

Modeling Recommender Ecosystems: Research Challenges at the Intersection of Mechanism Design, Reinforcement Learning and Generative Models

Craig Boutilier, Martin Mladenov, Guy Tennenholtz

Google Research

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

Abstract

Modern recommender systems lie at the heart of complex ecosystems that couple the behavior of users, content providers, advertisers, and other actors. Despite this, the focus of the majority of recommender research—and most practical recommenders of any import—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 recommenders could generate for its users. We argue that explicitly modeling the incentives and behaviors of all actors in the system—and the interactions among them induced by the recommender’s policy—is strictly necessary if one is to maximize the value the system brings to these actors and improve overall ecosystem “health.” Doing so requires: optimization over long horizons using techniques such as *reinforcement learning*; making inevitable tradeoffs in the utility that can be generated for different actors using the methods of *social choice*; reducing information asymmetry, while accounting for incentives and strategic behavior, using the tools of *mechanism design*; better modeling of both user and item-provider behaviors by incorporating notions from *behavioral economics* and *psychology*; and exploiting recent advances in *generative and foundation models* to make these mechanisms interpretable and actionable. We propose a conceptual framework that encompasses these elements, and articulate a number of research challenges that emerge at the intersection of these different disciplines.

Overview. This paper argues that a holistic *ecosystem-oriented* perspective on recommender systems is necessary to ensure recommender systems serve the long-term interests of users, item providers (e.g., content creators, product vendors) and society as a whole. Taking economic *mechanism design* as its main conceptual organizing framework, this paper outlines a collection of research challenges that we believe must be addressed to effectively model and optimize recommender ecosystems. The paper is organized in four main parts: Part I provides a brief introduction and a stylized formal model of *recommender ecosystems* to help ground the discussion in the remainder of the paper. Part II describes why mechanism design should prove to be an invaluable framework for modeling and optimizing recommender policies in such ecosystems. Part III outlines some of the research challenges required to bring mechanism design to bear on

recommender systems, while Part IV discusses the role of *social choice functions* in making the inevitable tradeoffs in the value the recommender systems generates for different participants (e.g., users, content providers, etc.)

Part I

Introduction

We begin with an overview of the problem, and describe a stylized model of an RS to help ground the discussion of key concepts and research problems that follow.

1 Introduction

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 the connection and communication between individuals or organizations. 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, and the need to integrate the resulting methods into practical RSs at scale.

The focus of the majority of RS research—and most practical recommenders of any import—is on the *local, myopic* optimization of the recommendations made to *individual users*. In other words, a recommendation made to one user does not consider the potential impact it might have on other users or other interested actors. This overlooks the fact that most RSs lie at the heart of complex *ecosystems*, 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 concretely below.

This leads to a set of questions of substantial import pertaining to such recommender ecosystems, or *recosystems*: (1) how should we analyze and model such recosystem interactions? (2) how do we optimize RS policies in the face of such dynamic interactions? (3) what criteria and objectives should be used to guide this optimization? In this position/research challenges paper, we outline a research agenda intended to develop a deeper understanding of recosystems, and encour-

age 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* (Hurwicz 1960; Hurwicz and Reiter 2006; Mas-Colell, Whinston, and Green 1995). Mechanism design has played at best a minor role in the design of RSs—of course, with some notable exceptions, some of which we point to later—in large part due to the local, single-user focus typical of current research and deployment. When moving to ecosystems, effective modeling and optimization of RS policies in the presence of interacting agents must account for: (i) the incentives and behaviors of the agents; (ii) the potential information asymmetry between these actors and the RS; (iii) the potential for strategic behavior; and (iv) tradeoffs in the value generated for different agents. The design of mechanisms—in our case, RS policies—in such settings is precisely the province of economic mechanism design. However, the complexity of the setting poses some unique challenges for traditional mechanism design, and we outline a number of them in this paper.

The remainder of the paper is organized as follows. We conclude Part I with Section 2 by outlining a stylized formalization of an RS—the *omniscient, omnipotent, benevolent recommender*—that makes unrealistic assumptions about the knowledge and control the recommender system has regarding user preferences, content creator skill, and other factors. This model is intended to allow the crisp articulation of some of the research challenges that follow, not to reflect the full set of considerations that go into modeling RSs in the literature. Over the remainder of the article, we describe the algorithmic and modeling challenges that emerge as we relax these assumptions.

Part II focuses on how mechanism design can be used to frame the design of RS policies in complex ecosystems. Using the stylized model introduced in Section 2, we begin in Section 3 with a concrete illustration of how multiagent interactions can negatively impact the value generated by a myopic, local RS, and offer a discussion of some of the considerations that can help overcome this impact. In Section 4, we introduce the classical formalization of mechanism design and outline its components can help address the issues that emerge as we move from modeling recommender *systems* to recommender *ecosystems*.

Part III describes a number of research challenges associated with using mechanism design for ecosystems, especially as we relax the assumptions underlying the stylized model. We outline a number of areas we think are ripe for new thinking, and make a few tentative suggestions along the way regarding specific approaches that might prove fruitful. Our research challenges touch on multiple areas of research in AI, ML, CS-Econ, and behavioral science, including: reinforcement learning, bandits and exploration, algorithmic mechanism design, incentive-compatible ML, user modeling, human-value alignment, generative models, ML fairness, privacy, among others.

Part III is broken down along the following lines: user preference modeling, elicitation and exploration is discussed in

Section 5. We outline the incentives and behaviors of content provider (e.g., creators, distributors, vendors) in Section 6. Section 7 details the inherent asymmetry present in typical ecosystems and how it can impact long-run ecosystem health, and in Section 8 we address the possibility of strategic behavior, focusing on that of content providers. We conclude this part of the paper in Section 9 by describing the challenges of optimization over extended horizons.

Part IV focuses on the inevitable tradeoffs that an RS must make when making decisions that influence the utility of multiple actors in the ecosystem. The *social choice* (or *social welfare*) function is the standard manner in which objectives are specified in mechanism design, and it plays a crucial role here. In Section 10, we examine tradeoffs associated with the “direct” utility generated for users and content/item providers by the RS, while in Section 11 we discuss how the social choice function can be used to encode broader “societal values” (e.g., minimization of filter bubbles or polarization) and how such social dynamics might be accounted for in RS policies. We offer some final remarks in Part V.

We note that this is a research challenges/position paper, and is not intended to serve as a comprehensive survey of research in this area. However, we do point to representative work that begins to address some of these challenges at various points in the paper. For a general overview of the multiagent perspective in RSs, see Abdollahpouri et al. (2020). There are a number of research themes in the literature that address specific aspects of multiagent modeling and optimization in RSs. Some of the more relevant for our purposes include (with representative samples of research cited):

- work on long-term optimization and game-theoretic analysis that models (usually relatively simple) multiagent interactions (Ben-Porat and Tennenholtz 2018; Ben-Porat et al. 2019; Mladenov et al. 2020b; Zhan et al. 2021; 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);
- work on fairness in RSs (Abdollahpouri et al. 2019; Akpınar et al. 2022; Asudeh et al. 2019; Barocas, Hardt, and Narayanan 2023; Basu et al. 2020; Biega, Gummadi, and Weikum 2018; Castera, Loiseau, and Pradelski 2022; Celis and Vishnoi 2017; Heuss, Sarvi, and de Rijke 2022; Wu et al. 2022; Ekstrand 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; DeGroot 1974; Ribeiro et al. 2020; Celma 2010).

2 The Omniscient, Omnipotent, Benevolent Recommender

We set the stage with an intentionally stylized model of an RS that abstracts away a number of subtleties. This will allow us to illustrate a few key challenges in modeling ecosystems while avoiding unnecessary complexity. It is important to keep in mind that adopting a *less* stylized model—one that incorporates additional details, real-world constraints and the many nuances of agent behaviors and incentives—only *exacerbates* the problems we describe, making them even more

challenging. At various points in this paper, we address relaxations of this stylized abstraction and the research problems so-induced.

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 equate affinity with user utility.¹

For expository purposes, we assume items embody some form of *content* (e.g., news, music, video) and that each item i is generated by some content provider or creator $c(i) \in \mathcal{C}$ (Konstan et al. 1997; Jacobson et al. 2016; Covington, Adams, and Sargin 2016), though everything that follows applies to other forms of item recommendation, such as product recommendation.² In general, any $c \in \mathcal{C}$ can produce multiple items, but for ease of exposition, we’ll assume each c creates a single item (which can be altered or replaced as discussed below). Hence, we equate creators with their items for the time being. Similarly, user affinity for items typically varies with their query/request, their context, etc. and embeddings are often dependent on these factors; but again for simplicity, we’ll treat each u as single point in latent space rather than as some collection or distribution.

Omniscience. Realistically, these embeddings are characterized by considerable uncertainty, nonstationarity and generalization error. But we start by assuming that the RS is *omniscient*. Specifically, its user and content embeddings are *perfect* in the following senses:

- There is no (resolvable) uncertainty:³ the RS’s estimated affinity for $\langle u, c \rangle$ accurately reflects u ’s (expected) utility for c . (This would hold approximately if, say, the models are built from sufficient amounts of data regarding both u and c .)
- The model is stationary: user preferences over item space do not change (change in the content corpus is discussed below).
- The model generalizes out-of-distribution: even though RS models are usually trained on the existing item corpus \mathcal{I} , we assume that the model generalizes to new items that may be embedded at points in latent space different than any existing $i \in \mathcal{I}$.⁴

¹This stylized model captures, for example, many forms of collaborative filtering, including matrix factorization (Salakhutdinov and Mnih 2007) and dual encoders/two-tower models (Yi et al. 2019; Yang et al. 2020).

²For instance, if the items are products for sale, one should read “vendor” for “creator,” “sales” for “engagement,” etc. The lessons below apply *mutatis mutandis* to all such settings.

³By “resolvable” we refer to *epistemic* uncertainty that can be eliminated with a sufficient amount of accessible data (e.g., asking the user for their preferences, or interacting with a user to gather enough data about those preferences). Irresolvable uncertainty is generally *aleatoric* w.r.t. the model granularity and data access available to the RS.

⁴We will not focus on new users here.

If the goal is to maximize total user utility w.r.t. some user distribution, under these assumptions the optimal RS policy is to *match* each user u to the creator/item $c^*(u) = \arg \max_{c \in \mathcal{C}} u \cdot c$ with which they have greatest affinity whenever they engage with the system (e.g., issue a query or request a recommendation). Note that expected user utility under this policy is $\mathbb{E}_{\mathcal{U}} u \cdot c^*(u)$, since we explicitly equate utility with item affinity. Importantly for practical RSs, the optimal matching is *fully local*: the recommendation decision for one user is independent of that for another.

Naturally, since none of these assumptions hold in practice, (non-stylized) RSs will need to take steps to: (i) acquire some of this information (e.g., explore or query users about their preferences); (ii) predict this information some of this information (e.g., learn models that generalize across users and items); and (iii) optimize its matching of users to content under conditions of uncertainty.

Omnipotence. Assuming the user population \mathcal{U} is fixed (or that new users are drawn from the same distribution), we can treat the collection of user embeddings as *perfect market research* w.r.t. content space. Specifically, one can potentially use this information to identify *new content* from which users will derive significant utility.

We can draw a simple analogy with how new products are designed (e.g., see the literature on product design or product line optimization (Talluri and van Ryzin 2004; Schön 2010)). Potential consumers are surveyed about their preferences over properties or features of some feasible space of product designs. These preferences—together with constraints on the feasibility of products exhibiting specific combinations of features, and information on production (and other) costs—are used to determine the optimal product, product line, or product differentiation best-suited to a specific audience. The product is then designed, produced, marketed and sold.

Of course, our stylized RS does not actually *create* content, it merely matches users to the content generated by creators.⁵ Thus, the utility it generates for users is fully reliant on the *content-generation decisions of creators*. However, we might imagine an *omnipotent* RS that can readily induce *any creator* to produce *any content* the RS proposes. We might visualize this as the RS having the ability to move any creator $c(i)$ from their existing point i in latent space and move them to some new point i' . This magically makes new content available for recommendation to users (for simplicity, we’ll assume i' *replaces* i , ignoring the persistence of digital content or other digital goods).

Under what circumstances would the RS do this? There are several such conditions under which moving creators in content space can have a positive effect—for example, if the distribution of users is such that this new point i' generates greater user affinity for some subset of users than $c(i)$ ’s current content i or that of any other creator; and if the total loss

⁵We will discuss below the use of generative models to support content creation (e.g., large-language models (Devlin et al. 2018; Radford et al. 2018; Thoppilan et al. 2022), image generation models (Vinyals et al. 2015; Esser, Rombach, and Ommer 2021; Ramesh et al. 2021)). It is not far-fetched to imagine the RS generating content itself in some cases.

in expected affinity due to c moving away from i is less than that gained by moving to i' .

Determining the placement of creators must be done *jointly* of course, since the net gain associated with the move of one creator depends on the positioning of others in embedding space. Let G be the space of (*creator*) configurations, where a configuration $g \in G$ is a vector of embedding points $(g_c)_{c \in \mathcal{C}}$ (somewhat abusing notation and treating c as the creator’s “identity” since its location g_c is variable). Given a fixed configuration g , a user u is again matched to the creator $c^*(u; g) = \arg \max_{c \in \mathcal{C}} u \cdot g_c$ with whom they have greatest affinity. This generates total expected user utility $EU(g) = \mathbb{E}_{\mathcal{U}} u \cdot g_{c^*(u; g)}$. The omnipotent RS will then shuffle creators to obtain an *optimal configuration* $g^* = \arg \max_{g \in G} EU(g)$.

Notice that this problem can be viewed as a simple form of uncapacitated k -facility location (with no opening costs, hence k -medians (see, e.g., Arora, Raghavan, and Rao 1998)), where each creator is a facility to be opened at a location that (collectively) best serves (or minimizes the distances across) the user population. Unlike the case above (omniscient, not omnipotent), optimization is now *non-local* since it requires making utility tradeoffs across users—this is a case where genuine *ecosystem interactions and tradeoffs* come into play. Few, if any, large-scale RS platforms engage in this form a global, ecosystem-level optimization.

Of course, no RS can simply pick up creators and place them in “content space” wherever it sees fit. Creators are independent actors that have both innate and learned (or dynamic) abilities, skills, interests; specific incentives and costs for producing certain types of content; knowledge, beliefs and risk attitudes; and particular strategic, semi-strategic, myopic, or arbitrary decision-making processes. All of these factor into their selection of content to produce. As a consequence, it is vital to design RS policies that appropriately encourage or incentivize creators to make moves in content space, given their own self-interests, to improve overall user utility.

Benevolence. In the stylized model above, we focused on the objective of maximizing a specific form of *user social welfare*, namely, the expected total user affinity for the content to which they are matched. We might be tempted to view such a RS as *benelovent*, in the sense that it is doing its best to generate utility for users, including (in the case of the omnipotent RS) moving creators around to best serve users.

As appealing as this objective might be, it is untenable for many reasons, of which we list just a small selection:

- Even assuming that matched-item affinity reflects user utility, simply maximizing expected affinity is inherently *utilitarian* (in the Rawlsian sense), and does not capture more general notions of social welfare reflecting, say, fairness (Ekstrand et al. 2022; Li et al. 2023). This problem is even more severe under certain types of dynamics, as we show in the next section.
- It completely ignores the costs, benefits and incentives of content creators, who themselves derive utility from engaging with the RS, a utility which plays no role in the RS behavior described above. As we see in the next section, this can have significant, negative impacts on the

long-term health of the ecosystem, including the ability to generate user utility.

- Utility, neither for users nor creators, is derived from a single matching, or the consumption of a single content item.⁶ Instead, it must reflect the ongoing, long-term engagement with the RS. For users, this includes a number of factors, and their dynamics, such as: contexts, preferences, histories, beliefs and discovery, etc. For creators, factors include skills, beliefs, costs, incentives, etc.
- As discussed above, the assumption of omniscience is untenable. The RS’s models are characterized by incomplete information—it can, at best, merely *estimate* the preferences of users, the incentives and policies of content providers, etc. Moreover, an effective RS must generally take *active* steps to reduce this uncertainty (e.g., through exploration or explicit preference elicitation) to improve the quality of the recommendations it makes.
- The converse of incomplete information on the part of the RS is the fact that both users and providers also lack information about the agents with whom they interact. For instance, the RS may have much greater information about the preferences of the entire population of RS users than any provider might have, since any given provider engages with only a small fraction of the population. This *information asymmetry* can limit the ability of providers to make effective decisions about the content they make available, potential decreasing the value the RS can generate for users and providers alike.
- Finally, the behaviors of users and providers will generally be self-interested and—in the case of providers moreso than users—*strategic*. An RS that fails to account for the incentives of its participants, and how these influence their behaviors, is unlikely to optimize its objectives effectively.

These factors compel us to develop a much richer modeling and optimization toolkit to manage ecosystems and the interactions among its constituent actors.

Part II

Why Mechanism Design?

We now turn to the the role of *mechanism design* in the construction of ecosystem-aware RSs, and specifically, why it is well-suited (at least conceptually) to address some of the characteristics identified above that challenge traditional approaches to RS design. We begin in Section 3 with an illustration of how ecosystem interactions complicate the design of RSs, using the stylized model above. In Section 4 we describe why mechanism design provides a valuable framework for modeling—and optimizing the design of—RSs in such ecosystems.

⁶Somewhat glibly, the model above would suggest that a music RS would maximize a user’s utility by playing their favorite song on repeat forever.




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

3 Ecosystem Interactions: An Illustration

Before elaborating on the role of mechanism design, we present a straightforward example to illustrate how even simple dynamics can induce complex multiagent interactions that require non-trivial optimization by the RS.

We consider a very simple, stylized scenario, based on an example of Mladenov et al. (2020b), to illustrate how ignoring ecosystem interactions can limit the ability of an RS to maximize user utility over the long run. In Fig. 1(a), we use (inverse) distance to reflect affinity between users (circles) and the content of specific creators (rectangles). We have two creators, P1 (red) and P2 (blue), and eight users, five of whom (red) 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* in the sense that it maximizes immediate affinity for the user, and *local* in that it does not consider impact that serving one user request might have on other users. If each user issues a single request during a specific time period, P1 engages with five users, and P2 three. However, suppose that, if a creator that does not attract four users per period, it abandons the RS—perhaps due to monetary, social or other incentives. In this case, P2 leaves 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 all blue users much worse off than had the blue creator stayed active. This clearly demonstrates that ignoring *creator behavior* induced by the RS policy can have a deleterious impact on *users*. By contrast, an alternative matching can be devised by anticipating the distribution of user requests—and recognizing the (incentive-induced) behavior of the creators—that matches requests to optimize overall user welfare in the *long run*. This matching (Fig. 1(c)) matches the red user u to P2 and is non-local and non-myopic (it considers the impact of any user match on the long-term utility of *all* users). Since u is nearly indifferent to P1 and P2, it incurs a relatively small sacrifice in utility that enables a large increase in the long-run utility of the three blue users.⁷

This equilibrium maximizes *user (utilitarian) social welfare*, and is “optimal” under this very specific criterion. It is also arguably more fair from an “egalitarian” perspective.

⁷See Mladenov et al. (2020b) for algorithms, analysis, and a deeper discussion of tradeoffs under realistic utility functions.

Note, however, that it does impose a small cost on user u and on creator P1.⁸ Indeed, *optimality* of the equilibrium (or any other outcome) depends in the choice of objective function—in the terms of economic mechanism design, the *social choice/welfare function* under consideration. This argues for applying mechanism design to the construction of recommender ecosystems.

Of course, not only is this example intentionally simplified for expository purposes, it is unrealistic in the many assumptions required to induce equilibrium behavior. 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 these assumptions are relaxed presents numerous research challenges, and is critical if RSs are to act in the best interests of users, creators, and society as a whole.

4 Mechanism Design to the Rescue

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 the preferences of different actors. 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 1960; Hurwicz and Reiter 2006; Mas-Colell, Whinston, and Green 1995) which essentially involves the design of 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

⁸See Mladenov et al. (2020b) for a discussion of ways in which to distribute that cost (e.g., using stochastic matchings) or minimize user regret relative to the myopic/local matching.

of those agents. We outline some of the ways in which a market/mechanism design perspective can help improve RS performance and play a central role in the design of RSs that lie at the center of complex ecosystems. Along the way, we suggest ways in which “classical” techniques need to be modified, and the required non-trivial research advances. We elaborate on a number of these research directions in the sections that follow.

Classical Mechanism Design. To set the stage, we provide a very (!) informal description of classical mechanism design, outlining its key conceptual ingredients.⁹ We then describe how these map to entities and concepts in our recomsystem setting. Concrete examples of these mappings will be provided in the subsequent sections.

Classical mechanism design assumes some action space A , where each $a \in A$ induces a distribution over some outcome space O . Each agent (or player) $g \in G$ has preferences over outcomes in the form of a utility function u_g , which is generally unknown to the mechanism M .¹⁰ The mechanism also embodies a *social choice function (SCF)* C that, given any vector $\bar{u} = \{u_g\}_{g \in G}$ of agent utility functions, determines the “social utility” $C(\bar{u}, o)$ of any $o \in O$.

The goal of mechanism M is to *implement* the SCF by taking the action $a_{\mathbf{u}}^* = \arg \max_{a \in A} \mathbb{E}_O C(\mathbf{u}, o)$.¹¹ To do so, M must take steps to observe, elicit, estimate or otherwise (possibly implicitly) determine \mathbf{u} by interacting with the agents. A mechanism M does so by (a) making available a set of *strategies* or “agent actions” to each agent; and (b) specifying a mapping from the strategies/actions selected by each agent to a choice of outcome $o \in O$. It is this strategy selection by the agents that may (indirectly) reveal something about their preferences. However, any agent g may act strategically in an attempt to manipulate the selected outcome o to maximize its own utility $u_g(o)$, often at the expense of C (and the utility of 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* w.r.t. the equilibrium notion in question.¹²

The form such mechanisms take are many and varied, and we cannot hope to cover them here. It is important to note that not all SCFs are implementable. But one of the foundational results of mechanism design is the *revelation principle*, which states (informally) that if an SCF is implementable, it can

⁹For an overview of mechanism design, please see (Mas-Colell, Whinston, and Green 1995; Nisan 2007).

¹⁰More generally, each agent has private information encoded in its *type*, which may reflect any information not known to the designer or other agents (e.g., we might think of a content creator’s “skill” as being unknown to the RS). However, since the role of private information is to determine the utility of an outcome to the agent, the type can be viewed as simply encoding the agent’s (expected) utility function over outcomes.

¹¹The SCF in this form is sometimes called a *social welfare function (SWF)*, which induces choices, hence the SCF, via the $\arg \max$.

¹²Common forms of implementation correspond to dominant strategy, ex-post and Bayes-Nash equilibria.

be implemented *truthfully* by a mechanism which asks all participants to reveal their types (i.e., their preferences) to the mechanism. Such *direct revelation mechanisms* are appealing because they restrict the actions that are made available to the agents, thus simplifying the design/policy space, and as such play an outsized role in classic mechanism design. As we detail below, however, direct revelation is completely impractical in typical recommender settings.

We now map our recomsystem setting into the elements of a “standard” mechanism design formalization.

Agents. The agents (or actors) are those parties impacted by the actions of the RS, and who derive value from interacting with the RS and (possibly indirectly) with each other. Actors who have an “interest” in the actions of the RS (or the induced outcomes), without directly interacting with it, may be also be considered as participants in the ecosystem. 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.

Actions, Outcomes and Agent Utility. In the recomsystem, we use the term *RS or system actions* to refer to those available to be taken by the RS. In the simplest settings above, these are merely the possible one-off recommendations to users. In such cases, the RS action (in the mechanism design sense) should be viewed as the *joint* set of recommendations made to all users, though we often refer to the “individual” recommendations as actions as well. We use the term *RS policy* when emphasizing the multiple-recommendation nature of the RS’s choices (e.g., across users or across time).

We distinguish RS actions from the *agent actions*, or “strategies” in the mechanism design sense, and often use the term “(user or creator) behavior” to refer to the latter.

The outcome space reflects the range of possible impacts that the RS actions can have on users and creators—again, the formal outcome in the mechanism design sense is the *joint* outcome over all actors, but we abuse the term to talk about the outcome for specific actors (i.e., those aspects of the outcome relevant to that individual). In one-off recommendation, this might simply be: for a user, the actual content consumed; and for a creator, the total user engagement generated by the RS for their content.

Finally, the agent utility functions reflect their preferences over the possible outcomes—generally, this assesses only parts of the outcome “relevant” to them—e.g., a user’s utility may depend only on the item recommended to or consumed by them, not those recommended to other users; and a creator’s utility might only depend on the engagement generated by the RS for their content, and not on the engagement of other creators.¹³

The full complexity of realistic RSs naturally demands that we enrich what we consider to be outcomes induced by the RS, hence the nature of agent utility, and even the action space of the RS itself. For instance, the utility-bearing outcomes for a user of a content RS will often comprise an extended *consumption stream*, or sequence of items consumed.

¹³In other words, we will often assume agent utilities embody no externalities.

User utility for the sequence may not be decomposable (say, additively) into scores/utilities of the items in isolation (e.g., consider a user who values some topic diversity in the content they consume). This in turn renders the RS action/decision space much more complex: the RS will generally adapt its recommendations to the user’s feedback or responses—hence the RS action space comprises the set of *policies* that map the current history of user interactions—and in general, the history of *all* users—to the choice of next recommendation to a specific user.¹⁴

For creators, outcomes may involve not just cumulative user engagement, but also temporal or other properties of that engagement over time. For instance, a creator may value some smoothness in the user traffic they serve over long periods, or user diversity across their audience.

In other domains, the (utility-bearing) outcomes of the recommendations may not be directly or immediately observable to the RS. For example, if a user accepts a product recommendation by purchasing the product, their ultimate utility will depend on the degree to which the product, its features and its functionality serves their needs over time.

The Social Choice Function. A mechanism’s SCF embodies the relevant tradeoffs between the interests of different agents. Our stylized example above makes clear that in any RS: (i) the interests of users and creators may not fully coincide; (ii) the interests of different creators often conflict, since they are competing for audience/market; and (iii) user interests also conflict, though often indirectly via creator behavior. The specification of an explicit SCF gives the RS a clear objective when determining how utility for one user or creator should be balanced against that of another. For example: how much user attention should be directed to one creator (at the expense of others); how user utility for one group should be balanced against that of another in the near-term or long-term; and so on.

Population-level effects may play a role in the SCF. For instance, fairness across different users or groups is usually critical, and cannot¹⁵ be defined without comparison of the outcomes or utility generated for different groups. Even social values can (and perhaps should) be encoded in an SCF; e.g., the emergence of filter bubbles or polarization or other social dynamics may be penalized by the SCF. Note that this cannot be done without defining population-level outcomes.¹⁶

The use of SCFs is valuable because it forces RS designers to examine, explore, and (at least, partially) articulate the

¹⁴Even one-off recommender settings can require sequential interaction if (say) preference elicitation, example critiquing or other forms of user feedback on proposed items is allowed.

¹⁵Perhaps we should say “should not.” A wide swath of AI/ML fairness research does not explicitly refer to the outcomes of *actions* or their utility, but rather focuses on properties of models, e.g., as in equalized odds (Barocas, Hardt, and Narayanan 2023).

¹⁶Moreover, social value judgements are necessary if population-level outcomes are embedded in the SCF: a “filter bubble” in which one group of users only consumes content for “hobby 1” and another group consumes “hobby 2” may not be viewed as problematic; but if the content difference is about a particular perspective on some consequential news event, this may be treated as troublesome (e.g., because of the externalities w.r.t. civic discourse.)

specific tradeoffs and population-level values the RS embodies and attempts to target. In complex ecosystems, RSs are unlikely to actually implement the target SCF (see discussion below of estimation, learning and optimization challenges), but the specification of core principles will undoubtedly be of tremendous value. That said, SCFs should be viewed as a “framework” for exploration and decision making w.r.t. value tradeoffs. They do not provide any guidance on what the right tradeoffs are—indeed, they cannot in general (see the impossibility results of Arrow (1950), Gibbard-Satterthwaite (1973; 1975), and Roberts (1979) among others).

Preference Revelation. A critical aspect of mechanism design is *preference revelation*—indeed, one of the primary aims of any mechanism is to induce agents to reveal private information that enables the implementation of the SCF to the greatest extent possible.¹⁷ In our setting, this manifests as the RS taking actions to elicit information from—or induce behavior by—users or creators that reveal something about their preferences.

While a substantial fraction of mechanism design research concentrates on *direct revelation*—mechanisms that attempt to have agents *explicitly* and *completely* state their preferences—such approaches are not generally feasible in RSs. Cognitive biases typically mean that users cannot reliably state their preferences accurately even in simple settings (Tversky and Kahneman 1974; Kahneman and Tversky 2000; Camerer, Loewenstein, and Rabin 2003). Moreover, the vast number of items in RSs precludes anything approaching complete revelation. When coupled with outcomes that may be dictated by combinatorial bundles or consumption streams, and the fact that user utility may be derived from some *indirect* effect of item consumption,¹⁸ we recognize that RSs will generally rely on *indirect and incomplete feedback* from users. Finally, communication costs (e.g., cognitive load, potential—actual or perceived—privacy concerns) and opportunity costs (e.g., consumption of less-than-ideal items in the course of RS exploration) must be traded off against the value of any information elicited from users or creators.

Some of these factors have been studied in economic mechanism design, e.g., dynamic mechanism design (Parkes 2007; Bergemann and Välimäki 2019), partial preference elicitation (Parkes 1999; Blum et al. 2004; Hyafil and Boutilier 2007; Lu and Boutilier 2020), noisy responses and machine-learned preferences (Lahaie and Parkes 2004; Beyeler et al. 2021), etc., but largely in rather stylized models. Preference revelation in RSs will look quite different, relying on a combination of partial, incremental elicitation (Boutilier 2013), attribute-based preference models to support natural (even conversational) user interaction (Burke 2002; Chen and Pu 2012; Christakopoulou, Radlinski, and Hofmann 2016), generalization of preference estimation across users given observed behaviors (e.g., as in CF) (Hu, Koren, and Volinsky 2008;

¹⁷As noted above, all private information can be encoded in the utility function.

¹⁸E.g., the utility of a content item for one user may be stochastically connected to its educational impact, while for another it may increase social connection. See Keeney’s (1992) discussion of means vs. ends objectives.

Koren, Bell, and Volinsky 2009), bandit-style exploration (Auer, Cesa-Bianchi, and Fischer 2002; Christakopoulou and Banerjee 2018; Yue et al. 2012), and much more.

A final important factor is that of *preference construction* (Lichtenstein and Slovic 2006; Warren, McGraw, and Van Boven 2011). Classic mechanism design generally assumes that preferences (and other private information) completely determine an agent’s utility and are unchanging. However, users often do not have sufficient familiarity with the item corpus—nor with novel items that may not yet exist—to have fully formed preferences. From a behavioral economics perspective, we might view this as preference construction, where a user’s interaction with the RS helps shape their preferences. Moreover, we should recognize that user preferences may change over time in ways that are influenced by the RS. Naturally, similar remarks apply to content creators (e.g., their skills, interests or decision-making processes may change), item vendors (e.g., costs and opportunities may change), and other actors.¹⁹

Information Sharing, Privacy and Control. While preference revelation is critical to RS design, we must also consider the *decisions* made by users and creators. These are autonomous decision makers that, in many cases, rely on the RS to provide them with information upon which they base their decisions (e.g., consumption or creation/distribution). As such, elicitation and (indirect) sharing of some aspects of the private *types* of other agents may well fall within the scope of RS’s policy, and will play a critical role in improving ecosystem outcomes and social welfare.²⁰ The information sharing of traditional RSs—recommendations to users, realized engagement to creators—offers far too narrow a communication channel for this purpose.

As one example, consider a creator (or vendor) whose item does not attain the desired engagement (or sales). In a complex RS ecosystem, the possible causes are myriad: (i) there is no demand for this type of item; (ii) item quality or other attributes make it unattractive to most users who might otherwise seek this item; (iii) numerous other creators offer similar items; and so on. The first two causes reflect user preferences while the third arises due to the “competitive landscape.” The most suitable creator response to each of these causes of low engagement is, of course, quite different. Unfortunately, the creator will have limited visibility into these causes, and generally be hampered in their decision making as a consequence. By contrast, the RS generally has much greater insight into both the specific preferences of the entire user population and the items of other providers that might be reducing the target item’s engagement. Designing information-sharing mechanisms that allow creators to make more informed decisions could dramatically increase the ecosystem value created by RSs.²¹

¹⁹Some of what is colloquially interpreted as “preference construction” may be recast as changing the state of knowledge (e.g., item familiarity) of the user, which allows them to “apply” their well-formed preferences to item consumption differently.

²⁰See Myerson (Myerson 1983) for a treatment of information sharing by the principal (RS in our setting).

²¹In some sense, this could be viewed as the RS providing “mar-

Mechanisms could include the creation of marketplaces with suggestions regarding under-served parts of the market/audience, or “nudging” specific creators with hints about potential improvements in their item offerings. Doing the latter effectively may require the RS to estimate (or elicit), say, the “skills” of the creator.²²

It is essential, of course, that such mechanisms preserve the privacy of sensitive data (e.g., specific user preferences, specific behaviors and private information of “competing” creators). The use of aggregate audience demands, differentially private item representations (Chien et al. 2021; Fang, Du, and Wu 2022; Ribero et al. 2022) and novel mechanisms to give users and creators fine-grained control over usage of their data are key enabling technologies that should be explored.

Strategic Behavior. Accounting for strategic behavior is a fundamental aspect of mechanism design. Classic mechanism design research typically assumes that agents are hyper-rational in two senses: (1) they reason about and plan their own behavior *optimally* under fixed environmental assumptions; and (2) they anticipate/determine equilibrium behavior in the presence of other strategic agents. Strategic behavior clearly plays a role in RS ecosystems, as hinted above, but creators and (especially) users will generally behave in a boundedly rational fashion, a topic that gets relatively sparse attention in mechanism design (Zhang and Levin 2017; Karaliopoulos, Koutsopoulos, and Spiliopoulos 2019).

Users will no doubt fall prey to behavioral biases that influence their choices in an RS, e.g., framing or endowment effects, hyperbolic discounting, recency bias (Tversky and Kahneman 1974; Camerer, Loewenstein, and Rabin 2003); and while they may exhibit some long-term planning capabilities—potentially even seeking to influence future recommendations (Guo et al. 2021)—they are unlikely to engage in strategic reasoning w.r.t. the behavior of other users, creators, or the RS itself. By contrast, we expect that creators will have greater incentive to reason strategically, especially in regard to “competition” for market share or audience in the RS. In either case, optimal RS policies (ignoring inherent non-stationarity) should generally aim to produce approximate equilibria—either in a dynamical systems sense, a game-theoretic sense, or some mix of the two (Tennenholtz and Kurland 2019; Ben-Porat and Tennenholtz 2018; Mladenov et al. 2020b; Akpınar et al. 2022).

Joint Optimization. Another daunting practical challenge in mechanism design is the computational demands of outcome selection. Even if information revelation and strategic issues are managed, outcome selection often requires solving difficult combinatorial optimization problems since outcomes must embody the features that are relevant (i.e.,

ket research” to creators to enable the design of better items, as in conjoint analysis for product design (Green and Srinivasan 1990) or product line optimization (Kohli and Sukumar 1990).

²²This is an example of private information that we don’t colloquially consider part of the creator’s utility function; but of course, formally, it can be encoded as such.

utility-bearing) to *every* agent in the system.²³ In a reecosystem, the RS must consider the impact of its policy on the outcomes experienced by all agents in the system.

Even in one-off recommendation settings, as discussed above, this can induce a complex (offline or online) matching problem—matching millions or billions of users to thousands or millions of items and their creators—when creator behavior is influenced by the RS. As suggested above, an RS that attempts to organize or influence creator content production faces a problem akin to facility location. When user outcomes are defined w.r.t. entire sequences of recommendations, the optimization of the RS policy from the perspective of a single user is a form of Markov decision process (or reinforcement learning problem). Moreover, most of these problems have to be solved under conditions of incomplete information, turning (say) user sequence optimization into a POMDP. Finally, an optimal RS policy will generally take steps to gather information to improve its ability to make good recommendations, which will require computation involving quantities such as information gain or expected value of information (Shachter 1986). Most of these constituent problems are themselves combinatorial in nature, with very large inputs: the number of users, the number of creators, the number of items, high-dimensional embedding representations, etc. Taken together, these problems require considerable research to develop approximate or heuristic solution techniques that scale to practical recommender ecosystems.

Part III

A Nonstandard Form of Mechanism Design

We now enumerate a selection of research directions that should improve our ability to design recommender systems that work effectively in complex ecosystems. Many of the research questions that follow can be posed and addressed without appeal to mechanism design; however, some aspects of these problems—and especially the connections across these problems—are placed in clearer focus when viewed through the mechanism design lens. For this reason, we use classical mechanism design as our starting point for much of the discussion. That said, these challenges make clear that the practical design of reecosystems requires that one take some license with the modelling, algorithmic and analytical techniques usually adopted in mechanism design, resulting in a decidedly “nonstandard” form of mechanism design.

²³For example, classical auction designs like VCG are simple in the single-item case (McAfee and McMillan 1987), since at most one bidder receives the item, so the outcome space is small. By contrast, *combinatorial auctions* present significant computational challenges due to the induced (set of) optimization problems that must be solved for mechanisms like VCG to be used (Cramton, Shoham, and Steinberg 2005).

5 User Preferences: Modeling, Elicitation and Exploration

Considerable RS research and practice deals with the prediction of user *affinities* for items, which are used to “score” candidate items for recommendation. To the extent that these reflect user utility—for instance, if the sum of the affinities of items consumed by a user reflects their utility—we can view such RSs as attempting to uncover user preferences to better optimize its action choices, in the spirit of mechanism design.²⁴

5.1 Indirect, Behavior-based Preference Assessment

A key distinction between the typical RS approach to preference assessment vs. classic mechanism design methods is the fact that a user’s affinities are generally estimated from the past behavior (e.g., clicks, views, or other consumption behavior) of both the user in question and *other users* (e.g., using methods such as collaborative filtering (Salakhutdinov and Mnih 2007; Koren, Rendle, and Bell 2011)). In mechanism design, the problem of incentivizing agents to reveal private information truthfully (or expend effort to support collective decision making) when this information is to be *aggregated* to generate a predictive model has been studied under the guise of *incentive-compatible machine learning* (Dekel, Fischer, and Procaccia 2010; Balcan et al. 2005). Adapting such models to the complexity of reecosystems would be fruitful, though strategic reporting on the part of users may be rare in typical RS settings.

That said, generalization error should be accounted for in any assessment of user preferences (including when trade-offs are to be made, as we discuss below). Moreover, the assumption that a user’s behavior (e.g., choice of item from a slate) is indicative of their preferences ignores the possibility of various cognitive biases and heuristics that drive that behavior, e.g., anchoring, framing or endowment effects (Tversky and Kahneman 1974; Camerer, Loewenstein, and Rabin 2003), position bias (Joachims et al. 2007), or even phenomena like popularity bias (Abdollahpouri 2019) that may be partly or wholly induced by the RS itself. Given the noise, biases and incompleteness of such affinity models, we may need richer optimization techniques than those usually considered in mechanism design.

5.2 Exploration: Indirect Revelation

Another characteristic of the classic RS approach to preference assessment described above is the fact that no explicit steps are taken to reduce uncertainty in the RS’s assessment of a user’s preferences—models are trained on the “organic” behavior of users interacting with their recommendations.

²⁴We note that this incorporates just one aspect of mechanism design, namely preference revelation, with no multiagent preference aggregation via an SCF, nor any accounting for strategic user behavior. We will use the term “in the spirit of mechanism design” to refer to methods that evoke one or more elements of classical mechanism design rather than the full framework.

This stands in contrast to mechanism design, where the mechanism takes explicit steps that induce (possibly indirect) revelation of preferences.

Of course, considerable work has addressed *exploration* in RSs, where, for example, (*contextual*) *bandit methods* are employed to present novel items to users to construct more refined models of their preferences (Li et al. 2010; Tang et al. 2014).²⁵ Active exploration of this form can be viewed as incremental, indirect preference revelation in the spirit of mechanism design.

Generalization across users remains an important feature of such models. This in itself raises questions of a multiagent nature, since the value of making an *exploratory recommendation* to one user—one which is not predicted to be best for that user in the current context—may lie largely in the ability of her response to improve estimates of *other* users’ preferences. Given that exploration imposes costs on users (at least an opportunity cost, if not annoyance, delay, etc.), interesting questions emerge regarding the potential impact too much exploration may have on user satisfaction or retention, and the equitable distribution of exploratory actions over a user population. Recent work has begun to consider strategic elements that emerge in exploration (e.g., Liu, Mania, and Jordan 2020).

5.3 Explicit Preference Elicitation: Direct Revelation

Methods for explicit *preference elicitation* have been studied extensively in the RS community, as well as in management science, operations research and many other research communities (White, Sage, and Dozono 1984; Salo and Hämmäläinen 2001; Rashid et al. 2002; Boutilier 2002; Toubia, Hauser, and Simester 2004; Boutilier et al. 2006; Pu and Chen 2008). The aim is to ask a user one or more questions about their preferences over items, which is akin to incremental, *direct* revelation in the mechanism design sense. A variety of query types can be used to assess preferences over items. Two broad classes are those that ask about items directly and those that ask about item attributes. The former class includes rating queries (“How much do you like item i ?”), choice or comparison queries (“Which item from set I do you prefer?”), and ranking queries (“Rank items from set I in order of preference.”), each of which can be framed using a variety of modalities (text, image, GUI, speech, and combinations thereof).

By contrast, *attribute queries* focus a user’s attention on specific properties of the items in the corpus, e.g., color, weight, price, the presence of various product features, etc. Attribute-based methods are frequently used in multiattribute item spaces, such as shopping domains (Burke 2002),²⁶ but are less common in content domains, despite the fact that content attributes are often used for content search, filtering and navigation (e.g., artists, genre in music recommendation;

²⁵We discuss exploration further in the context of items and creators in Sections 6 and 10.

²⁶Attributes are commonly used for example critiquing (Chen and Pu 2012) and faceted search (Koren, Zhang, and Liu 2008; Wei et al. 2013) which are somewhat related to preference elicitation.

topic or news source in news recommendation) (Vig, Sen, and Riedl 2012), and sometimes for *onboarding* (Rashid et al. 2002).

Critiquing methods (Burke 2002; Viappiani, Faltings, and Pu 2006) can also be viewed as a form of elicitation in which more control is given to the user: when an item i is recommended, the user can critique the item by requesting a different item with more, less or a different value of some attribute than i (e.g., a less expensive restaurant, a camera with 10X zoom, more upbeat music). Handling open-ended item critiquing will become increasingly important as *conversational recommenders* (Christakopoulou, Radlinski, and Hofmann 2016; Sun and Zhang 2018)—where users interact with the RS using natural-language dialogue—emerge as a dominant mode of interaction with continued improvements in foundational language modeling (Devlin et al. 2018; Radford et al. 2018; Thoppilan et al. 2022). This poses several new challenges for RSs in assessing user preferences, including: understanding how open-ended utterances reflect a user’s underlying preferences; and dealing with the soft, subjective nature of attribute usage (Radlinski et al. 2022; Göpfert et al. 2022).

Whether item- or attribute-based, *interpreting* user responses to direct elicitation queries—that is, reliably, albeit stochastically, relating them to a user’s underlying preferences—will have to account for the types of behavioral phenomena and cognitive biases mentioned above.

The revelation principle has meant that direct revelation mechanisms are the most widely studied within the mechanism design research community. However, considerable mechanism design research has considered the role of incremental and partial revelation, especially when the outcome space is complex or combinatorial. One example is the body of work on incremental elicitation in combinatorial auctions (see (Sandholm and Boutilier 2006) for a survey). The principles underlying such mechanisms should play a central role in the design of RSs for complex ecosystems. This work also accounts for the strategic behavior of mechanism participants, though we expect users of most RSs to largely act in a non-strategic fashion; see the discussion in Section 8.

Even if users are non-strategic, they *must perceive value* in the engagement they have with the RS, especially in conversational settings or when providing responses to an RS’s preference elicitation queries. Each interaction imposes a potential cost on the user, reflecting, say, cognitive burden, perceived intrusiveness, or time delay in receiving a recommendation. Many elicitation techniques can quantify the *expected value of information (EVOI)* of a query—how much it improves recommendation quality—which can be traded off against that cost (should the latter be quantifiable) (Boutilier 2002), which can be used by the RS to decide whether to elicit further or simply recommend. However, the user must also believe (e.g., through past experience or query relevance) that the information they provide is worthwhile.

Other complexities must be dealt with when understanding user preferences, for example, the fact that most preferences are inherently *contextual*—they depend on a user’s current context, e.g., location, activity, companions, mood, etc.—and *conditional*—the utility of one part of a recommendation

depends on other parts of the recommendation.²⁷ Sequential recommendations are an important special case discussed below.

Conversational recommenders (Christakopoulou, Radlinski, and Hofmann 2016; Sun and Zhang 2018) have been proposed as a means to allow more flexibility in a user’s interactions with an RS, including allowing more open-ended dialogue to support richer forms of steering, critiquing, preference elicitation, and user probing/exploration, and at the same time allowing the RS to develop a more nuanced understanding of a user’s preferences and context to support many of the considerations above. The rapid development of generative AI and foundation models (Devlin et al. 2018; Radford et al. 2018; Thoppilan et al. 2022; Vinyals et al. 2015; Esser, Rombach, and Ommer 2021; Ramesh et al. 2021) holds promise for building highly performative conversational RSs. For example, the use of large-language models (LLMs) to support multi-turn dialogue-based recommendation (Friedman et al. 2023; Gao, Gao, and Tang 2023) offers an immediate opportunity to allow more open-ended preference expression, critiquing, preference elicitation and recommendation explanation. At the same time, significant challenges remain, including developing models that are inherently personalized, and that seamlessly blend the rich, behavior-based models of users and items commonly used in CF-based RSs with the rich semantic understanding of items and users afforded by LLMs. Incorporating multi-modal interactions into RSs also represents rich opportunities, e.g., using text-to-image models (Ramesh et al. 2021) to synthesize new content or stylistic variations of products to allow more efficient user exploration or critiquing, or using audio models to generate new forms of preference elicitation queries.

5.4 Unobservable Outcomes

The discussion of preferences above focuses on the assessment, elicitation and revelation of preferences or utilities for individual items and makes several unrealistic assumptions. One of these is that fact that RSs usually *equate the item and the outcome*. In other words, they treat the consumption, purchase, or other engagement with the item by the user as the outcome over which preferences should be assessed.

There are a variety of situations in which this is not the case. In particular, often the utility-bearing outcomes for a user are merely *facilitated* by the item recommended. This embodies a distinction between what Keeney (1992) famously called *means objectives* vs. *ends objectives*. While an RS talks about items in terms of their features, user’s often derive value from the usage of the item facilitated by those features. For example, an RS might help a user navigate the purchase of camera by discussing various features and technical specifications, but its value may ultimately depend on whether its intended use is primarily for vacation pictures, nature photos, or great

²⁷We use the term “contextual” to refer to dependence on things beyond the control of the RS, e.g., a user may prefer more upbeat music when exercising, and more relaxing music when studying. We use “conditional” to denote a dependence on aspects of the decision or recommendation the RS can influence, e.g., a user may prefer a late reservation at one restaurant, but an earlier reservation at another.

action shots of kids on their sports teams. Understanding user preferences over outcomes (the ends) may be more natural than a focus on item features (the means), especially when the mapping between the two is complex and better understood by the RS.

This also implies that user utility is often influenced by exogenous factors beyond the control of the RS or the user. For example, a restaurant recommendation may be better or worse for a specific user depending on the availability of parking nearby and the nature of the user’s commute. Moreover, the value of a reservation at a specific time—should the restaurant have highly variable service times—might depend on whether the user has after-dinner theatre tickets. The outcomes are stochastic, and such factors require assessing user preferences and risk attitudes in *that outcome space*, as well as having an understanding of domain dynamics.

Finally, *latent factors* may play a dominant factor in user outcomes. For example, a user’s true satisfaction with (or enjoyment derived from) a music playlist will generally not be observable. Such latent factors may be reflective of (or even equated with) user utility, and some practical RSs have started incorporating regular user surveys to assess and predict factors like user satisfaction with recommended items (Goodrow 2021). Of course, it may be difficult for a user to express such factors with any degree of confidence or precision when asked. Other latent factors may be more directly outcome-related, e.g., the consumption of news content to stay current or informed on a specific issue, or entertainment content to facilitate social interactions with friends. Here the user’s *state of knowledge* may be latent, but changes in that state may be (part of) the outcome of interest.

Notice that each of these factors necessitates a reconception of the outcome space relative to traditional RS designs.

5.5 Sequential, Multi-item Recommendation

Finally, in most recommendation domains, typically users engage with a specific RS repeatedly for the recommendation of multiple items. This is especially true in many content domains (e.g., listening to music, reading news, watching video content), but in others as well (e.g., shopping, commerce, entertainment and dining).²⁸ In some cases, utility may derived over a single session (e.g., a music listening session), while in others it may extend over long, multi-session horizons (e.g., consumption of news or educational content, or numerous music sessions intended to broaden one’s musical appreciation). As such, user value derived from the collection or sequence of recommendations is not simply the sum of individual item utilities, and will generally require some form of sequential optimization using techniques such as reinforcement learning (RL) or other nonmyopic planning methods.

In some cases, sequence utility may be a function of individual “rewards,” but exhibit a certain *risk profile* w.r.t. the sum of rewards (Keeney and Raiffa 1976; Mihatsch and Neuneier 2002; Chow et al. 2017), diminishing returns, or even the classic sigmoidal utility form of prospect theory (Kahne-

²⁸Exceptions are largely confined to recommendations involving one-off or infrequent decisions (which tend to be higher-stakes).

man and Tversky 1979; Prashanth et al. 2016). For instance, user utility or satisfaction for video content may be convex in the region of low cumulative reward (a large amount of low reward content may be viewed as unsatisfying, or a waste of time) and concave in the region of high cumulative reward (saturation or diminishing returns).

In other cases, utility may not even be a function of the isolated scores or rewards of the individual items; it may instead depend on other properties of the items in the sequence and their relationship. A classic example is content diversity, where a user may value topic or stylistic diversity in a consumption stream. In this case, the quality or reward of individual items is not sufficient to assess the utility of the sequence; the distribution of diversity-relevant attributes across the sequence must also be considered.

Preference assessment and elicitation over sequences is generally much more involved than for single (even multi-attribute) items. While techniques such as elicitation of reward functions in Markov decision processes (MDPs) (Regan and Boutilier 2011) and preference-based RL (Wirth et al. 2017) may offer some insights, richer techniques will be needed in complex recommender domains. Accounting for behavioral biases (e.g., hyperbolic discounting Laibson 1997; Ericson and Laibson 2019) and risk preferences (Charness, Gneezy, and Imas 2013) in the assessment of long-term outcomes should prove critical as well. Elicitation and assessment over recommendation sequences is also relevant to the discussion of dynamic mechanism design in Section 9.

6 Creator Incentives, Skills and Beliefs

The providers of the items in the RS’s recommendable corpus (e.g., content creators, distributors, item vendors) have their own incentives for making these items available, which influence their behaviors. These incentives could be economic or monetary: product vendors have direct monetary interest in recommendations being made to users to whom their products appeal, while content creators may have a direct monetary interest in user engagement (e.g., product placement or ad revenue), or indirect economic interest in such (e.g., eventual monetization of reputation and brand). Incentives might also be driven by social, influence or reputational factors, e.g., a content creator who derives utility from the broader influence of their ideas. User engagement may serve as, at best, an indirect proxy for creator utility.

6.1 Creator Preferences and Incentives

In contrast to the detailed modeling of users (albeit still limited as discussed above), most RSs do much less detailed modeling of provider incentives and behaviors, and in many cases, none at all.²⁹ As our stylized example in Section 3 illustrates, failing to account for creator behavior can have a significant impact on the utility an RS can generate for its users. Moreover, a good understanding of creator preferences allows the RS to *better incentivize alternative behaviors* (e.g., provision of new products, generation of new content) that can improve overall user utility over time. Finally, given that

²⁹A notable exception is work in online advertising, which we discuss briefly below.

creators are themselves actors in the system, RSs should generally balance the utility creators derive from engaging with the RS against each other and against user utility (Abdollahpouri et al. 2019; Akpınar et al. 2022; Asudeh et al. 2019; Basu et al. 2020; Biega, Gummadi, and Weikum 2018; Celis and Vishnoi 2017; Singh and Joachims 2018; Heuss, Sarvi, and de Rijke 2022; Wu et al. 2022; Ekstrand et al. 2022; Mehrotra et al. 2018). We discuss each of these issues in turn.

Models of provider utility seem to be used infrequently in practical RSs and incorporated into RS research in fairly simple ways. To the extent that such models exist, they tend to be of a relatively simple form, for example, using product sales (Wang et al. 2015; Louca et al. 2019), user impressions/views or the like as a proxy for utility. For example, the body of work on *fairness of exposure* focuses on user impressions across different creators (Singh and Joachims 2018; Heuss, Sarvi, and de Rijke 2022), though without explicitly labeling this as creator utility (see Section 10). Different forms of creator utility are almost certainly at play in real recommender ecosystems. For example, a creator’s utility for engagement with their content may be a complex function of total, cumulative engagement, e.g., exhibiting specific risk profiles, diminishing returns, sigmoidal structure and saturation, etc. (Keeney and Raiffa 1976; Kahneman and Tversky 1979). This may be especially true when factoring in item production or sourcing costs, etc. Even factors such as user demographic mix, the distribution of user engagement across their audience, and related user properties can play a role (as can indirect outcomes, as discussed in Section 5.4 in the case of users). Finally, properties of the temporal sequencing of engagement or conversions may play a direct role (e.g., smoothness).

Given the potential complexity of creator utility, considerable research is needed to develop both indirect preference assessment and exploration techniques, as well as explicit preference elicitation methods. Indirect methods will of course require understanding how a creator’s preferences are manifest in their behavior (see Section 6.2).

As we discuss in Section 8, unlike users, we expect creators to be somewhat *strategic* in their decision making. This then requires designing (indirect or direct) elicitation mechanisms in the spirit of classic mechanism design: the mechanism itself should either account for the game-theoretic responses of creators or provide incentives that minimize the degree to which creators misrepresent or misreport their preferences. In the case of indirect signals (e.g., the type, quality or frequency of new content generation; pricing and type of new product offerings), connecting these to the creator’s underlying preferences in equilibrium becomes critical. This may be especially challenging when monetary compensation is not allowed, not only due to classic impossibility results in the theory of mechanism design (Gibbard 1973; Satterthwaite 1975; Roberts 1979), but also because of the need to learn and generalize behavioral models across creators.

One area in which the principles of mechanism design have been applied with great success is computational advertising (Mehta 2013; Dave and Varma 2014). If we (somewhat loosely) view ad servers as RSs—where the ads of specific advertisers are recommended to users depending on the rela-

tive affinity a user has for an ad (or the product and services they represent)—this serves as a successful example of creator (advertiser) incentive modeling and elicitation. Not only do classic auction mechanisms or their variants play a central role (first-price, second-price, generalized VCG, etc.), methods such as budget smoothing, ROI metrics w.r.t. conversions, and a variety of other mechanisms can be interpreted as adding at least some nuance to the usual “additive” clicks/conversions view of utility. However, apart from ad “relevance,” little attempt is made to model user utility to any degree; and most modeling—whether pCTR or conversion predictions, or even more subtle concepts such as ad blindness (Hohnhold, O’Brien, and Tang 2015)—tends to focus on user behavior rather than user utility. Finally, from a creator perspective, preference modeling and elicitation—including interpersonal comparisons of utility across creators—is greatly simplified by the use of quasi-linear utility, where money serves as the common numeraire. That said, the adaptation of techniques from online advertising to mechanism design for recommender ecosystems holds considerable promise.

6.2 Creator Behaviors

Understanding creator preferences is obviously important when these are incorporated in the SCF. However, even if the SCF inputs consist solely of user preferences, since creator behaviors are shaped by their incentives, creator preferences themselves impact the ability to optimize a “pure user” SCF as we showed in Section 3.

The behaviors or decisions of any rational actor are based not only on their preferences, but also on their available actions and their beliefs. Hence, modeling creator behaviors may, in many cases, be best served by modeling both:

- (a) their *action spaces*. For example, what skills and resources to creators have to generate new content, of specific type and specific quality? What abilities do vendors have to source, design and/or manufacture new products and at what capacity and cost?
- (b) their *beliefs*: For example, what do creators believe (e.g., are they able to predict) about the audience for a new content item, the time and cost required to generate attractive new content for an existing or new audience, or the anticipated quality of the results? What does a vendor believe the market will be for a new product as a function of its price?

Predictive models of creator behavior may be feasible without explicit modeling of skills and beliefs, at least in simple cases. For instance, our stylized example assumes creators abandon the RS if they do receive sufficient user engagement with their content. Note that this model of behavior does not explicitly capture any form of utility, but could be relatively simple to learn (given, say, sufficiently exploratory and generalizable data). However, the underlying explanation for abandonment w.r.t. creator incentives and costs, available actions/responses, and beliefs, allows for a much richer set of robust interventions by the RS with more predictable effects. For example, a creator might abandon the RS because they believe there is little audience for the type of content they produce, or because they believe other creators offer higher

quality content of the same type. How the RS responds to keep the creator engaged (either to improve utility for that creator, drive value for users, or both) depends on the rationale for their abandonment decision. We discuss this further in Section 7.

Developing models of creator skills and abilities, of creator beliefs, and of creator decision-making policies—especially those that generalize across creators—present a number of interesting research challenges. Among them are: developing models of the appropriate form; exploration and elicitation methods for indirect/direct assessment of these models; and appropriate incentivization for revealing this private information when creators deliberate strategically.

7 Information Asymmetry and Ecosystem Health

The “health” of a recommender ecosystem is often discussed without precise definition, but is typically evoked to refer to the ability of an RS to generate diverse recommendations to its users. We take a somewhat broader, though still informal, view, equating ecosystem health with *the ability of the RS to generate significant value for all of its participants over the long run*. This may encompass a variety of factors relating to the total utility or welfare it generates for users, creators, vendors, etc.; the distribution of this utility; the facilitation of user well-being; the ability to generate beneficial social dynamics; and much more. Many of these factors may, of course, incorporate diversity as one component (see Part IV).

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, for example, as user preferences and tastes evolve, as new content creation or product 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/provision decisions of item providers (e.g., content creators, vendors) and the consumption decisions of users, it can often impact value creation most directly by supporting this decision making, encouraging decisions that maintain ecosystem health and promote long-term social welfare. Here we focus primarily on content creation (or item production) decisions to illustrate some key insights and research challenges.

One of the impediments to ecosystem health is the non-trivial *information asymmetry* that exists between the RS and content creators. First, we note that the RS has rich models of user preferences over the existing item corpus. If we assume that this model generalizes *out-of-corpus* to some degree, we can treat this as a comprehensive model of *latent user demand* over content space which predicts affinity or utility for *hypothetical items* that do not (yet) exist. Second, the RS has a holistic view of the corpus itself, and potentially some insight into the abilities of providers to source or create new items (see discussion above of creator skill). This serves as a deep understanding of both the current and *potential* supply of items for recommendation and consumption. Un-

fortunately, no provider has the same breadth of insight into global demand or supply, since they tend to interact with only a small subset of the user population. This information asymmetry limits the ability of creators to make informed content generation decisions, and is a key source of economic inefficiency.

To illustrate (see discussion on Section 4), consider a creator whose content does not attain the desired user engagement: as discussed above, possible causes include: (i) no demand for this type of content; (ii) content quality makes it unattractive to most users; or (iii) numerous other creators offer similar items. While the RS has information that distinguishes these causes, the creator may not. In most RS settings, creators are left to their own devices to explore, design and test new offerings—a process that is generally inefficient, costly, and sometimes disincentivizing.

Breaking the information asymmetry through some form of direct or indirect information sharing can improve the decision-making capabilities of creators and drive significant improvements in both user and creator utility. For instance, if the competitive landscape is limiting a creator’s audience (cause (iii)), the RS could either communicate this directly, or provide more indirect guidance by “steering” the creator to produce content in a less well-supplied part of content space for which latent demand is predicted to be high.

We can interpret the RS demand and supply modeling as a form of “indirect” market research which can be shared with creators to enable the production of items that better meet untapped user demand. But a variety of design and research challenges must be resolved to make this happen. We list a few of these here:

- Direct information sharing may not be feasible in many domains, for example, 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., in our example above, revealing the predicted, aggregate audience for new content in an effort to support a creator’s content creation decisions. Of course, the leakage of private or sensitive information (Dwork et al. 2012) must be safeguarded in any such mechanism.
- Sharing information alone may not be sufficient to induce welfare-improving changes in content generation by a creator. The costs and uncertainty (hence risks) associated with major changes may serve as a disincentive for a creator to generate truly novel content. For example, the creator’s lack of experience, skill or access to resources, their uncertainty regarding the size/value of the new audience, and their potential lack of trust in the RS’s “advice” might all need to be overcome. Hence, appropriate mechanisms must be investigated that: incentivize new content production; de-risk creator exploration; facilitate the development of new skills; and open new audiences—all while maintaining trust in the RS’s suggestions and analytics. Many such mechanisms (e.g., gradual exploration, skill development, access to resources) may extend over lengthy horizons.
- Models of user preferences generally rely on behavioral

interactions with items in the corpus, content features of these items, or a mix of both. Proposing “hypothetical” items to creators thus presents some significant challenges. For example, a collaborative filtering (behavioral) model will generally embed users and items in some latent space (see Figure 1). If some point in content space is predicted to have high user demand, mechanisms for generating actionable descriptions or prompts for creators is critical. There may be a significant role for generative models, such as LLMs (Devlin et al. 2018; Radford et al. 2018; Thoppilan et al. 2022), or text-to-image models (Vinyals et al. 2015; Esser, Rombach, and Ommer 2021; Ramesh et al. 2021), to synthesize language descriptions, item features, evocative images and the like to serve as effective guidance for the creator.

- RSs may also play an active role in content creation. There is a clear trend toward placing generative AI and other tools in the hands of creators (Gillick et al. 2019; Yang, Srivastava, and Mandt 2022), and such approaches could be extended to support explicit product or content “design.” More than this, RSs are in a unique position to facilitate a full design-create-test-refine loop for creators or designers, coupling direct *experimentation and analytics* into the iterative design/creation process. In addition, such experimentation could involve preference elicitation with willing users to serve as a form of *direct* market research.
- While the discussion of information sharing and advising on content generation above has emphasized individual creators, the process is in fact one that has to be *coordinated* across the entire set of creators. Suppose the RS proposes generation of novel content to one creator that suggests a certain audience can be reached. A similar suggestion to a different creator might undercut the audience suggested to the first. As such, we must think of the problem of “nudging” creators to new parts of content space as a large-scale *joint optimization* problem (Prasad, Mladenov, and Boutilier 2023)—see our discussion of the Omnipotent Recommender in Section 2, where we loosely characterized this as a facility location problem. Of course, the problem is rendered even more complex by the recognition that prompting creators is itself a long-horizon, multi-stage process. Moreover, there are subtleties associated with the potential to promote *strategic competition*: unlike facility location, where the most effective solution tends to “spread out” facilities, in an RS ecosystem, the ideal configuration of production may actually keep several creators “near” each other, since some degree of competition for user engagement may actually incentivize the production of higher quality content.
- Once an RS takes steps to “nudge” providers to certain points in content space, the incentive for explicit strategic behavior may, in fact, increase. For example, one creator may refuse to change its content strategy when prompted if it believes that this refusal might induce a competing creator to make such a change (see discussion of Strategic Behavior by creators in Section 8.2).
- Rather than nudging specific creators—a process that

requires estimating (or eliciting) a creator’s incentives, skills and resources— another class of mechanisms for (indirect or direct) information sharing could include the creation of semi-open marketplaces that communicate under-served parts of the market/audience to *any* creator interested in filling that gap. This is itself an interesting market design and optimization problem that should target desirable (long-run) equilibrium behavior.

- Finally, any mechanism that engages in direct or indirect information sharing must take pains to avoid *sharing* sensitive or private information about users or other creators, or *leaking* such information when providing any aggregate data, metrics, prompts, suggestions, etc. (for instance, in the sense of differential privacy (Dwork et al. 2012)). This includes direct information about the preferences or behaviors of specific users that a creator does not already have access to, or revealing information about a competitor’s incentives, strategies, skills or cost structures. While differentially private methods have been explored in standard RS settings (Liu, Wang, and Smola 2015; Chien et al. 2021; Carranza et al. 2023), extending such techniques and analyses to complex ecosystems is needed.

8 Strategic Behavior

By *strategic behavior*, we adopt the standard informal economic/game-theoretic notion, referring to actions taken by an agent that anticipate the actions or reactions of other agents. Our discussion above has largely assumed that users and creators behave non-strategically. We turn to what strategic behavior might look like for both users and content creators (or other item providers).

8.1 User Strategic Behavior

We expect users, for the most part, to behave non-strategically. So when a user is presented with several content recommendations, they will simply select their most preferred option given their current context, possibly subject to random noise. In sequential settings, such a choice may be more challenging since the value of the selected item may depend on future recommendations/selections, and may influence the recommendations the user receives in the future. Accounting for the first factor requires that the user behave sequentially rationally (i.e., plan); but since the RS policy may be influenced by the initial selection, a user may also invoke some “mental model” of the RS policy and thus behave in a way to explicitly influence subsequent recommendations (Guo et al. 2021).

The latter could be viewed as simply planning in the environment “determined” by the RS policy, and thus remains non-strategic—this would be the case if the user believes the RS policy is stationary and/or attempting to act in the user’s best interests. For example, a user might select a (non-preferred) track by a specific artist in a music RS because they believe it will induce the RS to propose additional (preferred) tracks by that artist moving forward. If by contrast the user believed the RS were trying to promote specific content independent of the user’s interests, they might try to take

actions to prevent that, in which case, we might view such behavior as strategic.

Other forms of 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). Handling potential misrepresentation of preferences during explicit elicitation is the primary objective of classic mechanism design approaches, though as discussed in Section 5.3, any mechanism design approach must accommodate a very complex outcome/preference space and the incremental, incomplete, noisy and biased nature of practical user preference assessment. With regard to indirect preference assessment (see Sections 5.1 and 5.2), the somewhat nascent body of work on *incentive-compatible machine learning* (Dekel, Fischer, and Procaccia 2010; Balcan et al. 2005) is certainly relevant, though currently far from meeting the practical demands of complex ecosystems.

All this said, user strategic behavior, if existent at all, is likely to be relatively limited. Often it will involve trying to promote the interests of favorite item providers, and hence be generally somewhat “cooperative.” Moreover, we expect users to exhibit extreme bounded rationality w.r.t. the full complexity of the RS ecosystem—they will rarely consider the true equilibrium state of the system.

We note that the increased use of explicit *two-way* communication about a user’s preferences as advocated above—both in terms of users providing explicit preferences and critiques, and the RS offering explanations and exploration to users—is likely to increase transparency and user trust. This, we believe, will reduce the incentives for users to expend cognitive effort in attempts at “non-strategic” manipulation of RS behavior.

8.2 Provider Strategic Behavior

Strategic reasoning is more likely to play a role in the behavior choices of item providers, such as content creators or product vendors, than it is with users. Disregarding any direct intervention by the RS, a product vendor will generally choose to offer products for sale that it not only predicts to have reasonable demand, but that are sufficiently differentiated from the offerings of other vendors, including those that others may produce or source in response to the choice of the vendor in question. Price setting will likewise be strategic. Similar considerations will generally play a role in content generation by creators even if RS users do not (directly) pay for that content.

An RS that takes actions to maximize some form of social welfare must account for provider strategic behavior. The RS matching policy itself can induce strategic behavior on the part of content providers, a problem that has been analyzed by Ben-Porat and Tennenholtz (2018). Extensions of this modeling approach to richer user and creator preference models would have tremendous import.

From the perspective of more direct elicitation of private information, consider our illustrative example where each provider requires some minimum audience engagement to continue making content available. Directly eliciting this target engagement clearly gives the provider a *prima facie* incentive to overstate their target, but not by too much.³⁰ Direct revelation mechanisms of this sort are akin to auctions, where the matching should account for equilibrium behavior of providers (e.g., as in a first-price auction). But in many RSs (e.g., content RSs) monetary transfer cannot be considered, so this class of problems falls squarely within the area of mechanism design without money (Schummer and Vohra 2007; Procaccia and Tennenholtz 2009a), which can be especially challenging to deal with.

Next consider the setting described above in which the RS takes steps to induce providers to generate (or otherwise make available) items that will improve overall social welfare in the ecosystem. Here again, provider strategic reasoning is likely to emerge, and should be taken into account. For instance, if creators C_1 and C_2 are producing content well-suited to the same audience, they may split the audience, receiving less engagement than if their content were more distinctive. To ameliorate this, the RS might prompt C_1 to make slightly different content, to the benefit of the broader user population as well as both creators. However, C_1 might refuse to do so in the hopes that the RS might then request that C_2 move instead, thus sparing C_1 the potential cost and risk of this change.

Incentivizing welfare-improving behavior across a diverse set of strategic or semi-strategic providers will require significant effort involving many aspects of mechanism design. While direct mechanisms may work in relatively simple domains (e.g., auction-like mechanisms), in complex settings we expect that more detailed modeling of provider private information—utilities (e.g., w.r.t. induced engagement), skills and costs, beliefs, etc.—and how these elements shape their strategies and behaviors—especially when providers exhibit various forms of bounded rationality—should ultimately prove most effective. See also our discussion above of information sharing, which should take into account the potential strategic effect it has on these actors (Myerson 1983) (providers in our case).

9 Dynamic Mechanism Design and Reinforcement Learning

The *sequential* nature of an RS’s interactions with users, creators, vendors and other actors brings with it a host of challenges for the mechanism design perspective. This includes questions regarding preference elicitation (see Section 5.5), learning and optimization. First turning to learning and optimization, constructing optimal sequential interaction policies has attracted a fair amount of attention within the recommendations literature, where a significant body of work

³⁰Notice that while demanding a larger audience increases the odds of the RS recommending a provider’s content to more users, it also runs a greater risk of being deemed infeasible given the demands of other providers, and hence shutting out the provider from being matched at all (Mladenov et al. 2020b).

has formulated sequential recommendations using MDPs or partially observable MDPs (Shani, Heckerman, and Brafman 2005) or tackled the problem using reinforcement learning (Taghipour, Kardan, and Ghidary 2007; Choi et al. 2018; Gauci et al. 2018; Chen et al. 2018; Je et al. 2019). However, the bulk of this work focuses on “local” policies that optimize for one user in isolation. Extending RL methods to handle the large-scale interactions present in ecosystems will generally require joint optimization of the form studied in multiagent RL (Busoniu, Babuška, and De Schutter 2008; Lanctot et al. 2017). The study of such *multiagent sequential* optimization for ecosystems is currently in its infancy (Mladenov et al. 2020a; Zhan et al. 2021).

Turning to incentive issues and information revelation—two key ingredients of mechanism design—sequential interactions are the subject of *dynamic mechanism design (DMD)* (Parkes 2007; Athey and Segal 2013; Bergemann and Välimäki 2010, 2019; Sandholm, Conitzer, and Boutilier 2007; Curry et al. 2023), which provides a framework for mechanism design where different agents can engage with the mechanism at different points in time and possibly over varying horizons. Most of the work in DMD addresses problems involving pricing in dynamic markets (e.g., dynamic pricing for tickets, surge pricing for ride-share programs, wholesale energy markets), and adapt mechanisms like VCG to such settings.

Dynamic mechanisms of this sort may provide useful foundations for mechanism design in ecosystems, but cannot be applied directly for a variety of reasons. One of these is the fact that the focus on pricing problems allows the use of monetary transfers to incentivize suitable behavior and information revelation. But, as discussed above, RSs often fall within the realm of “mechanism design without money,” a problem that has been studied in the challenging dynamics setting only sporadically (Guo and Hörner 2015; Balseiro, Gurkan, and Sun 2019). A second is the fact in the dynamics of complex ecosystems is rarely known, and often must be learned. Recent work has explored the use of RL to learning optimal dynamic mechanisms (Lyu et al. 2022a,b) from repeated interactions with the environment. While directions such as these are promising, nontrivial research is needed to realize the potential of DMD for modeling ecosystems.

Part IV

Tradeoffs and the Social Choice Function

As discussed above, the ecosystem perspective makes it abundantly clear that an RS must make tradeoffs in the utility it generates for different actors in the system: users, creators, vendors, etc. Except in systems with the most trivial outcome spaces, the incentives of these parties will rarely be fully aligned. While there are many types of incentive conflicts, and a variety of ways to “carve them up,” we consider two classes of tradeoffs here: those involving the preferences of the individual actors (Section 10); and those that we refer to

as “social tradeoffs,” pertaining to the properties of the RS policy that may be less easily articulated in terms of an SCF over the preferences of individual participants (e.g., these may be interpreted as imposing externalities beyond the RS ecosystem itself). We elaborate on this distinction below.

10 Tradeoffs over Actor’s Utilities

One of the primary calling cards of mechanism design is the use of an SCF to encode the tradeoffs the designer is prepared to make between the utilities generated by the mechanism for its participants. For example, maximizing utilitarian social welfare generally requires that some agents receive greater utility than others.

Our stylized example in Sec. 3 illustrates this in the case of RS ecosystems, where there is a fundamental tradeoff in the utilities that can be generated for different users—we generate a large increase in utility for the blue users at a small cost in utility for the red users under the “user utilitarian” optimal equilibrium policy. Similar remarks apply to the providers themselves: the “user utilitarian” optimal policy changes the utility mix among the providers, by more evenly distributing engagement across the red and blue providers. If provider utility is equated with engagement, we note that this too improves total provider utility (hence “provider utilitarian” welfare). Of course, under different utility functions, user utility and provider utility may not coincide in this fashion.

The use of mechanism design in the design of RS policies requires the specification of an SCF. This 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. A variety of factors can be brought to bear in the construction of an SCF. Several broad classes of social welfare functions include:

- Utilitarian SCFs, which maximize the sum of (expected) utilities generated over all agents (e.g., users, vendors), or some subclass (users only). One criticism of a pure utilitarian approach is that it can generate wide disparities in utilities across agents (e.g., users or providers), which some may deem unfair (we discuss fairness further below). For instance, in our stylized example, if none of the red users were sufficiently close to the blue provider, the utilitarian-optimal policy would not have “subsidized” the blue provider, thus leading to very low utility for the blue users (and no utility at all in steady state for the blue provider).
- Maxmin SCFs, which maximize the utility of the worst-off individual (with least utility), either across the population or within some subclass. While seemingly more egalitarian, maxmin SCFs can sacrifice a significant amount of utility for a large number of agents to accommodate the worst-off agents.
- SCFs with fairness and/or regret constraints or objectives. These can optimize some primary objective (say, sum of utilities), while constraining the solution in one or more respects. Fairness constraints tend to encourage something akin to maxmin outcomes (ensuring outcome utilities are somewhat less disparate), while regret constraints can

be used to penalize policies that impose too large a cost (loss in utility) on any individual in service of the primary objective (Mladenov et al. 2020b). Fairness is discussed further below, while regret measures require one to define baseline behavior (e.g., a default RS policy) against which to measure the “loss” incurred by an individual.

Naturally, there are many other factors that can be incorporated into the design of an SCF.

The use of an SCF presents numerous challenges for specification, elicitation, assessment, measurement, and optimization as discussed in the preceding sections. For instance, most RSs rarely take steps to assess true user, creator or vendor utility, instead relying on proxy measurements (e.g., return engagement). From an optimization perspective, the fact that agent utility should generally be measured over extended horizons itself presents challenges, since RS rarely optimize policies non-myopically.

Another challenge pertains to interpersonal comparison of utilities (Harsanyi 1955). Abstractly, an SCF simply maps preferences into RS action/policy choices. In practice, however, one generally assumes *quantitative agent utilities* over outcomes, defines a social welfare function that *scores* outcomes as a function of these utilities, and selects the choice that is score (or welfare) maximizing. As such, the tradeoffs in utility of one agent (or group) against that of another embodied in the SWF must assume that these utilities (as estimated or elicited by the RS) are comparable.

The need for quantitative utility functions is especially acute in RSs, where learning methods are commonly used to assess quantitative user affinities, and expectations must be taken over uncertain outcomes. Moreover, boiling things down into individual utilities and assuming comparability provides significant traction w.r.t. optimization. But it remains to be seen what one loses in terms of generality. We note that in much of mechanism design, the use of monetary transfers and the assumption of quasi-linear utility obviates this concern (e.g., this assumption underlies most of auction theory). As discussed above, however, mechanism design within RSs will often fall within the realm of *mechanism design without money* (Schummer and Vohra 2007; Procaccia and Tennenholtz 2009b).

Questions about disparate treatment of actors are often treated within the domain of ML fairness. Classical notions of fairness are concerned with equalising the performance of a machine learning model amongst groups or individuals by correcting for various data or modelling biases. In RSs, these concerns are reflected both on the user and creator side (Ekstrand et al. 2022; Li et al. 2023), where the goal is generally to achieve more equitable model quality of user preferences/item features across various protected groups. Another form of fairness is that of fairness of exposure, which deals with the fair distribution of attention to creators, usually based on their estimated quality, by manipulating the ranking of recommended items. All of these approaches can be seen as implicit attempts to promote certain social outcomes over others, but the relationship between adjusting models and actual utility of outcomes can sometimes be hard to quantify. Therefore, SCFs can be seen as a more powerful language to address these issues at the level of outcomes, rather than

interventions on models.

11 Social Objectives

Some elements of the SCF may not be readily definable in terms of the utilities or preferences of the individual actors. This may be the case, for instance, when the SCF involves properties of the *joint* outcome over, say, users, but where each user is concerned only about their “local” part of that outcome. One such phenomenon with this property, commonly discussed in relation to RSs, is that of *filter bubbles*, and the potential for induced social or intellectual fragmentation across the user population (Pariser 2011; Aridor, Goncalves, and Sikdar 2020). This may emerge when the preferences of individual users do not explicitly value diversity w.r.t. some types of content.³¹ Thus, different user subpopulations, when engaging with an RS that can effectively maximize user utility, may end up consuming different collections of content.

Whether such filter bubbles are problematic, of course, depends on the nature of the content in question. A “filter bubble” in which one groups of users largely consumes music from the sub-genre electro-house, while another listens to indie folk will should not be considered problematic. But if the content, say, involves diverse perspectives on social, scientific, civic or political issues of significant social import, the fact that large groups of users have no exposure to each other’s perspectives may be of greater concern. This might be best viewed as a societal value, where these consumption patterns have a negative societal externality, e.g., preventing meaningful civic discourse. Penalizing such patterns (or assigning utility to at least some degree of overlap in consumption) to minimize these externalities within the SCF (i.e., as part of the RS’s long-term objective) provides one means for encoding such considerations. Other phenomena of this sort include polarization and radicalization (Ribeiro et al. 2020; Musco, Musco, and Tsourakakis 2018), personal diversity of consumption, “echo chambers,” and the like.³²

Finally, it is worth drawing attention to the interaction between social objectives and fairness. For example, rich-get-richer feedback loops—which can easily lead to heavy-tailed distributions of content/creator popularity, influence or income—are typically seen as undesirable from the point of view of fairness. Yet, economies of scale often have the property that high production-value content can be more easily produced by a single creator with access to greater resources or budget than by ten creators, each with one-tenth of the budget, under conditions of competition. Moreover, having a few items of content with out-sized popularity may be inherently socially beneficial as “cornerstone content” for building

³¹This preference attitude may be due to general indifference, or may be an active social preference, e.g., as in the case of a user’s perception of being an influential member of a community, and the associated sense of belonging/identity (Klein 2020).

³²Of course, if one believes that these are harmful externalities for the RS’s users themselves, in principle, one could encode these considerations in the utility functions of the users; but doing so would require a deeply sophisticated model of long-term outcome dynamics. Hence, simpler proxies involving overlap and/or diversity of content consumption should prove more actionable in practical reecosystems.

genres and communities. It is easy to construct hypothetical scenarios where fairness and other societal values are at odds. This too suggests that SCFs should be taken seriously as a means for addressing fundamental tradeoffs in reecosystems.

12 The Elephant in the Room

One of the most fundamental questions in the use of mechanism design to design and optimize reecosystems is the selection of the SCF. While akin to the selection of the objective in any optimization problem, or loss function in any machine learning problem—especially those of a multi-objective or multi-task nature—deciding on an SCF carries with it additional “baggage,” considerations that (usually) play no role in ML or traditional multi-objective problems.

One distinction is that tradeoffs are being made over the utility generated for different individuals. This stands in contrast to single-party multi-objective optimization, where tradeoffs across objectives lie in the hands, and only impact, a single decision making entity. The self-interested agents impacted by the RS policy may not agree with the tradeoffs being made (see more below). This ultimately raises important questions regarding who should help shape the nature of the SCF, and who should make the final determination? Such considerations are rendered more complex still if one is to account for social objectives of the form discussed in Section 11.

A second factor are the seminal impossibility results which make the considerations above even more difficult. For instance, Arrow’s (1950) impossibility theorem makes clear that appeal to some “universal” SCF that embodies incontrovertible principles is pointless.³³ Even seemingly uncontroversial conditions like Pareto optimality can sometimes be called into question on philosophical grounds (Sen 1970).

Apart from deciding how to aggregate preferences to make collective choices, the question of accurately eliciting or estimating the preferences of the participants in the RS ecosystem is yet a third distinguishing factor. Indeed, the self-interested nature of the reecosystem participants renders this especially challenging. The classic Gibbard-Satterthwaite (Gibbard 1973; Satterthwaite 1975) impossibility result states that no mechanism can incentivize all participants to reveal their preferences truthfully (except, again, in relatively simple settings). This result may be of marginally less practical import than Arrow’s, since strategic misrepresentation or misreporting of preferences can be both computationally and informationally difficult, and can often have a relatively limited impact on outcome quality (Lu et al. 2012); but the possibility of strategic misreporting should still play a role in the selection of a suitable SCF.

As a result, determining a suitable SCF may be one of the most daunting tasks in the application of mechanism de-

³³Arrow’s result states that, except in some simple settings, no SCF (i.e., preference aggregation scheme) satisfies all of a set of (what are generally viewed as) reasonable conditions. This mathematical impossibility does not mean aggregation doesn’t work; it simply means that, for any (non-dictatorial) SCF, there exists *some* configuration of preferences where the SCF will do something considered to be unreasonable.

sign principles to the design of recoms. For *any* given SCF, simply predicting its ultimate impact on the *realized* outcomes of an RS policy, and the *realized* expected utility generated for ecosystem players (and their tradeoffs), is tremendously complex. The development of analytical, simulation, visualization and scenario-analysis tools seems vital to help in exploring the space of SCFs. This should also play a critical role in the transparency of RS designs and their consequences for individuals and societal groups.

Part V

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, and on societal dynamics and values.

References

Abdollahpouri, H. 2019. Popularity Bias in Ranking and Recommendation. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 529–530.

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.

Akpınar, 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.

Arora, S.; Raghavan, P.; and Rao, S. 1998. Approximation Schemes for Euclidean k-medians and Related Problems. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing (STOC-98)*, 106–113.

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.

Athey, S.; and Segal, I. 2013. An Efficient Dynamic Mechanism. *Econometrica*, 81(6): 2463–2485.

Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47: 235–256.

Balcan, M.-F.; Blum, A.; Hartline, J. D.; and Mansour, Y. 2005. Mechanism Design via Machine Learning. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS-05)*, 605–614.

Balseiro, S. R.; Gurkan, H.; and Sun, P. 2019. Multiagent Mechanism Design Without Money. *Operations Research*, 67(5): 1417–1436.

Barocas, S.; Hardt, M.; and Narayanan, A. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. Cambridge, MA: MIT Press.

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.; Goren, G.; Rosenberg, I.; and Tennenholtz, M. 2019. From Recommendation Systems to Facility Location Games. In *Proceedings of the Thirty-third AAAI Conference on Artificial Intelligence (AAAI-19)*, 1772–1779. Honolulu.

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.

Bergemann, D.; and Välimäki, J. 2010. The Dynamic Pivot Mechanism. *Econometrica*, 78(2): 771–789.

Bergemann, D.; and Välimäki, J. 2019. Dynamic Mechanism Design: An Introduction. *Journal of Economic Literature*, 57(2): 235–274.

Beyeler, M.; Brero, G.; Lubin, B.; and Seuken, S. 2021. iMLCA: Machine Learning-powered Iterative Combinatorial Auctions with Interval Bidding. In *Proceedings of the Twenty-second ACM Conference on Economics and Computation (EC’21)*, 136–136.

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.

Blum, A.; Jackson, J. C.; Sandholm, T.; and Zinkevich, M. 2004. Preference Elicitation and Query Learning. *Journal of Machine Learning Research*, 5: 649–667.

- 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. 2013. Computational Decision Support: Regret-based Models for Optimization and Preference Elicitation. In Crowley, P. H.; and Zentall, T. R., eds., *Comparative Decision Making: Analysis and Support Across Disciplines and Applications*, 423–453. Oxford: Oxford University Press.
- Boutilier, C.; Patrascu, R.; Poupart, P.; and Schuurmans, D. 2006. Constraint-based Optimization and Utility Elicitation using the Minimax Decision Criterion. *Artificial Intelligence*, 170(8–9): 686–713.
- Burke, R. 2002. Interactive Critiquing for Catalog Navigation in E-Commerce. *Artificial Intelligence Review*, 18(3–4): 245–267.
- Busoni, L.; Babuška, R.; and De Schutter, B. 2008. Multi-Agent Reinforcement Learning: An Overview. *IEEE Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, 38(2): 156–172.
- Camerer, C. F.; Loewenstein, G.; and Rabin, M., eds. 2003. *Advances in Behavioral Economics*. Princeton, New Jersey: Princeton University Press.
- Carranza, A. G.; Farahani, R.; Ponomareva, N.; Kurakin, A.; Jagielski, M.; and Nasr, M. 2023. Privacy-Preserving Recommender Systems with Synthetic Query Generation using Differentially Private Large Language Models. ArXiv:2305.05973.
- Castera, R.; Loiseau, P.; and Pradelski, B. S. 2022. Statistical Discrimination in Stable Matchings. In *Proceedings of the Twenty-third ACM Conference on Economics and Computation (EC’22)*, 373–374. Boulder, CO.
- Celis, L. E.; and Vishnoi, N. K. 2017. Fair Personalization. ArXiv:1707.02260.
- Celma, Ò. 2010. The Long Tail in Recommender Systems. In *Music Recommendation and Discovery*, 87–107. Springer.
- 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*.
- Charness, G.; Gneezy, U.; and Imas, A. 2013. Experimental Methods: Eliciting Risk Preferences. *Journal of Economic Behavior and Organization*, 87: 43–51.
- Chen, L.; and Pu, P. 2012. Critiquing-based Recommenders: Survey and Emerging Trends. *User Modeling and User-Adapted Interaction*, 22(1): 125–150.
- Chen, M.; Beutel, A.; Covington, P.; Jain, S.; Belletti, F.; and Chi, E. 2018. Top-K Off-Policy Correction for a REINFORCE Recommender System. In *12th ACM International Conference on Web Search and Data Mining (WSDM-19)*, 456–464. Melbourne, Australia.
- 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.
- Choi, S.; Ha, H.; Hwang, U.; Kim, C.; Ha, J.-W.; and Yoon, S. 2018. Reinforcement Learning-based Recommender System using Biclustering Technique. ArXiv:1801.05532.
- Chow, Y.; Ghavamzadeh, M.; Janson, L.; and Pavone, M. 2017. Risk-constrained Reinforcement Learning with Percentile Risk Criteria. *Journal of Machine Learning Research*, 18(1): 6070–6120.
- Christakopoulou, K.; and Banerjee, A. 2018. Learning to Interact with Users: A Collaborative-Bandit Approach. In *Proceedings of the 2018 SIAM International Conference on Data Mining (ICDM-18)*, 612–620. San Diego.
- 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.
- Covington, P.; Adams, J.; and Sargin, E. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, 191–198. Boston.
- Cramton, P.; Shoham, Y.; and Steinberg, R., eds. 2005. *Combinatorial Auctions*. Cambridge: MIT Press.
- Curry, M.; Thoma, V.; Chakrabarti, D.; McAleer, S. M.; Kroer, C.; Sandholm, T.; He, N.; and Seuken, S. 2023. Automated Design of Affine Maximizer Mechanisms In Dynamic Settings. In *Sixteenth European Workshop on Reinforcement Learning (EWRL-23)*. Brussels.
- Dave, K.; and Varma, V. 2014. Computational Advertising: Techniques for Targeting Relevant Ads. *Foundations and Trends in Information Retrieval*, 8(4–5): 263–418.
- DeGroot, M. H. 1974. Reaching a Consensus. *Journal of the American Statistical Association*, 69(345): 118–121.
- 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.
- Ericson, K. M.; and Laibson, D. 2019. Intertemporal Choice. In Bernheim, B. D.; DellaVigna, S.; and Laibson, D., eds., *Handbook of Behavioral Economics: Applications and Foundations*, volume 2, chapter 1, 1–67. Elsevier.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming Transformers for High-resolution Image Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-21)*, 12873–12883.
- Fang, L.; Du, B.; and Wu, C. 2022. Differentially Private Recommender System with Variational Autoencoders. *Knowledge-Based Systems*, 250. Article 109044.

- Friedman, L.; Ahuja, S.; Allen, D.; Tan, T.; Sidahmed, H.; Long, C.; Xie, J.; Schubiner, G.; Patel, A.; Lara, H.; et al. 2023. Leveraging Large Language Models in Conversational Recommender Systems. ArXiv:2305.07961.
- Gao, W.; Gao, X.; and Tang, Y. 2023. Multi-Turn Dialogue Agent as Sales' Assistant in Telemarketing. In *2023 International Joint Conference on Neural Networks (IJCNN)*. Gold Coast, Australia.
- Gauci, J.; Conti, E.; Liang, Y.; Virochsiri, K.; He, Y.; Kaden, Z.; Narayanan, V.; and Ye, X. 2018. Horizon: Facebook's Open Source Applied Reinforcement Learning Platform. ArXiv:1811.00260.
- Gibbard, A. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4): 587–601.
- Gillick, J.; Roberts, A.; Engel, J.; Eck, D.; and Bamman, D. 2019. Learning to Groove with Inverse Sequence Transformations. In *Proceedings of the Thirty-sixth International Conference on Machine Learning (ICML-19)*, 2269–2279. Long Beach, CA.
- Goodrow, C. 2021. On YouTube's Recommendation System.
- Göpfert, C.; Chow, Y.; Hsu, C.-W.; Vendrov, I.; Lu, T.; Ramachandran, D.; and Boutilier, C. 2022. Discovering Personalized Semantics for Soft Attributes in Recommender Systems Using Concept Activation Vectors. In *Proceedings of the 31st International Web Conference (WWW-22)*, 2411–2421. Lyon. (Also arXiv 2202.02830).
- Green, P. E.; and Srinivasan, V. 1990. Conjoint Analysis in Marketing: New Developments with Implications for Research and Practice. *Journal of Marketing*, 54(4): 3–19.
- 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.
- Guo, Y.; and Hörner, J. 2015. Dynamic Mechanisms without Money. IHS Economics Series, No. 310, Institute for Advanced Studies (IHS), Vienna.
- 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.
- Hohnhold, H.; O'Brien, D.; and Tang, D. 2015. Focusing on the Long-term: It's Good for Users and Business. In *Proceedings of the Twenty-first ACM International Conference on Knowledge Discovery and Data Mining (KDD-15)*, 1849–1858. Sydney.
- Hron, J.; Krauth, K.; Jordan, M. I.; Kilbertus, N.; and Dean, S. 2022. Modeling Content Creator Incentives on Algorithm-curated Platforms. ArXiv:2206.13102.
- Hu, Y.; Koren, Y.; and Volinsky, C. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *Proceedings of the Eighth IEEE International Conference on Data Mining (ICDM-08)*, 263–272. Pisa, Italy.
- 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.
- Hyafil, N.; and Boutilier, C. 2007. Mechanism Design with Partial Revelation. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, 1333–1340. Hyderabad, India.
- Ie, E.; Jain, V.; Wang, J.; Narvekar, S.; Agarwal, R.; Wu, R.; Cheng, H.-T.; Chandra, T.; and Boutilier, C. 2019. SlateQ: A Tractable Decomposition for Reinforcement Learning with Recommendation Sets. In *Proceedings of the Twenty-eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*, 2592–2599. Macau.
- 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.
- Joachims, T.; Granka, L.; Pan, B.; Hembrooke, H.; Radlinski, F.; and Gay, G. 2007. Evaluating the Accuracy of Implicit Feedback from Clicks and Query Reformulations in Web Search. *ACM Transactions on Information Systems*, 25(2). Article 7.
- Kahneman, D.; and Tversky, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263–292.
- Kahneman, D.; and Tversky, A., eds. 2000. *Choices, Values, and Frames*. Cambridge: Cambridge University Press.
- Karaliopoulos, M.; Koutsopoulos, I.; and Spiliopoulos, L. 2019. Optimal User Choice Engineering in Mobile Crowdsensing with Bounded Rational Users. In *IEEE Conference on Computer Communications (INFOCOM-19)*, 1054–1062.
- Keeney, R. L. 1992. *Value-Focused Thinking: A Path to Creative Decisionmaking*. Cambridge, MA: Harvard University Press.
- Keeney, R. L.; and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. New York: Wiley.
- Klein, E. 2020. *Why We're Polarized*. Simon and Schuster.
- Kohli, R.; and Sukumar, R. 1990. Heuristics for Product-line Design using Conjoint Analysis. *Management Science*, 1464–1478.
- 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, J.; Zhang, Y.; and Liu, X. 2008. Personalized Interactive Faceted Search. In *Proceedings of the 17th International World Wide Web Conference (WWW-08)*, 477–486. Beijing.

- Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix Factorization Techniques for Recommender Systems. *IEEE Computer*, 42(8): 30–37.
- 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.
- Lahaie, S.; and Parkes, D. 2004. Applying Learning Algorithms to Preference Elicitation. In *ACM Conference on Electronic Commerce*, 180–188. New York.
- Laibson, D. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(2): 443–478.
- Langlotz, M.; Zambaldi, V.; Gruslys, A.; Lazaridou, A.; Tuyls, K.; Pérolat, J.; Silver, D.; and Graepel, T. 2017. A Unified Game-Theoretic Approach to Multiagent Reinforcement Learning. In *Advances in Neural Information Processing Systems 30 (NIPS-17)*, 4193–4206. Long Beach, CA.
- 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.
- Liu, Z.; Wang, Y.-X.; and Smola, A. 2015. Fast Differentially Private Matrix Factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys15)*, 171–178. Boston.
- Louca, R.; Bhattacharya, M.; Hu, D.; and Hong, L. 2019. Joint Optimization of Profit and Relevance for Recommendation Systems in E-commerce. In *RMSE Workshop, RecSys-19*. Copenhagen.
- Lu, T.; and Boutilier, C. 2020. Preference Elicitation and Robust Winner Determination for Single- and Multi-winner Social Choice. *Artificial Intelligence*, 279. Article:103203.
- Lu, T.; Tang, P.; Procaccia, A. D.; and Boutilier, C. 2012. Bayesian Vote Manipulation: Optimal Strategies and Impact on Welfare. In *Proceedings of the Twenty-eighth Conference on Uncertainty in Artificial Intelligence (UAI-12)*, 543–553. Catalina, CA.
- Lyu, B.; Meng, Q.; Qiu, S.; Wang, Z.; Yang, Z.; and Jordan, M. I. 2022a. Learning Dynamic Mechanisms in Unknown Environments: A Reinforcement Learning Approach. ArXiv:2202.12797.
- Lyu, B.; Wang, Z.; Kolar, M.; and Yang, Z. 2022b. Pessimism meets VCG: Learning Dynamic Mechanism Design via Offline Reinforcement Learning. In *Proceedings of the Thirty-ninth International Conference on Machine Learning (ICML-22)*, 14601–14638. Baltimore.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. New York: Oxford University Press.
- McAfee, R. P.; and McMillan, J. 1987. Auctions and Bidding. *Journal of Economic Literature*, 25: 699–738.
- 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.
- Mihatsch, O.; and Neuneier, R. 2002. Risk-sensitive Reinforcement Learning. *Machine learning*, 49: 267–290.
- Mladenov, M.; Creager, E.; Ben-Porat, O.; Swersky, K.; Zemel, R.; and Boutilier, C. 2020a. Optimizing Long-term Social Welfare in Recommender Systems: A Constrained Matching Approach. ArXiv:2008.00104.
- Mladenov, M.; Creager, E.; Swersky, K.; Ben-Porat, O.; Zemel, R. S.; and Boutilier, C. 2020b. 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.
- Musco, C.; Musco, C.; and Tsourakakis, C. E. 2018. Minimizing Polarization and Disagreement in Social Networks. In *The 2018 Web Conference (WWW 2018)*. Lyon.
- Myerson, R. B. 1983. Mechanism Design by an Informed Principal. *Econometrica*, 1767–1797.
- Nisan, N. 2007. Introduction to Mechanism Design (for Computer Scientists). In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V., eds., *Algorithmic Game Theory*, 209–241. Cambridge University Press.
- 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. 1999. iBundle: An efficient ascending price bundle auction. In *Proceedings of the First ACM Conference on Electronic Commerce (EC'99)*, 148–157. Denver.
- 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.
- Prasad, S.; Mladenov, M.; and Boutilier, C. 2023. Content Prompting: Modeling Content Provider Dynamics to Improve User Welfare in Recommender Ecosystems. In *RecSys-23 Workshop on Causality, Counterfactuals & Sequential Decision-Making (CONSEQUENCES-23)*. Singapore.
- Prashanth, L.; Jie, C.; Fu, M.; Marcus, S.; and Szepesvári, C. 2016. Cumulative Prospect Theory meets Reinforcement

- Learning: Prediction and Control. In *Proceedings of the Thirty-third International Conference on Machine Learning (ICML-16)*, 1406–1415. New York.
- Procaccia, A. D.; and Tennenholtz, M. 2009a. Approximate Mechanism Design without Money. In *Proceedings of the Tenth ACM Conference on Electronic Commerce (EC'09)*, 177–186. Stanford, CA.
- Procaccia, A. D.; and Tennenholtz, M. 2009b. Approximate Mechanism Design without Money. In *Proceedings of the Tenth ACM Conference on Electronic Commerce (EC'09)*, 177–186. Cambridge, MA.
- 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.
- Radlinksi, 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.
- Regan, K.; and Boutilier, C. 2011. Eliciting Additive Reward Functions for Markov Decision Processes. In *Proceedings of the Twenty-second International Joint Conference on Artificial Intelligence (IJCAI-11)*, 2159–2164. Barcelona.
- 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.
- Ribero, M.; Henderson, J.; Williamson, S.; and Vikalo, H. 2022. Federating Recommendations using Differentially Private Prototypes. *Pattern Recognition*, 129. Article 108746.
- Roberts, K. 1979. The Characterization of Implementable Choice Rules. In Laffont, J.-J., ed., *Aggregation and Revelation of Preferences*, 321–349. Amsterdam: North-Holland.
- 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 Trans. 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.
- Sandholm, T.; Conitzer, V.; and Boutilier, C. 2007. Automated Design of Multistage Mechanisms. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, 1500–1506. Hyderabad, India.
- 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.
- Schön, C. 2010. On the Optimal Product Line Selection Problem with Price Discrimination. *Management Science*, 56(5): 896–902.
- Schummer, J.; and Vohra, R. V. 2007. Mechanism Design without Money. In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V., eds., *Algorithmic Game Theory*, 243–265. Cambridge University Press.
- Sen, A. 1970. The Impossibility of a Paretian Liberal. *Journal of Political Economy*, 78(1): 152–157.
- Shachter, R. D. 1986. Evaluating Influence Diagrams. *Operations Research*, 33(6): 871–882.
- Shani, G.; Heckerman, D.; and Brafman, R. I. 2005. An MDP-based Recommender System. *Journal of Machine Learning Research*, 6: 1265–1295.
- Singh, A.; and Joachims, T. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD-18)*, 2219–2228. London: Association for CompMachinery.
- Sun, Y.; and Zhang, Y. 2018. Conversational recommender system. In *41st International ACM/SIGIR Conference on Research & Development in Information Retrieval (SIGIR-18)*, 235–244.
- Taghipour, N.; Kardan, A.; and Ghidary, S. S. 2007. Usage-based Web Recommendations: A Reinforcement Learning Approach. In *Proceedings of the 1st ACM Conference on Recommender Systems (RecSys07)*, 113–120. Minneapolis: ACM.
- Talluri, K.; and van Ryzin, G. 2004. Revenue Management under a General Discrete Choice Model of Consumer Behavior. *Management Science*, 50(1): 15–33.
- Tang, L.; Jiang, Y.; Li, L.; and Li, T. 2014. Ensemble Contextual Bandits for Personalized Recommendation. In *Proceedings of the 8th ACM Conference on Recommender Systems (RecSys14)*, 73–80. Foster City, CA.
- Tennenholtz, M.; and Kurland, O. 2019. Rethinking Search Engines and Recommendation Systems: A Game Theoretic Perspective. *Communications of the ACM*, 62(12): 66–75.
- 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.

- Tversky, A.; and Kahneman, D. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157): 1124–1131.
- Viappiani, P.; Faltings, B.; and Pu, P. 2006. Preference-based Search using Example-Critiquing with Suggestions. *Journal of Artificial Intelligence Research*, 27: 465–503.
- Vig, J.; Sen, S.; and Riedl, J. 2012. The Tag Genome: Encoding Community Knowledge to Support Novel Interaction. *ACM Transactions on Interactive Intelligent Systems*, 2(3): 13:1–13:44.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and Tell: A Neural Image Caption Generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-15)*, 3156–3164.
- Wang, C.; Kalra, A.; Borcea, C.; and Chen, Y. 2015. Revenue-optimized Webpage Recommendation. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, 1558–1559. Atlantic City.
- Warren, C.; McGraw, A. P.; and Van Boven, L. 2011. Values and Preferences: Defining Preference Construction. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(2): 193–205.
- Wei, B.; Liu, J.; Zheng, Q.; Zhang, W.; Fu, X.; and Feng, B. 2013. A Survey of Faceted Search. *Journal of Web Engineering*, 12(1 & 2): 41–64.
- White, C. C., III; Sage, A. P.; and Dozono, S. 1984. A Model of Multiattribute Decision-making and Trade-off Weight Determination under Uncertainty. *IEEE Transactions on Systems, Man and Cybernetics*, 14(2): 223–229.
- Wirth, C.; Akrou, R.; Neumann, G.; and Fürnkranz, J. 2017. A Survey of Preference-based Reinforcement Learning Methods. *Journal of Machine Learning Research*, 18(136): 1–46.
- 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.
- Yang, R.; Srivastava, P.; and Mandt, S. 2022. Diffusion Probabilistic Modeling for Video Generation. ArXiv:2203.09481.
- 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.
- Yue, Y.; Broder, J.; Kleinberg, R.; and Joachims, T. 2012. The k-armed Dueling Bandits pProblem. *Journal of Computer and System Sciences*, 78(5): 1538–1556.
- Zhan, R.; Christakopoulou, K.; Le, Y.; Ooi, J.; Mladenov, M.; Beutel, A.; Boutilier, C.; Chi, E.; and Chen, M. 2021. Towards Content Provider Aware Recommender Systems: A Simulation Study on the Interplay Between User and Provider Utilities. In *Proceedings of the Web Conference 2021 (WWW-21)*, 3872–3883.
- Zhang, L.; and Levin, D. 2017. Bounded Rationality and Robust Mechanism Design: An Axiomatic Approach. *American Economic Review*, 107(5): 235–239.