Tailored Recommendations*

Eric Danan[†]

Thibault Gajdos[‡]

Jean-Marc Tallon[§]

October 1, 2020

Abstract

Many popular internet platforms use so-called collaborative filtering systems to give personalized recommendations to their users, based on other users who provided similar ratings for some items. We propose a novel approach to such recommendation systems by viewing a recommendation as a way to extend an agent's expressed preferences, which are typically incomplete, through some aggregate of other agents' expressed preferences. These extension and aggregation requirements are expressed by an Acceptance and a Pareto principle, respectively. We characterize the recommendation systems satisfying these two principles and contrast them with collaborative filtering systems, which typically violate the Pareto principle.

Keywords. Recommendation systems; incomplete preferences; extension; aggregation; Pareto principle; collaborative filtering.

JEL Classification. D71, D81.

1 Introduction

The digitalization of our societies has enabled the personalization of advice to an extent never seen before. We now routinely receive recommendations or advertisements not for the “best” possible product, but, rather, for the one that is “best fit” for us. Netflix, for instance, adjusts the movies and shows it suggests to each user based on her profile, what she has seen and how much she appreciated it. Different users receive different recommendations. As for many other internet platforms using so-called “collaborative filtering” recommendation systems (Facebook, Twitter, Amazon, Spotify, Last.fm, LinkedIn, ...), the recommendation for a given user will hinge upon other users that are “similar” to her in one way or another. How this notion of similarity between users is modeled is crucial for the properties of the recommendation system.

The present paper approaches the issue of personalized, tailored recommendation from a normative perspective. Abstracting from the sophistication of actual recommendation systems, we explore the implications of imposing basic axiomatic principles in a simple formal framework. More specifically, we

*We thank Franz Dietrich, Georgios Gerasimou, Itzhak Gilboa, Philippe Jehiel, Antonin Macé, Tigran Melkonyan, Efe Ok, Marcus Pivato, Daniel Read, Régis Renault, Zvi Safra, two anonymous referees, and seminar participants at Düsseldorf U., Warwick Business School, PSE, U. St Andrews as well as participants to the conference Times, Uncertainty and Strategy (Dec. 2018) and the CHOp Workshop, for useful feedback. This work is part of the project “Coping with Heterogeneous Opinions”, which benefits from the support of the grant ANR-17-CE26-0003. Danan acknowledges financial support from the grants ANR-11-LBX-0023-01 and ANR-16-IDEX-0008. Tallon acknowledges financial support from the grant ANR-17-EURE-001.

[†]CY Cergy Paris Université, CNRS, THEMA, F-95000 Cergy, France (eric.danan@cyu.fr).

[‡]LPC, CNRS, Université d'Aix-Marseille (thibault.gajdos@univ-amu.fr).

[§]Paris School of Economics, CNRS (jean-marc.tallon@psemail.eu).

view a recommendation system as a way to “extend” an agent’s preferences to alternatives she has not yet rated, on the basis of some aggregate of other agents’ preferences. Preferences here encode the information available to the recommendation system – typically based on a user’s ratings or viewing behavior. Agents might actually be able to rank more alternatives than those they have rated, but such additional rankings are not modeled as the recommendation system cannot take them into account. We implicitly interpret the available information as being truthful, thereby leaving aside strategic considerations. We allow all agents’ preferences to be incomplete, i.e. not rank all alternatives – and actually expect these rankings to be quite sparse in practice, given the vast amount of alternatives the above platforms cope with.

Formally, we consider a society of $N + 1$ agents, $n = 0, 1, \dots, N$. Agent 0 is the one to whom we want to give a recommendation. Each agent is endowed with a possibly incomplete preference relation and we look for a recommendation in the form of another possibly incomplete preference relation. The actual list of alternatives suggested to agent 0 can then be generated from this recommendation ranking in various ways – e.g. suggesting the top-ranked alternatives among those available at a given time, or the top-ranked alternatives agent 0 has rated along with the top-ranked ones she has not rated. We do not model this stage and instead focus on desirable properties of the recommendation ranking. We consider two basic axioms:

Acceptance principle Whenever agent 0 has expressed a preference for an alternative over another, the former is recommended over the latter.

Pareto principle Whenever agents $1, \dots, N$ have expressed a unanimous preference for an alternative over another, the former is recommended over the latter.

The Acceptance principle captures the requirement that the recommendation to agent 0 must extend her preferences whereas the Pareto principle captures the requirement that it must aggregate the other agents’ preferences. As we first show through illustrative examples, this formalization of the recommendation problem makes it apparent that existing collaborative filtering systems violate the Pareto principle.

We then provide a simple and tractable characterization of the recommendations satisfying these two principles. To this end, we assume all agents have expected utility preferences over lotteries. Such preferences, when complete, can be represented by a (cardinal) expected utility function ([von Neumann and Morgenstern, 1944](#)) and, when incomplete, by a set of “possible” expected utility functions in the sense that a lottery is preferred to another if and only if all functions in the set unanimously assign a higher expected utility to the former than the latter ([Baucells and Shapley, 2008](#); [Dubra et al., 2004](#)). When there is no such unanimity in favor of one of the two lotteries, they are unranked. While in practice the alternatives at stake are usually sure prospects rather than lotteries, one can always take as primitive agents’ ratings over sure prospects and derive their preferences over lotteries “parsimoniously” under the expected utility axioms ([Dubra and Ok, 2002](#)).

Our main result establishes that the recommendations satisfying the Acceptance and Pareto principles are those that are based on possible utilities for agent 0 that are also positive linear combinations of possible utilities for agents $1, \dots, N$. Roughly speaking, such recommendations hinge upon tailored “virtual guides” who have complete preferences exactly agreeing with those of agent 0 on the alternatives she has ranked, while (linearly) aggregating those of agents $1, \dots, N$. The class of such recommendations contains, at one extreme, complete recommendations based on a single virtual guide and, at the

other extreme, the most incomplete recommendation based on the whole set of virtual guides. As we illustrate below, even this most incomplete recommendation generally ranks alternatives that agent 0 has not ranked.

The paper is built as follows. We end this introduction with a very brief literature review. Section 2 presents two stylized examples illustrating how the virtual guide recommendations we will later characterize work and contrasting them with typical collaborative filtering systems, in particular with respect to the Pareto principle. Section 3 presents the formal model and results. Besides the main characterization mentioned above, we also characterize a more general class of “signed” virtual guide recommendations and analyze the computation of (signed) virtual guides as well as how to select among them. Section 4 concludes and highlights avenues for future research to go beyond the illustrative examples and simple model presented here. Proofs appear in the Appendix.

There is obviously a huge literature in computer science on recommendation systems that we will not attempt to summarize (see e.g. the survey by [Chen et al., 2018](#)). Applications of the axiomatic approach based on social choice concepts seem scarce in this literature ([Pennock et al., 2000](#); [Altman and Tennenholtz, 2007](#)). There is, on the other hand, a large axiomatic literature in economics on preference aggregation (for expected utility preferences in particular, see e.g. [Harsanyi, 1953, 1955](#); [Weymark, 1991](#); [Mongin and Pivato, 2016](#)). A few papers in this literature allow for incomplete preferences ([Pivato, 2011, 2013, 2014](#); [Danan et al., 2013, 2015, 2016](#)). The standard preference aggregation problem can be seen as the special case of our tailored recommendation problem in which the preferences of the agent to which the recommendation is made are “void” – rank no alternatives. Non-tailored recommendations have been explored axiomatically by e.g. [Demange \(2014, 2017\)](#). [Eliaz and Spiegler \(2019\)](#) consider the problem of providing a tailored recommendation to a user through statistical models estimated on other users rather than aggregation of their preferences. They take the models as given and focus on strategic issues that we ignore here.

2 Examples

2.1 Rich ratings

Consider five alternatives a, b, c, d, e and three agents 0, 1, 2. Agents’ preferences are represented by the following utility functions u_0, u_1, u_2 , that one can for instance assimilate to ratings on a scale between 0 and 10 given on an internet platform:

	a	b	c	d	e
u_0	3	4	6	-	-
u_1	5	2	10	5	8
u_2	1	6	2	9	1

We want to give a recommendation to agent 0 by extending her preferences to alternatives d and e that she has not yet rated, based on the preferences of agents 1 and 2 who have rated all alternatives. The collaborative filtering systems used by popular internet platforms typically fill agent 0’s missing rating

for alternative $z = d, e$ with:

$$\bar{u}_0 + \frac{s_{0,1}(u_1(z) - \bar{u}_1) + s_{0,2}(u_2(z) - \bar{u}_2)}{|s_{0,1}| + |s_{0,2}|},$$

where $\bar{u}_n$ is the average of the ratings expressed by agent $n = 0, 1, 2$ – $\bar{u}_0 \approx 4.33, \bar{u}_1 = 6, \bar{u}_2 = 3.75$ – and $s_{0,n}$ is some measure of the similarity between the ratings of agent 0 and those of agent $n = 1, 2$ on the commonly-rated alternatives a, b, c . A typical similarity measure is the Pearson correlation coefficient – $s_{0,1} \approx 0.76, s_{0,2} = 0$ – leading to fill agent 0’s missing ratings for alternatives d and e with 3.33 and 6.33, respectively. This yields the recommendation ranking $e > c > b > d > a$.

The “virtual guide” recommendations characterized by the Acceptance and Pareto principles, on the other hand, are based on positive linear combinations of agents 1 and 2’s ratings that agree with agent 0’s ratings on the commonly-rated alternatives a, b, c . The unique such combination – or virtual guide – is $0.5u_1 + 0.5u_2$. Hence there is a unique such recommendation, which fills agent 0’s missing ratings for d and e with $0.5u_1(d) + 0.5u_2(d) = 7$ and $0.5u_1(e) + 0.5u_2(e) = 4.5$, respectively. This yields the recommendation ranking $d > c > e > b > a$. In particular, d is recommended over e , whereas collaborative filtering systems recommend e over d . Also, the recommendation follows agents 1 and 2’s unanimous preference for c over e , whereas collaborative filtering systems violate the Pareto principle by recommending e over c .

2.2 Sparse ratings

The above example has two very peculiar features: first, all agents except 0 have rated all alternatives – their preferences are complete – and, second, agent 0 has rated enough alternatives to pin down a *unique* linear combination of the other agents’ ratings. To illustrate how collaborative filtering and virtual guide recommendations cope with sparser ratings, we now consider the same example but taken at an earlier stage when only the following ratings have been expressed:

	a	b	c	d	e
u_0	3	4	-	-	-
u_1	5	2	10	-	8
u_2	1	6	2	9	-

We again want to give a recommendation to agent 0 based on these ratings. For collaborative filtering systems, computing the average ratings and the Pearson correlation coefficients on the commonly-rated alternatives a, b leads to fill agent 0’s missing rating for alternative c with 0.375. For alternative d , the same approach, discarding agent 1 who has not rated this alternative, leads to filling agent 0’s missing rating with $\bar{u}_0 + u_2(d) - \bar{u}_2 = 8$. Finally, for alternative e , we similarly discard agent 2 and fill agent 0’s missing rating with $\bar{u}_0 - u_1(e) + \bar{u}_1 = 1.75$. This yields the recommendation ranking $d > b > a > e > c$, which violates the Pareto principle since both agents 1 and 2 have rated c above a .

For virtual guide recommendations, we instead look for positive linear combinations $\theta_1 u_1 + \theta_2 u_2 + \kappa$ that agree with u_0 on the commonly-rated alternatives a, b , where u_1 and u_2 fill agent 1 and 2’s missing ratings for d and e with any numbers δ_1 and ϵ_2 between 0 and 10, respectively. Such combinations are

characterized by the system:

$$\begin{aligned}
3 &= 5\theta_1 + \theta_2 + \kappa, & 0 &\leq 10\theta_1 + 2\theta_2 + \kappa \leq 10, & 0 &\leq \theta_1, \\
4 &= 2\theta_1 + 6\theta_2 + \kappa, & 0 &\leq \delta_1\theta_1 + 9\theta_2 + \kappa \leq 10, & 0 &\leq \theta_2, \\
& & 0 &\leq 8\theta_1 + \epsilon_2\theta_2 + \kappa \leq 10, & 0 &\leq \delta_1 \leq 10, \\
& & & & 0 &\leq \epsilon_2 \leq 10,
\end{aligned}$$

which is equivalent to:

$$\begin{aligned}
0 \leq \theta_1 \leq \frac{17}{14}, & \quad \theta_2 = 0.2 + 0.6\theta_1, & \quad 0 \leq \delta_1 \leq \min \left\{ 10, 0.2 + \frac{5.4}{\theta_1} \right\}, \\
& \quad \kappa = 2.8 - 5.6\theta_1, & \quad 0 \leq \epsilon_2 \leq \min \left\{ 10, \frac{36 - 12\theta_1}{1 + 3\theta_1} \right\}.
\end{aligned}$$

Any such vector $(\theta_1, \theta_2, \kappa, \delta_1, \epsilon_2)$ defines a virtual guide filling agent 0's missing ratings for c , d , and e with $3.2 + 5.6\theta_1$, $4.6 + (\delta_1 - 0.2)\theta_1$, and $2.8 + 0.2\epsilon_2 + (0.6\epsilon_2 + 2.4)\theta_1$, respectively. Any set of such virtual guides defines a virtual guide recommendation. All these recommendations have in common that $c > a$ and $d > b > a$. At one extreme, the most incomplete of these recommendations is based on unanimity among all virtual guides and makes no further ranking – note though that it does rank some alternatives that agent 0 has not ranked. At the other extreme, complete recommendations are based on a single virtual guide and fully rank the five alternatives – i.e. also rank c with respect to b and d as well as e with respect to the other four alternatives). For instance, selecting the single virtual guide $0.5u_1 + 0.5u_2$ with $\delta_1 = 5$ and $\epsilon_2 = 1$ yields the same recommendation ranking $d > c > e > b > a$ as in the rich ratings example, whereas selecting the single virtual guide $0.1u_1 + 0.26u_2 + 2.24$ with $\delta_1 = 5$ and $\epsilon_2 = 8$ yields the recommendation ranking $e > d > b > c > a$.

3 Model

3.1 Expected multi-utility

We let A be a finite set of alternatives and consider the set X of lotteries – probability distributions – over A . A *preference relation* on X is a binary relation $\succsim$ on X , where $x \succsim y$ is interpreted as lottery x being weakly preferred to lottery y . A preference relation $\succsim'$ on X *extends* a preference relation $\succsim$ on X if for all $x, y \in X$, $x \succsim y$ implies $x \succsim' y$. A preference relation $\succsim$ on X is *complete* if for all $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both), and *incomplete* otherwise.

Definition 1. A preference relation $\succsim$ on X is an *expected multi-utility* preference relation if there exists a non-empty, compact, convex set U of utility functions $u : A \rightarrow \mathbb{R}$ such that for all $x, y \in X$:

$$x \succsim y \text{ if and only if } \mathbb{E}u(x) \geq \mathbb{E}u(y) \text{ for all } u \in U.$$

Expected multi-utility preferences are generally incomplete and are axiomatized by [Baucells and Shapley \(2008\)](#) and [Dubra et al. \(2004\)](#). They generalize standard expected utility preferences by allowing the set U to contain more than one utility function or, equivalently, by allowing $\succsim$ to be incomplete. A lottery x is weakly preferred to a lottery y if and only if all functions in the set U assign a weakly higher

expected utility to x . If x has a strictly higher expected utility for one of these functions whereas y has a strictly higher expected utility for another one, the two lotteries are not ranked. The interpretation is that the agent is unsure about her utility function and considers all functions in U as possible. The utility functions in U are cardinal in the sense of being unique up to positive affine transformations.¹

For technical convenience, we restrict attention to expected multi-utility preferences that can be represented as above with U being the convex hull – set of convex combinations – of finitely many utility functions.² Such preferences are axiomatized in Dubra and Ok (2002). They naturally arise, in particular, if $\succsim$ is derived from some primitive cardinal ratings of some alternatives in A as per Definition 1 with U consisting of all possible ways to fill in the missing ratings.³ For instance, in the rich ratings example, denoting by $u_0^{\delta,\epsilon}$ the utility function filling agent 0's missing ratings for d and e with δ and ϵ , respectively, we take $U_0 = \{u_0^{\delta,\epsilon} : \delta, \epsilon \in [0, 10]\}$, $U_1 = \{u_1\}$, $U_2 = \{u_2\}$ and we note that U_0 is then the convex hull of the four utility functions $u_0^{0,0}, u_0^{0,10}, u_0^{10,0}, u_0^{10,10}$.

We say as usual that two lotteries x and y are *indifferent*, denoted, $x \sim y$, if both $x \succsim y$ and $y \succsim x$, and that x is *strictly preferred* to y , denoted $x \succ y$, if $x \succsim y$ but not $y \succsim x$. We further say that x is *strongly preferred* to y , denoted $x \gg y$, if $x' \succ y'$ for every lotteries x' and y' “sufficiently close” to x and y , respectively. Formally, with the usual notation for mixed lotteries, $x \gg y$ if for all $x', y' \in X$, there exists $\lambda \in (0, 1)$ such that $\lambda x + (1 - \lambda)x' \succ \lambda y + (1 - \lambda)y'$. We have that $x \sim y$ if and only if $Eu(x) = Eu(y)$ for all $u \in U$, and $x \succ y$ if and only if $Eu(x) \geq Eu(y)$ for all $u \in U$ with strict inequality for at least one $u \in U$. We further have, without loss of generality, that $x \gg y$ if and only if $Eu(x) > Eu(y)$ for all $u \in U$.⁴ Strong preference is thus equivalent to strict preference for complete preferences but more demanding for incomplete preferences. For instance, in the rich ratings example, we have $c \gg_0 b \gg_0 a$ and $b \sim_0 (a, \frac{2}{3}; c, \frac{1}{3})$ – the lottery yielding a and c with probabilities $\frac{2}{3}$ and $\frac{1}{3}$, respectively – as well as $b \succ_0 (a, \frac{6}{7}; e, \frac{1}{7})$ but not $b \gg_0 (a, \frac{6}{7}; e, \frac{1}{7})$.

3.2 The recommendation problem

We consider a finite set $\{0, 1, \dots, N\}$ of at least two agents. Each agent $n = 0, 1, \dots, N$ is endowed with an expected multi-utility preference relation $\succsim_n$ on X . We want to give a recommendation to agent 0, in the form of another expected multi-utility preference relation $\succsim$ on X satisfying the following two simple properties:

Axiom (Acceptance principle). For all $x, y \in X$, if $x \succsim_0 y$ then $x \succsim y$ and if $x \gg_0 y$ then $x \gg y$.

Axiom (Pareto principle). For all $x, y \in X$, if $x \succsim_n y$ for all $n = 1, \dots, N$ then $x \succsim y$.

Acceptance means that the recommendation must extend agent 0's preferences, more specifically preserve her weak and strong preferences. In the rich ratings example, we must therefore have $c \gg b \gg a$ and $b \sim (a, \frac{2}{3}; c, \frac{1}{3})$ as well as $b \succ (a, \frac{6}{7}; e, \frac{1}{7})$ but not necessarily $b \gg (a, \frac{6}{7}; e, \frac{1}{7})$. The Pareto principle, on the other hand, expresses the requirement that the recommendation must be based on the preferences

¹More precisely, the closure of the set $\{\theta u + \kappa : u \in U, \theta \in \mathbb{R}_+, \kappa \in \mathbb{R}\}$ is unique.

²This allows us to dispense with the closure operator in the uniqueness statement above as well as the closedness assumption of Danan et al. (2015).

³While in the examples above we consider possible any rating within the allowed scale, our framework can accommodate alternative specifications (for instance, one might consider intermediate ratings only).

⁴It is without loss of generality in the sense that we can always choose a representation U such that this holds.

of agents $1, \dots, N$. If those agents have a unanimous weak preferences then the recommendation must follow it. In the rich ratings example, we must therefore have $c \succsim e \succsim a$, $d \succsim a$, and $d \succsim b$.

We will see below that the Acceptance and Pareto principles together imply that the recommendation must also follow unanimous strong preferences from agents $1, \dots, N$ in most cases – and in the two examples above in particular, so that we must have $c \gg e$, $c \gg a$, and $d \gg b$ in the first example. We will also explain there why we generally do not expect the recommendation to preserve agent 0's strict preferences or follow unanimous strict preferences from agents $1, \dots, N$.

Definition 2. The preference relations $\succsim_0, \succsim_1, \dots, \succsim_N$ are *coherent* if there are no lotteries x and y such that $x \gg_0 y$ whereas $y \succsim_n x$ for all $n = 1, \dots, N$.

Coherence is obviously necessary for the Acceptance and Pareto principles to be mutually compatible.⁵ We will show below that it is sufficient as well.⁶ We will also consider below a weakening of the Pareto principle that is more often compatible with the Acceptance principle.

3.3 Virtual guide recommendations

We now provide a characterization of the recommendations satisfying the Acceptance and Pareto principles within the expected multi-utility model. Let $U_0, U_1, \dots, U_N$ be representations of $\succsim_0, \succsim_1, \dots, \succsim_N$ as per Definition 1, respectively.

Definition 3. A utility function $u : A \rightarrow \mathbb{R}$ is a *virtual guide* (for agent 0) if $u \in U_0$ and $u = \sum_{n=1}^N \theta_n u_n + \kappa$ for some $u_1 \in U_1, \dots, u_N \in U_N, \theta_1, \dots, \theta_N \in \mathbb{R}_+, \kappa \in \mathbb{R}$.

A virtual guide is thus a possible utility function for agent 0 that is also a positive linear combination of possible utility functions for agents $1, \dots, N$. It corresponds to a complete preference relation agreeing with that of agent 0 on the alternatives she has ranked while aggregating complete preference relations agreeing with those of agents $1, \dots, N$ on the alternatives they have ranked, respectively.

We let V denote the set of all virtual guides. Because of the cardinal uniqueness of expected multi-utility representations, V is essentially determined by the preference relations $\succsim_0, \succsim_1, \dots, \succsim_N$, independently of their particular representations $U_0, U_1, \dots, U_N$. More precisely, choosing different representations $U_1, \dots, U_N$ leaves V unchanged and only affects the coefficients $\theta_1, \dots, \theta_N$ and κ corresponding to a given virtual guide. Choosing a different representation U_0 , on the other hand, causes the virtual guides in V to be rescaled accordingly but leaves the corresponding expected utility preferences unchanged.

Proposition 1. A recommendation $\succsim$ satisfies the Acceptance and Pareto principles if and only if it can be represented as per Definition 1 by some set $U \subseteq V$ of virtual guides.

The recommendations satisfying the Acceptance and Pareto principles are thus exactly those that are based on virtual guides. Any recommendation that is not of this form violates at least one of the two principles – in particular, as we illustrated above, collaborating filtering systems violate the Pareto principle. In the rich ratings example, there is a unique virtual guide and, hence, the corresponding (complete) preference relation is the only possible recommendation. In the sparse ratings example, there

⁵Indeed, if there are lotteries x and y such that $x \gg_0 y$ whereas $y \succsim_n x$ for all $n = 1, \dots, N$, then the Acceptance principle implies $x \gg y$ whereas the Pareto principle implies $y \succsim x$, a contradiction.

⁶This is not trivial because of our restriction to expected multi-utility preferences.

are multiple virtual guides and, hence, any subset of V corresponds to a possible recommendation. It is possible, in particular, to select a single virtual guide, yielding a complete recommendation as in the first example, but it is also possible to select multiple virtual guides, yielding an incomplete recommendation. Selecting the whole set V yields the most incomplete such recommendation.

Our recommendation problem generalizes the classical preference aggregation problem, which corresponds to the particular case where $\succsim_0$ is void – or agent 0 has not rated more than one alternative. In this case the Acceptance principle is trivially satisfied and Proposition 1 boils down to the following statement: $\succsim$ satisfies the Pareto principle with respect to $\succsim_1, \dots, \succsim_N$ if and only if each utility function in U is a positive linear combination of some utility functions in $U_1, \dots, U_N$. This generalization of Harsanyi (1955)’s aggregation theorem was proved in Danan et al. (2015). Assuming further that $\succsim$ as well as $\succsim_1, \dots, \succsim_N$ are complete, we recover Harsanyi’s theorem itself. When $\succsim_0$ is not void, the Acceptance principle requires each utility function in U to be a positive affine transformation of some utility function in U_0 . Note that this generally places restrictions on the possible weights entering the linear combinations – sometimes pinning down unique weights as in the rich ratings example – whereas these weights are arbitrary in the pure aggregation setting.

When $\succsim_0$ contains at least one strong preference – or agent 0 has given different ratings to at least two alternatives – a corollary to Proposition 1 is that the recommendations satisfying the Acceptance and Pareto principles also follow unanimous strong preferences from agents $1, \dots, N$: if $x \succ_n y$ for all $n = 1, \dots, N$ then $x \succ y$.⁷ This principle generalizes to incomplete preferences the “Weak Pareto” or “P₂” principle analyzed by Weymark (1993, 1995) and de Meyer and Mongin (1995) in the aggregation context. On the other hand, recommendations satisfying the Acceptance and Pareto principles generally neither preserve agent 0’s strict preferences nor follow unanimous strict preferences from agents $1, \dots, N$. For instance, in the sparse ratings example, we have $b \succ_0 (a, \frac{6}{7}; e, \frac{1}{7})$ but selecting a single virtual guide filling agent 0’s missing rating for alternative e with 10 leads to $b \sim (a, \frac{6}{7}; e, \frac{1}{7})$. We also have $(c, \frac{1}{2}; e, \frac{1}{2}) \succ_1 a$ and $(c, \frac{1}{2}; e, \frac{1}{2}) \succ_2 a$ but selecting a single virtual guide filling agent 0’s missing ratings for c and e with 3.2 and 2.8, respectively, leads to $(c, \frac{1}{2}; e, \frac{1}{2}) \sim a$.⁸

3.4 Signed virtual guides

A limitation of the concept of virtual guide is that the weights in the linear combination cannot be negative. Collaborative filtering systems, on the other hand, allow for negative weights, so that the recommendation can possibly draw information from agents whose ratings are negatively correlated with the ratings of agent 0. In order to characterize virtual guide recommendations with possibly negative weights, we now weaken the Pareto principle as follows:

Axiom (Pareto-Indifference principle). For all $x, y \in X$, if $x \sim_n y$ for all $i = 1, \dots, N$ then $x \sim y$.

The Pareto-Indifference principle merely requires the recommendation to follow unanimous indifference from agent $1, \dots, N$. It does not require the recommendation to follow unanimous weak preferences from agents $1, \dots, N$ when some of these weak preferences are strict or strong, opening the possibility to weigh some of these agents negatively.

⁷Indeed, in that case, U_0 can not contain constant utility functions and, hence, any virtual guide $u = \sum_{n=1}^N \theta_n u_n + \kappa \in U_0$ must be such that $\theta_n > 0$ for at least one $n = 1, \dots, N$.

⁸In this example, the recommendation could possibly preserve agent 0’s strict preferences and/or follow unanimous strict preferences from agents $1, \dots, N$, depending on which virtual guides are selected. Other examples can be constructed where that is not possible.

Definition 4. The preference relations $\succsim_0, \succsim_1, \dots, \succsim_N$ are *weakly coherent* if there are no lotteries x and y such that $x \succ_0 y$ whereas $x \sim_n y$ for all $n = 1, \dots, N$.

Weak coherence is substantially weaker than coherence and we expect it to hold in most practical cases. It is obviously necessary for the Acceptance and Pareto-Indifference principles to be mutually compatible and we will show below that it is also sufficient.

Definition 5. A utility function $u : A \rightarrow \mathbb{R}$ is a *signed virtual guide* (for agent 0) if $u \in U_0$ and $u = \sum_{n=1}^N (\theta_n u_n - \omega_n v_n) + \kappa$ for some $u_1, v_1 \in U_1, \dots, u_N, v_N \in U_N, \theta_1, \omega_1, \dots, \theta_N, \omega_N \in \mathbb{R}_+, \kappa \in \mathbb{R}$.

A signed virtual guide is thus a possible utility function for agent 0 that is also the difference of two positive linear combinations of possible utility functions for agents $1, \dots, N$. In the sparse ratings example, the utility function $u = 3u_2^0 - 2.8u_2^1 + 2.8$, where u_2^0 and u_2^1 fill agent 2's missing rating for alternative e with 0 and 1, respectively, is a signed virtual guide but not a virtual guide.⁹ We let W denote the set of all signed virtual guides. Like the set V of virtual guides, W is essentially determined by the preference relations $\succsim_0, \succsim_1, \dots, \succsim_N$, independently of their particular representations $U_0, U_1, \dots, U_N$.

Proposition 2. A recommendation $\succsim$ satisfies the Acceptance and Pareto-Indifference principles if and only if it can be represented as per Definition 1 by some set $U \subseteq W$ of signed virtual guides.

3.5 Computing (signed) virtual guides

We now show that the set of (signed) virtual guides can be computed simply by solving a finite linear system. This allows us to deduce some properties of this set and, in particular, to identify the cases where it is empty – or, equivalently, where (signed) virtual guide recommendations do not exist. To this end, we enumerate as $\{a_1, \dots, a_K\}$ the finite set A of alternatives, where $K \geq 1$ denotes the number of alternatives. For all $n = 0, 1, \dots, N$, we also enumerate as $\{u_{n,1}, \dots, u_{n,I(n)}\}$ the finite set of utility functions of which U_n is the convex hull, where $I(n) \geq 1$ denotes the number of such functions. A utility function $u : A \rightarrow \mathbb{R}$ then belongs to U_n if and only if $u = \sum_{i=1}^{I(n)} \lambda_{n,i} u_{n,i}$ for some $\lambda_{n,1}, \dots, \lambda_{n,I(n)} \in \mathbb{R}_+$ such that $\sum_{i=1}^{I(n)} \lambda_{n,i} = 1$. Virtual guides then correspond to solutions to the finite system:

$$\begin{aligned} \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) &= \sum_{n=1}^N \sum_{i=1}^{I(n)} \theta_n \lambda_{n,i} u_{n,i}(a_k) + \kappa, & k &= 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(n)} \lambda_{n,i}, & n &= 0, 1, \dots, N, \\ 0 &\leq \lambda_{n,i}, & n &= 0, 1, \dots, N, i = 1, \dots, I(n), \\ 0 &\leq \theta_n, & n &= 0, 1, \dots, N. \end{aligned}$$

⁹Indeed, u fills agent 0's missing ratings for alternatives c, d , and e with 3.2, 4.6, and 0, respectively, whereas all virtual guides fill agent 0's missing rating for e with at least 2.8. Note that u cannot be obtained as a linear combination $\theta_1 u_1 + \theta_2 u_2 + \kappa$ with $u_1 \in U_1, u_2 \in U_2$, and $\theta_1, \theta_2, \kappa \in \mathbb{R}$.

Defining $\phi_{n,i} = \theta_n \lambda_{n,i}$ for all $n = 1, \dots, N$ and $i = 1, \dots, I(n)$, it is sufficient to solve the finite linear system:

$$\begin{aligned} \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) &= \sum_{n=1}^N \sum_{i=1}^{I(n)} \phi_{n,i} u_{n,i}(a_k) + \kappa, & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}, \\ 0 &\leq \lambda_{0,i}, & i = 1, \dots, I(0), \\ 0 &\leq \phi_{n,i}, & n = 1, \dots, N, i = 1, \dots, I(n). \end{aligned} \quad (1)$$

Proposition 3. V is non-empty if and only if $\succsim_0, \succsim_1, \dots, \succsim_N$ are coherent. Moreover, V is then the convex hull of finitely many utility functions.

Similarly, signed virtual guides correspond to solutions to the finite system:

$$\begin{aligned} \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) &= \sum_{n=1}^N \sum_{i=1}^{I(n)} (\theta_n \lambda_{n,i} - \omega_n \mu_{n,i}) u_{n,i}(a_k) + \kappa, & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}, & n = 0, 1, \dots, N, \\ 1 &= \sum_{i=1}^{I(n)} \mu_{n,i}, & n = 1, \dots, N, \\ 0 &\leq \lambda_{n,i}, & n = 0, 1, \dots, N, i = 1, \dots, I(n), \\ 0 &\leq \mu_{n,i}, & n = 1, \dots, N, i = 1, \dots, I(n), \\ 0 &\leq \theta_n, & n = 0, 1, \dots, N, \\ 0 &\leq \omega_n, & n = 0, 1, \dots, N. \end{aligned}$$

Defining $\psi_{n,i} = \theta_n \lambda_{n,i} - \omega_n \mu_{n,i}$ for all $n = 1, \dots, N$ and $i = 1, \dots, I(n)$, it is sufficient to solve the finite linear system:

$$\begin{aligned} \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) &= \sum_{n=1}^N \sum_{i=1}^{I(n)} \psi_{n,i} u_{n,i}(a_k) + \kappa, & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}, \\ 0 &\leq \lambda_{0,i}, & i = 1, \dots, I(0). \end{aligned} \quad (2)$$

Proposition 4. W is non-empty if and only if $\succsim_0, \succsim_1, \dots, \succsim_N$ are weakly coherent. Moreover, W is then the convex hull of finitely many utility functions.

3.6 Selecting among (signed) virtual guides

As the sparse ratings example illustrates, which (signed) virtual guides are selected can critically impact the recommendation. Axiomatically characterizing particular subclasses of virtual guide recommendations – or even a single recommendation – would therefore provide useful guidance in this regard. To this end, it seems natural to extend the present framework by considering a *recommendation rule* f associating a recommendation $\succsim = f(\succsim_0, \succsim_1, \dots, \succsim_N)$ to each preference profile $(\succsim_0, \succsim_1, \dots, \succsim_N)$ in some domain D . We shall more specifically consider the domains D_c and D_{wc} of coherent and weakly coherent profiles, respectively, and say that f satisfies the Acceptance (resp. Pareto, Pareto-Indifference) principle if $f(\succsim_0, \succsim_1, \dots, \succsim_N)$ satisfies this principle for all $(\succsim_0, \succsim_1, \dots, \succsim_N) \in D$. This extended framework makes it possible to express additional principles, such as the following ones.

Axiom (Anonymity principle). For all $(z_0, z_1, \dots, z_N), (z_0, z'_1, \dots, z'_N) \in D$, if $(z'_1, \dots, z'_N)$ is a permutation of $(z_1, \dots, z_N)$ then $f(z_0, z'_1, \dots, z'_N) = f(z_0, z_1, \dots, z_N)$.

Axiom (Internal-Robustness principle). For all $(z_0, z_1, \dots, z_N), (z'_0, z_1, \dots, z_N) \in D$, if z'_0 extends z_0 then $f(z'_0, z_1, \dots, z_N)$ extends $f(z_0, z_1, \dots, z_N)$.

Axiom (External-Robustness principle). For all $(z_0, z_1, \dots, z_N), (z_0, z'_1, \dots, z'_N) \in D$, if z'_n extends z_n for all $n = 1, \dots, N$ then $f(z_0, z'_1, \dots, z'_N)$ extends $f(z_0, z_1, \dots, z_N)$.

Anonymity means that no agent $n = 1, \dots, N$ has more “intrinsic” influence on the recommendation than another. The Internal-Robustness and External-Robustness principles express the idea that as agent 0 or agents $1, \dots, N$ rate more alternatives, and provided the preference profile remains coherent or weakly coherent, the recommendation becomes more complete without reversing previous rankings. For instance, z_n from the rich ratings example extends z_n from the sparse ratings example for $n = 0, 1, 2$. If all virtual guides are kept in the sparse ratings example, then the rich ratings recommendation ranking extends the sparse ratings recommendation ranking, in accordance with the two robustness principles. If, on the other hand, the single virtual guide $0.1u_1^5 + 0.26u_2^8 + 2.24$ is selected in the sparse ratings example, then the sparse ratings recommendation ranking is partly reversed by the rich ratings recommendation ranking, violating at least one of these two principles.¹⁰

Proposition 5. There exists a unique recommendation rule on D_c satisfying the Acceptance, Pareto, and Internal-Robustness principles, which consists in selecting all virtual guides for all profiles. This rule also satisfies the Anonymity and External-Robustness principles.

Proposition 6. There exists a unique recommendation rule on D_{wc} satisfying the Acceptance, Pareto-Indifference, and Internal-Robustness principles, which consists in selecting all signed virtual guides for all profiles. This rule also satisfies the Anonymity and External-Robustness principles.

The most incomplete (signed) virtual guide recommendation rule is thus the only one that is compatible with the Internal-Robustness principle. In view of this result, we may call this rule the *robust (signed) virtual guide recommendation rule*. Note that the statements in Propositions 6 and 7 are well defined because, as explained above, the set of (signed) virtual guides is essentially determined by the preference profile independently of the chosen representations. One may therefore choose arbitrary representations in order to compute the recommendation ranking. A feature of this rule that might be of practical interest is that the recommendation ranking can be computed “incrementally” by dropping (signed) virtual guides as new ratings arrive.

Although the Internal-Robustness principle is clearly a very demanding requirement, the above results suggest that a certain degree of incompleteness in the recommendation ranking might help obtaining more “stable” recommendations. We note in this regard that the External-Robustness principle is weaker than the Internal-Robustness principle in our framework.¹¹ It is nevertheless incompatible with complete recommendation rules.

¹⁰Collaborative filtering systems, on the other hand, satisfy the anonymity principle but violate the two robustness principles.

¹¹For instance, any recommendation rule that selects a single (signed) virtual guide whenever $z_1, \dots, z_N$ are complete but keeps all (signed) virtual guides otherwise satisfies the Acceptance, Pareto (resp. Pareto-Indifference) and External-Robustness principles but not the Internal-Robustness principle.

Proposition 7. If A contains at least three alternatives then there exists no recommendation rule on D_c (resp. D_{wc}) satisfying the Acceptance, Pareto (resp. Pareto-Indifference), and External-Robustness principles that yields a complete recommendation for all profiles.

4 Conclusion

We view the present paper’s contribution as mainly conceptual: the illustrative examples and simple model we presented demonstrate how an axiomatic analysis of recommendation systems in terms of preference aggregation and extension can provide important insights into their normative properties. We showed in particular that collaborative filtering systems – used by popular internet platforms – typically violate the Pareto principle and characterized a new class of recommendation rules – based on virtual guides – satisfying this principle. Finally, we extended our formal framework to identify a particular rule within this class – the robust virtual guide recommendation rule. This rule has interesting stability properties but is arguably quite extreme. A natural avenue for future research is therefore to characterize more complete rules in this extended setting. It also seems interesting to go beyond this setting and do away with lotteries or allow for strategic behavior.

On the practical side, one may want to go beyond illustrative examples to assess how frequently do collaborative filtering systems violate the Pareto principle and how different are collaborative filtering and virtual guide recommendations from each other. Doing so would require an empirically relevant dataset, perhaps stochastically generated to avoid the endogeneity that some real datasets are based on the use of a particular recommendation system. We may wonder, in particular, whether virtual recommendation rules help alleviate some of the difficulties typically associated with collaborative filtering systems.

A first such difficulty is that collaborative filtering systems generally struggle to produce recommendations for “gray sheep” users, i.e. agents whose ratings are not highly correlated – positively or negatively – with those of any other agent. One may expect this issue to be less prevalent for virtual guide recommendations, because they rely on virtual agents whose ratings exactly match those of agent 0 rather than on real agents whose ratings are “close” to those of agent 0. These virtual agents are aggregates of real agents, but these real agents can as well be “close” or “remote” to agent 0.

A second difficulty is that collaborative filtering systems tend to forgo “diversity” by focusing on alternatives that have been rated by many agents, at the expense of less rated – perhaps newer – ones. Virtual guide recommendations seem to suffer less from this bias: even if an alternative has received no or few ratings, there will generally be virtual guides assigning a range of ratings to it. Opting for an incomplete recommendation rather than a complete one may in fact be useful in this regard, because more such alternatives will be at the top of the recommendation ranking.

A third difficulty is that collaborative filtering systems can be manipulated by “shilling attacks”, i.e. by injecting a large number of fake ratings in order to promote or demote some alternatives. Although the present model leaves strategic considerations aside and considers a fixed set of agents, we may note that virtual recommendations are not immune to the injection of a fake agent either, as this may enlarge the set of virtual guides and, hence, affect the recommendation. The set of virtual guides is, however, immune to injecting additional “clones” of this fake agent, as it does not depend on how many agents with a given preference relation there are. Furthermore, injecting additional agents that are “close” to the first one generally does not affect this set substantially. Whether and how the recommendation is affected

by the injection of additional clones then depends on which virtual guides are selected.

Appendix: Proofs

We derive Propositions 1 and 2 from slightly more general results. To this end we consider the following weakening of the Acceptance principle.

Axiom (Weak Acceptance principle). For all $x, y \in X$, if $x \succeq_0 y$ then $x \succeq y$.

Lemma 1. Let U be a representation of $\succeq$ as per Definition 1.

(a) $\succeq$ satisfies the Weak Acceptance and Pareto principles if and only if, for all $u \in U$, there exist $u_0 \in U_0, u_1 \in U_1, \dots, u_N \in U_N, \theta_0, \theta_1, \dots, \theta_N \in \mathbb{R}_+, \kappa_0, \kappa \in \mathbb{R}$ such that:

$$u = \theta_0 u_0 + \kappa_0 = \sum_{n=1}^N \theta_n u_n + \kappa. \quad (3)$$

(b) $\succeq$ satisfies the Weak Acceptance and Pareto-Indifference principles if and only if, for all $u \in U$, there exist $u_0 \in U_0, u_1, v_1 \in U_1, \dots, u_N, v_N \in U_N, \theta_0, \theta_1, \omega_1, \dots, \theta_N, \omega_N \in \mathbb{R}_+, \kappa_0, \kappa \in \mathbb{R}$ such that:

$$u = \theta_0 u_0 + \kappa_0 = \sum_{n=1}^N (\theta_n u_n - \omega_n v_n) + \kappa. \quad (4)$$

Proof. (a) It is obvious that if (3) holds then $\succeq$ satisfies the Weak Acceptance and Pareto principles. Conversely, assume (3) does not hold. Then at least one $u \in U$ must lie outside at least one of the two sets

$$S = \{\theta_0 u_0 + \kappa_0 : u_0 \in U_0, \theta_0 \in \mathbb{R}_+, \kappa_0 \in \mathbb{R}\} \text{ and } \\ T = \left\{ \sum_{n=1}^N \theta_n u_n + \kappa : u_1 \in U_1, \dots, u_N \in U_N, \theta_1, \dots, \theta_N \in \mathbb{R}_+, \kappa \in \mathbb{R} \right\}.$$

Note that both S and T are closed, convex cones in $\mathbb{R}^A$.

Case 1. If $u \notin S$ then, by the separating hyperplane theorem, there exists $h \in \mathbb{R}^A, h \neq 0$, such that $\sum_{a \in A} h(a)u(a) > 0 \geq \sum_{a \in A} h(a)s(a)$ for all $s \in S$. Since $s + \beta \in S$ for all $s \in S$ and $\beta \in \mathbb{R}$, this can only be the case if $\sum_{a \in A} h(a) = 0$, for otherwise we could make $\sum_{a \in A} h(a)(s(a) + \beta) = \sum_{a \in A} h(a)s(a) + \beta \sum_{a \in A} h(a)$ arbitrarily large by taking β sufficiently close to $+\infty$ or $-\infty$. Hence, letting $\xi = \sum_{a \in A} |h(a)|$ and defining $x, y \in \mathbb{R}^A$ by $x(a) = \frac{\max\{h(a), 0\}}{\xi}$ and $y(a) = \frac{\max\{-h(a), 0\}}{\xi}$, we have $\xi > 0, x, y \in X$, and $h = \xi(x - y)$. Since $\sum_{a \in A} h(a)u(a) > 0$, it follows that $E u(x) > E u(y)$, so it cannot be the case that $y \succeq x$ by Definition 1. But since $U_0 \subset S$, it also follows that $0 \geq \sum_{a \in A} h(a)u_0(a)$ and, hence, $E u_0(y) \geq E u_0(x)$ for all $u_0 \in U_0$, so that $y \succeq_0 x$ by Definition 1. This contradicts the Weak Acceptance principle.

Case 2. If $u \notin T$ then, by the separating hyperplane theorem, there exists $h \in \mathbb{R}^A, h \neq 0$, such that $\sum_{a \in A} h(a)u(a) > 0 \geq \sum_{a \in A} h(a)t(a)$ for all $t \in T$. By the same argument as in Case 1, we must have $\sum_{a \in A} h(a) = 0$ and, hence, we can write $h = \xi(x - y)$ with $\xi > 0$ and $x, y \in X$. Since $\sum_{a \in A} h(a)u(a) > 0$, it follows that $E u(x) > E u(y)$, so it cannot be the case that $y \succeq x$ by Definition 1. But for all $n = 1, \dots, N$, since $U_n \subset T$, it also follows that $0 \geq \sum_{a \in A} h(a)u_n(a)$ and, hence, $E u_n(y) \geq E u_n(x)$ for all $u_n \in U_n$, so that $y \succeq_n x$ by Definition 1. This contradicts the Pareto principle.

(b) It is obvious that if (4) holds then $\succsim$ satisfies the Weak Acceptance and Pareto-Indifference principles. Conversely, assume (4) does not hold. Then at least one $u \in U$ must lie outside at least one of the two sets S and:

$$T' = \left\{ \sum_{n=1}^N (\theta_n u_n - \omega_n v_n) + \kappa : u_1, v_1 \in U_1, \dots, u_N, v_N \in U_N, \theta_1, \omega_1, \dots, \theta_N, \omega_N \in \mathbb{R}_+, \kappa \in \mathbb{R} \right\}.$$

We then proceed as in the proof of Part (a), except that in Case 2, since T' is a linear subspace of $\mathbb{R}^A$, the separating hyperplane theorem now implies that $\sum_{a \in A} h(a)u(a) > 0 = \sum_{a \in A} h(a)t(a)$ for all $t \in T'$. Hence for all $n = 1, \dots, N$, we now have $E u_n(x) = E u_n(y)$ for all $u_n \in U_n$, so that $x \sim_n y$ by Definition 1, which contradicts the Pareto-Indifference principle. $\square$

Proof of Proposition 1. It is obvious that if $\succsim$ can be represented as per Definition 1 by some set $U \subseteq V$ of virtual guides, then it satisfies the Acceptance and Pareto principles. Conversely, assume $\succsim$ satisfies these two principles and fix some representation U of $\succsim$ as per Definition 1. Then for all $u \in U$, (3) holds by Lemma 1(a). It is sufficient to show that (3) actually holds with $\theta_0 > 0$ for some $u_0 \in U_0$, since then the closed convex hull of the set $\left\{ \frac{u - \kappa_0}{\theta_0} : u \in U \right\}$ also represents $\succsim$ as per Definition 1 and is a subset of V . So suppose (3) does not hold with $\theta_0 > 0$ for any $u_0 \in U_0$. Note that this implies that u is a constant function whereas U_0 contains no constant function. Hence, by the separating hyperplane theorem, there exists $h \in \mathbb{R}^A$, $h \neq 0$, such that $\sum_{a \in A} h(a)u_0(a) > \kappa \sum_{a \in A} h(a)$ for all $u_0 \in U_0$ and $\kappa \in \mathbb{R}$. By the same argument as in Case 1 of the proof of Lemma 1(a), we must have $\sum_{a \in A} h(a) = 0$ and, hence, we can write $h = \xi(x - y)$ with $\xi > 0$ and $x, y \in X$. For all $u_0 \in U_0$, since $\sum_{a \in A} h(a)u_0(a) > 0$, it follows that $E u_0(x) > E u_0(y)$ and, hence, $x \gg_0 y$. The Acceptance principle then implies that $x \gg y$, contradicting the fact that U contains a constant function. $\square$

Proof of Proposition 2. Similar to the proof of Proposition 1. $\square$

Proof of Proposition 3. Letting $\kappa = \kappa^+ - \kappa^-$ with $\kappa^+, \kappa^- \in \mathbb{R}_+$ in (1), V is non-empty if and only if the following finite linear system has a non-negative solution:

$$\begin{aligned} 0 &= \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) - \sum_{n=1}^N \sum_{i=1}^{I(n)} \phi_{n,i} u_{n,i}(a_k) - (\kappa^+ - \kappa^-), & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}. \end{aligned}$$

By Farkas' Lemma, this system has a non-negative solution if and only if there exist no $h \in \mathbb{R}^K$ and $\zeta \in \mathbb{R}$ such that:

$$\begin{aligned} 0 &\leq \sum_{k=1}^K h_k u_{0,i}(a_k) + \zeta, & i = 1, \dots, I(0), \\ 0 &\leq - \sum_{k=1}^K h_k u_{n,i}(a_k), & n = 1, \dots, N, i = 1, \dots, I(n), \\ 0 &\leq \sum_{k=1}^K h_k, \\ 0 &\leq - \sum_{k=1}^K h_k, \\ 0 &> \zeta, \end{aligned}$$

or, equivalently, if there exists no $h \in \mathbb{R}^K$ such that:

$$\begin{aligned} 0 &< \sum_{k=1}^K h_k u_{0,i}(a_k), & i = 1, \dots, I(0), \\ 0 &\geq \sum_{k=1}^K h_k u_{0,i}(a_k), & n = 1, \dots, N, i = 1, \dots, I(n), \\ 0 &= \sum_{k=1}^K h_k. \end{aligned}$$

By the same argument as in Case 1 of the proof of Lemma 1(a), this is equivalent to the non-existence of lotteries $x, y \in X$ such that $x \gg_0 y$ and $y \succeq_n x$ for all $n = 1, \dots, N$. The fact that V is then the convex hull of finitely many functions simply follows from the finiteness and linearity of the above system as well as the fact that V is a subset of the compact set U_0 . $\square$

Proof of Proposition 4. Letting $\psi_{n,i} = \psi_{n,i}^+ - \psi_{n,i}^-$ and $\kappa = \kappa^+ - \kappa^-$ where $\psi_{n,i}^+, \psi_{n,i}^-, \kappa^+, \kappa^- \in \mathbb{R}_+$ in (2), W is non-empty if and only if the following finite linear system has a non-negative solution:

$$\begin{aligned} 0 &= \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) - \sum_{n=1}^N \sum_{i=1}^{I(n)} (\psi_{n,i}^+ - \psi_{n,i}^-) u_{n,i}(a_k) - (\kappa^+ - \kappa^-), & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}. \end{aligned}$$

By Farkas' Lemma, this system has a non-negative solution if and only if there exist no $h \in \mathbb{R}^K$ and $\zeta \in \mathbb{R}$ such that:

$$\begin{aligned} 0 &\leq \sum_{k=1}^K h_k u_{0,i}(a_k) + \zeta, & i = 1, \dots, I(0), \\ 0 &\leq - \sum_{k=1}^K h_k u_{n,i}(a_k), & n = 1, \dots, N, i = 1, \dots, I(n), \\ 0 &\leq \sum_{k=1}^K h_k u_{n,i}(a_k), & n = 1, \dots, N, i = 1, \dots, I(n), \\ 0 &\leq \sum_{k=1}^K h_k, \\ 0 &\leq - \sum_{k=1}^K h_k, \\ 0 &> \zeta, \end{aligned}$$

or, equivalently, if there exists no $h \in \mathbb{R}^K$ such that:

$$\begin{aligned} 0 &< \sum_{k=1}^K h_k u_{0,i}(a_k), & i = 1, \dots, I(0), \\ 0 &= \sum_{k=1}^K h_k u_{0,i}(a_k), & n = 1, \dots, N, i = 1, \dots, I(n), \\ 0 &= \sum_{k=1}^K h_k. \end{aligned}$$

By the same argument as in Case 1 of the proof of Lemma 1(a), this is equivalent to the non-existence of lotteries $x, y \in X$ such that $x \gg_0 y$ and $x \sim_n y$ for all $n = 1, \dots, N$. The fact that W is then the convex hull of finitely many functions simply follows from the finiteness and linearity of the above system as well as the fact that W is a subset of the compact set U_0 . $\square$

Proof of Proposition 5. It is obvious that the recommendation rule on D_c consisting in selecting all virtual guides for all profiles satisfies the Acceptance, Pareto, Internal-Robustness, Anonymity, and External-Robustness principles. Conversely, let f be a recommendation rule on D_c satisfying the Ac-

ceptance, Pareto, and Internal-Robustness principles. Consider a profile $(\succsim_0, \succsim_1, \dots, \succsim_N) \in D_c$ and let $\succsim = f(\succsim_0, \succsim_1, \dots, \succsim_N)$. Fixing arbitrary representations $U_0, U_1, \dots, U_N$ of $\succsim_0, \succsim_1, \dots, \succsim_N$, respectively, let V denote the corresponding set of virtual guides. Then by Proposition 1, $\succsim$ can be represented as per Definition 1 by some set $U \subseteq V$.

Suppose $\succsim$ cannot be represented as per Definition 1 by V . By the uniqueness of expected multi-utility representation and our finiteness assumption, we must then have $\{\theta u + \kappa : u \in U, \theta \in \mathbb{R}_+, \kappa \in \mathbb{R}\} \neq \{\theta v + \kappa : v \in V, \theta \in \mathbb{R}_+, \kappa \in \mathbb{R}\}$ and, hence, there must exist some $u \in U$ that does not belong to $\{\theta v + \kappa : v \in V, \theta \in \mathbb{R}_+, \kappa \in \mathbb{R}\}$. By the same argument as in Case 1 of the proof of Lemma 1(a), it follows that there exist $x, y \in X$ such that $Eu(x) > Eu(y)$ whereas $Ev(y) \geq Ev(x)$ for all $v \in V$. Hence, letting $\succsim'_0$ be represented as per Definition 1 by $\{u_0\}$, we have $x \succ'_0 y$. On the other hand, letting $\succsim' = f(\succsim'_0, \succsim_1, \dots, \succsim_N)$, we have $y \succ' x$ by the Internal-Robustness principle. This contradicts the Acceptance principle. $\square$

Proof of Proposition 6. Similar to the proof of Proposition 5. $\square$

Proof of Proposition 7. Assume A contains at least three alternatives a, b, c , let $w, w' : A \rightarrow \mathbb{R}$ be such that $w(a) = w'(b) = 1$ and $w(b) = w'(a) = w(c) = w'(c) = 0$, and let U denote the convex hull of $\{w, w'\}$. Note that for all $u, v \in U$, if $u \neq v$ then there exist $x, y \in X$ such that $Eu(x) > Eu(y)$ whereas $Ev(y) > Ev(x)$. Now let f be a recommendation rule on D_c (resp. D_{wc}) satisfying the Acceptance and Pareto (resp. Pareto-Indifference) principles that yields a complete recommendation for all profiles. Let $\succsim$ be represented as per Definition 1 by U and $\succsim' = f(\succsim, \succsim, \dots, \succsim)$. Then by Proposition 1 (resp. 2), $\succsim'$ can be represented as per Definition 1 by $\{u\}$ for some $u \in U$. Let $\succsim''$ be represented by some $v \in U \setminus \{u\}$ and $\succsim''' = f(\succsim, \succsim'', \dots, \succsim'')$. Again by Proposition 1 (resp. 2), $\succsim'''$ can be represented as per Definition 1 by $\{v\}$. Hence $\succsim'''$ does not refine $\succsim'$, so f does not satisfy the External-Robustness principle. $\square$

References

- ALTMAN, A. AND M. TENNENHOLTZ (2007): “An Axiomatic Approach to Personalized Ranking Systems,” in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., IJCAI’07, 1187–1192.
- BAUCELLS, M. AND L. S. SHAPLEY (2008): “Multiperson utility,” *Games and Economic Behavior*, 62, 329–347.
- CHEN, R., Q. HUA, Y.-S. CHANG, B. WANG, L. ZHANG, AND X. KONG (2018): “A Survey of Collaborative Filtering-Based Recommender Systems: from Traditional Methods to Hybrid Methods Based on Social Networks,” *IEEE Access*, 6, 64301–64320.
- DANAN, E., T. GAJDOS, B. HILL, AND J.-M. TALLON (2016): “Robust Social Decisions,” *American Economic Review*, 106, 2407–2425.
- DANAN, E., T. GAJDOS, AND J.-M. TALLON (2013): “Aggregating sets of von Neumann-Morgenstern utilities,” *Journal of Economic Theory*, 148, 663–688.

- (2015): “Harsanyi’s Aggregation Theorem with Incomplete Preferences,” *American Economic Journal: Microeconomics*, 7, 61–69.
- DE MEYER, B. AND P. MONGIN (1995): “A note on affine aggregation,” *Economics Letters*, 47, 177–183.
- DEMANGE, G. (2014): “A ranking method based on handicaps,” *Theoretical Economics*, 9, 915–942.
- (2017): “Mutual rankings,” *Mathematical Social Sciences*, 90, 35–42.
- DUBRA, J., F. MACCHERONI, AND E. A. OK (2004): “Expected utility theory without the completeness axiom,” *Journal of Economic Theory*, 115, 118–133.
- DUBRA, J. AND E. A. OK (2002): “A Model of Procedural Decision Making in the Presence of Risk,” *International Economic Review*, 43, 1053–1080.
- ELIAZ, K. AND R. SPIEGLER (2019): “The Model Selection Curse,” *American Economic Review: Insights*, 1, 127–140.
- HARSANYI, J. (1955): “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility,” *Journal of Political Economy*, 63, 309–321.
- HARSANYI, J. C. (1953): “Cardinal Utility in Welfare Economics and in the Theory of Risk-taking,” *Journal of Political Economy*, 61, 434–435.
- MONGIN, P. AND M. PIVATO (2016): “Social Preference and Social Welfare Under Risk and Uncertainty,” in *Handbook of Well-being and Public Policy*, ed. by M. D. Adler and M. Fleurbaey, Oxford: Oxford University Press, chap. 24, 711–745.
- PENNOCK, D. M., E. HORVITZ, AND C. L. GILES (2000): “Social Choice Theory and Recommender Systems: Analysis of the Axiomatic Foundations of Collaborative Filtering,” in *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, AAAI Press, 729–734.
- PIVATO, M. (2011): “Risky social choice with incomplete or noisy interpersonal comparisons of well-being,” *Social Choice and Welfare*, 40, 123–139.
- (2013): “Social welfare with incomplete ordinal interpersonal comparisons,” *Journal of Mathematical Economics*, 49, 405–417.
- (2014): “Social choice with approximate interpersonal comparison of welfare gains,” *Theory and Decision*, 79, 181–216.
- VON NEUMANN, J. AND O. MORGENSTERN (1944): *Theory of games and economic behavior*, Princeton University Press.
- WEYMARK, J. A. (1991): “A reconsideration of the Harsanyi-Sen debate on utilitarianism,” in *Interpersonal Comparisons of Well-Being*, ed. by J. Elster and J. E. Roemer, Cambridge: Cambridge University Press.
- (1993): “Harsanyi’s Social Aggregation Theorem and the Weak Pareto Principle,” *Social Choice and Welfare*, 10, 209–221.

——— (1995): “Further remarks on Harsanyi’s social aggregation theorem and the weak Pareto principle,” *Social Choice and Welfare*, 12, 87–92.