cognitive biases (Camerer, Loewenstein, and Rabin 2003), position or popularity bias, etc. The noise, biases and incompleteness of affinity models demands richer optimization techniques than those usually considered in MD.

**Active Exploration.** Passive preference assessment above takes no explicit steps to reduce RS uncertainty about a user’s preferences; models are trained on “organic” user behavior. This stands in contrast to MD, where explicit actions induce (possibly indirect) revelation of preferences. Active exploration is common in RSs, e.g., using (*contextual*) *bandit methods* to present novel items to users to refine models of their preferences (Li et al. 2010). Active exploration can be viewed as incremental, indirect preference revelation in the MD sense. Generalization across users remains important, which raises questions of a multiagent nature, since the value of making an *exploratory recommendation* to one user—one not predicted to be best for that user in the moment—may depend on how it improves estimates of *other users’* preferences. Recent work has begun to consider strategic elements that emerge in exploration (e.g., Liu, Mania, and Jordan 2020).

**Explicit Preference Elicitation.** Methods for explicit *preference elicitation*, studied in RSs and other areas (Salo and Hämäläinen 2001; Rashid et al. 2002; Boutilier 2002; Toubia, Hauser, and Simester 2004; Pu and Chen 2008), question a user about their preferences for items, akin to incremental, *direct* revelation in MD. *Critiquing methods* (Burke 2002) can also be viewed as a form of elicitation with more direct user control: when an item is recommended, a user can critique one of its attributes (e.g., a less expensive restaurant, more upbeat music). Handling open-ended item critiquing has become increasingly important as *conversational recommenders* become more prominent (see below). This poses new challenges for RSs in assessing user preferences, including: understanding how open-ended utterances reflect a user’s underlying preferences; and dealing with the subjective nature of attribute usage (Radlinksi et al. 2022).

*Interpreting* user responses to direct elicitation queries must account for the types of cognitive biases mentioned above. While the revelation principle has meant that direct revelation is most widely studied in MD, much research considers the role of incremental and partial revelation, especially in complex outcome spaces (e.g., combinatorial auctions (Sandholm and Boutilier 2006)). The principles underlying such mechanisms should play a central role in MD for RESs. Another challenge is the fact that most preferences are *contextual* (i.e., depending on a user’s current context, e.g., location, activity, companions, mood) and *conditional*; sequential recommendations are an important special case.

*Conversational recommenders* (Christakopoulou, Radlinksi, and Hofmann 2016) allow more flexibility in a user’s interactions with an RS, including open-ended dialogue; richer forms of steering, critiquing, preference elicitation; and user probing/exploration. This enables the RS to develop a more nuanced understanding of a user’s preferences and context. While generative and foundation models (Devlin et al. 2018; Radford et al. 2018; Thoppilan et al. 2022) hold promise for highly performative conversational RSs, significant challenges remain, including developing models that are inherently personalized by blending the rich, behavior-based models of users and items commonly used in CF-based RSs with the semantic understanding of items and

users afforded by large language models (LLMs). Incorporating multi-modal interactions into RSs also offers new opportunities, e.g., using text-to-image models (Ramesh et al. 2021) to synthesize new content or stylistic variations of products for more efficient user exploration or critiquing.

*Preference construction* (Lichtenstein and Slovic 2006) may play a role in MD, since users (and providers) often do not have fully formed preferences (e.g., due to unfamiliarity with the item corpus). Handling dynamic construction of preferences, especially when these are influenced by the RS itself, is an important challenge.

**Creator Incentives.** While we focused on user preferences above, similar arise when assessing or eliciting the preferences (and other private information) of providers (e.g., creators, vendors). Moreover, engaging with providers is more likely to involve both strategic revelation and information sharing, topics we address below.

## Information & Effective Decision Making

The “health” of an RES is typically evoked to refer to the ability of an RS to generate diverse recommendations for its users. We take a broader view, equating RES health with *the ability of the RS to generate significant value for all of its participants over the long run*. Importantly, we think of health not as a snapshot of utility generation, but rather as the RS’s ability to anticipate and respond to fundamental changes in the underlying ecosystem (e.g., as user preferences/tastes evolve, as new creation or production capabilities emerge, or as user/creator communities form and disband). Moreover, we expect an RS to take actions to promote, or at least facilitate, beneficial changes.

Since the value created by the RS depends on both the production decisions of creators and the consumption decisions of users, the RS can impact ecosystem health directly by supporting this decision making. One impediment to effective decision making is the substantial *information asymmetry* that exists between the RS and creators (or vendors, etc.). The RS has rich models of user preferences over the existing item corpus that often generalize *out-of-corpus* to some degree, serving as model of *latent user demand*. The RS also has a holistic view of the corpus itself, and insight into the abilities of providers to source/create new items, giving a deep understanding of both current and *potential* supply of content. No provider has the same breadth of insight into global demand or supply; this information asymmetry limits a creator’s ability to make informed content generation decisions, and is a key source of economic inefficiency. For example, a creator whose content does not attain the desired user engagement may not be able to determine the cause, such as: (i) no demand for this type of content; (ii) content quality that makes it unattractive to most users; or (iii) numerous other creators offering similar items. By contrast the RS can distinguish these causes.

Breaking the information asymmetry through some form of direct or indirect information sharing can improve creator decision making and drive significant improvements in both user and creator utility. For instance, if the competitive landscape limits a creator’s audience (cause (iii) above), the RS

could communicate this *directly*, or provide more *indirect* guidance by “steering” the creator to a less well-supplied part of content space with high predicted demand. Such information sharing poses numerous research challenges:

- Direct information sharing may not be feasible, e.g., if it reveals personal user data, strategically important information about “competitors,” or sensitive information about RS policies. Indirect or *implicit* sharing may be more acceptable (e.g., predicted aggregate audience for new content). Leakage of private information must be safeguarded of course, even in indirect methods (Dwork et al. 2012; Chien et al. 2021).
- Sharing information alone may not suffice to induce welfare-improving changes by a creator. Costs and uncertainty (e.g., due to lack of experience, access to resources) may discourage a creator from generating truly novel content. Mechanisms must be investigated that incentivize new content production, de-risk creator exploration and facilitate development of new skills.
- Generative models, such as LLMs (Thoppilan et al. 2022), or text-to-image models (Ramesh et al. 2021), may be used to synthesize language descriptions, item features, evocative images, etc. to guide the creator.
- RS steering or advice must be *coordinated* across the entire set of creators (e.g., an RS may not want to propose generation of the *same* novel content to multiple creators).
- Strategic creator responses may arise (see below).

Information sharing is tied directly to the revelation component of MD, since the RS may be best served by eliciting/assessing private information about a creator’s skills, beliefs, costs, and decision-making processes to better decide how to encourage welfare-improving decisions.

## Strategic Behavior

By *strategic behavior*, we refer to actions taken by an agent that anticipate the actions/reactions of others. Handling, or obviating the need for, strategic manipulations of the type below is a vital component of MD in RECs.

**User Strategic Behavior.** We largely expect users to behave non-strategically. When presented with several recommendations, a user will (perhaps noisily) select their most preferred option. In sequential settings, this may be trickier since an item’s value may depend on future recommendations, and may influence future recommendations that user receives. This requires the user be sequentially rational (i.e., plan)—but not strategic—and possibly invoke a “mental model” of the RS policy to explicitly influence subsequent recommendations (Guo et al. 2021); e.g., a user selecting a (non-preferred) music track by some artist to induce future recommendations of (preferred) tracks by that artist.

Strategic behavior might involve spam-like activity to promote the popularity of a favorite musical artist, news outlet, or content creator; or to provide excessive ratings or glowing product reviews to increase the odds of certain items being recommended to others. Likewise, responses to

preference elicitation queries may intentionally be inaccurate to manipulate the RS’s future recommendations to the user in question, to other users, or to impact the providers of the recommended items (either positively or negatively). That said, user strategic behavior, if existent, is likely to exhibit a degree of bounded rationality w.r.t. full RES complexity. We note that the increased use of explicit *two-way* communication about a user’s preferences, as advocated above, should increase transparency and user trust, and in turn reduce incentives for users to expend cognitive effort for “non-strategic” manipulation of an RS policy.

**Provider Strategic Behavior.** Strategic reasoning is more likely to play a role in the behavior of item providers. Disregarding direct intervention by the RS, a product vendor offers products that it predicts to have reasonable demand *and* to be sufficiently differentiated from those of other vendors (in equilibrium). Price setting will likewise be strategic. Similar considerations arise in content RSs even if users do not (directly) pay for content. An RS that tries to maximize some form of social welfare must account for provider strategic behavior; indeed, the RS matching policy itself can induce strategic behavior (Ben-Porat and Tennenholtz 2018). Extending such MD approaches to richer user and creator preferences is of great import.

Direct elicitation of private information may also induce manipulation. To wit, in our example, where a provider requires a minimum audience, directly eliciting this target gives the provider a *prima facie* incentive to overstate their target, but not by too much (else they risk being shut out completely), much like a first-price auction. While any direct revelation should account for equilibrium provider behavior, content RSs often cannot use monetary transfer. This falls within the area of *MD without money* (Procaccia and Tennenholtz 2009), and can be especially challenging.

Strategic considerations may also emerge when an RS takes steps to induce providers to generate new items that will improve overall social welfare (see above). For instance, consider two creators producing content designed to appeal to the same audience, which they split. The RS might prompt the first creator to make slightly different content, to the benefit of the user population as well as both creators. However, the first creator might refuse in the hopes that the RS might then suggest that the *second* creator “move” instead, thus sparing it the potential cost and risk of this change.

Incentivizing welfare-improving behavior across a diverse set of strategic or semi-strategic providers requires significant effort involving many aspects of MD. While direct (e.g., auction-like) mechanisms may work in simple domains, complex settings need detailed modeling of provider private information (utilities, costs, skills, beliefs, etc.) and how these elements shape their strategies, especially when providers exhibit bounded rationality.

## Tradeoffs and the Social Choice Function

One of the thorniest problems in the use of MD for RESs is the adoption of a particular SCF. The ecosystem perspective makes clear that an RS must trade off the utility it generates for different actors (users, creators, vendors, etc.), whose in-

centives will not be fully aligned. We consider two classes of tradeoffs: those involving preferences of the individual actors, and those that we refer to as “social tradeoffs.”

**Actor Utility Tradeoffs.** A key principle of MD is the use of an SCF to encode the tradeoffs the designer makes over the utilities of its participants. Our stylized example illustrates this for RESs, where even maximizing simple utilitarian user welfare requires such tradeoffs (the large increase in blue-user utility imposes a small cost on red users; likewise for the providers). The use of MD and an SCF to optimize RS policies forces the RS designer to explicitly articulate the tradeoffs they are prepared to make—which outcomes, across all actors in the system, they consider more or less desirable—rather than leaving them to chance.

While we used (*user*) *utilitarian welfare* (or sum of user utilities) in our motivating example, this is only for illustration. Naturally, creator utility can be incorporated. But more broadly, a variety of factors can be used in the construction of an SCF: overall welfare; distribution of utility (worst-case, in expectation); fairness w.r.t. welfare, opportunity, etc.; regret considerations (e.g., relative to a local baseline policy); and many others. This raises the question of who should engineer the SCF, and how. Unfortunately, there is no “universal” SCF that embodies incontrovertible principles for preference aggregation (Arrow 1950). Even seemingly uncontroversial conditions like Pareto optimality can be called into question on philosophical grounds (Sen 1970).

The use of an SCF poses challenges for specification, elicitation, assessment, measurement, and optimization. For instance, most RSs rarely take steps to assess true user or creator utility, instead relying on proxy measurements (e.g., return engagement). The self-interested nature of the RES participants renders explicit assessment or elicitation especially challenging, since it may not be possible to incentivize participants to reveal their preferences (or other private information) truthfully (Gibbard 1973; Satterthwaite 1975).

The fact that agent utility should be measured over extended horizons presents MDP and RL-style optimization challenges, as well as important questions w.r.t. *dynamic MD* (Parkes 2007). Other challenges pertain to interpersonal comparison of utilities (Harsanyi 1955). Reducing preferences to individual *comparable* utilities provides significant traction, but it remains to be seen what one loses in terms of generality. While monetary transfers and quasi-linearity obviates this concern in many MD settings (e.g., auctions), RESs fall within the realm of MD design without money.

Fairness is an area where tradeoffs across individuals and groups has gotten much attention in RSs (see, e.g., Ekstrand et al. 2022; Li et al. 2023), both from the perspective of users (e.g., equitable quality of models of user preferences) and creators/vendors (e.g., fairness of exposure). These methods can be seen to promote certain social outcomes without an explicit SCF. However, an explicit MD approach at the level of outcomes, rather than (say) interventions on models, may prove to be more powerful.

**Social Objectives.** Some elements of the SCF may not be readily or naturally definable w.r.t. individual actor utilities, e.g., when the SCF uses properties of *joint* outcomes

over, say, users, but where each user is concerned only with their “local” outcome. For instance, an SCF might try to minimize the emergence of *filter bubbles* due the potential for induced social fragmentation across the user population (Pariser 2011; Aridor, Goncalves, and Sikdar 2020).<sup>5</sup> Other phenomena of this sort include polarization (Ribeiro et al. 2020), personal diversity of consumption, “echo chambers,” and the like. We also note the potential interaction between social objectives and fairness; e.g., rich-get-richer feedback loops can induce popularity bias, which is often seen as undesirable w.r.t. fairness. Yet, economies of scale induced by this effect may actually offer utilitarian benefits, especially if well-managed by the RS (e.g., via information sharing).

Determining a suitable SCF may be one of the most daunting tasks in the application of MD to the design of RESs. For *any* given SCF, simply predicting its ultimate impact on the *realized* outcomes of an RS policy, and the *realized* expected utility generated for RES players, is tremendously complex. The development of analytical, simulation, visualization and scenario-analysis tools should prove vital in helping the designer(s) explore the space of SCFs. SCFs will also play a crucial role in the transparency of RS designs and their consequences for individuals and societal groups.

Finally, while we have focused on the design of RSs that maximize (say) user utility and well-being, the RS itself typically derives value from its matching of users to items over time. This utility can be included in the SCF by treating the RS as an actor/agent, but conflicts might emerge if the RS acts as the mechanism designer as well. This may require the introduction of both policy constraints in the SCF, and more subtle modeling of ecosystems dynamics (e.g., the influence of recommendations on user preference construction or evolution (Lichtenstein and Slovic 2006; Carroll et al. 2022).

## Concluding Remarks

We have defined an ambitious research program, consisting of a set of challenging—but potentially immensely impactful—research problems in the arena of recommender ecosystems. Given the pervasive nature of recommender systems, and their ever-increasing scope and influence on our daily lives, success in addressing these research challenges will have broad implications for scientists studying RSs, practitioners who deploy RSs, the alignment of future RSs with the needs and preferences of its users, creators, providers, etc., and ultimately broad societal dynamics and values. We refer to the extended paper for a more in-depth treatment (Boutilier, Mladenov, and Tennenholtz 2023).

**Acknowledgements:** We thank the reviewers for their helpful comments on the original submission.

<sup>5</sup>Whether filter bubbles are problematic depends on the nature of the content in question: fragmentation in music consumption may be unproblematic; but fragmentation involving diverse perspectives on key social or civic issues of significant social import, may be of greater concern, and be viewed as a societal value, where these consumption patterns have a negative societal externality, e.g., preventing meaningful civic discourse.

## References

- Abdollahpouri, H.; Adomavicius, G.; Burke, R.; Guy, I.; Jannach, D.; Kamishima, T.; Krasnodebski, J.; and Pizzato, L. 2020. Multistakeholder Recommendation: Survey and Research Directions. *User Modeling and User-Adapted Interaction*, 30(1): 127–158.
- Abdollahpouri, H.; Mansoury, M.; Burke, R.; and Mobasher, B. 2019. The Unfairness of Popularity Bias in Recommendation. ArXiv:1907.13286.
- Abebe, R.; Kleinberg, J.; Parkes, D.; and Tsourakakis, C. E. 2018. Opinion Dynamics with Varying Susceptibility to Persuasion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD-18)*, 1089–1098. London.
- Akpinar, N.-J.; DiCiccio, C.; Nandy, P.; and Basu, K. 2022. Long-term Dynamics of Fairness Intervention in Connection Recommender Systems. In *Proceedings of the Fifth AAAI/ACM Conference on AI, Ethics, and Society (AIES-22)*, 22–35. Oxford, UK.
- Amelkin, V.; Bullo, F.; and Singh, A. K. 2017. Polar Opinion Dynamics in Social Networks. *IEEE Transactions on Automatic Control*, 62(11): 5650–5665.
- Aridor, G.; Goncalves, D.; and Sikdar, S. 2020. Deconstructing the Filter Bubble: User Decision-making and Recommender Systems. In *Proceedings of the 14th ACM Conference on Recommender Systems (RecSys20)*, 82–91.
- Arrow, K. J. 1950. A Difficulty in the Concept of Social Welfare. *Journal of Political Economy*, 58: 328–346.
- Asudeh, A.; Jagadish, H.; Stoyanovich, J.; and Das, G. 2019. Designing Fair Ranking Schemes. In *Proceedings of the 2019 International Conference on Management Of Data (SIGMOD-19)*, 1259–1276. Amsterdam.
- Basu, K.; DiCiccio, C.; Logan, H.; and Karoui, N. E. 2020. A Framework for Fairness in Two-Sided Marketplaces. ArXiv:2006.12756.
- Ben-Porat, O.; Cohen, L.; Leqi, L.; Lipton, Z. C.; and Mansour, Y. 2022. Modeling Attrition in Recommender Systems with Departing Bandits. In *Proceedings of the Thirty-sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, 6072–6079. Vancouver.
- Ben-Porat, O.; and Tennenholtz, M. 2018. A Game-Theoretic Approach to Recommendation Systems with Strategic Content Providers. In *Advances in Neural Information Processing Systems 31 (NeurIPS-18)*, 1118–1128. Montreal.
- Biega, A. J.; Gummadi, K. P.; and Weikum, G. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *41st International ACM Conference on Research & Development in Information Retrieval (SIGIR-18)*, 405–414. Ann Arbor, MI.
- Boutilier, C. 2002. A POMDP Formulation of Preference Elicitation Problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)*, 239–246. Edmonton.
- Boutilier, C.; Mladenov, M.; and Tennenholtz, G. 2023. Modeling Recommender Ecosystems: Research Challenges at the Intersection of Mechanism Design, Reinforcement Learning and Generative Models. ArXiv:2309.06375.
- Burke, R. 2002. Interactive Critiquing for Catalog Navigation in E-Commerce. *Artificial Intelligence Review*, 18(3–4): 245–267.
- Camerer, C. F.; Loewenstein, G.; and Rabin, M., eds. 2003. *Advances in Behavioral Economics*. Princeton, New Jersey: Princeton University Press.
- Carroll, M. D.; Dragan, A.; Russell, S.; and Hadfield-Menell, D. 2022. Estimating and Penalizing Induced Preference Shifts in Recommender Systems. In *Proceedings of the Thirty-ninth International Conference on Machine Learning (ICML-22)*, 2686–2708. Baltimore: PMLR.
- Cen, S. H.; and Ilyas, A. 2022. A Game-Theoretic Perspective on Trust in Recommendation. In *ICML Workshop on Responsible Decision-Making in Dynamic Environments*.
- Chien, S.; Jain, P.; Krichene, W.; Rendle, S.; Song, S.; Thakurta, A.; and Zhang, L. 2021. Private Alternating Least Squares: Practical Private Matrix Completion with Tighter Rates. In *Proceedings of the Thirty-eighth International Conference on Machine Learning (ICML-21)*, 1877–1887.
- Christakopoulou, K.; Radlinski, F.; and Hofmann, K. 2016. Towards Conversational Recommender Systems. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-16)*, 815–824. San Francisco.
- Dekel, O.; Fischer, F.; and Procaccia, A. D. 2010. Incentive Compatible Regression Learning. *Journal of Computer and System Sciences*, 76(8): 759–777.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS-12)*, 214–226. Association for Computing Machinery.
- Ekstrand, M. D.; Das, A.; Burke, R.; and Diaz, F. 2022. Fairness in Recommender Systems. In *Recommender Systems Handbook*, 679–707. Springer.
- Gibbard, A. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4): 587–601.
- Guo, W.; Krauth, K.; Jordan, M.; and Garg, N. 2021. The Stereotyping Problem in Collaboratively Filtered Recommender Systems. In *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO-21)*, 1–10.
- Harsanyi, J. C. 1955. Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility. *Journal of Political Economy*, 63(4): 309–321.
- Heuss, M.; Sarvi, F.; and de Rijke, M. 2022. 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 (SIGIR-22)*, 759–769. Madrid.

- Hron, J.; Krauth, K.; Jordan, M. I.; Kilbertus, N.; and Dean, S. 2022. Modeling Content Creator Incentives on Algorithm-curated Platforms. ArXiv:2206.13102.
- Hurwicz, L. 1960. Optimality and Informational Efficiency in Resource Allocation Processes. In Arrow, K.; Karlin, S.; and Suppes, P., eds., *Mathematical Methods in the Social Sciences*, 27–46. Stanford, CA: Stanford University Press.
- Hurwicz, L.; and Reiter, S. 2006. *Designing Economic Mechanisms*. Cambridge University Press.
- Jacobson, K.; Murali, V.; Newett, E.; Whitman, B.; and Yon, R. 2016. Music Personalization at Spotify. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys16)*, 373–373. Boston, Massachusetts, USA.
- Jagadeesan, M.; Garg, N.; and Steinhardt, J. 2022. Supply-Side Equilibria in Recommender Systems. ArXiv:2206.13489.
- Keeney, R. L. 1992. *Value-Focused Thinking: A Path to Creative Decisionmaking*. Cambridge, MA: Harvard University Press.
- Konstan, J. A.; Miller, B. N.; Maltz, D.; Herlocker, J. L.; Gordon, L. R.; and Riedl, J. 1997. GroupLens: Applying Collaborative Filtering to Usenet News. *Communications of the ACM*, 40(3): 77–87.
- Koren, Y.; Rendle, S.; and Bell, R. 2011. Advances in Collaborative Filtering. In *Recommender Systems Handbook*, 91–142. New York: Springer.
- Kurland, O.; and Tennenholtz, M. 2022. Competitive Search. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-22)*, 2838–2849. Madrid.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. 2010. A Contextual-Bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web (WWW-10)*, 661–670. Raleigh, NC.
- Li, Y.; Chen, H.; Xu, S.; Ge, Y.; Tan, J.; Liu, S.; and Zhang, Y. 2023. Fairness in Recommendation: Foundations, Methods and Applications. *ACM Transactions on Intelligent Systems and Technology*.
- Lichtenstein, S.; and Slovic, P. 2006. *The Construction of Preference*. Cambridge University Press.
- Liu, L. T.; Mania, H.; and Jordan, M. 2020. Competing Bandits in Matching Markets. In *Proceedings of the Twenty-third International Conference on Artificial Intelligence and Statistics (AISTats-20)*, 1618–1628.
- Mehrotra, R.; McInerney, J.; Bouchard, H.; Lalmas, M.; and Diaz, F. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the Trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM-18)*, 2243–2251. Turin.
- Mehta, A. 2013. Online Matching and Ad Allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4): 265–368.
- Mladenov, M.; Creager, E.; Swerksy, K.; Ben-Porat, O.; Zemel, R. S.; and Boutilier, C. 2020. Optimizing Long-term Social Welfare in Recommender Systems: A Constrained Matching Approach. In *Proceedings of the Thirty-seventh International Conference on Machine Learning (ICML-20)*, 6987–6998. Vienna.
- Pariser, E. 2011. *The Filter Bubble: How the New Personalized Web is Changing What We Read and How We Think*. Penguin Random House.
- Parkes, D. C. 2007. Online Mechanisms. In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds., *Algorithmic Game Theory*, 411–439. Cambridge: Cambridge University Press.
- Procaccia, A. D.; and Tennenholtz, M. 2009. Approximate Mechanism Design without Money. In *Proceedings of the Tenth ACM Conference on Electronic Commerce (EC'09)*, 177–186. Stanford, CA.
- Pu, P.; and Chen, L. 2008. User-involved Preference Elicitation for Product Search and Recommender Systems. *AI Magazine*, 29(4): 93–103.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving Language Understanding by Generative Pre-training. OpenAI.
- Radlinks, F.; Boutilier, C.; Ramachandran, D.; and Vendrov, I. 2022. Discovering Personalized Semantics for Soft Attributes in Recommender Systems Using Concept Activation Vectors. In *Proceedings of the Thirty-sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, 12287–12293. Vancouver.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot Text-to-image Generation. In *Proceedings of the Thirty-eighth International Conference on Machine Learning (ICML-21)*, 8821–8831.
- Rashid, A. M.; Albert, I.; Cosley, D.; Lam, S. K.; McNee, S. M.; Konstan, J. A.; and Riedl, J. 2002. Getting to Know You: Learning New User Preferences in Recommender Systems. In *International Conference on Intelligent User Interfaces (IUI-20)*, 127–134. San Francisco.
- Ribeiro, M. H.; Ottoni, R.; West, R.; Almeida, V. A. F.; and Jr., W. M. 2020. Auditing Radicalization Pathways on YouTube. In *FAT\* '20: Conference on Fairness, Accountability, and Transparency*, 131–141. Barcelona.
- Salakhutdinov, R.; and Mnih, A. 2007. Probabilistic Matrix Factorization. In *Advances in Neural Information Processing Systems 20 (NIPS-07)*, 1257–1264. Vancouver.
- Salo, A.; and Hämäläinen, R. P. 2001. Preference Ratios in Multiattribute Evaluation (PRIME)—Elicitation and Decision Procedures under Incomplete Information. *IEEE Transactions on Systems, Man and Cybernetics*, 31(6): 533–545.
- Sandholm, T.; and Boutilier, C. 2006. Preference Elicitation in Combinatorial Auctions. In Crampton, P.; Shoham, Y.; and Steinberg, R., eds., *Combinatorial Auctions*, 233–264. Cambridge, MA: MIT Press.
- Satterthwaite, M. A. 1975. Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory*, 10: 187–217.

- Sen, A. 1970. The Impossibility of a Paretian Liberal. *Journal of Political Economy*, 78(1): 152–157.
- Su, Y.; Bayoumi, M.; and Joachims, T. 2022. Optimizing Rankings for Recommendation in Matching Markets. In *Proceedings of the 31st International Web Conference (WWW-22)*, 328–338. Lyon.
- Thoppilan, R.; De Freitas, D.; Hall, J.; Shazeer, N.; Kulshreshtha, A.; Cheng, H.-T.; Jin, A.; Bos, T.; Baker, L.; and Du, Y. 2022. Lamda: Language Models for Dialog Applications. ArXiv:2201.08239.
- Toubia, O.; Hauser, J. R.; and Simester, D. I. 2004. Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis. *Journal of Marketing Research*, 41(1): 116–131.
- Vorobeychik, Y. 2023. The Many Faces of Adversarial Machine Learning. In *Proceedings of the Thirty-seventh AAAI Conference on Artificial Intelligence (AAAI-23)*, 15402–15409. Washington, DC.
- Wu, H.; Mitra, B.; Ma, C.; Diaz, F.; and Liu, X. 2022. Joint Multisided Exposure Fairness for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-22)*, 703–714. Madrid.
- Yang, J.; Yi, X.; Zhiyuan Cheng, D.; Hong, L.; Li, Y.; Xiaoming Wang, S.; Xu, T.; and Chi, E. H. 2020. Mixed Negative Sampling for Learning Two-tower Neural Networks in Recommendations. In *Proceedings of the Web Conference (WWW-20)*, 441–447. Taipei.
- Yi, X.; Yang, J.; Hong, L.; Cheng, D. Z.; Heldt, L.; Kumthekar, A.; Zhao, Z.; Wei, L.; and Chi, E. 2019. Sampling-bias-corrected Neural Modeling for Large Corpus Item Recommendations. In *Proceedings of the Thirteenth ACM Conference on Recommender Systems (RecSys19)*, 269–277. Copenhagen.