

Tailored Recommendations*

Eric Danan[†]

Thibault Gajdos[‡]

Jean-Marc Tallon[§]

December 24, 2019

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, axiomatic perspective. We view a recommendation system as a way to “extend” an agent’s preferences

*We thank Franz Dietrich, Itzhak Gilboa, Tigran Melkonyan, Efe Ok, Marcus Pivato, Régis Renault, Zvi Safra, and seminar participants at Düsseldorf U., Warwick Business School, and PSE 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 grant ANR-11-LBX-0023-01. Tallon acknowledges financial support from the grant ANR-17-EURE-001.

[†]THEMA, Université de Cergy-Pontoise, CNRS, 33 boulevard du Port, 95000 Cergy-Pontoise, France (eric.danan@u-cergy.fr).

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

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

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 illustrate below, our formalization of the recommendation problem makes it apparent that existing collaborative filtering systems violate the Pareto principle.

In order to provide a simple and tractable characterization of the recommendations satisfying these two principles, 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 ([Shapley and Baucells, 1998](#); [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 characterize work and contrasting them with typical collaborative filtering systems. Section 3 develops the formal model and contains the main results. Section 4 discusses informally how virtual guide recommendations may help alleviate some common issues affecting collaborative filtering systems and outlines venues for future research. An Appendix gathers all proofs as well as additional material.

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 [Groh et al., 2012](#)). 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. the survey by [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 \(2014, 2017\)](#).

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 fills agent 1 and 2’s missing ratings for d and e with any numbers δ_1 and ϵ_2 within the scale, 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:

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

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).

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 is *complete* if for all $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both), and *incomplete* otherwise. 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 $x, y \in X$:

$$x \succsim y \text{ if and only if } \mathbb{E}u(x) \geq \mathbb{E}u(y) \text{ for all } u \in U. \quad (1)$$

Expected multi-utility preferences are generally incomplete and are axiomatized by [Shapley and Baucells \(1998\)](#) 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 (1) with U consisting of all possible ways to fill in the missing ratings. For instance, in the rich ratings example,

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. \(2015\)](#).

denoting by $u_0^{\delta,\epsilon}$ the utility function filling agent 0's missing 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 $x \succeq y$ and $y \succeq x$, and that x is *strictly preferred* to y , denoted $x \succ y$, if $x \succeq y$ but not $y \succ x$. 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 $\tau \in (0, 1)$ such that $\lambda x + (1 - \lambda)x' \succ \mu y + (1 - \mu)y'$ for all $\lambda, \mu \in [\tau, 1]$. 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 $u \in U$.³ 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 \succeq_n on X . We want to give a recommendation to agent 0, in the form of another expected multi-utility preference relation \succeq on X satisfying the following two simple properties:

Axiom (Acceptance principle). For all $x, y \in X$, if $x \succeq_0 y$ then $x \succeq y$ and if $x \gg_0 y$ then $x \gg y$.

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

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 $b \succeq (a, \frac{6}{7}; e, \frac{1}{7})$ but not necessarily $b \succ (a, \frac{6}{7}; e, \frac{1}{7})$. The Pareto principle, on the other hand, expresses the requirement that the recommendation must be based on the preferences 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 $c \succeq e \succeq a$, $d \succeq a$, and $d \succeq b$.

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 $1, \dots, N$.

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 $\succeq_0, \succeq_1, \dots, \succeq_N$ as per (1), respectively. We say that a utility function $u : A \rightarrow \mathbb{R}$ is a *virtual guide* (for agent 0) if:

$$u \in U_0 \text{ and}$$

³It is without loss of generality in the sense that we can always choose a representation U such that this holds.

$$u = \sum_{n=1}^N \theta_n u_n + \kappa \text{ for some } u_1 \in U_1, \dots, u_N \in U_N, \theta_1, \dots, \theta_N \in \mathbb{R}_+, \kappa \in \mathbb{R}.$$

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.

Theorem 1. A recommendation \succsim satisfies the Acceptance and Pareto principles if and only if it can be represented as per (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 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.

A corollary to Theorem 1 is that whenever agent 0 has expressed at least one strong preference – or given different ratings to at least two alternatives –, 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 $x \succ y$.⁴ 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 $(c, \frac{1}{2}; e, \frac{1}{2}) \sim a$.⁵

Our recommendation problem generalizes the classical preference aggregation problem, which corresponds to the particular case where \succsim_0 is void – no two distinct alternatives are ranked by agent 0. In this case the Acceptance principle is trivially satisfied and Theorem 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

⁴Indeed, 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 $n = 1, \dots, N$.

⁵In 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.

(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.

Finally, we show in Appendix A.1 that the set V of 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 virtual guide recommendations do not exist. We say that the preference relations $\succsim_0, \succsim_1, \dots, \succsim_N$ are *compatible* if there exist no $x, y \in X$ such that $x \succ_{\succsim_0} y$ and $y \succ_{\succsim_n} x$ for all $n = 1, \dots, N$.

Theorem 2. V is non-empty if and only if $\succsim_0, \succsim_1, \dots, \succsim_N$ are compatible. Moreover, V is then the convex hull of finitely many utility functions.

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 indifferences 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.

We say that a utility function $u : A \rightarrow \mathbb{R}$ is a *signed virtual guide* (for agent 0) if:

$u \in U_0$ and

$$u = \sum_{n=1}^N (\theta_n u_n - \omega_n v_n) + \kappa \text{ for some } u_1, v_1 \in U_1, \dots, u_N, v_N \in U_N, \theta_1, \omega_1, \dots, \theta_N, \omega_N \in \mathbb{R}_+, \kappa \in \mathbb{R}.$$

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.⁶ 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 .

Theorem 3. A recommendation \succsim satisfies the Acceptance and Pareto-Indifference principles if and only if it can be represented as per (1) by some set $U \subseteq W$ of signed virtual guides.

⁶Indeed, 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 $\theta_1, \theta_2, \kappa \in \mathbb{R}$.

We show in Appendix A.1 that the set W of signed virtual guides can also be computed simply by solving a finite linear system, and this allows in particular to characterize the cases where signed virtual guide recommendations do not exist. We say that the preference relations $\succsim_0, \succsim_1, \dots, \succsim_N$ are *weakly compatible* if there exist no $x, y \in X$ such that $x \succ_0 y$ and $x \sim_n y$ for all $n = 1, \dots, N$.

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

3.5 Sketch of proofs

The proofs of Theorems 1–4 appear in Appendix A.2 and are quite straightforward. We actually derive Theorems 1 and 3 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 $x \succ_0 y$ then $x \succ y$.

Proposition 1. Let U be a representation of \succ as per (1).

(a) \succ 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:

$$u = \theta_0 u_0 + \kappa_0 = \sum_{n=1}^N \theta_n u_n + \kappa. \quad (2)$$

(b) \succ 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:

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

Each part of Proposition 1 could be simply deduced from a double application of the aggregation theorem in Danan et al. (2015) – once with respect to agent 0 alone and once with respect to agents $1, \dots, I$. Appendix A.2 instead presents self-contained proofs based on standard separation arguments. The strong preference clause in the Acceptance principle then implies that θ_0 can be taken strictly positive in both parts of Proposition 1, and the two theorems readily follow from a suitable normalization of the utility functions in U .

To establish Theorems 2 and 4, we start from the finite systems of linear equalities and inequalities determining the sets of virtual guides and signed virtual guides, which are given in Appendix A.1. Applying Farkas' Lemma to these systems – after some simple changes of variables – yields the first statement in these two theorems. The second statement directly follows from the finiteness and linearity of the systems together with the fact that U_0 is compact.

4 Discussion

As explained above, an important difference between virtual guide and collaborative filtering recommendations is that the former always satisfy the Pareto principle whereas the latter generally do not. Besides this difference, virtual recommendations may 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.

As the latter two issues illustrate, the selected virtual guides can critically affect the recommendations’s properties. 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 model by considering *recommendation rules* associating a recommendation \succsim to each compatible preference profile $\succsim_0, \succsim_1, \dots, \succsim_N$ – either for fixed or variable N . Then, beyond requiring the recommendation to satisfy the Acceptance and Pareto principle for each profile, axioms constraining how the recommendation \succsim varies with the profile could be expressed. One could for instance require the recommendation to treat agents $1, \dots, N$ “anonymously”, or to remain stable when some of these agents are “duplicated”. Identifying well-behaved recommendations in this extended setting, as well as testing them on actual data, are avenues for future research.

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.
- 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.
- 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.
- GROH, G., S. BIRNKAMMERER, AND V. KÖLLHOFER (2012): *Social Recommender Systems*, Berlin, Heidelberg: Springer Berlin Heidelberg, 3–42.
- HARSANYI, J. (1955): “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility,” *Journal of Political Economy*, 63, 309–321.
- 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. 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.
- SHAPLEY, L. S. AND M. BAUCCELLS (1998): “Multiperson utility,” Tech. Rep. 779, UCLA Department of Economics.
- VON NEUMANN, J. AND O. MORGENSTERN (1944): *Theory of games and economic behavior*, Princeton University Press.

A Appendix

A.1 Computing (signed) virtual guides

We show that the set V of virtual guides can be computed by solving a finite linear system. 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 of the finite system:

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

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:

$$\begin{aligned} \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) &= \sum_{n=1}^N \sum_{i=1}^{I(n)} \phi_{n,i} u_{n,i}(a_k) + \kappa, & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}, \\ 0 &\leq \lambda_{0,i}, & i = 1, \dots, I(0), \\ 0 &\leq \phi_{n,i}, & n = 1, \dots, N, i = 1, \dots, I(n). \end{aligned} \tag{4}$$

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

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

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:

$$\begin{aligned} \sum_{i=1}^{I(0)} \lambda_{0,i} u_{0,i}(a_k) &= \sum_{n=1}^N \sum_{i=1}^{I(n)} \psi_{n,i} u_{n,i}(a_k) + \kappa, & k = 1, \dots, K, \\ 1 &= \sum_{i=1}^{I(0)} \lambda_{0,i}, \\ 0 &\leq \lambda_{0,i}, & i = 1, \dots, I(0). \end{aligned} \tag{5}$$

A.2 Proofs

Proof of Proposition 1. (a) It is obvious that if (2) holds then \succsim satisfies the Weak Acceptance and Pareto principles. Conversely, assume (2) 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 $y \succsim x$ by (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 $y \succsim_0 x$ by (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 $y \succsim x$ by (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 $y \succsim_n x$ by (1). This contradicts the Pareto principle.

(b) It is obvious that if (3) holds then \succsim satisfies the Weak Acceptance and Pareto-Indifference 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 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 (1), which contradicts the Pareto-Indifference principle. \square

Proof of Theorem 1. It is obvious that if \succsim can be represented as per (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 (1). Then for all $u \in U$, (2) holds by Proposition 1(a). It is sufficient to show that (2) 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 (1) and is a subset of V . So suppose (2) 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 Proposition 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 Theorem 3. Similar to the proof of Theorem 1. \square

Proof of Theorem 2. As explained in Appendix A.1, V is non-empty if and only if the system (4) has a solution. Equivalently, letting $\kappa = \kappa^+ - \kappa^-$ where $\kappa^+, \kappa^- \in \mathbb{R}_+$, 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:

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

By the same argument as in Case 1 of the proof of Proposition 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 $n = 1, \dots, N$, i.e. to compatibility of $\succeq_0, \succeq_1, \dots, \succeq_N$. 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 Theorem 4. As explained in Appendix A.1, W is non-empty if and only if the system (5) has a solution. Equivalently, letting $\psi_{n,i} = \psi_{n,i}^+ - \psi_{n,i}^-$ and $\kappa = \kappa^+ - \kappa^-$ where $\psi_{n,i}^+, \psi_{n,i}^-, \kappa^+, \kappa^- \in \mathbb{R}_+$, 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 Proposition 1(a), this is equivalent to the non-existence of lotteries $x, y \in X$ such that $x \succcurlyeq_0 y$ and $x \sim_n y$ for all $n = 1, \dots, N$, i.e. to weak compatibility of $\succcurlyeq_0, \succcurlyeq_1, \dots, \succcurlyeq_N$. 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 . □