

Multi-attribute proportional representation [☆]

Jérôme Lang ^a, Piotr Skowron ^b

^a Université Paris-Dauphine, Paris, France

^b University of Warsaw, Warsaw, Poland



ARTICLE INFO

Article history:

Received 10 January 2017

Received in revised form 16 July 2018

Accepted 23 July 2018

Available online 25 July 2018

Keywords:

Proportional representation

Diversity

Multiwinner elections

Apportionment

Recommendation systems

Algorithms

Computational complexity

Approximation algorithms

ABSTRACT

We consider the following problem in which a given number of items has to be chosen from a predefined set. Each item is described by a vector of attributes and for each attribute there is a desired distribution that the selected set should have. We look for a set that fits as much as possible the desired distributions on all attributes. An example of application is the choice of members for a representative committee, where candidates are described by attributes such as gender, age and profession, and where we look for a committee that for each attribute offers a certain representation, i.e., a single committee that contains a certain number of young and old people, certain number of men and women, certain number of people with different professions, etc. Another example of application is the selection of a common set of items to be used by a group of users, where items are labelled by attribute values. With a single attribute the problem collapses to the apportionment problem for party-list proportional representation systems (in such a case the value of the single attribute would be a political affiliation of a candidate). We study the properties of the associated subset selection rules, as well as their computational complexity.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Consider the following example. A research department has to choose k members for a recruiting committee. A selected committee should be gender-balanced, ideally containing 50% of male and 50% of female. Additionally, a committee should represent different research areas in certain proportions: ideally it should contain 55% of researchers specialising in area 1, 25% of experts in area 2, and 20% in area 3. Another requirement is that the committee should contain 30% of junior and 70% of senior researchers, and finally, the repartition between local and external members should be kept in proportions 30% to 70%. The pool of candidates from which the department can select members of such a committee is as shown in the table below.

In the given example, if the department wants to select $k = 3$ members, then it is easy to see that there exists no committee that would satisfy all the criteria perfectly. Nevertheless, some committees are better than others: intuitively we feel that in the selected committee the ratio of the numbers of members representing different genders should be either equal to 2:1 or to 1:2, the ratio of the numbers of members representing areas 1, 2 and 3, should be equal to 2:1:0. Further, the selected committee should contain one junior and two senior members, and exactly one member of the selected

[☆] The preliminary version of this paper was presented at the 30th Conference on Artificial Intelligence (AAAI-2016).

E-mail address: p.skowron@mimuw.edu.pl (P. Skowron).

Name	Gender	Group	Age	Affiliation
Ann	F	1	J	L
Bob	M	1	J	E
Charlie	M	1	S	L
Donna	F	2	S	E
Ernest	M	1	S	L
George	M	1	S	E
Helena	F	2	S	E
John	M	2	J	E
Kevin	M	3	J	E
Laura	F	3	J	L

committee should have local affiliation. Such relaxed criteria can be achieved by selecting Ann, Donna, and George. Now, let us consider the above example for the case when $k = 4$. In such a case, the ideal ratios between the numbers of members for each of the four attributes should be equal to 1:1, 2:1:1, 1:3, and 1:3, respectively. Observe, however, that there exists no committee satisfying such relaxed criteria. According to different criteria, in this case the best committee can be for instance {Ann, Charlie, Donna, George}, with two externals instead of three, or {Charles, Donna, George, Kevin}, with males being over-represented.

In this paper we formalise the intuition given in the above example and we define what it means for a committee to be optimal, with respect to multi-attribute proportional representation. In our approach we leverage classical tools from political and social sciences, in particular we adapt the concept of *proportional apportionment* from the political science literature [3] to the case of multiple attributes. The central question of the apportionment problem is how to distribute parliament seats between political parties, given the numbers of votes cast for each party. Indeed, we can consider our multi-attribute problem, with the single attribute being a political affiliation of a candidate, and the desired distributions being the proportions of votes cast for different parties. In such a case we can see that selecting a committee in our multi-attribute proportional representation system boils down to selecting a parliament according to some apportionment criterion.

To emphasise the analogy between our model and the apportionment methods, we should provide some discussion on where the desired proportions for attributes come from. Typically, but not always, they come from *votes*. For instance, each voter might give her preferred value for each attribute, and the ideal proportions coincide with the observed frequencies. For instance, out of 20 voters, 10 would have voted for a male and 10 for a female, 13 for a young person and 7 for a senior one, etc.¹ It is worth mentioning that the voters might cast approval ballots, that is for each attribute they might define a set of approved values rather than pointing out the single most preferred one. On the other hand, sometimes, instead of votes, there are “global” preferences on the composition of the committee, expressed directly by the group, imposed by law, or by other constraints that should be respected as much as possible independently of voters’ preferences.

There is a variety of apportionment methods considered in the literature (we refer the reader to the survey of Balinski and Young [3]). They are evaluated by means of properties; among those that are deemed important and have been extensively studied in the literature, we find *non-reversal*, *respect of quota*, *population monotonicity*, and *house monotonicity* (see [2]). We define the analogs of these properties for the multi-attribute domain. These properties give us some insights into the nature of multi-attribute committee selection mechanisms; in particular, their analysis allows us to view certain selection methods as generalisations of the appropriate apportionment rules. Specifically, following this approach, in this paper we define multi-attribute variants of the Hamilton rule and of the d’Hondt rule of apportionment, hereinafter referred to as the *multi-attribute Hamilton rule* and the *multi-attribute d’Hondt rule*.

The multi-attribute case, however, is also substantially different from the single-attribute one. In particular, multi-attribute proportional representation systems exhibit computational problems that do not appear in the single-attribute setting. Indeed, in the second part of our paper we show that finding an optimal committee is often NP-hard. However, we show that this challenge can be addressed by designing efficient approximation and fixed-parameter tractable algorithms. In particular, the core technical contribution of this paper lies in the analysis of approximation guarantees provided by the local-search algorithm for the problem of finding an optimal committee, with respect to a certain measure of multi-attribute proportional representation.

We believe that the model formalised in this paper has broad applications. As an example, consider a political system where the voters do not vote for the candidates directly, but rather for their opinions on various issues. For instance, quoting Lang and Xia [32], in 2012, voters in California had to decide in simultaneous multiple referenda whether to adopt each of the given eleven propositions²; a similar vote also took place in Florida. Given that the voters vote on propositions, our algorithms can be used to find a set of candidates that, in some sense, best represents opinions of voters about propositions. The number of propositions can be even larger: for instance, political parties have usually quite elaborate programs in which they refer to tens or hundreds of issues.

¹ How to aggregate in a consistent way ideal proportions specified by different voters is a nontrivial problem addressed in [15].

² http://en.wikipedia.org/wiki/California_elections,_November_2012.

Further, our algorithms can be useful for selecting diversified groups of people. For instance, assume that our goal is to prepare an advertisement campaign. In such a case it is often desirable to depute this task to a team where men, women, people with different age and different education level are well represented. Similarly, when we select a jury we would like it to be representative according to different criteria, such as ethnicity, gender, age, religious beliefs, education level and wealth. Admitting PhD students is another example, where we would like to have a diversity with respect to ethnicity, gender, nationality, but also with respect to skills, education background, or disciplines of interest.

As another example, consider a library offering a set of movies to buy. In IMDb³ movies can be described by many attributes, such as genre, country, language, year, actors, directors, awards, etc. Users often look for movies by their attributes. Our algorithms can help such library to find a representative collection of movies that fits the collective will as much as possible. Finding a representative collective set of attribute-value items can be also used as a tool for implementing *group recommendations* [1], where the goal is to recommend a set of items for a group of agents, based on their (possibly conflicting) preferences: in some recent approaches to group recommendation (see [27] for a survey and comparison of four approaches), each item is seen as a set of features, users' preferences over features are elicited, and the aim of the system is to suggest a few representative items, such as set of movies or a set of tourist activities that comply as much as possible with the users' preferences over features.

This paper is organised as follows. In Section 2 we recall some useful concepts and definitions relating to methods of apportionment. We present our model in Section 3 and in Section 4 we introduce two different optimisation criteria and define multi-attribute committee selection rules optimising these criteria. In Section 5 we show that, although computation of optimal committees is generally NP-hard, there exist good approximation and fixed-parameter tractable algorithms for finding them. We position our work with respect to related areas in Section 6. In Section 7 we give a detailed discussion on the model and some of its possible extensions. Finally, in Section 8 we conclude and point to further research issues.

2. Preliminaries: methods of apportionment

For each integer $i \in \mathbb{N}$, by $[i]$ we denote the set of the first i natural numbers, $[i] = \{1, \dots, i\}$.

Consider a sequence of t political parties, denoted as P_1, \dots, P_t . For each $i \in [t]$, let v_i denote the number of votes given to party P_i . An apportionment rule is a method that given a distribution of votes among parties, $v = (v_1, \dots, v_t)$ where v_i denotes the number of votes cast for party P_i , and the number of seats h (the size of the house), returns a distribution of the h seats among the t parties. We denote the number of seats allocated to party P_i by r_i .

As is often the case in social choice, ties may occur and we have to choose between resoluteness and neutrality between parties: a resolute apportionment rule returns a single solution by sacrificing neutrality in case a tie occurs, and an irresolute apportionment rule returns all tied apportionments. In the rest of the paper we focus on resolute rules, and assume that ties are broken by an exogenous priority relation between parties. All our results are easily adaptable to irresolute rules.

Formally, an apportionment rule is a function $\mathcal{A}: \mathbb{N}^t \times \mathbb{N} \rightarrow \mathbb{N}^t$ that for each $v \in \mathbb{N}^t$ and each $h \in \mathbb{N}$ returns a vector $\mathcal{A}(v, h) = (r_1, \dots, r_t)$ satisfying the following two conditions: (i) $\sum_{i \in [t]} r_i = h$, (ii) $r_i \in \mathbb{N} \cup \{0\}$ for each $i \in [t]$. We will use the symbol v_+ to denote the sum of all votes, $v_+ = \sum_{i=1}^t v_i$.

There are numerous apportionment rules considered in the literature. The two most commonly-used classes of apportionment rules are the *largest remainder* and the *divisor* methods [3], which we briefly describe below.

2.1. Largest remainder methods

The following definition describes one of the most prominent classes of apportionment methods.

Definition 1 (*The largest remainder methods*). Let $q \in \mathbb{Q}$ be a rational number. The largest remainder method with quota q works in two steps. In the first step, each party P_i is allocated $\lfloor v_i/q \rfloor$ seats (the quota value must be chosen in such a way that the number of seats allocated in the first step is guaranteed to be between $h - t$ and h). In the second step, the remaining seats are allocated to the parties so that each party is allocated either one or zero additional seats. The parties which are allocated an additional seat are the ones with the largest values of the remainders $v_i/q - \lfloor v_i/q \rfloor$ (using the tie-breaking priority relation if necessary).

The most common choice of a quota is the *Hare quota*, defined as $q_{\text{Hare}} = v_+/h$; the method based on the Hare quota is called the *Hamilton method* (also known as the *largest remainder method* or *Hare–Niemeyer method*).⁴ The Hamilton method was one of the first methods used in the contemporary democracies. Its definition dates back to the 18th century and it was first used to select the members of the U.S. House of Representatives between 1852 and 1900. Currently, with slight modifications, it is used in parliamentary elections in Russia, Ukraine, Tunisia, Namibia, and Hong Kong. Below we provide an example illustrating the Hamilton method.

³ <http://www.imdb.com/>.

⁴ Other common choices are the *Droop quota* $1 + \frac{v_+}{1+h}$, the *Hagenbach–Bischoff quota* $\frac{v_+}{1+h}$ and the *Imperiali quota* $\frac{v_+}{2+h}$.

Example 1. Consider the instance with four parties and 100 voters. Assume that 4, 12, 33, and 51 votes were cast for parties $P_1, P_2, P_3,$ and $P_4,$ respectively. Let us set $h = 10,$ thus $q_{\text{Hare}} = v_+/h = 10.$ In the first step the parties $P_1, P_2, P_3,$ and P_4 are allocated 0, 1, 3, and 5 seats, respectively. The remainders for the four parties equal to $4/10, 2/10, 3/10,$ and $1/10,$ respectively. In the second step, the single remaining seat goes to the party with the highest remainder, i.e., to $P_1.$ Consequently, the allocation of the seats returned by the Hamilton method is given by the vector $(1, 1, 3, 5).$ □

2.2. Divisor methods

Divisor methods (also known as *highest average methods*) constitute another class of common and important apportionment methods.

Definition 2 (Divisor methods). Let $d = (d_1, d_2, \dots)$ be a nondecreasing sequence of positive values. The divisor method defined by sequence d starts with an empty allocation $(0, \dots, 0),$ and in each of the h consecutive steps assigns one additional seat to some party. Let $s_i(j)$ denote the number of seats allocated to party P_j just before step $i.$ In the i -th step the party P_j with the highest ratio $v_j/d_{s_i(j)+1}$ is allocated an additional seat (using the tie-breaking priority relation if necessary). We denote this party as $A(v, h, i).$

The most commonly used sequences of divisors are $d_{\text{DHondt}} = (1, 2, 3, \dots)$ and $d_{\text{SL}} = (1, 3, 5, \dots).$ The divisor method based on the sequence d_{DHondt} is called the *d'Hondt method* (it is also known as the *Jefferson method* or the *Hagenbach-Bischoff method*). The definition of the d'Hondt method dates back to the 18th century as well, and it is currently used for apportionment in more than 40 countries. The divisor method based on the sequence d_{SL} is known as the *Sainte-Laguë method* (sometimes referred to as the *Webster method, Schepers method,* or the *method of major fractions*) and is currently used in several countries.

Example 2. Consider the instance from Example 1. The below table shows the computation of the d'Hondt method. In the i -th iteration the i -th highest value from the table is selected and a seat is allocated to the party that corresponds to this value. For instance, the first seat will be allocated to party $P_4,$ which corresponds to the highest value of 51. The highest 10 values are shown in bold font: these are the values that correspond to the 10 seats allocated to parties.

	v_1	v_2	v_3	v_4
$v_i/1$	4	12	33	51
$v_i/2$	2	6	16.5	25.5
$v_i/3$	1.33	4	11	17
$v_i/4$	1	3	8.25	12.75
$v_i/5$	0.8	2.4	6.6	10.2
$v_i/6$	0.66	2	5.5	8.5
$v_i/7$	0.57	1.71	4.71	7.28

According to the d'Hondt method the following parties will be allocated consecutive seats: we start by giving a seat to P_4 (that is, $A(v, 10, 1) = P_4,$) because the largest value in the table is $v_4/1 = 51;$ then a seat to $P_3,$ because the second largest value is $v_3/1 = 33;$ then a second seat to $P_4,$ because the third largest value is $v_4/2 = 25.5;$ then a third seat to $P_4,$ and then $P_3, P_4, P_2, P_3, P_4, P_4.$ In the end, parties $P_1, P_2, P_3,$ and P_4 will get 0, 1, 3, and 6 seats, respectively. □

2.3. Properties of methods of apportionment

Several properties of apportionment methods have been studied, starting with Balinski and Young [2]. Below, we recall the definitions of the several of them, which will be useful in our further discussion. Recall that v denotes the vector of votes, v_+ denotes the total number of all votes, h is the number of available seats and that $(r_1, \dots, r_t) = \mathcal{A}(v, h).$

Non-reversal. The rule \mathcal{A} is said to satisfy *non-reversal* if for each parties $P_i, P_j, r_i \geq r_j$ holds whenever $v_i > v_j.$

Respect of quota. The rule \mathcal{A} is said to *respect quota* if for each party P_i it holds that $\lfloor v_i h / v_+ \rfloor \leq r_i \leq \lceil v_i h / v_+ \rceil.$

Party population monotonicity. Consider two vectors of votes $v = (v_1, \dots, v_t)$ and $v' = (v'_1, \dots, v'_t)$ and a party P_i such that: (i) $v_i / v_+ > v'_i / v'_+,$ and (ii) $v_j / v_\ell = v'_j / v'_\ell$ for each $j, \ell \neq i.$ The rule \mathcal{A} satisfies *party population monotonicity* if for each such vectors of votes v and $v',$ it holds that $r_i \geq r'_i,$ where $(r_1, \dots, r_t) = \mathcal{A}(v, h)$ and $(r'_1, \dots, r'_t) = \mathcal{A}(v', h).$ In other words, if the relative number of votes of a party increases *ceteris paribus,* then this party cannot receive less seats. Conditions (i) and (ii) are satisfied in particular if v is obtained from v' by adding more votes for $P_i.$

House monotonicity. The rule \mathcal{A} satisfies *house monotonicity* if for each two numbers of available seats, h and $h',$ with $h' > h,$ and for each party P_i it holds that $r'_i \geq r_i,$ where $(r_1, \dots, r_t) = \mathcal{A}(v, h)$ and $(r'_1, \dots, r'_t) = \mathcal{A}(v, h').$

The property that we call party population monotonicity is sometimes called population monotonicity (for instance, this is often the case in the literature on fair allocation, and sometimes in the literature on apportionment [19]). However, most

commonly in the context of apportionment, the term “population monotonicity” is used to refer to a stronger property, which covers cases when voters migrate between parties [2] (intuitively, party population monotonicity describes only cases when the population of one party grows while the populations of others remain unchanged). In particular, it is known that only divisor methods satisfy population monotonicity [3]. Party population monotonicity is more interesting for our study, since we will show that it is satisfied by the Hamilton method, and so this property will be useful in understanding the relation between the Hamilton method and its multi-attribute counterpart that we introduce in this paper.

One assumption that is often implicitly made in the analysis of the apportionment methods is that each party has at least h members, i.e., that there will always be enough candidates in each party to be given the allocated seats. This assumption will be very relevant in our further discussion. We will refer to it as to the *full supply property*. It is commonly known that under full supply property the Hamilton method satisfies non-reversal and respect of quota, and that it fails house monotonicity (this failure of house monotonicity is better known under the name *Alabama paradox*). It is also known that the Hamilton method fails population monotonicity, and we will show that it satisfies its weaker variant – the party population monotonicity. On the other hand, the d’Hondt method satisfies all four properties except the respect of quota.

Proposition 1. *Under full supply property the Hamilton method satisfies party population monotonicity.*

The proof of Proposition 1, as all proofs omitted from the main text, is relegated to the Appendix A.

We note that there are also other properties of the apportionment methods considered in the literature, such as consistency, or the properties that deal with strategyproofness issues, such as resistance to party merging or to party splitting. We selected the above four properties for our analysis as the most basic ones, and perhaps the most often referred to in the literature. Moreover, as we shall soon see, they are relevant for our multi-attribute generalisation of apportionment, which does not seem to be the case for other properties listed above.

3. The multi-attribute model

In this section we give a formal description of our model and discuss its specific elements. We explain that our model can be viewed as a generalisation of the apportionment setting to the case of multiple attributes and we discuss how the properties of the apportionment methods from Section 2.3 can be formulated in such a generalised model.

3.1. The formal setting

Let $X = \{X_1, \dots, X_p\}$ be a set of p attributes, each with a finite domain $D_i = \{x_i^1, \dots, x_i^{q_i}\}$. We say that X_i is binary if $|D_i| = 2$. We set $D = D_1 \times \dots \times D_p$, and let $C = \{c_1, \dots, c_m\}$ be a set of candidates (also referred to as items); C is also referred to as the candidate database. Each candidate c_i is represented as a vector of attribute values $(X_1(c_i), \dots, X_p(c_i)) \in D$.⁵

For each $i \in [p]$, by π_i we denote a target distribution $\pi_i = (\pi_i^1, \dots, \pi_i^{q_i})$ with $\sum_{j=1}^{q_i} \pi_i^j = 1$. We set $\pi = (\pi_1, \dots, \pi_p)$. Typically, n voters have cast a ballot expressing their preferred value on every attribute X_i , and π_i^j is the fraction of voters who have x_i^j as their preferred value for X_i , but the results presented in the paper are independent from where the values π_i^j come from (see the discussion in the introduction).

A multi-attribute committee selection rule⁶ is a function \mathcal{R} that for each database of candidates C , each vector of target distributions π and each committee size $k \in [m]$, select a set of k candidates $\mathcal{R}(C, \pi, k)$ from C .⁷ Again, we focus on the *resolute* version of such rules, using a tie-breaking mechanism whenever necessary.

Intuitively, a good multi-attribute committee selection rule should select such candidates that the distribution of attribute values in the selected set is as close as possible to π . Let $S_k(C)$ denote the set of all subsets of C of cardinality k . Given $A \in S_k(C)$, the representation vector for A is defined as $r(A) = (r_1(A), \dots, r_p(A))$, where $r_i(A) = (r_i^j(A) \mid j \in [q_i])$ for each $i \in [p]$, and $r_i^j(A) = \frac{|\{c \in A \mid X_i(c) = x_i^j\}|}{k}$.

Example 3. Consider the example from the introduction. In this example there are $p = 4$ four attributes, $X = \{\text{Gender, Group, Age, Affiliation}\}$. There are two possible values for the attribute “Gender”, thus $D_1 = \{F, M\}$ ($x_1^1 = F$ and $x_1^2 = M$). Similarly, $D_2 = \{A, B, C\}$, with $x_2^1 = A$, $x_2^2 = B$, and $x_2^3 = C$, etc. For example, Ann can be represented in our model as a tuple (F, A, J, L) . The target distributions used in the introduction can be formulated in our model as:

⁵ By writing $X_j(c_i)$, we slightly abuse notation, that is, we consider X_j both as an attribute name and as a function that maps a candidate to an attribute value, yet this will not lead to any ambiguity.

⁶ We will stick to the terminology “committee” although the meaning of subsets of candidates has sometimes nothing to do with electing a committee.

⁷ Observe that the outputs of the committee selection rules for the single-attribute and the multi-attribute cases are different. In the single-attribute model the rule returns a vector while in the multi-attribute one, it returns a set of k candidates. This is because in the multi-attribute case it is often not reasonable to assume the full supply property, hence the exact structure of the candidate database is important – in particular, for our problem, it will be important which candidates exist in the database (and then which ones will be selected from it).

$$\begin{array}{cccc} \pi_1^1 = 0.5 & \pi_2^1 = 0.55 & \pi_3^1 = 0.3 & \pi_4^1 = 0.3 \\ \pi_1^2 = 0.5 & \pi_2^2 = 0.25 & \pi_3^2 = 0.7 & \pi_4^2 = 0.7 \\ & \pi_2^3 = 0.2 & & \end{array}$$

Consider a three-element committee $A = \{\text{Ann, Donna, George}\}$. For this committee, the values in appropriate representation vectors are the following:

$$\begin{array}{cccc} r_1^1(A) = 2/3 & r_2^1(A) = 2/3 & r_3^1(A) = 1/3 & r_4^1(A) = 1/3 \\ r_1^2(A) = 1/3 & r_2^2(A) = 1/3 & r_3^2(A) = 2/3 & r_4^2(A) = 2/3 \\ & r_2^3(A) = 0 & & \square \end{array}$$

The following definition formalises our intuition regarding which committees are considered ideal.

Definition 3. A committee $A \in S_k(C)$ is *perfect* for π if $r_i(A) = \pi_i$ for all i .

Thus, a perfect committee matches exactly the target distributions. Clearly, there is no perfect committee if for some i, j , π_i^j is not an integer multiple of $1/k$. In some of our results we will focus on target distributions such that for each i, j the value $k\pi_i^j$ is an integer. We will refer to such target distributions as to *natural* distributions. Further, we see that the number of possible combinations of the attribute values grows exponentially with the number of nontrivial attributes (attributes which have at least two values). Consequently, even for natural distributions, finding a perfect committee cannot always be possible, simply because there are not sufficiently many appropriate candidates in the database. This observation suggests that the multi-attribute analog of the full supply property might be harder to satisfy in the multi-attribute setting (especially when the number of attributes is large). Below, we generalise the definition of the full supply property to the case of multiple attributes.

Definition 4. A candidate database C satisfies the *full supply property* with respect to k if for any $\vec{x} \in D$ there are at least k candidates in C who have the value of each attribute X_i equal to $\vec{x}[i]$.

An alternative interpretation of the full supply property is that it is always possible to *create* any number of candidates corresponding to a specific vector of attribute values.

While in some cases, specifically when the number of attributes is very low and the number of candidates is very large, it is reasonable to expect that the database will satisfy the full supply property, it is a much less realistic assumption when the number of attributes is large and/or when the database is small. In this paper we show how to deal with such cases, and in particular, how to extend two standard methods of apportionment, the Hamilton method and the d'Hondt method, to the case of multiple attributes and to the case where the full supply property is violated. Nevertheless, the analysis of the full supply property for multiple attributes can also give us some interesting insights into the nature of the analysed multi-attribute committee selection rules, and in particular, it can allow us to view some of these rules as extensions of the classic methods of apportionment.

We can observe that there exists a straightforward polynomial-time algorithm for checking if the candidate database satisfies the full supply property. Indeed, first we need to check if the size of the database is at least equal to $k|D|$, and if this is the case, we should additionally check if for each $\vec{x} \in D$ there exists at least k appropriate candidates in the database.

3.2. Properties of multi-attribute committee selection rules

We now generalise the properties discussed in Section 2.3 to multi-attribute committee selection rules (which we simply refer to as “rules”).

Non-reversal. Intuitively, non-reversal says that if one value x of a certain attribute has a target value higher than that of another value y , then x should be represented in the resulting committee at least as well as y (in short, values that deserve more get more). Formally, a rule \mathcal{R} satisfies non-reversal if for each triple (C, π, k) , if $\mathcal{R}(C, \pi, k) = A$, then for all $i \in [p]$ and $j, j' \in [q_i]$, $\pi_i^j > \pi_i^{j'}$ implies $r_i^j(A) \geq r_i^{j'}(A)$.

Respect of quota. A rule respects quota if its results match target distributions “almost” exactly, i.e., if they are allowed not to match them exactly, only because of the rounding issues. Formally, a rule \mathcal{R} respects quota if for each triple (C, π, k) , if $\mathcal{R}(C, \pi, k) = A$, then for all $i \in [p]$ and $j \in [q_i]$, either $r_i^j(A) = \lfloor k\pi_i^j \rfloor$ or $r_i^j(A) = \lceil k\pi_i^j \rceil$.

Value monotonicity (with respect to attribute X_i). Informally, value monotonicity says that if we increase the demand for a certain value x of a certain attribute X_i while not changing other demands, then in the new committee x should be at least as well represented as in the old one. Formally, consider a candidate database C , an integer k , and two vectors of target distributions π and ρ , such that there exist i, j with: (i) $\pi_i^j > \rho_i^j$, (ii) $\pi_i^{j''}/\pi_i^{j'} = \rho_i^{j''}/\rho_i^{j'}$ for all

$j', j'' \neq j$, and (iii) $\rho_{i'}^j = \pi_{i'}^j$ for all $i' \neq i$ and all $j \in [q_{i'}]$. \mathcal{R} satisfies value monotonicity with respect to attribute X_i if for each such C, k, π and ρ , if $\mathcal{R}(C, k, \pi) = A$ and $\mathcal{R}(C, k, \rho) = B$, then $r_i^j(A) \geq r_i^j(B)$.

House monotonicity. Informally, house monotonicity says that if we increase the number of available seats, then in the new committee each value of each attribute will be at least as well represented as in the old one. Formally, a rule \mathcal{R} satisfies house monotonicity if for any candidate database C , target distribution π , and two integers k and k' with $k' > k$, if $\mathcal{R}(C, \pi, k) = A$ and $\mathcal{R}(C, \pi, k') = B$ then $r_i^j(B) \geq r_i^j(A)$ holds for all i, j .

Clearly, these four definitions generalise the classical definitions of the properties when there is a single attribute (in particular, value monotonicity generalises party population monotonicity). Even though there could be other generalisations, those are arguably natural ones. In what follows, we will use these properties to argue that a certain class of multi-attribute committee selection rules can be viewed as extensions of the appropriate methods of apportionment.

4. Multi-attribute extensions of methods of apportionment

As we argued in the previous section, finding perfect committees in many cases might not be feasible, either because the target distributions are not natural, or because the candidate database does not satisfy the full supply property. These two observations lead us to define two metrics measuring how well a committee fits a target distribution. These two metrics induce two different methods of finding committees, which can be viewed as extensions of the Hamilton rule and of the d'Hondt rule to the multi-attribute domains. Other metrics will be briefly considered in Section 7.3.

4.1. Multi-attribute Hamilton rule

We start by defining what we believe to be the most natural metric measuring how close a given committee is to a target distribution and by arguing that such metric induces a rule which can be viewed as an extension of the Hamilton method: the L^1 metric.

Definition 5 (Multi-attribute Hamilton rule). The multi-attribute Hamilton rule is the function \mathcal{R}_H that given a candidate database C , a vector of target distributions π , and an integer k , finds a committee $A \in S_k(C)$ minimising $\sum_{i,j} |r_i^j(A) - \pi_i^j|$.

In other words, Multi-Attribute Hamilton Rule minimises the total variation distance between π/p and r/p viewed as probability distributions (we normalise π and r , multiplying them by $1/p$, so that their sums are equal to 1, and so that they could be viewed as probability distributions). The above definition is illustrated by the following example.

Example 4. For the example from the introduction, we have $X = \{\text{Gender, Group, Age, Affiliation}\}$, $D = \{F, M\} \times \{A, B, C\} \times \{J, S\} \times \{L, E\}$, and $X_1(\text{Ann}) = F$, $X_1(\text{Bob}) = M$ etc. Further, we have $\pi_1 = (0.5, 0.5)$, $\pi_2 = (0.55, 0.25, 0.2)$, $\pi_3 = (0.3, 0.7)$, and $\pi_4 = (0.3, 0.7)$. For $k = 4$, there are eight different committees which minimise our expression – let us show the calculation for one of them: $A = \{\text{Ernest, George, Helena, Laura}\}$. We have $\sum_{i,j} |r_i^j(A) - \pi_i^j| = 0 + 0 + 0.05 + 0 + 0.05 + 0.05 + 0.05 + 0.2 + 0.2 = 0.6$. \square

Now, let us argue that Definition 5 can be viewed as an extension of the Hamilton method of apportionment. We start by considering the case of a single-attribute ($p = 1$). Without loss of generality, let us assume that the single attribute is a party affiliation. Moreover, let us for a moment assume the full supply property, i.e., that for each value x_1^j there are at least k candidates with value x_1^j (this is typically the case in party-list elections; in fact we need to use this assumption only to ensure that the Hamilton method is well defined). After defining $\pi_i^j = \frac{v_i}{v_+}$ for all i , we obtain the following result:

Proposition 2. When $p = 1$ and assuming there are at least k candidates for each value of the unique attribute, then \mathcal{R}_H coincides with the Hamilton apportionment rule.

Definition 5 is inspired by the idea of *distance rationalisation of voting rules* (see [23] for one of the most recent works on the topic). In distance rationalisation, one first defines the outcome of a voting rule for elections where there exist an obvious, non-controversial winning candidate; such elections are called *consensus elections*. Second, in order to define an outcome of a voting rule for an arbitrary election E we determine the closest consensus election E' according to some distance (for instance, one can use the swap distance to measure the closeness between preferences of voters expressed as rankings over candidates), and we set the outcome of the rule for E to the same set of winners as the outcome for E' . When viewed from this perspective, Proposition 2 can be viewed as giving a distance rationalisation of the Hamilton rule – the consensus apportionments are those where there exist perfect committees, and the distance is the Manhattan (or L^1) distance.

Under the full supply assumption, a very similar result to Proposition 2 holds for multiple attributes.

Proposition 3. Consider a candidate database that satisfies the full supply property. For any attribute X_i , any committee A that minimises $\sum_{i,j} |r_i^j(A) - \pi_i^j|$ is a Hamilton committee for the single-attribute problem $(\{X_i\}, D^{\downarrow X_i}, \pi_i, k)$, where $D^{\downarrow X_i}$ is the projection of D on $\{X_i\}$.

Proof. Let us fix an attribute X_i and towards a contradiction, let us assume that there exists a committee A that minimises $\sum_{i,j} |r_i^j(A) - \pi_i^j|$ and that is not a Hamilton committee for the single-attribute problem $(\{X_i\}, D^{\downarrow X_i}, \pi_i, k)$. By Proposition 2 we infer that there exists another committee B such that $\sum_j |r_i^j(B) - \pi_i^j| < \sum_j |r_i^j(A) - \pi_i^j|$. We will show that it is possible to construct a committee D from A and B such that $\sum_{i,j} |r_i^j(D) - \pi_i^j| < \sum_{i,j} |r_i^j(A) - \pi_i^j|$. This will contradict the assumption that A that minimises $\sum_{i,j} |r_i^j(A) - \pi_i^j|$.

We construct D as follows. We sort A and B in some fixed arbitrary way; let $A = \{a_1, \dots, a_k\}$ and $B = \{b_1, \dots, b_k\}$. For each $i \in [k]$ we take $a_i \in A$ and $b_i \in B$, and define d_i as a candidate which has the value of the i -th attribute the same as b_i and the value of all other attributes the same as a_i . We add d_i to D ; the full supply assumption guarantees that it will always be possible to find such a candidate. We have that:

$$\sum_j |r_i^j(D) - \pi_i^j| = \sum_j |r_i^j(B) - \pi_i^j| < \sum_j |r_i^j(A) - \pi_i^j|$$

and that for each $\ell \neq i$ it holds that:

$$\sum_j |r_\ell^j(D) - \pi_\ell^j| = \sum_j |r_\ell^j(A) - \pi_\ell^j|.$$

Consequently, $\sum_{i,j} |r_i^j(D) - \pi_i^j| < \sum_{i,j} |r_i^j(A) - \pi_i^j|$, which gives a contradiction and completes the proof. \square

Note that the construction given in the proof gives us a practical way of constructing an optimal committee under the full supply assumption.

Further, below we also show that under the full supply assumption, the multi-attribute Hamilton rule satisfies the same from the four (multi-attribute variants) of the properties considered in Section 2.3 as the Hamilton method of apportionment. We start by noticing that if a property fails to be satisfied in the single-attribute case, *a fortiori* it is not satisfied in the multi-attribute case. As a consequence, house monotonicity is not satisfied, even under the full supply assumption.

Proposition 4. Under the full supply assumption, non-reversal, respect of quota, and value monotonicity with respect to every attribute are all satisfied by the multi-attribute Hamilton rule. In the general case, non-reversal, and respect of quota are not satisfied. If X_i is a binary variable, then value monotonicity with respect to X_i is satisfied; however it is not satisfied in the general case.

Importantly, if a perfect committee A exists for C , π and k , then $\mathcal{R}_H(C, \pi, k) = A$.

We close this subsection by a short discussion on the metric that is minimised in the definition of the multi-attribute Hamilton rule. It is perfectly reasonable to consider other metrics such as $\max_{i,j} |r_i^j(A) - \pi_i^j|$ (max-max) or $\sum_i \max_j |r_i^j(A) - \pi_i^j|$ (sum-max). In this paper we focus on $\sum_{i,j} |r_i^j(A) - \pi_i^j|$ because we believe this is the most natural choice. For a discussion on other metrics we refer the reader to the conference version of this paper [31]. (Note that Proposition 3 does not hold with the max-max metric.)

Finally, note that similar extensions to other largest remainder methods can be defined in the same way, after changing the value of the quota.

4.2. Multi-attribute d'Hondt rule

In this subsection we extend the idea implemented in the d'Hondt apportionment method to the multi-attribute setting. This generalisation would work for any other divisor method; for the sake of brevity, and also because the d'Hondt apportionment method is the most often used divisor method, throughout the rest of the paper we focus on this specific method. We first observe that the result of the d'Hondt apportionment can be formulated equivalently, as the solution to an optimisation problem.

Lemma 1. Consider the apportionment problem and an allocation $r = (r_1, \dots, r_t)$. If for each $i, j \in [t]$ it holds that $v_i/r_{i+1} \leq v_j/r_j$, then r is a d'Hondt apportionment.

Proof. For the sake of contradiction let us assume that r is not a d'Hondt apportionment. Let us run the d'Hondt method on v (breaking ties arbitrarily), outputting r^* . Let u be the last step such that $s_u(i) \leq r_i$ for all i : for some i we have $s_{u+1}(i) = r_i + 1$ (such a step exists because $r^* \neq r$). By the pigeonhole principle, for some j we have $s_{u+1}(j) = s_u(j) < r_j$. By definition of the d'Hondt procedure at step u , $\frac{v_j}{s_u(j)+1} \leq \frac{v_i}{s_u(i)+1} = \frac{v_i}{r_i+1}$. Since $s_u(j) < r_j$ we have $\frac{v_i}{r_i+1} \geq \frac{v_j}{s_u(j)+1} > \frac{v_j}{r_j}$, which contradicts the condition in the statement of the lemma. \square

This leads us to the following equivalent formulation of the d'Hondt rule. The following proposition states an analogous result for the d'Hondt method as Proposition 2 for the Hamilton rule.

Proposition 5. *In the classical apportionment setting, an allocation (r_1, \dots, r_t) maximises the value of the expression $\sum_{i \in [t]} \sum_{j \in [r_i]} v_i/j$ if and only if it is one of the outcomes of the irresolute version of the d'Hondt apportionment rule.*

Proof. Define $S(r, v) = \sum_{i \in [t]} \sum_{j \in [r_i]} v_i/j$. Also, given an apportionment r and two attribute values (parties) i, j such that $r_i \neq k$ and $r_j \neq 0$, let $r[i^+j^-]$ be the apportionment obtained from r by giving one more seat to i and one less to j from r .

First, we show that if allocation $r = (r_1, \dots, r_t)$ maximises $S(r, v)$, then it is a d'Hondt apportionment for v . For the sake of contradiction let us assume that r is not a d'Hondt apportionment. By Lemma 1 we know that in such a case there exist two parties i, j such that $v_i/r_{i+1} > v_j/r_j$. Let $r' = r[i^+j^-]$. We have $S(r', v) = S(r, v) + v_i/r_{i+1} - v_j/r_j > S(r, v)$, therefore r does not maximise $S(r, v)$.

Next, we prove that if r is a d'Hondt apportionment for v , then it maximises $S(r, v)$. For the sake of contradiction let us assume that there exists r' such that $S(r', v) > S(r, v)$. Let $r^{(0)} = r$. We define the following sequences of apportionments: for an integer $s \geq 0$, if $r^{(s)} \neq r'$ then, since $\sum_{i \in [t]} r_i^{(s)} = \sum_{i \in [t]} r'_i = k$, there must exist two indices i and j such that $r_i^{(s)} > r'_i$ and $r_j^{(s)} < r'_j$. Let $r^{(s+1)} = r^{(s)}[j^+i^-]$. We have $S(r^{(s+1)}, v) - S(r^{(s)}, v) = v_j/r_j^{(s)+1} - v_i/r_i^{(s)}$. Consider the step in the execution of the d'Hondt method when the r_i -th seat has been allocated to party P_i . At this step, party P_j is allocated x seats with $x \leq r_j \leq r'_j - 1$. Since the d'Hondt method allocated the seat to party P_i instead of P_j it must hold that $v_i/r_i^{(s)} \geq v_j/r_j \geq v_j/x+1 \geq v_j/r_j+1 \geq v_j/r_j^{(s)+1}$. Thus, each transformation does not increase the value of the expression. Yet, after a number of such transformations we reach (r'_1, \dots, r'_t) which has a higher value of the expression than (r_1, \dots, r_t) . This gives a contradiction and completes the proof. \square

Observe that $\sum_{j \in [r_i]} v_i/j = v_i H(r_i)$, where $H(n) = \sum_{i=1}^n 1/i$ denotes the n -th harmonic number. Proposition 5 leads us to the following extension of the d'Hondt method to multi-attribute scenarios.

Definition 6 (Multi-attribute d'Hondt rule). The multi-attribute d'Hondt rule is the function \mathcal{R}_{dHondt} that given a candidate database C , a vector of target distributions π , and an integer k , finds a committee $A \in S_k(C)$ maximising $\sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k)$.

Since for each $x \in \mathbb{N}$ we have that $\log(x+1) \leq H(x) \leq \log(x+1) + 1$, the maximisation of $\sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k)$ is intuitively a very close objective to the maximisation of $\sum_{i,j} \pi_i^j \log(r_i^j(A) \cdot k)$, which is equivalent to the maximisation of $\sum_{i,j} \pi_i^j \log\left(\frac{r_i^j(A)}{\pi_i^j}\right)$, and so, to the minimisation of $\sum_{i,j} \pi_i^j \log\left(\frac{\pi_i^j}{r_i^j(A)}\right)$, which is the Kullback–Leibler divergence from r/p to π/p viewed as probability distributions.⁸

Example 5. Consider again our running example. For $k = 4$ there are two optimal committees {Bob, Donna, Ernest, Helena} and {Bob, Charlie, Donna, Helena}. The value of the optimised function for the first committee can be computed as $0.5 \cdot H(2) + 0.5 \cdot H(2) + 0.55 \cdot H(2) + 0.25 \cdot H(2) + 0.3 \cdot H(1) + 0.7 \cdot H(3) + 0.3 \cdot H(1) + 0.7 \cdot H(3) = 0.6 \cdot H(1) + 1.8 \cdot H(2) + 1.4 \cdot H(3) \approx 5.866$. \square

We can formulate a result for the multi-attribute d'Hondt rule that is analogous to Proposition 3 for the multi-attribute Hamilton rule.

Proposition 6. *Consider a candidate database that satisfies the full supply property. For any attribute X_i , any committee A that maximises $\sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k)$ is a d'Hondt committee for the single-attribute problem $(\{X_i\}, D^{\downarrow X_i}, \pi_i, k)$, where $D^{\downarrow X_i}$ is the projection of D on $\{X_i\}$.*

Let us now examine properties of the multi-attribute d'Hondt method. It is known that for a single-attribute case the d'Hondt method satisfies non-reversal, house monotonicity, and party population monotonicity, yet it does not respect quota.⁹ Consequently, respect of quota is not satisfied by the multi-attribute d'Hondt method even under the full supply assumption.

⁸ Thanks to one of the anonymous reviewers for this observation.
⁹ It is known that the d'Hondt method satisfies a weaker form of respect of quota – it respects lower quota, i.e., for each party P_i it holds that $r_i \geq \lfloor v_i/h/v_+ \rfloor$.

Proposition 7. *Under the full supply assumption, non-reversal, house monotonicity, and value monotonicity with respect to every attribute are all satisfied by the multi-attribute d'Hondt method. In the general case, non-reversal and house monotonicity are not satisfied. If X_i is a binary variable, then value monotonicity with respect to X_i is satisfied; however it is not satisfied in the general case.*

Finally, let us observe that if a perfect committee A exists for C, π and k , then $\mathcal{R}_{\text{dHondt}}(C, \pi, k) = A$. This follows from Proposition 5 and from the fact that in the single-attribute case a committee that exactly matches the target distributions is always selected by the d'Hondt method. This property, which, as we have seen, also holds for \mathcal{R}_H , will be useful in our further discussion on computational properties of our multi-attribute rules.

5. Computing multi-attribute rules

Now, we are ready to formally define the main computational problems that we address in this paper.

Problem 1. We are given X, C, π , and k . In the `OPTIMALHAMILTONREPRESENTATION` we look for a committee $A \in S_k(C)$ that minimises the expression $\sum_{i,j} |r_i^j(A) - \pi_i^j|$. In the `OPTIMALDHONDTRERESENTATION` problem our goal is to find a committee $A \in S_k(C)$ maximising $\sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k)$.

In this section we investigate the computational complexity of the problem of finding optimal committees. We start with observing that the problem of deciding whether there is a perfect committee for a given instance is NP-complete.

Proposition 8. *Given set of attributes X , a set of candidates C , a vector of target distributions π , an integer k , deciding whether there exists a perfect committee is NP-complete.*

Proof. Membership is straightforward. Hardness follows by reduction from the NP-complete problem `EXACT COVER WITH 3-SETS`, or `x3c` [28]. Let $I = (X, \mathcal{S})$ with $X = \{x_1, \dots, x_{3k}\}$ and $\mathcal{S} = \{S_1, \dots, S_n\}$ with $|S_i| = 3$ for each i . I is a positive instance of `x3c` iff there is a collection $\mathcal{S}' \subseteq \mathcal{S}$ with $|\mathcal{S}'| = k$ and $\cup\{S \mid S \in \mathcal{S}'\} = X$. Define the following instance of perfect committee: let X_1, \dots, X_{3k} be $3k$ binary attributes, and let C consist of m candidates c_1, \dots, c_m with $X_i(c_j) = 1$ if $x_i \in S_j$ and $X_i(c_j) = 0$ if $x_i \notin S_j$. Finally, for each i , $\pi_i^0 = \frac{k-1}{k}$ and $\pi_i^1 = \frac{1}{k}$. We want a committee of size k . $A = \{c_{i_1}, \dots, c_{i_k}\}$ is perfect for π if for each X_i , there is exactly one $j \in \{1, \dots, k\}$ such that $X_i(c_{i_j}) = 1$, which is equivalent to saying that for each x_i , there is exactly one $S_j \in \{S_{i_1}, \dots, S_{i_k}\}$ such that $x_i \in S_j$. Thus, there is a perfect committee for π and C if and only if I is a positive instance. \square

Since the multi-attribute Hamilton and d'Hondt methods always find a perfect committee if there exists one, this simple result implies that the decision problem associated with finding an optimal committee is NP-hard. In the next subsections we will explore two natural approaches to alleviate the NP-hardness of the problem: we will ask if the problem can be computed efficiently when certain natural parameters are small, and we will ask whether it can be well approximated.

In this paper we mostly present computational results for binary domains. However, this assumption is not as restrictive as it may seem – every instance of the `OPTIMALHAMILTONREPRESENTATION` problem can be transformed to a new instance with binary domains in the following way:

- $X_{\text{new}} = \{X_{i,j} \mid i \in [p], j \in [|D_i|]\}$; for each i, j we set $D_{i,j} = \{0, 1\}$.
- $C_{\text{new}} = \{c'_l \mid l = 1, \dots, m\}$; for each ℓ, i, j we have $X_{i,j}(c'_l) = 0$ iff $X_i(c_l) = x_i^j$.
- $\pi_{\text{new}} = (\pi_{i,j} \mid i \in [p], j \in [|D_i|])$, where for all $i = [p]$ and $j = [|D_i|]$, $\pi_{i,j}^0 = \pi_i^j$ and $\pi_{i,j}^1 = 1 - \pi_i^j$.

The following proposition establishes the relation between the optimal committees for the original problem, and for the problem transformed to binary domains.

Proposition 9. *For a given committee A and target distribution π , let A_{new} and π_{new} denote the committee and target distributions obtained as above. The following holds:*

$$\sum_{i,j} |r_i^j(A_{\text{new}}) - \pi_{i,j}^j| = 2 \sum_{i,j} |r_i^j(A) - \pi_i^j|.$$

Proof.

$$\sum_{i,j} |r_i^j(A) - \pi_i^j| = \sum_{i,j} \left| \frac{|\{c \in A : X_i(c) = x_i^j\}|}{k} - \pi_i^j \right| = \sum_{i,j} \left| \frac{|\{c \in A_{\text{new}} : X_{i,j}(c) = 0\}|}{k} - \pi_{i,j}^0 \right|$$




Fig. 1. The functions h_1, \dots, h_k used in the proof of Theorem 2 for $k = 4$. Observe that for each $x \in [k]$ the highest value v such that $v \leq h_i(x)$ for each $i \in [k]$ is equal to $v = H(x)$.

$$\begin{aligned}
 &= \frac{1}{2} \sum_{i,j} \left(\left| \frac{|\{c \in A_{\text{new}} : X_{i,j}(c) = 0\}|}{k} - \pi_{i,j}^0 \right| + \left| \frac{|\{c \in A_{\text{new}} : X_{i,j}(c) = 1\}|}{k} - \pi_{i,j}^1 \right| \right) \\
 &= \frac{1}{2} \sum_{i,j} \sum_{\ell \in \{0,1\}} |r_{i,j}^\ell(A_{\text{new}}) - \pi_{i,j}^\ell| = \frac{1}{2} \sum_{i,j} |r_i^j(A_{\text{new}}) - \pi_i^j|. \quad \square
 \end{aligned}$$

Proposition 9 has interesting implications – first, it shows that the transformed instance has the same perfect committees, and the same optimal Hamilton committees as the original instance; then it shows how to obtain approximation guarantees for OPTIMALHAMILTONREPRESENTATION for arbitrary domains having guarantees for the problem restricted to binary domains, which will be useful in Section 5.3.

5.1. Parameterised complexity

In this section, we study the parameterised complexity of the problem of finding optimal Hamilton and d’Hondt committees. We are specifically interested whether for some natural parameters there exist fixed parameter tractable (FPT) algorithms. We recall that the problem is FPT for a parameter P if its each instance I can be solved in time $O(f(P) \cdot \text{poly}(|I|))$ for some computable function f . From the point of view of parameterised complexity, FPT is seen as the class of easy problems. There is also a whole hierarchy of hardness classes, $\text{FPT} \subseteq W[1] \subseteq W[2] \subseteq \dots$ (for details, we point the reader to appropriate overviews [18,21,25,40]).

Obviously, the problem admits an FPT algorithm for the parameter m (we can enumerate all k -element subsets of the set of candidates and select the best one). Now, we present a negative result for parameter k (committee size) and a positive result for the parameter p (number of attributes).

Theorem 1. *The problem of deciding whether there exists a perfect committee is W[1]-hard for the parameter k , even for binary domains.*

Proof. By reduction from the W[1]-complete PERFECTCODE problem [13]. Let I be an instance of PERFECTCODE that consists of a graph $G = (V, E)$ and a positive integer k . We ask whether there exists $V' \subseteq V$ with $|V'| = k$ such that each vertex in V is adjacent to exactly one vertex from V' (by convention, a vertex is adjacent to itself). From I we construct the following instance I' of the problem of deciding whether there exists a perfect committee. For each $v \in V$ there is a binary attribute X_v and a candidate c_v . For each $u, v \in V$, $X_v(c_u) = 1$ if and only if u and v are adjacent in G . We look for a committee of size k . For each v , $\pi_v^1 = 1 - \pi_v^0 = \frac{1}{k}$. It is easy to see that perfect codes in I correspond to perfect committees in I' . \square

Theorem 2. *For binary domains, there exists an FPT algorithm for OPTIMALHAMILTONREPRESENTATION and for OPTIMALDHONDTRERESENTATION for parameter p .*

Proof. We will show a linear integer program for each of the two problems, OPTIMALHAMILTONREPRESENTATION and for OPTIMALDHONDTRERESENTATION, with the number of integer variables bounded by a function of p . Such a linear program, by the result of Lenstra [34, Section 5], can be solved in FPT time for parameter p . Currently, the best known running time of algorithms solving integer linear programs is $O(n^{2.5n+o(n)} \cdot \text{poly}(|I|))$, where n is the number of integer variables and $|I|$ is the length of encoding of the input instance [26,29]. We will start from describing the set of variables and constraints which are common for the two problems that we consider.

Each candidate can be viewed as a vector of values indexed with the attributes; there are 2^p such possible vectors: v_1, \dots, v_{2^p} . For each v_i , let a_i denote the number of candidates that correspond to v_i . For each $i \in [2^p]$ we introduce an integer variable b_i , which intuitively denotes the number of candidates corresponding to v_i in an optimal committee.

Further, for each $i \in [p]$ and each $j \in [q_i]$ we introduce a variable r_i^j , which in the optimal solution will be equal to $k \cdot r_i^j(A)$, where A is the optimal committee returned by our integer linear program. Consider the following set of linear constraints:

$$\left. \begin{array}{l} \text{(a): } b_i \in \mathbb{Z} \\ \text{(b): } b_i \geq 0 \\ \text{(c): } b_i \leq a_i \end{array} \right\} \quad i \in [2^p]$$

$$\text{(d): } \sum_{i=1}^{2^p} b_i = k$$

$$\text{(e): } r_i^j = \sum_{\ell: v_\ell[i]=x_i^j} b_\ell \quad i \in [p], j \in [q_i]$$

These constraints ensure that for a committee A described by the variables $(b_i)_{i \in [2^p]}$ it holds that $r_i^j = k \cdot r_i^j(A)$, for each $i \in [p]$ and $j \in [q_i]$. Now, for OPTIMALHAMILTONREPRESENTATION we additionally introduce two real variables, x_i^j and y_i^j , for each $i \in [p]$ and $j \in [q_i]$, and the following constraints:

$$\left. \begin{array}{l} \text{(f): } x_i^j \geq 0 \\ \text{(g): } x_i^j \geq r_i^j - k \cdot \pi_i^j \\ \text{(h): } y_i^j \geq 0 \\ \text{(i): } y_i^j \geq k \cdot \pi_i^j - r_i^j \end{array} \right\} \quad i \in [2^p], j \in [q_i]$$

These constraints ensure that for each $i \in [p]$ and $j \in [q_i]$ it holds that $k|r_i^j(A) - \pi_i^j| \leq x_i^j + y_i^j$. Thus, we can find an optimal Hamilton committee by minimising the objective function $\sum_{i \in [p]} \sum_{j \in [q_i]} (x_i^j + y_i^j)$ subject to constraints (a)–(i).

Finding an optimal d'Hondt committee requires an additional construction. The idea of this construction is similar to the one used by Caragiannis et al. [10]. Let us construct k linear functions: h_1, \dots, h_k , as follows. For each $i \in [k]$ we define h_i as a linear function such that $h_i(i - 1) = H(i - 1)$ and $h_i(i) = H(i)$, where $H(i)$ denotes the i -th harmonic number. The functions $(h_i)_{i \in [k]}$ for $k = 4$ are depicted in Fig. 1. Now, for each $i \in [p]$ and $j \in [q_i]$ we introduce one real variable z_i^j and the following constraints:

$$\text{(j): } z_i^j \leq h_\ell(r_i^j) \quad \ell \in [k]$$

The highest value of z_i^j which satisfies constraints (j) is equal to $H(r_i^j)$. Thus, to find an optimal d'Hondt committee we need to maximise the expression $\sum_{i \in [p]} \sum_{j \in [q_i]} z_i^j$ subject to constraints (a)–(e) and (j).

Of course, when there is no candidate corresponding to a given vector of values of the attributes v_i , then we can skip the respective variable b_i . This can make our ILPs more practical when the size of the candidate database is small. \square

Example 6. Let $p = 2$, $k = 5$, and let the candidate database C consists of 4 candidates with value vector $v_1 = (0, 0)$, 2 with value vector $v_2 = (1, 0)$, 2 candidates with value vector $v_3 = (0, 1)$ and 2 candidates with value vector $v_4 = (1, 1)$. Let $\pi = ((0.2, 0.8), (0.6, 0.4))$. The constraints (a)–(e) of the integer linear program are:

$$\begin{array}{ll} \text{(a): } b_i \in \mathbb{Z} & 1 \leq i \leq 4 \\ \text{(b): } b_i \geq 0 & 1 \leq i \leq 4 \\ \text{(c): } b_1 \leq 4; b_2 \leq 2; b_3 \leq 2; b_4 \leq 2 \\ \text{(d): } b_1 + b_2 + b_3 + b_4 = 5 \\ \text{(e): } r_1^0 = b_1 + b_3; r_1^1 = b_2 + b_4; r_2^0 = b_1 + b_2; r_2^1 = b_3 + b_4; \end{array}$$

and a solution is $(b_1 = 1, b_2 = 2, b_3 = 0, b_4 = 2)$: a perfect committee is obtained by taking one candidate with value vector $(0, 0)$, two candidates with value vector $(1, 0)$, and two with value vector $(1, 1)$. Thus, this is an optimal d'Hondt and Hamilton committee. \square

As a corollary of Theorem 2 we get that the problem of checking whether there exists a perfect committee is in FPT for the parameter p . We can see that the proof of Theorem 2 easily extends to the case where the size of each domain D_i is bounded by a constant.

Table 1

Running times (average and maximal) of the ILP for the problem of finding optimal Hamilton and d'Hondt committees. The entry "> 10 h" means that none of the ten experiments finished before the deadline of 10 hours.

m	p	k	average	maximal	m	p	k	average	maximal
50	20	5	0.06 s	0.07 s	50	20	10	0.07 s	0.09 s
50	20	10	0.33 s	1.17 s	50	20	10	1.05 s	1.73 s
100	50	10	55 s	98 s	100	50	10	59 s	89 s
100	50	25	2.7 h	9.5 h	100	50	25	1.9 h	4.3 h
300	80	40	> 10 h	> 10 h	300	80	40	> 10 h	> 10 h

a. Finding optimal Hamilton committees

b. Finding optimal d'Hondt committees

```

A ← ∅;
for i ← 1 to k do
  c ← argmaxc' ∈ C \ A ∑i,j πij H(rij(A ∪ {c'}) · k);
  A ← A ∪ {c};
return A;

```

Fig. 2. Greedy approximation algorithm for the OPTIMALDHONDTREPRESENTATION problem.

We conclude this section by a short discussion. Finding an optimal committee is likely to be difficult if the candidate database C is large, and the number of attributes not small. Assume $|C|$ is large compared to the size of the domain $\prod_{i=1}^p |D_i|$, that each attribute value appears often enough in C and that there is no strong correlation between attributes in C : then, the larger $|C|$, the more likely C satisfies the full supply property, in which case finding an optimal committee is easy. The really difficult cases are when $|C|$ is not significantly larger than the domain, or when C shows a high correlation between attributes.

We have run a set of simple experiments to better understand the limitations of the ILP-based approach presented in the proof of Theorem 2. For several different values of the parameters m (the number of candidates), p (the number of binary attributes) and k (the size of the committee) we run the following simulations. We selected the value of each attribute for each candidate independently, following a uniform i.i.d. distribution. For each $i \in [p]$ we set the target distribution for the i -th attribute to $\pi_i^0 = \pi_i^1 = 1/2$. For each combination of the parameter values (m, p, k) we run 10 experiments and computed the average and the worst-case time that the appropriate ILP required to solve the respective instance. For our experiments we used the Gurobi ILP solver and a desktop machine with 4 processors Intel Core i5-4200U, 1.62 GHz, 3072 KB of cache. The running times of our ILPs are summarised in Table 1.

We observe that for a small number of attributes, we can efficiently apply our FPT algorithms. Nevertheless, for large instances, with more than a hundred attributes and more than a few hundreds candidates, the ILP-based approach is no longer feasible. Following this observation, in the next two subsection we will discuss an alternative approach, which uses the concept of approximation. This approach is suitable when the number of attributes is too large for the ILP-based algorithms.

One more reason for studying approximation algorithms for finding a committee is that an approximation algorithm can be viewed as a new rule, which might or might not inherit some good properties of the original rule that we aim to approximate (this view of approximation algorithms was taken first in [9]). This new rule is not only simpler to compute but also may be easier to understand by humans. Moreover, for low-stake domains where optimality is not crucial, it often does not matter whether we apply the initial rule of one of its approximations.

5.2. Approximating the multi-attribute d'Hondt method

Let us first consider the problem of approximating the multi-attribute d'Hondt method. We will use the following standard definition of approximation.

Definition 7. An algorithm \mathcal{A} is an α -approximation algorithm for OPTIMALDHONDTREPRESENTATION if for each instance I of OPTIMALDHONDTREPRESENTATION it holds that

$$\sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k) \geq \alpha \sum_{i,j} \pi_i^j H(r_i^j(A^*) \cdot k),$$

where A is the committee returned by \mathcal{A} for I , and is A^* an optimal committee.

For the OPTIMALDHONDTREPRESENTATION problem we show that a simple greedy algorithm from Fig. 2 achieves an approximation ratio of $1 - 1/e$.

Theorem 3. *The greedy algorithm from Fig. 2 is a $(1 - 1/e)$ -approximation algorithm for OPTIMALHAMILTONREPRESENTATION.*

Proof. Consider the set function that we optimise $f(A) = \sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k)$. It can be expressed as a linear combination of functions $f(A) = \sum_{i,j} f_{i,j}(A)$, where $f_{i,j}(A) = H(r_i^j(A))$. We will now show that for each i and j the function $f_{i,j}$ is submodular. For that we need to show that for each A, B with $A \subset B \subset C$, and for each $c \in C \setminus B$ it holds that:

$$f_{i,j}(A \cup \{c\}) - f_{i,j}(A) \geq f_{i,j}(B \cup \{c\}) - f_{i,j}(B).$$

Now, if $X_i(c) \neq x_i^j$ then $r_i^j(A \cup \{c\}) = r_i^j(A)$ and $r_i^j(B \cup \{c\}) = r_i^j(B)$, thus:

$$f_{i,j}(A \cup \{c\}) - f_{i,j}(A) = 0 = f_{i,j}(B \cup \{c\}) - f_{i,j}(B).$$

On the other hand, if $X_i(c) = x_i^j$ then $r_i^j(A \cup \{c\}) = r_i^j(A) + 1$ and $r_i^j(B \cup \{c\}) = r_i^j(B) + 1$. Since $A \subset B$, we have $r_i^j(B) \geq r_i^j(A)$ and so:

$$f_{i,j}(A \cup \{c\}) - f_{i,j}(A) = \frac{1}{r_i^j(A) + 1} \geq \frac{1}{r_i^j(B) + 1} = f_{i,j}(B \cup \{c\}) - f_{i,j}(B).$$

Thus, each function $f_{i,j}$ is submodular. Also, f as a linear combination of submodular functions is submodular. The thesis follows from the famous result of Nemhauser et al. [39] which established the $(1 - 1/e)$ -approximation bound for the greedy algorithm for the problem of maximising a submodular function. \square

5.3. Approximating the multi-attribute Hamilton rule

Now, we move to the problem of approximating the multi-attribute Hamilton method. Before proceeding to presentation of our approximation algorithms for this problem, we define the notion of approximability used in our analysis. First, we observe that there is no hope for a polynomial time approximation algorithm according to the notion of multiplicative approximation, perhaps the most commonly used definition of approximation.

Proposition 10. *Unless $P = NP$, for each computable function $\alpha : \mathbb{N} \rightarrow \mathbb{N}$ there exists no polynomial-time algorithm that for each instance I of OPTIMALHAMILTONREPRESENTATION returns a committee A such that:*

$$\sum_{i,j} |r_i^j(A) - \pi_i^j| \leq \alpha(|I|) \cdot \sum_{i,j} |r_i^j(A^*) - \pi_i^j|,$$

where A^* is an optimal committee for I .

Proof. For the sake of contradiction, let us assume that such a polynomial-time algorithm exists. Then for each instance I for which there exists a perfect committee A^* , we have $\sum_{i,j} |r_i^j(A^*) - \pi_i^j| = 0$, and thus our algorithm would need to find a committee A with $\sum_{i,j} |r_i^j(A) - \pi_i^j| = 0$. This means that we could use our algorithm to find a perfect committee, whenever such exists. Yet, by Proposition 8 deciding whether there exists a perfect committee is NP-hard. \square

Given this strong negative result, we move to analysing the *additive* approximation of the problem.

Definition 8. An algorithm \mathcal{A} is an α -additive-approximation algorithm for OPTIMALHAMILTONREPRESENTATION if for each instance I of OPTIMALHAMILTONREPRESENTATION it holds that

$$\left| \sum_{i,j} |r_i^j(A) - \pi_i^j| - \sum_{i,j} |r_i^j(A^*) - \pi_i^j| \right| \leq \alpha,$$

where A is the committee returned by \mathcal{A} for I , and A^* an optimal committee.

Now, we are ready to show an approximation algorithm for the OPTIMALHAMILTONREPRESENTATION problem. The algorithm is given in Fig. 3 and is parameterised by an integer value ℓ . It starts with a random collection of k samples and, in each step, it looks whether it is possible to replace some ℓ candidates from the current solution with some other ℓ candidates to obtain a better solution (if there exist many choices for replacing such ℓ candidates the algorithm can pick an arbitrary of them). The algorithm continues until it cannot find any pair of sets of ℓ candidates that would improve the current solution.

Let $q = \max_{i \in [p]} q_i$. For the sake of simplicity, assume that target distributions are natural (this is almost without loss of generality, as there we can always find target distributions that are close to the initial distributions; see the end of Section 5). We show that in that case, the running time of the local search algorithm is $O(p^2 m^\ell k^{\ell+1} q^\ell)$. For instance, the simplest variant of the algorithm (that is, the algorithm for $\ell = 1$) for binary domains works in time $O(mp^2 k^2)$. Indeed,

Parameters:
 $\pi = (\pi_1, \dots, \pi_p)$ – input target distributions.
 ℓ – the parameter of the algorithm.
 $A \leftarrow k$ random candidates from C ;
while there exist $C_\ell \subset C$ and $A_\ell \subset A$ such that $|C_\ell| = |A_\ell| \leq \ell$, and
 $\sum_{i,j} |r_i^j(A) - \pi_i^j| > \sum_{i,j} |r_i^j((A \setminus A_\ell) \cup C_\ell) - \pi_i^j|$ **do**
 $A \leftarrow (A \setminus A_\ell) \cup C_\ell$;
return A ;

Fig. 3. Local search approximation algorithm for the OPTIMALHAMILTONREPRESENTATION problem.

the algorithm starts with a random committee A ; the worst case distance to the target distributions can be upper bounded by $\sum_{i,j} |r_i^j(A) - \pi_i^j| \leq \sum_{i,j} 1 \leq pq$. In each iteration of the *while* loop the solution improves: since the distributions are natural, the distance must improve by at least $1/k$. This is because for each $i \in [p]$, $j \in [q_i]$ the values $r_i^j(A)$ and π_i^j are integral multiples of $1/k$, and so is the optimised value. Thus, there will be at most pqk iterations of the *while* loop. In each iteration we check all ℓ -element subsets of the set of candidates and compare each such a subset with all ℓ -element subsets of the current best committee – thus, there are at most $m^\ell k^\ell$ such comparisons. For a single comparison we need to check all the attributes of the selected candidates in order to verify if replacing the appropriate subsets gives an improvement, which results in $p\ell$ operations.

As we show now, the approximation guarantees depend on the value of the parameter ℓ .

Theorem 4. Recall that $p = |X|$. For binary domains and natural distributions the local search algorithm from Fig. 3 with $\ell = 1$ is a p -additive-approximation algorithm for OPTIMALHAMILTONREPRESENTATION.

Proof. Let A^* denote an optimal solution for a given instance I of the OPTIMALHAMILTONREPRESENTATION problem. Let $A \in S_k(C)$ denote the set returned by the local search algorithm from Fig. 3. From the condition in the “while” loop, we know that there exist no $c \in C$ and $a \in A$ such that $\sum_{i,j} |r_i^j(A) - \pi_i^j| > \sum_{i,j} |r_i^j((A \setminus \{a\}) \cup \{c\}) - \pi_i^j|$. Now, let $X_{\text{ex}} \subseteq X$ denote the set of all attributes for which A achieves exact match with π , that is, such that for each $X_i \in X_{\text{ex}}$, we have that $r_i^1(A) = \pi_i^1$ and $r_i^2(A) = \pi_i^2$.

Let us consider the procedure consisting in taking the candidates from $A \setminus A^*$ and, one by one, replacing them with arbitrary candidates from $A^* \setminus A$. This procedure, in $|A \setminus A^*|$ steps, transforms A into an optimal solution A^* . We now estimate the total gain g induced by this procedure. For each candidate $a \in A \setminus A^*$, by $a' \in A^* \setminus A$ we denote the candidate which was taken to replace a in the procedure. For each attribute $X_i \in X$ we define the gain $g_i(a, a')$ of replacing a by a' as:

$$g_i(a, a') = \sum_{j \in \{1,2\}} (|r_i^j(A) - \pi_i^j| - |r_i^j(A \setminus \{a\} \cup \{a'\}) - \pi_i^j|).$$

We now extend this definition to sets of k candidates:

$$g_i(B, B') = \sum_{j \in \{1,2\}} (|r_i^j(A) - \pi_i^j| - |r_i^j((A \setminus B) \cup B') - \pi_i^j|).$$

If $X_i \in X_{\text{ex}}$, then $r_i(A) = \pi_i$, and so the replacement cannot improve the quality of the solution relatively to X_i , hence

$$\sum_{i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) \leq 0. \tag{1}$$

Since the distribution is natural, we have that $g_i(a, a') \in \left\{-\frac{2}{k}, 0, \frac{2}{k}\right\}$. This is because replacing a single candidate in A can change the value of $|\{c \in A : X_i(c) = x_i^j\}|$ by $-1, 0$, or 1 , and so, it can change the value of $r_i^j(A)$ by $-\frac{1}{k}, 0$, or $\frac{1}{k}$. Moreover, for each attribute $X_i \notin X_{\text{ex}}$ there are two possible cases:

1. $r_i^j(A) > \pi_i^j$ and each exchange of candidate that results in a negative gain increases $r_i^j(A)$.
2. $r_i^j(A) < \pi_i^j$ and each exchange that results in a negative gain decreases $r_i^j(A)$.

Intuitively, 1. and 2. mean that for attributes outside of X_{ex} , the negative gains cumulate. Formally, for each $X \notin X_{\text{ex}}$:

$$g_i(A \setminus A^*, A^* \setminus A) \leq \sum_{a \in A \setminus A^*} g_i(a, a'). \tag{2}$$

From the condition in the “while” loop, we have that for each $a \in A \setminus A^*$: $\sum_i g_i(a, a') \leq 0$, and so:

$$\sum_i \sum_{a \in A \setminus A^*} g_i(a, a') \leq 0. \tag{3}$$

We now give the following sequence of inequalities:

$$\begin{aligned} g &= \sum_i g_i(A \setminus A^*, A^* \setminus A) = \sum_{i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) + \sum_{i \notin X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) \\ &\leq \sum_{i \notin X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) \leq \sum_{i \notin X_{\text{ex}}} \sum_{a \in A \setminus A^*} g_i(a, a') \leq - \sum_{i \in X_{\text{ex}}} \sum_{a \in A \setminus A^*} g_i(a, a') \end{aligned}$$

The last transition in the above sequence is due to Inequality (3). Consequently, we get that:

$$g \leq \left| \sum_{i \in X_{\text{ex}}} \sum_{a \in A \setminus A^*} g_i(a, a') \right| \leq |X_{\text{ex}}| \cdot k \cdot \frac{2}{k} = 2|X_{\text{ex}}|. \tag{4}$$

Finally, for each attribute $X_i \notin X_{\text{ex}}$ the loss relative to X_i , i.e., $|r_i^0(A) - \pi^0| + |r_i^1(A) - \pi^1|$, is at most 2. Thus, we get $g \leq 2(|X| - |X_{\text{ex}}|)$, which leads to $g \leq |X|$. \square

One way to interpret the bound $|X|$ is to observe that a solution that for half of the attributes gives exact match, and for other half is arbitrarily bad, is an $|X|$ -approximate solution. We do not know whether the bound $|X|$ is reached, but below we show a lower bound of $\frac{2}{3}|X|$ on the error made by the algorithm with $\ell = 1$.

Example 7. Consider $3p$ binary attributes X_1, \dots, X_{3p} , 4ℓ candidates $C = \{a_1, \dots, a_{2\ell}, b_1, \dots, b_{2\ell}\}$, and let $k = 2\ell$. For each $i \leq p$, we have: for $j \leq \ell$, $X_i(a_j) = 1$ and $X_i(b_j) = 1$; for $j > \ell$, $X_i(a_j) = 0$ and $X_i(b_j) = 0$. For each i such that $p < i \leq 2p$ we have: for $j \leq \ell$, $X_i(a_j) = 1$ and $X_i(b_j) = 0$; for $j > \ell$, $X_i(a_j) = 0$ and $X_i(b_j) = 1$. For $i > 2p$ we have: for each j , $X_i(a_j) = 1$ and $X_i(b_j) = 0$. Finally, for $i \leq 2p$ let $\pi_i^0 = \pi_i^1 = \frac{1}{2}$, and for $i > 2p$ let $\pi_i^0 = 1 - \pi_i^1 = 1$. It can be easily checked that $B = \{b_1, \dots, b_{2\ell}\}$ is a perfect committee. Now, $A = \{a_1, \dots, a_{2\ell}\}$ is locally optimal. To check this, we consider two cases: in the first case, where $(r \leq \ell$ and $q \leq \ell)$ or $(r > \ell$ and $q > \ell)$, replacing a_r with b_q does not change the distance to the target distribution on each of the first p attributes, increases the distance on each of the next p attributes and decreases the distance on each of the last p attributes. For the second case, where $(r \leq \ell$ and $q > \ell)$ or $(r > \ell; q \leq \ell)$, the line of reasoning is similar. Finally, $\sum_{i,j} |r_i^j(A) - \pi_i^j| = 2p = \frac{2}{3}|X|$. \square

A better approximation bound can be obtained with $\ell = 2$, yet the analysis of this case is much more involved.

Lemma 2. Consider n buckets X_1, \dots, X_n , such that in the i -th bucket X_i there are x_i white balls and y_i black balls. Let A denote the number of pairs of balls such that both balls in the pair belong to the same bucket and are of different colour. Let us consider the procedure in which one iteratively selects a bucket and takes out two balls with different colours from the selected bucket. The procedure ends after B steps, when no further steps are possible (in each bucket, either there are no balls anymore, or all balls have the same colour). It holds that $A \geq \frac{B^2}{n}$.

Proof. Without loss of generality let us assume that for each i : $x_i \leq y_i$. Thus, $B = \sum_i x_i$ and $A = \sum_i x_i y_i \leq \sum_i x_i^2$. The inequality $\sum_i x_i^2 \geq \frac{(\sum_i x_i)^2}{n}$ follows from Jensen’s inequality [17] applied to the quadratic function. \square

Lemma 3. Let $x_i, y_i, A_i, 1 \leq i \leq n$, be real values satisfying the following constraints:

1. $x_i \geq \frac{A_i}{2n-2(i-1)}$, for each $1 \leq i \leq n$,
2. $A_i \geq A_{i-1} - 2x_{i-1}$, for each $2 \leq i \leq n$,
3. $y_i \geq \frac{x_i}{2n-2(i-1)-1}$, for each $1 \leq i \leq n$.

Then:

$$\sum_{i=1}^n y_i \geq \frac{|A_1| \ln n}{4n}.$$

Proof. We can view the set of above inequalities 1, 2, 3 as a linear program with $(3n - 1)$ variables (all x_i and y_i for $1 \leq i \leq n$ and A_i for $2 \leq i \leq n$; we treat A_1 as a constant) and $(3n - 1)$ constraints. Thus, we know that $\sum_i y_i$ achieves the minimum when each from the above constraints is satisfied with equality.

We show by induction that the values $x_i = \frac{A_1}{2n}$ and $A_i = \frac{2n-2(i-1)}{2n} A_1$ constitute the solution to the set of equalities that is derived by taking constraints 1, and 2, and treating them as equalities. We can show that by induction: First, consider the base step, i.e., the case when $i = 1$. Since constraint 2 is defined only for $i \geq 2$, we need to check only constraint 1. This constraint written in the form of equality gives us:

$$x_1 = \frac{A_1}{2n - 2(i - 1)} = \frac{|A_1|}{2n},$$

which proves that our hypothesis holds for $i = 1$. Next, let us assume that from the equalities 1 and 2 taken for $i < j$, it follows that $x_i = \frac{A_1}{2n}$ and $A_i = \frac{2n-2(i-1)}{2n} A_1$, for $i < j$. We will show that from equalities 1 and 2 for $i = j$ it follows that $x_j = \frac{A_1}{2n}$ and $A_j = \frac{2n-2(j-1)}{2n} A_1$:

$$x_j = \frac{A_j}{2n - 2(j - 1)} = \frac{1}{2n - 2(j - 1)} \cdot \frac{2n - 2(j - 1)}{2n} A_1 = \frac{|A_1|}{2n},$$

$$A_j = A_{j-1} - 2x_{j-1} = \frac{2n - 2((j - 1) - 1)}{2n} A_1 - 2 \frac{|A_1|}{2n} = \frac{2n - 2(j - 1)}{2n} A_1.$$

From constraint 3, treated as equality, we get:

$$y_i = \frac{x_i}{2n - 2(i - 1) - 1} = \frac{|A_1|}{2n(2n - 2(i - 1) - 1)}.$$

Thus, we infer that $\sum_{i=1}^n y_i$ is minimised when $y_i = \frac{|A_1|}{2n(2n-2(i-1)-1)}$. We recall that H_n denotes the n -th harmonic number ($H_n = \sum_{i=1}^n \frac{1}{i}$), and that $\ln(n + 1) < H_n \leq 1 + \ln(n)$. As a result we get:

$$\sum_{i=1}^n y_i \geq \frac{A_1}{2n} \sum_{i=1}^n \frac{1}{(2n - 2(i - 1) - 1)} \geq \frac{A_1}{2n} \sum_{i=1}^n \frac{1}{2n - 2(i - 1)} \tag{5}$$

$$= \frac{A_1}{4n} \sum_{i=1}^n \frac{1}{(n - i + 1)} = \frac{A_1}{4n} H_n \geq A_1 \frac{\ln n}{4n}. \quad \square \tag{6}$$

Theorem 5. For binary domains and natural distributions the local search algorithm from Fig. 3 with $\ell = 2$ is a $\frac{\ln(k/2)}{2 \ln(k/2) - 1} (|X| + \frac{6|X|}{k})$ -additive-approximation algorithm for OPTIMALHAMILTONREPRESENTATION.

Proof. In this proof we use similar idea to the proof of Theorem 4, but the proof is technically more involved. As before, by A^* and A we denote the optimal solution and the solution returned by the local search algorithm, respectively. Similarly to the previous proof, by $X_{\text{ex}} \subset X$ we denote the set of all attributes for which A achieves exact match with π , i.e.,

$$X_{\text{ex}} = \{X_i \in X : r_i^1(A) = \pi_i^1\}.$$

We also define the set $X_{\text{aex}} \subset X$ of all attributes for which A achieves almost exact match with π , i.e.,

$$X_{\text{aex}} = \left\{ X_i \in X : |r_i^1(A) - \pi_i^1| \leq \frac{1}{k} \right\}.$$

Let $q_f = \frac{|A \setminus A^*|}{2}$ and $q = \lfloor q_f \rfloor$. Let us rename the candidates from $A \setminus A^*$ so that $A \setminus A^* = \{a_1, a_2, \dots, a_{2q_f}\}$, and the candidates from $A^* \setminus A$, so that $A^* \setminus A = \{a'_1, a'_2, \dots, a'_{2q_f}\}$. Hereinafter, we follow a convention in which the elements from $A^* \setminus A$ are marked with primes. Renaming of the candidates that we described above, allows us to define the following sequence of pairs $(a_1, a'_1), \dots, (a_{2q_f}, a'_{2q_f})$ in which each element from $A \setminus A^*$ is paired with (assigned to) exactly one element from $A^* \setminus A$.

For each pair (a_j, a'_j) and for each attribute X_i we consider what happens if we replace a_i in $A \setminus A^*$ with a'_i . One of three scenarios can happen, after such a replacement:

1. The value $r_i^0(A)$ can increase by $\frac{1}{k}$ (in this case $r_i^1(A)$ decreases by $\frac{1}{k}$), which we denote by $X_i(a_j \leftrightarrow a'_j) = 1$,
2. The value $r_i^0(A)$ can decrease by $\frac{1}{k}$ (in this case $r_i^1(A)$ increases by $\frac{1}{k}$), which we denote by $X_i(a_j \leftrightarrow a'_j) = -1$, or
3. The value $r_i^0(A)$ can remain unchanged (in this case $r_i^1(A)$ also remains unchanged), which we denote by $X_i(a_j \leftrightarrow a'_j) = 0$.

We follow a procedure which, in q consecutive steps, replaces pairs of candidates from $A \setminus A^*$, with the pairs of candidates from $A^* \setminus A$. A pair (a_i, a_j) is always replaced with (a'_i, a'_j) . In other words, when looking for a pair from $A^* \setminus A$ to replace (a_i, a_j) we follow the assignment rule induced by renaming, as described above. The way in which we create pairs

Table 2

An example illustrating the concept of annihilating pairs. In this example we have $X_{\text{ex}} = \{X_1, X_2, X_3, X_4, X_5, X_6, X_7\}$ and $\bar{A}_1 = \{a_1, a_2, a_3, a_4\}$. The cell in row “ $X_i(a_j \leftrightarrow a'_j)$ ” for $j \in [4]$ and in column “ $X_i = X_\ell$ ” for $\ell \in [7]$ denotes the value of $X_\ell(a_j \leftrightarrow a'_j)$. We recall that $X_i(a_i \leftrightarrow a'_i) = 1$ if replacing a_i with a'_i moves A further from the target distribution in one direction and $X_i(a_i \leftrightarrow a'_i) = -1$ if replacing a_i with a'_i moves A further from the target distribution in the other direction. Here, we have $W_1(X_1) = \{(a_1, X_1), (a_2, X_1), ((a_1, X_1), (a_4, X_1))\}$, $W_1(X_2) = \{(a_1, X_2), (a_2, X_2), ((a_1, X_2), (a_3, X_2))\}$, $W_1(X_3) = \{(a_1, X_3), (a_3, X_3), ((a_1, X_3), (a_4, X_3)), ((a_2, X_3), (a_3, X_3)), ((a_2, X_3), (a_4, X_3))\}$, etc. Further, $W_1 = W_1(X_1) \cup W_1(X_2) \cup W_1(X_3) \cup W_1(X_4) \cup W_1(X_5) \cup W_1(X_6) \cup W_1(X_7)$. There are many choices for the set W , but it must hold that $P = |W| = 6$; we give the following example: $W = \{(a_1, X_1), (a_2, X_1), ((a_1, X_2), (a_2, X_2)), ((a_1, X_3), (a_3, X_3)), ((a_2, X_3), (a_4, X_3)), ((a_1, X_4), (a_4, X_4)), ((a_1, X_7), (a_3, X_7))\}$.

	$X_i = X_1$	$X_i = X_2$	$X_i = X_3$	$X_i = X_4$	$X_i = X_5$	$X_i = X_6$	$X_i = X_7$
$X_i(a_1 \leftrightarrow a'_1)$	1	1	1	1	0	0	-1
$X_i(a_2 \leftrightarrow a'_2)$	-1	-1	1	0	0	1	0
$X_i(a_3 \leftrightarrow a'_3)$	0	-1	-1	0	1	0	1
$X_i(a_4 \leftrightarrow a'_4)$	-1	1	-1	-1	1	0	-1

within $A \setminus A^*$ for replacement (the way how (a_i, a_j) is selected in each of q consecutive steps) will be described later. After this whole procedure A can differ from A^* with at most one element, hence, having distance to the optimal distribution at most equal to $|X| \frac{2}{k}$. Let us define the sequence of sets $\bar{A}_1, \bar{A}_2, \dots, \bar{A}_q$ in the following way: we define $\bar{A}_1 = A \setminus A^*$, and we define \bar{A}_{j+1} as \bar{A}_j after removing the pair from $A \setminus A^*$ that was used in replacement in the j -th step of our procedure.

As before, for each $B \subseteq A \setminus A^*$ and $B' \subseteq A^* \setminus A$, and for each attribute $X_i \in X$ we define the gain $g_i(B, B')$:

$$g_i(B, B') = \sum_{j \in \{1, 2\}} (|r_i^j(A) - \pi_i^j| - |r_i^j((A \setminus B) \cup B') - \pi_i^j|).$$

Similarly as in the proof of Theorem 4, we observe that for $X_i \notin X_{\text{aex}}$ the negative gains cumulate: i.e., that for each sequence of disjoint sets B_1, B_2, \dots, B_s and B'_1, B'_2, \dots, B'_s such that for every $1 \leq j \leq s$, $B_j \subseteq A \setminus A^*$, $B'_j \subseteq A^* \setminus A$, and $|B_j| = |B'_j| \leq 2$ we have that:

$$g_i\left(\bigcup_j B_j, \bigcup_j B'_j\right) \leq \sum_j g_i(B_j, B'_j). \tag{7}$$

Why is this the case? If $X_i \notin X_{\text{aex}}$, then the distance between A and the target distribution on attribute X_i is at least equal to $2 \cdot \frac{2}{k}$. In other words: $|r_i^0(A) - \pi_i^0| \geq \frac{2}{k}$ and $|r_i^1(A) - \pi_i^1| \geq \frac{2}{k}$. Without loss of generality let us assume that $r_i^0(A) - \pi_i^0 \geq \frac{2}{k}$. Since each set B_j and each set B'_j has at most two elements, replacing B_j with B'_j can change the distance between A and the target distribution, for each attribute, by at most $\frac{2}{k}$. Consequently, if $g_i(B_j, B'_j)$ is negative, then it means that replacing B_j with B'_j makes the difference $r_i^0(A) - \pi_i^0$ even greater. Thus, each such replacement with the negative gain g causes A to move further from the target distribution by the value g . Naturally, each replacement with the positive gain g causes A to move closer to the target distribution by at most g . Consequently, after the sequence of replacement $\cup_j B_j \leftrightarrow B'_j$ the distance on the attribute X_i cannot improve by more than $\sum_j g_i(B_j, B'_j)$.

In contrast to the proof of Theorem 4, we note that here we require that $X_i \notin X_{\text{aex}}$ instead of $X_i \notin X_{\text{ex}}$ – the above observation is not valid if $X_i \in X_{\text{aex}}$ even if $X_i \notin X_{\text{ex}}$.¹⁰

Next, for each \bar{A}_j , and each attribute $X_i \in X_{\text{ex}}$, we define a set W_j of annihilating pairs as:

$$W_j(X_i) = \{((a_x, X_i), (a_y, X_i)) : a_x \in \bar{A}_j; a_y \in \bar{A}_j; x < y; X_i(a_x \leftrightarrow a'_x) = -X_i(a_y \leftrightarrow a'_y)\}.$$

Intuitively, if $((a_x, X_i), (a_y, X_i)) \in W_j$, then both replacing a_x with a'_x and replacing a_y with a'_y move the original set A (i.e., the set before any of the replacements) further from the target distribution for the attribute X_i , but replacing $\{a_x, a_y\}$ with $\{a'_x, a'_y\}$ does not change the distance of A from the target distribution for the attribute X_i .

For each j , we set $W_j = \cup_{i \in X_{\text{ex}}} W_j(X_i)$. Let us denote by P the number of annihilated pairs of candidates considered in the process of replacing candidates from $A \setminus A^*$ with candidates from $A^* \setminus A$. Formally, P is the size of the maximal subset $W \subseteq W_1$ composed of disjoint annihilating pairs, i.e., for each $i \leq p$, for each a_x , and for each a_y , if $((a_x, X_i), (a_y, X_i)) \in W$ then there exists no $b \neq a_y$ such that $((a_x, X_i), (b, X_i)) \in W$ or $((b, X_i), (a_x, X_i)) \in W$. From Lemma 2, after defining each bucket X_i as containing x_i white balls and y_i black balls, where x_i (respectively, y_i) is the number of candidates $a_j \in \bar{A}_1$ with the value $X_i(a_j \leftrightarrow a'_j)$ equal to 1 (respectively, -1), it follows that $W_1 \geq \frac{p^2}{|X_{\text{ex}}|}$. The concept of annihilating pairs is explained on example in Table 2.

¹⁰ Consider an example in which $\pi_i^1 = \frac{1}{k}$ and $r_i^1(A) = \frac{2}{k}$. Let us consider sets $B = \{b_1, b_2\}$, $B' = \{b'_1, b'_2\}$, $C = \{c_1, c_2\}$, $C' = \{c'_1, c'_2\}$ such that: $X_i(c_1) = X_i(c_2) = X_i(b'_1) = X_i(b'_2) = d_1^1$, and $X_i(c'_1) = X_i(c'_2) = X_i(b_1) = X_i(b_2) = d_2^1$. Thus, we have that:

- Replacing B with B' results with $r_i^1(A) = \frac{4}{k}$.
- Replacing C with C' results with $r_i^1(A) = 0$.
- Replacing $B \cup C$ with $B' \cup C'$ results with $r_i^1(A) = \frac{2}{k}$.

We can repeat this reasoning for $r_i^2(A)$, thus having, $g_i(B, B') = -\frac{4}{k}$, $g_i(C, C') = 0$ and $g_i(B \cup C, B' \cup C') = 0$.




Fig. 4. Figure illustrating that for $X_i \in X_{ex}$, $g_i(\{a_1, a_2\}, \{a'_1, a'_2\})$ is greater than $(g_i(a_1, a'_1) + g_i(a_2, a'_2))$ if and only if $((a_1, X_i), (a_2, X_i))$ is an annihilating pair. The figure presents 3 scenarios: a) $((a_1, X_i), (a_2, X_i))$ is an annihilating pair. Both replacing a_1 with a'_1 and replacing a_2 with a'_2 moves us further from the target distribution for attribute X_i (the target distribution is marked as a black dot), thus $g_i(a_1, a'_1) = -\frac{2}{k}$ and $g_i(a_2, a'_2) = -\frac{2}{k}$. However these changes annihilate, and $g_i(\{a_1, a_2\}, \{a'_1, a'_2\}) = 0$. b) $g_i(a_1, a'_1) = -\frac{2}{k}$ and $g_i(a_2, a'_2) = -\frac{2}{k}$, but these changes do not annihilate, and thus: $g_i(\{a_1, a_2\}, \{a'_1, a'_2\}) = -\frac{4}{k}$. c) $g_i(a_1, a'_1) = -\frac{2}{k}$ and $g_i(a_2, a'_2) = 0$, if at least one change does not move the solution against the target distribution, the changes do not annihilate, and $g_i(\{a_1, a_2\}, \{a'_1, a'_2\}) = g_i(a_1, a'_1) + g_i(a_2, a'_2)$.

We are now ready to describe the way in which we select pairs from $A \setminus A^*$ in our procedure. In each step j , the pair $(a_{j,1}, a_{j,2})$ from $A \setminus A^*$ is selected in the following way. For each candidate a let $s_{j,1}(a)$ be the number of pairs p in W_j such that $p = ((a, \cdot), (\cdot, \cdot))$ or $p = ((\cdot, \cdot), (a, \cdot))$, let $a_{j,1}$ be such that $s_{j,1}(a_{j,1}) = \max_{a \in \bar{A}_j} s_{j,1}(a)$, and let $s_{j,1} = s_{j,1}(a_{j,1})$. Next, for each candidate b let $s_{j,2}(b)$ be the number of pairs p in W_j such that $p = ((a_{j,1}, \cdot), (b, \cdot))$ or $p = ((b, \cdot), (a_{j,1}, \cdot))$, let $a_{j,2}$ be such that $s_{j,2}(a_{j,2}) = \max_{b \in \bar{A}_j} s_{j,2}(b)$, and let $s_{j,2} = s_{j,2}(a_{j,2})$.

Let us consider the procedure described above on the example from Table 2. The candidate a_1 belongs to 8 pairs in W_1 (a_1 belongs to 2 pairs for attribute X_1, X_2 , and X_3 , and to one pair for attributes X_4 and X_7), thus: $s_{1,1}(a_1) = 8$. Moreover, $s_{1,1}(a_2) = 5$, $s_{1,1}(a_3) = 6$, and $s_{1,1}(a_4) = 7$. Consequently, a_1 will be the candidate that will be replaced with a'_1 in the first step: $a_{j,1} = a_1$ and $s_{j,1} = 8$. Further, $s_{1,2}(a_2) = 2$ (there are two annihilating pairs including a_1 and a_2 , i.e.: $((a_1, X_1), (a_2, X_1))$ and $((a_1, X_2), (a_2, X_2))$); similarly: $s_{1,2}(a_3) = 3$, and $s_{1,2}(a_4) = 3$. Thus, an arbitrary of the two candidates, a_3 and a_4 , say a_3 , will be the second candidate that will be replaced with a'_3 in the first step. In the second step only two candidates, a_2 and a_4 , are left, so both will be replaced with a'_2 and a'_4 in the second step. Nevertheless, let us illustrate our definitions also in the second step of the replacement procedure. The set \bar{A}_2 consists of two remaining candidates: a_2 and a_4 . We have $W_2 = \{((a_2, X_2), (a_4, X_2)), ((a_2, X_3), (a_4, X_3))\}$. Naturally, $s_{2,1}(a_2) = s_{2,1}(a_4) = s_{2,2}(a_2) = s_{2,2}(a_4) = 2$.

We want now to derive bounds on the values $s_{j,1}$ and $s_{j,2}$. The following inequalities hold:

$$1. s_{j,1} \geq \frac{2|W_j|}{2q_f - 2(j-1)} \text{ for each } 1 \leq j \leq q.$$

W_j contains pairs of candidates belonging to \bar{A}_j . \bar{A}_1 has $2q_f$ candidates, and \bar{A}_{j+1} is obtained from \bar{A}_j by removing two candidates. Consequently, \bar{A}_j has $2q_f - 2(j-1)$ candidates, and thus, W_j contains pairs of $2q_f - 2(j-1)$ different candidates. From the pigeonhole principle it follows that there exists a candidate that belongs to at least $\frac{2|W_j|}{2q_f - 2(j-1)}$ pairs. Naturally, we also get the weaker constraint: $s_{j,1} \geq \frac{|W_j|}{2q_f - 2(j-1)}$.

$$2. |W_j| \geq |W_{j-1}| - 2s_{j-1,1} \text{ for each } 2 \leq j \leq q.$$

Each candidate in W_{j-1} belongs to at most $s_{j-1,1}$ pairs (this follows from the definition of $s_{j-1,1}$). W_j contains all pairs that W_{j-1} contained, except for the pairs involving $a_{j-1,1}, a_{j-2,2}$ (to obtain \bar{A}_j , we removed these two candidates from \bar{A}_{j-1}). Consequently, W_j is obtained from W_{j-1} by removing at most $2s_{j-1,1}$ pairs of candidates.

$$3. s_{j,2} \geq \frac{s_{j,1}}{2q_f - 2(j-1) - 1} \text{ for each } 1 \leq j \leq q.$$

In W_j , there are $s_{j,1}$ pairs of candidates involving $a_{j,1}$. As we noted before, W_j contains pairs of $2q_f - 2(j-1)$ different candidates. Thus, in W_j , $a_{j,1}$ is paired with at most $2q_f - 2(j-1) - 1$ candidates. From the pigeonhole principle it follows that $a_{j,1}$ must be paired with some candidate at least $\frac{s_{j,1}}{2q_f - 2(j-1) - 1}$ times.

From Lemma 3 we get that:

$$\sum_{j=1}^q s_{j,2} \geq \frac{|W_1| \ln q}{4q}. \tag{8}$$

Before we proceed further let us make three observations regarding annihilating pairs. First, we note that for each $X_i \in X_{ex}$, and each a_x and a_y , if the value $g_i(\{a_x, a_y\}, \{a'_x, a'_y\})$ is different from $(g_i(a_x, a'_x) + g_i(a_y, a'_y))$ then it is greater from $(g_i(a_x, a'_x) + g_i(a_y, a'_y))$ by $\frac{4}{k}$. We also note that $g_i(\{a_x, a_y\}, \{a'_x, a'_y\})$ is greater than $(g_i(a_x, a'_x) + g_i(a_y, a'_y))$ if and only if the changes $X_i(a_x \leftrightarrow a'_x)$ and $X_i(a_y \leftrightarrow a'_y)$ annihilate (this is illustrated in Fig. 4). Further, we recall that the value $s_{j,2}$ counts all attributes for which $a_{j,1}$ and $a_{j,2}$ constitute an annihilating pair. Thus, for each $1 \leq j \leq q$:




Fig. 5. Figure illustrating the effect of replacing 10 candidates for an attribute $X_i \in X_{\text{ex}}$. Each replacement imposes a negative gain: $g_i(a_j, a'_j) = -\frac{2}{k}$ for $1 \leq j \leq 10$. Thus, $\sum_{a \in A \setminus A^*} g_i(a, a') = -\frac{20}{k}$. In this example four pairs annihilated, and, consequently, $g_i(A \setminus A^*, A^* \setminus A) = -\frac{4}{k}$.

$$\sum_{i \in X_{\text{ex}}} g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) = \sum_{i \in X_{\text{ex}}} \left(g_i(a_{j,1}, a'_{j,1}) + g_i(a_{j,2}, a'_{j,2}) \right) + s_{j,2} \frac{4}{k} \tag{9}$$

Our second observation is similar in spirit to the first one. We note that for each $X_i \in X_{\text{ex}}$:

$$g_i(A \setminus A^*, A^* \setminus A) - \sum_{a \in A \setminus A^*} g_i(a, a') = \text{the number of pairs that annihilated for } X_i \times \frac{4}{k}.$$

The above equality is illustrated in Fig. 5. As a consequence, we get that:

$$\sum_{X_i \in X_{\text{ex}}} \left(g_i(A \setminus A^*, A^* \setminus A) - \sum_{a \in A \setminus A^*} g_i(a, a') \right) = \text{the number of pairs that annihilated} \times \frac{4}{k}.$$

We recall that after the replacement procedure A can differ from A^* with at most one element, hence, having distance to the optimal distribution at most equal to $|X| \frac{2}{k}$. Thus:

$$\sum_{X_i \in X_{\text{ex}}} \left(g_i(A \setminus A^*, A^* \setminus A) - \sum_{j=1}^q \left(g_i(a_{j,1}, a'_{j,1}) + g_i(a_{j,2}, a'_{j,2}) \right) \right) \leq P \cdot \frac{4}{k} + |X| \frac{2}{k}. \tag{10}$$

Our third observation says that:

$$\sum_{X_i \in X_{\text{aex}} \setminus X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) - \sum_{X_i \in X_{\text{aex}} \setminus X_{\text{ex}}} \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) \leq |X_{\text{aex}} \setminus X_{\text{ex}}|. \tag{11}$$

Where does Inequality (11) come from? Let us use the geometric interpretation, like the one from Fig. 5. Let us consider an $X_i, X_j \in X_{\text{aex}}$. For X_i, A lies in a distance of $\frac{2}{k}$ on the left or on the right from the target distribution. Without loss of generality, let us assume it lies on the right. Now, if $g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) < 0$ then replacing $(a_{j,1}, a_{j,2})$ with $(a'_{j,1}, a'_{j,2})$ moves the current solution right. If $g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) = \frac{2}{k}$, then replacing $(a_{j,1}, a_{j,2})$ with $(a'_{j,1}, a'_{j,2})$ moves the current solution by $\frac{2}{k}$ on left. If $g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) = 0$, then replacing $(a_{j,1}, a_{j,2})$ with $(a'_{j,1}, a'_{j,2})$ either does not move the solution or moves it by $\frac{4}{k}$ on left.

Let us define $y_i = g_i(A \setminus A^*, A^* \setminus A) - \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\})$. If the solution moves q times to the right, then the total gain $-\sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\})$ will be maximised, achieving $q \frac{4}{k}$. In such a case however, the value $g_i(A \setminus A^*, A^* \setminus A)$ will be equal to $-q \frac{4}{k}$, and thus the value y_i will be equal to 0. After some consideration, the reader will see that the value y_i is maximised if the current solution moves $\frac{q}{2}$ times right and $\frac{q}{2}$ times left, each time by the value of $\frac{4}{k}$. This way, the moves to the right induce the total gain of $\frac{q}{2} \cdot \frac{4}{k}$, the moves to the left induce the zero gain, but as a consequence, the current solution for X_i does not change ($g_i(A \setminus A^*, A^* \setminus A) = 0$). Thus, for each $X_i \in X_{\text{aex}}$, y_i is upper bounded by $\frac{q}{2} \cdot \frac{4}{k} \leq 1$, which proves Inequality (11).

We can further proceed with the proof by observing that from the condition in the “while” loop we get that for each $1 \leq j \leq q$:

$$\begin{aligned} 0 &\geq \sum_i g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) \\ &\geq \sum_{i \in X_{\text{ex}}} g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) + \sum_{i \notin X_{\text{ex}}} g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) \end{aligned}$$

From (9):

$$\geq \sum_{i \in X_{\text{ex}}} \left(g_i(a_{j,1}, a'_{j,1}) + g_i(a_{j,2}, a'_{j,2}) \right) + s_{j,2} \frac{4}{k} + \sum_{i \notin X_{\text{ex}}} g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}).$$

Thus, we get:

$$- \sum_{i \in X_{\text{ex}}} \left(g_i(a_{j,1}, a'_{j,1}) + g_i(a_{j,2}, a'_{j,2}) \right) - \frac{4}{k} s_{j,2} > + \sum_{i \notin X_{\text{ex}}} g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}). \tag{12}$$

Next, we give the following sequence of inequalities:

$$\begin{aligned} g &= \sum_i g_i(A \setminus A^*, A^* \setminus A) \\ &= \sum_{X_i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) + \sum_{X_i \in X_{\text{aex}} \setminus X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) + \sum_{X_i \notin X_{\text{aex}}} g_i(A \setminus A^*, A^* \setminus A) \end{aligned}$$

From Inequality (7), for all $i \notin X_{\text{aex}}$, we have $g_i(A \setminus A^*, A^* \setminus A) \leq \sum_{a \in A \setminus A^*} g_i(a, a')$. Since the set $A \setminus A^*$ and $\cup_{j=1}^q \{a_{j,1}, a_{j,2}\}$ can differ by at most one candidate (which induces distance $\frac{2|X|}{k}$ to the optimal solution), we have that

$$\sum_{X_i \notin X_{\text{aex}}} g_i(A \setminus A^*, A^* \setminus A) \leq \sum_{X_i \notin X_{\text{aex}}} \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) + \frac{2|X|}{k}.$$

And, as a consequence:

$$\begin{aligned} g &\leq \sum_{X_i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) + \sum_{X_i \in X_{\text{aex}} \setminus X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) \\ &\quad + \sum_{X_i \notin X_{\text{aex}}} \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) + \frac{2|X|}{k} \\ &\leq \sum_{X_i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) + \sum_{X_i \in X_{\text{aex}} \setminus X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) \\ &\quad + \sum_{X_i \notin X_{\text{ex}}} \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) - \sum_{X_i \in X_{\text{aex}} \setminus X_{\text{ex}}} \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) + \frac{2|X|}{k}. \end{aligned}$$

From Inequality (11) we get:

$$g \leq \sum_{X_i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) + \sum_{X_i \notin X_{\text{ex}}} \sum_{j=1}^q g_i(\{a_{j,1}, a_{j,2}\}, \{a'_{j,1}, a'_{j,2}\}) + \frac{2|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}|.$$

From Inequality (12):

$$\begin{aligned} g &\leq \frac{2|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| + \sum_{X_i \in X_{\text{ex}}} g_i(A \setminus A^*, A^* \setminus A) \\ &\quad - \sum_{X_i \in X_{\text{ex}}} \sum_{j=1}^q \left(g_i(a_{j,1}, a'_{j,1}) + g_i(a_{j,2}, a'_{j,2}) \right) - \frac{4}{k} \sum_j s_{j,2}. \end{aligned}$$

From Inequality (8):

$$\begin{aligned} g &\leq \frac{2|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| - \frac{|W_1| \ln q}{4q} \cdot \frac{4}{k} \\ &\quad + \sum_{i \in X_{\text{ex}}} \left(g_i(A \setminus A^*, A^* \setminus A) - \sum_{j=1}^q \left(g_i(a_{j,1}, a'_{j,1}) + g_i(a_{j,2}, a'_{j,2}) \right) \right) \end{aligned}$$

From Inequality (10):

$$g \leq \frac{4|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| - \frac{|W_1| \ln q}{kq} + P \frac{4}{k}.$$

As we noted before, from Lemma 2, we have that $W_1 \geq \frac{P^2}{|X_{\text{ex}}|}$. Thus:

$$g \leq \frac{4|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| + \frac{4}{k} \left(P - \frac{P^2 \ln q}{4|X_{\text{ex}}|q} \right).$$

Since $q \leq \frac{k}{2}$, and since the function $\frac{\ln x}{x}$ is decreasing for $x \geq 1$:

$$g \leq \frac{4|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| + \frac{4}{k} \left(P - \frac{P^2 \ln(k/2)}{2|X_{\text{ex}}|k} \right)$$

The function $f(P) = P - \frac{P^2 \ln(k/2)}{2|X_{\text{ex}}|k}$ takes its maximum for $P = \frac{|X_{\text{ex}}|k}{\ln(k/2)}$. Thus:

$$g \leq \frac{4|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| + \frac{4}{k} \cdot \frac{|X_{\text{ex}}|k}{2 \ln(k/2)} = \frac{4|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| + \frac{2|X_{\text{ex}}|}{\ln(k/2)}.$$

Since our local-search algorithm for $\ell = 2$ also tries to perform local swaps on single candidates, we can repeat the analysis from the proof of Theorem 4. Thus, using Inequality (4) from there, we get that $g \leq 2|X_{\text{ex}}|$, and as a consequence: $\left(\frac{1}{2} - \frac{1}{\ln(k/2)}\right)g \leq |X_{\text{ex}}| - \frac{2|X_{\text{ex}}|}{\ln(k/2)}$.

For each attribute $X_i \in X \setminus X_{\text{aex}}$ the distance from A and the target distribution is bounded by 2. For $X_i \in X_{\text{aex}}$ this distance is bounded by $\frac{2}{k}$. Thus, we get that $g \leq 2(|X| - |X_{\text{ex}}| - |X_{\text{aex}} \setminus X_{\text{ex}}|) + |X|\frac{2}{k}$, and so:

$$\begin{aligned} g + \left(\frac{1}{2} - \frac{1}{\ln(k/2)}\right)g + \frac{1}{2}g &\leq \frac{4|X|}{k} + |X_{\text{aex}} \setminus X_{\text{ex}}| + \frac{2|X_{\text{ex}}|}{\ln(k/2)} \\ &\quad + |X_{\text{ex}}| - \frac{2|X_{\text{ex}}|}{\ln(k/2)} \\ &\quad + (|X| - |X_{\text{ex}}| - |X_{\text{aex}} \setminus X_{\text{ex}}|) + |X|\frac{2}{k} \\ &= |X| + \frac{6|X|}{k} \end{aligned}$$

Finally, we get:

$$g \leq \frac{\ln(k/2)}{2 \ln(k/2) - 1} \left(|X| + \frac{6|X|}{k} \right)$$

which completes the proof. \square

Since a brute-force algorithm can be used to compute an optimal solution for small values of k , Theorem 5 implies that for every $\epsilon > 0$ we can achieve an additive approximation of $\frac{1}{2}(|X| + \epsilon)$. That is, we can guarantee that the solution returned by our algorithm will be at least 4 times better than a solution that is arbitrarily bad on each attribute. A natural open question is whether the local search algorithm achieves even better approximation guarantees for larger values of ℓ .

One may argue that the restriction to natural target distributions is quite strong. However, for a given vector of target distributions π , we can easily find a vector ρ of target natural distributions such that $\sum_{i,j} |\rho_i^j - \pi_i^j| \leq \frac{2|X|}{k}$. For instance for $k = 5$ and $p = |X| = 3$ the distribution

$$\pi = \left(\left(\frac{2}{5} + \frac{1}{10}, \frac{3}{5} - \frac{1}{10} \right), \left(\frac{1}{5} + \frac{1}{7}, \frac{4}{5} - \frac{1}{7} \right), \left(\frac{1}{6}, 1 - \frac{1}{6} \right) \right)$$

is not natural, yet there exists a natural distribution

$$\rho = \left(\left(\frac{2}{5}, \frac{3}{5} \right), \left(\frac{1}{5}, \frac{4}{5} \right), (0, 1) \right)$$

such that $\sum_{i,j} |\rho_i^j - \pi_i^j| \leq \frac{2|X|}{k}$. Thus, the results from Theorem 4 and Theorem 5 can be modified by providing approximation ratios that are worse by an additive value of $\frac{2|X|}{k}$ but valid for arbitrary target distributions. Again, since an optimal solution can easily be computed for small values of k , we can get approximation guarantees arbitrarily close to the ones given by Theorem 4 and Theorem 5, even for non-natural target distributions.

Below we show a lower bound of $\frac{2|X|}{7}$ for the approximation ratio of the local search algorithm from Fig. 3 with $\ell = 2$.

Example 8. Consider 7 binary attributes X_1, \dots, X_7 , and the set of $12p$ candidates $C = \{a_1, \dots, a_{2p}, a'_1, \dots, a'_{2p}, b_1, \dots, b_{2p}, b'_1, \dots, b'_{2p}, c_1, \dots, c_{2p}, c'_1, \dots, c'_{2p}\}$. For each $i \in [k]$, we have:

	X_1	X_2	X_3	X_4	X_5	X_6	X_7
a_i	1	0	1	1	0	0	1
a'_i	0	1	0	0	1	1	1
b_i	0	0	0	0	0	0	0
b'_i	0	0	1	1	1	1	0
c_i	1	1	1	1	0	0	0
c'_i	1	1	0	0	1	1	0

We note that for each candidate the value of the attribute X_3 is the same as of X_4 and the value of the attribute X_5 is the same as of X_6 . For $i \in \{1, 2, 3, 4, 5, 6\}$ let $\pi_i^0 = \pi_i^1 = \frac{1}{2}$, and let $\pi_7^0 = 1 - \pi_7^1 = 1$.

Let us fix $k = 4p$. It can be easily checked that the set consisting of p copies of candidates b_i, b'_i, c_i, c'_i is a perfect committee. On the other hand, the set A consisting of $2p$ copies of candidates a_i and a'_i is locally optimal. Indeed, replacing candidate a_i or a'_i with b_i or b'_i moves the solution closer to the target distribution on X_7 , but the further from the target distribution on X_1 or X_2 . The same situation happens if we replace candidates a_i or a'_i with c_i or c'_i . If we replace two a -candidates with the pair consisting of one b -candidate (b_i or b'_i) and one c -candidate (c_i or c'_i), then such a replacement will move the solution closer by $4/k$ to the target distribution on X_7 , but will move the solution further by $2/k$ on two attributes from $\{X_3, X_4, X_5, X_6\}$.

Finally, $\sum_{i,j} |r_i^j(A) - \pi_i^j| = 2p = \frac{2}{7}|X|$. \square

6. Related work

Our model is related to the following research areas:

6.1. Apportionment for party-list representation systems

As we already pointed out, classical apportionment methods correspond to the restriction of our model to a single attribute (albeit with a different motivation). See the work of Balinski and Young [3] for a survey. While voting on multi-attribute domains and multiwinner elections have led to significant research effort in computational social choice, this is less the case for party-list representation systems. Ding and Lin [20] studied a game-theoretic model for a party-list proportional representation system under specific assumptions, and show that computing the Nash equilibria of the game is NP-hard.

6.2. Biapportionment

The biapportionment setting [4,5] has some similarities with our multi-attribute proportional representation setting (MAPR). In biapportionment we are given two attributes, one corresponding to parties and the other one to voting districts. The input consists of (1) hard constraints expressing lower and upper bounds on the number of candidates to be elected in each district, and similarly, bounds on the number of candidates to be elected from each party; (2) for each district i and party P_j , a value p_{ij} corresponding to the number of votes for party P_j in district i . (2) induces a soft proportionality constraint: the number of elected candidates from party P_j in district i should be as much as possible proportional to p_{ij} .

There are however substantial differences between biapportionment and MAPR. First, we do not have anything that corresponds to the values p_{ij} : while in biapportionment the target composition of the committee consists of a target number of seats for each combination of the two attributes, in MAPR, on the other hand, we have a smaller input consisting of a target number for each value of each attribute.¹¹ The second (and most important) difference between biapportionment and MAPR is that in MAPR we have a limited supply of available candidates characterised each by a tuple of attribute values: in our words, we focus on the case when the full supply assumption is not satisfied, which not only corresponds to the practical cases we have in mind, but is required in practice when the number of attributes is large. On the other hand, in biapportionment, it is implicitly assumed that there are enough candidates so that there always exists a solution satisfying given (often restrictive) hard constraints. Note finally that the computation of biapportionment methods has been investigated in a few recent papers [33,43,44].

6.3. Constrained approval voting

Constrained approval voting (CAP) [7,41] is also close to MAPR. In CAP there are also multiple attributes, candidates are represented by tuples of attribute values, there is a target composition of the committee and we try to find a committee close to this target. However, there are also substantial differences between MAPR and CAP. First, in CAP, like in

¹¹ Yet, a target number of seats for each combination of two or more attributes could be incorporated to our model by providing an attribute that corresponds to the Cartesian product of the given attributes; but, for combinatorial reasons, this ceases to be realistic for more than two or three attributes.

biapportionment, the target composition of the committee, exogenously defined, consists of a target number of seats for each combination of attributes (called a cell), that is, for each $\vec{z} \in D_1 \times \dots \times D_p$, we have a value $s(\vec{z})$; while in MAPR, as we said above, we have a smaller input consisting of a target number for each value of each attribute. Note that the input in CAP is exponentially large in the number of attributes, which makes it infeasible in practice as soon as this number exceeds a few units (probably CAP was designed for very small numbers of attributes). Second, in CAP, the selection criterion of an optimal committee is made in two consecutive steps: first a set of *admissible committees* is defined, and the choice between these admissible committees is made by using approval ballots, and the chosen committee is the admissible committee maximising the sum, over all voters, of the number of candidates approved (there are no target fractions as in MAPR). A simple translation of CAP into an integer linear programming problem is given in [41,48].

6.4. Voting on multi-attribute domains and judgement aggregation

Another interesting degenerated case is when $k = 1$, i.e., when we must select a single candidate from the database. The ideal case is when there exists a candidate in the database whose value on each attribute i coincides with the attribute value x_i^j . In this case, this candidate should certainly be selected; otherwise, the most representative candidate should be selected, for some measure of representativity.

This problem relates to voting in multi-attribute (or combinatorial) domains (cf. the recent survey chapter [32]). There, the aim is to output a single winning combination of attributes given the preferences of voters over combinations of attribute values, generally expressed in some compact form. When $k = 1$, our model can be viewed as a voting problem in a *constrained* multi-attribute domain (constrained because not all combinations are feasible). Another important difference is that in voting in multi-attribute domains, the focus is generally on the way of dealing with nonseparable preferences; here, the issue is avoided, as throughout our paper preferences are assumed to be separable.¹²

Our model also relates to judgement aggregation (see [24] for a recent survey). In judgement aggregation, there is a set of propositions $\{\varphi_1, \dots, \varphi_p\}$; the set of consistent (and complete) judgement sets is a subset \mathcal{J} of $\times_{i=1}^p \{\varphi_i, \neg\varphi_i\}$; a judgement aggregation profile $V = (V_1, \dots, V_n)$ is a collection of judgement sets from \mathcal{J} ; an irresolute judgement aggregation rule F maps a judgement aggregation profile to a nonempty subset of \mathcal{J} ; such a rule is said to be based on the weighted majoritarian judgement set if its output can be computed from the vector $\alpha_V = (\alpha_1, \dots, \alpha_p)$, where α_j is the proportion of judgement sets in V which contain φ_j .

Now, consider a multi-attribute proportional representation (MAPR) setting where all attributes are binary; we can view each attribute X_i as a proposition φ_i . Next, for each database candidate $c \in C$, the judgement set J_c is defined by $J_c = \{\varphi_i : i \in [p], X_i(c) = 1\} \cup \{\neg\varphi_i : i \in [p], X_i(c) = 0\}$. The set of consistent judgement sets \mathcal{J}_C is defined as $\mathcal{J}_C = \{J_c \mid c \in C\}$; in other words, $J \in \mathcal{J}_C$ is consistent if and only if there is a candidate c in C such that $(X_1(c), \dots, X_p(c))$ corresponds to J .

Finally, let $(\alpha_1, \dots, \alpha_p) = (\pi_1^1, \dots, \pi_1^p)$. Let \mathcal{R} be an irresolute MAPR rule; \mathcal{R} induces an irresolute judgement aggregation rule $F_{\mathcal{R}}$, based on the weighted majoritarian judgement set, defined by $F_{\mathcal{R}}(V) = \{J_c \mid \{c\} \in \mathcal{R}(C, \alpha_V, 1)\}$. Conversely, from a judgement aggregation rule F based on the weighted majoritarian judgement set we can define a MAPR rule \mathcal{R}_F restricted to $k = 1$, by $\mathcal{R}_F(C, \alpha, 1) = \{\{c\}, J_c \in F(V_\alpha)\}$. It is interesting to see which judgement aggregation rules correspond to the two MAPR rules we have defined when $k = 1$.

The *median* judgement aggregation rule¹³ is defined as follows: given a weighted majoritarian judgement set $\alpha_V = (\alpha_1, \dots, \alpha_p)$ and $J \in \mathcal{J}$, let $(J|\alpha) = \sum_{i, \varphi_i \in J} \alpha_i + \sum_{i, \neg\varphi_i \in J} (1 - \alpha_i)$. Then $median(\alpha) = \operatorname{argmax}_{J \in \mathcal{J}} (J|\alpha)$. Now, let C be a candidate database over a domain of binary attributes, and $k = 1$. Given α_V , we have $\{c\} \in \mathcal{R}_H(C, \alpha_V, 1)$ if $\sum_{i=1}^p \sum_{j=1,2} |\pi_i^j - X_i(c)|$ is minimum; now,

$$\begin{aligned} \sum_{i=1}^p \sum_{j=1,2} |\pi_i^j - X_i(c)| &= \sum_{i \in [p], X_i(c)=1} (1 - \alpha_i) + \sum_{i \in [p], X_i(c)=0} \alpha_i \\ &= p - \left(\sum_{i \in [p], X_i(c)=1} \alpha_i + \sum_{i \in [p], X_i(c)=0} (1 - \alpha_i) \right) = p - (J_c|\alpha), \end{aligned}$$

therefore $\mathcal{R}_H(\alpha_V)$ contains $\{c\}$ if $(J_c|\alpha)$ is maximum, that is, if $J_c \in median(\alpha)$.

¹² Extending our model to nonseparable preferences would consist in expressing preferences such as *if the gender ratio is 50–50 then the ideal group ratio is 40–30–30, otherwise 50–25–25*, or else *we want a gender ratio 50–50 or a seniority ratio 50–50*. We are not sure whether it is worth developing this generalisation.

¹³ This rule has been introduced independently in several different papers under different names, and it is probably not relevant to cite them here. A recent paper on the median rule, together with an axiomatisation, is [38].

The calculations for the multi-attribute d'Hondt rule are similar:

$$\begin{aligned} \sum_{i \in [p]} \sum_{j=1,2} \pi_i^j H(r_i^j(\{c\})) &= \sum_{i \in [p], X_i(c)=1} \pi_i^1 + \sum_{i \in [p], X_i(c)=0} \pi_i^0 \\ &= \sum_{i \in [p], X_i(c)=1} \alpha_i + \sum_{i \in [p], X_i(c)=0} 1 - \alpha_i = (J_c | \alpha). \end{aligned}$$

Therefore, $\mathcal{R}_{dHondt}(\alpha_V)$ contains $\{c\}$ if $J_c \in \text{median}(\alpha)$. In summary:

Observation 1. $F_{\mathcal{R}_H}$ and $F_{\mathcal{R}_{dHondt}}$ coincide with the median judgement aggregation rule.

6.5. Multiwinner (or committee) elections

In multiwinner elections the voters vote directly for candidates and do not consider attributes that characterise them. Thus, in this literature, the term “proportional representation” [14,37] has a different meaning: these methods are ‘representative’ because each voter feels represented by some member of the elected committee. The computational aspects of full proportional representation and its extensions have raised a lot of attention lately [42,6,16,47,36]. Our study of the properties of multi-attribute proportional representation is close in spirit to the work of Elkind et al. [22], who gives a normative study of multiwinner election rules. *Budgeted social choice* [35,46] is technically close to committee elections, but it has a different motivation: the aim is to make a collective choice about a set of objects to be consumed by the group (perhaps, subject to some constraints) rather than about the set of candidates to represent voters.

There exists an interesting line of research on multiwinner voting [11,12,30,8,45,32], where it is assumed that the elected committee runs a sequence of independent ballots on various issues – for instance consider a parliament voting on issues such as monetary politics, changes to the national-health care system, or educational reforms. Each issue can be represented by an attribute; in this setting our MAPR methods can be used to find a representative committee with respect to its collective views on a certain set of issues.

7. Discussion of the model and of its possible extensions

In this section we discuss several other approaches to the problem of achieving proportional representation with respect to multiple attributes, and we compare them with the model discussed so far.

7.1. Lower and upper quotas for attributes

In Section 3 we assumed that the input contains a vector of target distributions which describe desirable proportions of values for different attributes in an ideal committee – such an ideal committee might not exist, e.g., because there is not enough diversity within the candidate database (in particular, the candidate database might not satisfy the full supply property), or even if an ideal committee exists it might be computationally infeasible to find one. For this reason we formulated two optimisation metrics which, intuitively, allow one to assess how good are certain committees, and to find committees which are good enough, though not necessary ideal. Thus, it is natural to consider another approach: instead of getting a vector of ideal target distributions, we could assume that for each value of each attribute we are given a lower and an upper bound (also referred to as lower and upper quota, respectively) on the number of committee members with such a value of the respective attribute. For instance, instead of specifying that we would like to have 50% of men and 50% of women in a committee, we could ask for a committee with at least 40% of women and at least 40% of men.

Having lower and upper quotas for attributes gives more flexibility and makes it more likely that a committee satisfying the constraints exists. However, if the number of attributes is large (for instance, hundreds or thousands) and the size of the candidate database is moderate, it is still likely that a committee satisfying all the constraints does not exist. Further, coming up with the constraints which, on the one hand are restrictive enough to implement multi-attribute proportionality to the extent that would be satisfactory, and on the other hand are liberal enough to ensure that a committee satisfying the constraints exists, is much less straightforward and requires more cognitive effort than simply providing a vector of ideal distributions.

Interestingly, our results from Section 5.1 can be extended to the model with lower and upper quotas. Indeed, for the hardness it suffices to observe that the problem of finding a perfect committee can be easily formulated in the model with lower and upper quotas – it suffices to set the upper and lower quotas to the same value, equal to the value of the respective target distribution. For the positive result from Theorem 2 it suffices to change constraints (g) and (i) in the proof of the theorem so that variables x_i^j and y_i^j are compared against specific quotas instead of $k\pi_i^j$. Also the analysis from the proofs of Theorem 4 and Theorem 5 carries over to the case with quotas: in the proofs of these theorems one needs to define X_{ex} as the set of all attributes for which the analysed committee A does not exceed the lower and upper bounds. Specifically, our local search algorithm would treat lower and upper quotas as soft constraints and would approximate the total violation of the constraints:

$$\sum_{r_i^j(A) < \underline{\pi}_i^j} (\underline{\pi}_i^j - r_i^j(A)) + \sum_{r_i^j(A) > \overline{\pi}_i^j} (r_i^j(A) - \overline{\pi}_i^j),$$

where $\underline{\pi}_i^j$ and $\overline{\pi}_i^j$ denote the lower and upper quotas, respectively (cf. Definition 8).

7.2. Dependent attributes

In our model we assume that the attributes are independent, which sometimes may lead to undesirable outcomes. For example, consider an instance where the goal is to select a committee consisting of 50% of men and 50% of women and of 50% of junior and 50% of senior people. In this instance, for $k = 10$, a committee A that consists of 5 junior men and 5 senior women is a perfect committee. However, junior women and senior men are clearly underrepresented in A . Another example is when our goal is to select a set of k movies and when half of the population likes drama movies with Meryl Streep starring the main role and the other half likes action movies with Dwayne Johnson. A set of $k/2$ action movies with Meryl Streep and $k/2$ drama movies with Dwayne Johnson would form a perfect committee, even though it is incompatible with the voters' preferences.¹⁴

This phenomenon is known as the separability dilemma:

- either preferences are assumed to be separable: in that case, they are cheap to communicate (and computing the outcome is generally easy); but it is a strong domain restriction. In our example, if we assumed that the preferences of the society expressed by the target distributions were separable, then a set with $k/2$ action movies with Meryl Streep and $k/2$ drama movies with Dwayne Johnson would form an excellent solution.
- or we don't make such an assumption and allow preferential dependencies between attributes. This increases the cost of communication exponentially in the worst case, and makes computation harder.

Both approaches are often seen as too extreme, and the usual trade-off consists in allowing a reasonable amount of preferential dependencies. We can for instance introduce an artificial attribute combining some dependent attributes. For instance, in our first example we could introduce a combined attribute (Gender, Age) and we could require that there are 25% of committee members representing each of the four values: (male, junior), (male, senior), (female, junior), and (female, senior). Since combining the attributes leads to an exponential growth of the length of the representation of the target distributions, this approach is only possible when the number of dependent attributes is relatively small (see also the discussion below Definition 3 in Section 3, and the discussion on Constrained Approval Voting in Section 6.3).

7.3. Other metrics measuring the distance to the target distributions

In Section 4 we defined the multi-attribute d'Hondt rule and the multi-attribute Hamilton rule in terms of minimisation or maximisation of sums of expressions. Another possibility is to define the L^∞ -multi-attribute d'Hondt rule as the one which outputs a committee A maximising $\min_i \sum_j \pi_i^j H(r_i^j(A) \cdot k)$ and the L^∞ -multi-attribute Hamilton rule which outputs a committee A minimising $\max_i \sum_j |r_i^j(A) - \pi_i^j|$. Both approaches have their advantages and shortcomings. For instance, with the L^1 metric it may happen that there exists an optimal committee which is far from the target distributions for half of the attributes while there exists another committee which violates the target distribution for each attribute, but to a significantly lower extent. Such a committee seems more appropriate in the context of proportional representation. On the other hand, if we follow an L^∞ -optimisation approach, it may happen that among a large number of attributes there exists a single "outlier" attribute X_i with the target distribution set to $\pi_i^1 = 1$ in spite of the fact that all candidates in the database have the value of this attribute equal to x_i^0 . In such case a rule would select any committee, in particular it could select a committee which is far from the target distributions for every attribute, even though there might exist a committee which would be perfect for all attributes except for X_i . Naturally, there exist intermediate approaches – for instance, one could aim at maximising/minimising the L^p norms of the appropriate expressions.

The results from Section 5.1 easily extend to the case of optimising L^∞ -aggregate. For instance, the ILPs from the proof of Theorem 2 can be naturally extended to the L^∞ -optimisation case, by using the standard constructions for implementing the "max" operator in the objective function. A natural question which remains open is whether the L^∞ -variants of our problems can be well approximated.

8. Conclusion

In this paper we have defined and studied multi-attribute generalisations of a well-known class of apportionment methods, in particular of the Hamilton and the d'Hondt methods of apportionment, albeit with motivations that go far beyond party-list elections (such as the selection of a collective set of candidates). We have formulated several axioms, commonly

¹⁴ We thank the anonymous AIJ reviewer for suggesting this example.

considered in the political science literature in the context of apportionment, for multi-attribute committee selection rules. Motivated with this axiomatic approach we have identified two multi-attribute committee selection rules that can be considered as extensions of the Hamilton and d'Hondt methods to multi-attribute scenarios.

We have studied the computational complexity of the problem of finding committees that, in some sense, best fit some given distribution of attribute values. We have found out that the problem is in general NP-hard, but that it can be handled efficiently if the number of attributes is small. We have shown that the multi-attribute extensions of the Hamilton and d'Hondt methods can be well approximated. In particular, we have provided an interesting involved analysis of the local-search algorithm in the context of our multi-attribute setting.

Acknowledgements

We thank Eunjung Kim for giving us the initial idea of the paper, and for fruitful discussions. We further thank Katarína Cechlárová for her comments on the relation between population and party population monotonicity. Jérôme Lang was supported by the ANR project CoCoRiCo-CoDec, project number ANR-14-CE24-0007. Piotr Skowron was supported by the European Research Council grant ERC-StG 639945 (ACCORD) and by the Foundation for Polish Science within the Homing/2017-4/40 programme (Project title: “Normative Comparison of Multiwinner Election Rules”).

Appendix A. Proofs omitted from the main text

Proposition 1. *Under full supply property the Hamilton method satisfies party population monotonicity.*

Proof. Consider an instance I of the apportionment problem, and let I' be an instance obtained from I by increasing the quota v_i/v_+ for one party P_i , but leaving the ratios of quotas between the other parties unchanged. Let us consider the Hamilton method as the process that in steps allocates seats to parties (the first step is rounding down the quotas, and the next steps correspond to allocating seats to the parties in the descending order of their remainders). We want to prove that if this is the case that the Hamilton method assigns an x -th seat to P_i before assigning a y -th seat to $P_j \neq P_i$ in I , then it is also the case in I' . This will show that the number of seats assigned to P_i in I' is at least as large as in I .

We know that the quota of P_i in I' is higher than in I . Also, for any other party $P_j \neq P_i$, we know that the quota of P_j in I' is lower than in I (this is because the ratios of the quotas of the other parties remain unchanged; note that this argument would not work if we used population monotonicity instead of party population monotonicity). Thus, in the phase of rounding quotas down P_i will get at least the same number of seats in I' as in I . Also, if P_j got the same number of seats after rounding in I' as it got in I , then the remainder of P_i is higher than the remainder of P_j in I' whenever it is the case that it was higher in I . Thus, if the Hamilton method assigned a seat to P_i before P_j in I , then it must also happen in I' . \square

Proposition 2. *When $p = 1$ and assuming there are at least k candidates for each value of the unique attribute, then \mathcal{R}_H coincides with the Hamilton apportionment rule.*

Proof. Let s_j^* denote the ideal number of seats for party P_j , i.e., $s_j^* = \pi^j k$. Let A be a committee of size k and let $R^j(A) = k r^j(A)$ be the number of members of A that belong to party P_j . Since $|R^j(A) - s_j^*| = k|r^j(A) - \pi^j|$, we need to show that the following two assertions are equivalent:

1. A minimises $\sum_j |R^j(A) - s_j^*|$.
2. A is a Hamilton committee.

We first show $1 \Rightarrow 2$. Assume A is not a Hamilton committee: then there exists an attribute value (party) that receives strictly more or strictly less seats than it would receive according to the Hamilton method. Naturally, there must also exist an attribute that receives strictly less or strictly more seats, respectively. Formally, this means that there are two attribute values (parties), say 1 and 2, such that the target number of seats for parties 1 and 2 are $s_1^* = p + \alpha_1$ and $s_2^* = q + \alpha_2$, with p, q being integers and $1 > \alpha_2 > \alpha_1 \geq 0$, and such that $R^1(A) \geq p + 1$ and $R^2(A) \leq q$. We have

$$\begin{aligned} \sum_j |R^j(A) - \pi^j| &= \sum_{j \neq 1,2} |R^j(A) - s_j^*| + |R^1(A) - s_1^*| + |R^2(A) - s_2^*| \\ &\geq \sum_{j \neq 1,2} |R^j(A) - s_j^*| + (1 - \alpha_1) + \alpha_2. \end{aligned}$$

Consider the committee A' obtained from A by giving one less seat to 1 and one more to 2 and consider the following three cases

Case 1: If $R^1(A) > p + 1$ then:

$$\begin{aligned} & \sum_j |R^j(A) - s_j^*| - \sum_j |R^j(A') - s_j^*| \\ &= |R^1(A) - s_1^*| - |R^1(A') - s_1^*| + |R^2(A) - s_2^*| - |R^2(A') - s_2^*| \\ &\geq 1 + (1 - \alpha_2) - \alpha_2 > 0. \end{aligned}$$

Case 2: If $R^2(A) < q$ then similarly, $\sum_j |R^j(A) - s_j^*| - \sum_j |R^j(A') - s_j^*| > 0$.

Case 3: If $R^1(A) = p + 1$ and $R^2(A) = q$ then we have:

$$\sum_j |R^j(A) - s_j^*| = \sum_{j \neq 1,2} |R^j(A) - s_j^*| + (1 - \alpha_1) + \alpha_2$$

and

$$\sum_j |R^j(A') - s_j^*| = \sum_{j \neq 1,2} |R^j(A') - s_j^*| + (1 - \alpha_2) + \alpha_1$$

Hence:

$$\sum_j |R^j(A) - s_j^*| - \sum_j |R^j(A') - s_j^*| = 2(\alpha_2 - \alpha_1) > 0.$$

In all three cases, A does not minimise $\sum_j |R^j(A) - s_j^*|$, which gives a contradiction.

It remains to be shown that $2 \Rightarrow 1$, i.e., that if A is a Hamilton committee then it minimises $\sum_j |R^j(A) - s_j^*|$. If there is a unique Hamilton committee then this follows immediately from $1 \Rightarrow 2$. Assume there are several Hamilton committees A_1, \dots, A_q . Then there are q parties, w.l.o.g., let us call them P_1, \dots, P_q , with equal remainders $\alpha \in [0, 1)$, that is, $s_1^* = p_1 + \alpha, \dots, s_q^* = p_q + \alpha$, and these committees differ only with respect to whether they get an extra seat or not. We easily check that for any two A, A' of these committees we have $\sum_j |R^j(A) - s_j^*| = \sum_j |R^j(A') - s_j^*|$. \square

Proposition 4. *Under the full supply assumption, non-reversal, respect of quota, and value monotonicity with respect to every attribute are all satisfied by the multi-attribute Hamilton rule. In the general case, non-reversal, and respect of quota are not satisfied. If X_i is a binary variable, then value monotonicity with respect to X_i is satisfied; however it is not satisfied in the general case.*

Proof. Under the full supply assumption, the result easily comes from Proposition 3 and the fact that the property holds in the single-attribute case.

In the general case, we give counterexamples. For respect of quota, we have two binary attributes, and two candidates a, b with value vectors (x_1^2, x_2^2) and (x_1^1, x_2^1) , $k = 1$, π defined as $\pi_1^1 = 0, \pi_1^2 = 1, \pi_2^1 = 1, \pi_2^2 = 0$. The committee minimising our metric is either $\{a\}$ or $\{b\}$, and does not respect quota even though all values $k\pi_i^j$ are integers.

For non-reversal we have two binary attributes and six candidates: a, b, c , each with vector (x_1^1, x_2^1) and d, e, f , each with vector (x_1^2, x_2^2) . We have a target distribution π defined as follows: $\pi_1^1 = 0.35, \pi_1^2 = 0.65, \pi_2^1 = 1, \pi_2^2 = 0$. We set $k = 3$. The committees minimising our metric are $\{a, b, c\}$ and all triples made up from two candidates out of $\{d, e, f\}$ and one out of $\{d, e, f\}$. In all cases, we have $r_1^1(A) > r_1^2(A)$ even though $\pi_1^1 < \pi_1^2$.

Now, we prove that value monotonicity holds for binary domains. In the following we will use notation $\|r(A) - \pi\| = \sum_{i,j} |r_i^j(A) - \pi_i^j|$. Consider a binary attribute X_i , with $D_i = \{x_i^0, x_i^1\}$. Assume that $\rho_i^0 > \pi_i^0$ (and so $\rho_i^1 < \pi_i^1$), and that for all $i' \neq i$ we have $\rho_{i'} = \pi_{i'}$. Let A be an committee minimising our metric for π and, for the sake of contradiction, assume that for all committees B minimising our metric for ρ we have $r_i^0(B) < r_i^0(A)$. Let B be such a committee. The proof is a case by case study, with six cases to be considered: (C1) $r_i^0(B) \leq \pi_i^0 < \rho_i^0 \leq r_i^0(A)$; (C2) $\pi_i^0 \leq r_i^0(B) \leq \rho_i^0 \leq r_i^0(A)$; (C3) $\pi_i^0 < \rho_i^0 \leq r_i^0(B) < r_i^0(A)$; (C4) $r_i^0(B) \leq \pi_i^0 \leq r_i^0(A) \leq \rho_i^0$; (C5) $\pi_i^0 \leq r_i^0(B) < r_i^0(A) \leq \rho_i^0$; and (C6) $r_i^0(B) < r_i^0(A) \leq \pi_i^0 < \rho_i^0$.

- Case 1: $r_i^0(B) \leq \pi_i^0 < \rho_i^0 \leq r_i^0(A)$. In this case we have $r_i^1(B) \geq \pi_i^1 > \rho_i^1 \geq r_i^1(A)$ and the following holds:

$$\|r(B) - \pi\| = \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \pi_{i'}^j| + (\pi_i^0 - r_i^0(B)) + (r_i^1(B) - \pi_i^1) \tag{1}$$

$$= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(B)) + (r_i^1(B) - \rho_i^1) + \pi_i^0 - \pi_i^1 - \rho_i^0 + \rho_i^1 \tag{2}$$

$$= \|r(B) - \rho\| + 2(\pi_i^0 - \rho_i^0) \tag{3}$$

$$< \|r(A) - \rho\| + 2(\pi_i^0 - \rho_i^0) \tag{4}$$

$$= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \rho_i^0) + (\rho_i^1 - r_i^1(A)) + 2(\pi_i^0 - \rho_i^0) \tag{5}$$

$$= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \pi_i^0) + (\pi_i^1 - r_i^1(A)) + \pi_i^0 - \pi_i^1 - \rho_i^0 + \rho_i^1 + 2(\pi_i^0 - \rho_i^0) \tag{6}$$

$$= \|r(A) - \pi\| + 4(\pi_i^0 - \rho_i^0) \tag{7}$$

$$\leq \|r(A) - \pi\| \tag{8}$$

(4) comes from the fact that A does not minimise f for ρ . Since, there is one strong inequality in the sequence, we imply that A does not minimise f for π , a contradiction.

- Case 2: $\pi_i^0 \leq \rho_i^0(B) \leq \rho_i^0 \leq r_i^0(A)$.

$$\begin{aligned} \|r(B) - \pi\| &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \pi_{i'}^j| + (r_i^0(B) - \pi_i^0) + (\pi_i^1 - r_i^1(B)) \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(B)) + (r_i^1(B) - \rho_i^1) + 2r_i^0(B) - \pi_i^0 - \rho_i^0 - 2r_i^1(B) + \pi_i^1 + \rho_i^1 \\ &= \|r(B) - \rho\| + 4r_i^0(B) - 2\pi_i^0 - 2\rho_i^0 \\ &< \|r(A) - \rho\| + 4r_i^0(B) - 2\pi_i^0 - 2\rho_i^0 \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \rho_i^0) + (\rho_i^1 - r_i^1(A)) + 4r_i^0(B) - 2\pi_i^0 - 2\rho_i^0 \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \pi_i^0) + (\pi_i^1 - r_i^1(A)) + \pi_i^0 - \rho_i^0 - \pi_i^1 + \rho_i^1 + 4r_i^0(B) \\ &\quad - 2\pi_i^0 - 2\rho_i^0 \\ &= \|r(A) - \pi\| + 4r_i^0(B) - 4\rho_i^0 \\ &\leq \|r(A) - \pi\| \end{aligned}$$

Again we obtain a contradiction.

- Case 3: $\pi_i^0 < \rho_i^0 \leq r_i^0(B) < r_i^0(A)$.

$$\begin{aligned} \|r(B) - \pi\| &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \pi_{i'}^j| + (r_i^0(B) - \pi_i^0) + (\pi_i^1 - r_i^1(B)) \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \rho_{i'}^j| + (r_i^0(B) - \rho_i^0) + (\rho_i^1 - r_i^1(B)) - \pi_i^0 + \rho_i^0 + \pi_i^1 - \rho_i^1 \\ &= \|r(B) - \rho\| - 2\pi_i^0 + 2\rho_i^0 \\ &< \|r(A) - \rho\| - 2\pi_i^0 + 2\rho_i^0 \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \rho_i^0) + (\rho_i^1 - r_i^1(A)) - 2\pi_i^0 + 2\rho_i^0 \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \pi_i^0) + (\pi_i^1 - r_i^1(A)) + \pi_i^0 - \rho_i^0 - \pi_i^1 + \rho_i^1 - 2\pi_i^0 + 2\rho_i^0 \\ &= \|r(A) - \pi\| \end{aligned}$$

- Case 4: $r_i^0(B) \leq \pi_i^0 \leq r_i^0(A) \leq \rho_i^0$.

$$\begin{aligned} \|r(B) - \pi\| &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \pi_{i'}^j| + (\pi_i^0 - r_i^0(B)) + (r_i^1(B) - \pi_i^1) \\ &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(B)) + (r_i^1(B) - \rho_i^1)\pi_i^0 - \rho_i^0 - \pi_i^1 + \rho_i^1 \\ &= \|r(B) - \rho\| + 2\pi_i^0 - 2\rho_i^0 \\ &< \|r(A) - \rho\| + 2\pi_i^0 - 2\rho_i^0 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(A)) + (r_i^1(A) - \rho_i^1) + 2\pi_i^0 - 2\rho_i^0 \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \pi_i^0) + (\pi_i^1 - r_i^1(A)) - 2r_i^0(A) + 2r_i^1(A) + \pi_i^0 + \rho_i^0 - \pi_i^1 - \rho_i^1 \\
 &\quad + 2\pi_i^0 - 2\rho_i^0 \\
 &= \|r(A) - \pi\| - 4r_i^0(A) + 4\pi_i^0 \\
 &\leq \|r(A) - \pi\|
 \end{aligned}$$

- Case 5: $\pi_i^0 \leq r_i^0(B) < r_i^0(A) \leq \rho_i^0$.

$$\begin{aligned}
 \|r(B) - \pi\| &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \pi_{i'}^j| + (r_i^0(B) - \pi_i^0) + (\pi_i^1 - r_i^1(B)) \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(B)) + (r_i^1(B) - \rho_i^1) + 2r_i^0(B) - 2r_i^1(B) - \pi_i^0 - \rho_i^0 + \pi_i^1 + \rho_i^1 \\
 &= \|r(B) - \rho\| + 4r_i^0(B) - 2\pi_i^0 - 2\rho_i^0 \\
 &< \|r(A) - \rho\| + 4r_i^0(B) - 2\pi_i^0 - 2\rho_i^0 \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(A)) + (r_i^1(A) - \rho_i^1) + 4r_i^0(B) - 2\pi_i^0 - 2\rho_i^0 \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (r_i^0(A) - \pi_i^0) + (\pi_i^1 - r_i^1(A)) + 4r_i^0(B) - 2r_i^0(A) + 2r_i^1(A) + \pi_i^0 + \rho_i^0 \\
 &\quad - \pi_i^1 - \rho_i^1 - 2\pi_i^0 - 2\rho_i^0 \\
 &= \|r(A) - \pi\| + 4r_i^0(B) - 4r_i^0(A) \\
 &\leq \|r(A) - \pi\|
 \end{aligned}$$

- Case 6: $r_i^0(B) < r_i^0(A) \leq \pi_i^0 < \rho_i^0$.

$$\begin{aligned}
 \|r(B) - \pi\| &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \pi_{i'}^j| + (\pi_i^0 - r_i^0(B)) + (r_i^1(B) - \pi_i^1) \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(B) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(B)) + (r_i^1(B) - \rho_i^1) + \pi_i^0 - \rho_i^0 - \pi_i^1 + \rho_i^1 \\
 &= \|r(B) - \rho\| + 2\pi_i^0 - 2\rho_i^0 \\
 &< \|r(A) - \rho\| + 2\pi_i^0 - 2\rho_i^0 \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (\rho_i^0 - r_i^0(A)) + (r_i^1(A) - \rho_i^1) + 2\pi_i^0 - 2\rho_i^0 \\
 &= \sum_{i' \neq i} \sum_j |r_{i'}^j(A) - \rho_{i'}^j| + (\pi_i^0 - r_i^0(A)) + (r_i^1(A) - \pi_i^1) - \pi_i^0 + \rho_i^0 + \pi_i^1 - \rho_i^1 + 2\pi_i^0 - 2\rho_i^0 \\
 &= \|r(A) - \pi\|
 \end{aligned}$$

Finally, we give an example showing that value monotonicity does not hold in the general case. First, we describe the set of attributes. We have one distinguished attribute X_1 with 5 possible values $x_1^1, x_1^2, x_1^3, x_1^4,$ and x_1^5 and 64 groups of binary attributes, indexed with the pairs of integers $i, j \in \{1, 2, 3, 4\}$. These groups of attributes are denoted as $X_{(1,1)}, X_{(1,2)}, \dots, X_{(1,8)}, X_{(2,1)}, \dots, X_{(8,8)}$. Each group contains some large number λ of indistinguishable attributes, each having the same set of possible values $\{x_2^1, x_2^2\}$. We have 16 alternatives $A_1, A_2, \dots, A_8,$ and $B_1, B_2, \dots, B_8,$ and our goal is to select a subset of $k = 8$ of them.

We start with describing these alternatives on binary attributes: each alternative A_i has the value x_2^1 on all attributes $X_{(i,\cdot)}$ and the value x_2^2 on all the remaining ones; each alternative B_i has the value x_2^1 on all attributes $X_{(\cdot,i)}$ and the value x_2^2 on all the remaining ones. For the binary attributes we set the target probabilities to $\pi_2^1 = 1/8$ and $\pi_2^2 = 7/8$. Due to this construction, we see that the only two subsets that perfectly agree with target distributions on each of binary attributes are $A = \{A_1, A_2, \dots, A_8\}$ and $B = \{B_1, B_2, \dots, B_8\}$. Indeed, every subset S including A_i and $B_j,$ would have $r(S) \geq 1/4$ at least for one group of attributes $X_{(i,j)}$. Since λ is large, we infer that, independently of what happens on the distinguished attribute $X_1,$ the only possible winning committee is either $A = \{A_1, A_2, \dots, A_8\}$ or $B = \{B_1, B_2, \dots, B_8\}$.

Next, let us describe what happens on the attribute X_1 . The vector $\langle r_1^j(A) \rangle$ is equal to $\langle r_1^j(A) \rangle = (1/2, 0, 1/2, 0, 0)$. For the committee B , we have $\langle r_1^j(B) \rangle = (1/4, 1/4, 1/4, 1/8, 1/8)$, and the vector of target distributions for X_1 is equal $\pi_1 = (0, 0, 3/8 + \epsilon, 5/8 - \epsilon, 0)$. We can see that $\|r(A) - \pi\| = 1/2 + 1/8 - \epsilon + 5/8 - \epsilon = 1.25 - 2\epsilon$. Since, $\|r(B) - \pi\| = 1/4 + 1/4 + 1/8 + \epsilon + 4/8 - \epsilon + 1/8 = 1.25$, we get that A is a winning committee. However, if we modify the target fractions so that $\rho_1 = (1/4, 0, 9/32 + \epsilon_1, 15/32 - \epsilon_2, 0)$, we will get $\|r(A) - \rho\| = 1/4 + 7/32 - \epsilon_1 + 15/32 - \epsilon_2 = 30/32 - \epsilon_1 - \epsilon_2$ and $\|r(B) - \rho\| = 1/4 + 1/32 + \epsilon_1 + 11/32 - \epsilon_2 + 1/8 = 24/32 + \epsilon_1 - \epsilon_2$, thus, B is winning according to ρ . However, B has lower representation of x_1^1 than A , and ρ was obtained from π , by increasing the fraction of π_1^1 . This completes the proof. \square

Proposition 6. Consider a candidate database that satisfies the full supply property. For any attribute X_i , any committee A that maximises $\sum_{i,j} \pi_i^j H(r_i^j(A) \cdot k)$ is a d'Hondt committee for the single-attribute problem $(\{X_i\}, D^{\downarrow X_i}, \pi_i, k)$, where $D^{\downarrow X_i}$ is the projection of D on $\{X_i\}$.

Proof. The idea from the proof of Proposition 3 works also for this proposition. If there exists a committee A which maximises $\sum_i \sum_j \pi_i^j H(r_i^j(A) \cdot k)$ and which is not a d'Hondt committee for the single-attribute problem $(\{X_i\}, D^{\downarrow X_i}, \pi_i, k)$, then by Proposition 5 there exists a committee B such that $\sum_j \pi_i^j H(r_i^j(B) \cdot k) > \sum_j \pi_i^j H(r_i^j(A) \cdot k)$. Similarly as in the proof of Proposition 3, it is possible to build a committee D from A and B such that $\sum_i \sum_j \pi_i^j H(r_i^j(D) \cdot k) > \sum_i \sum_j \pi_i^j H(r_i^j(A) \cdot k)$, which gives a contradiction and completes the proof. \square

Proposition 7. Under the full supply assumption, non-reversal, house monotonicity, and value monotonicity with respect to every attribute are all satisfied by the multi-attribute d'Hondt method. In the general case, non-reversal and house monotonicity are not satisfied. If X_i is a binary variable, then value monotonicity with respect to X_i is satisfied; however it is not satisfied in the general case.

Proof. Similarly as in the proof of Proposition 4 we infer that the result for full supply assumption follows from Proposition 6 and from the fact that the respective properties holds in the single-attribute case.

In the general case, we give counterexamples. For non-reversal, the same example as in the proof of Proposition 4 works also for the case of the multi-attribute d'Hondt method.

For house monotonicity we have two binary attributes and three candidates: a with vector (x_1^1, x_2^2) , b with vector (x_1^2, x_2^1) , and c with vector (x_1^2, x_2^2) . We have a target distribution π defined as follows: $\pi_1^1 = \pi_2^1 = 0.5 - \epsilon$ and $\pi_1^2 = \pi_2^2 = 0.5 + \epsilon$, for some small positive ϵ . For $k = 1$ candidate c should be selected, while for $k = 2$ committee $\{a, b\}$ is optimal.

Now, we prove that value monotonicity holds for binary domains. Consider a binary attribute X_i , with $D_i = \{x_i^0, x_i^1\}$. Assume that $\rho_1^0 > \pi_1^0$, and that for all $i' \neq i$ we have $\rho_{i'} = \pi_{i'}$. Let A be an committee maximising our metric for π and, for the sake of contradiction, assume that for all committees B maximising our metric for ρ we have $r_i^0(B) < r_i^0(A)$. Let B be such a committee.

$$\begin{aligned}
& \sum_i \sum_j \pi_i^j H(r_i^j(B) \cdot k) = \sum_{i \neq 1} \sum_j \pi_i^j H(r_i^j(B) \cdot k) + \pi_1^0 H(r_1^0(B) \cdot k) + (1 - \pi_1^0) H(k - r_1^0(B) \cdot k) \\
& = \sum_{i \neq 1} \sum_j \rho_i^j H(r_i^j(B) \cdot k) + \rho_1^0 H(r_1^0(B) \cdot k) + (1 - \rho_1^0) H(k - r_1^0(B) \cdot k) \\
& \quad + (\rho_1^0 - \pi_1^0) \cdot (H(k - r_1^0(B) \cdot k) - H(r_1^0(B) \cdot k)) \\
& > \sum_i \sum_j \rho_i^j H(r_i^j(A) \cdot k) + (\rho_1^0 - \pi_1^0) \cdot (H(k - r_1^0(B) \cdot k) - H(r_1^0(B) \cdot k)) \\
& \geq \sum_i \sum_j \rho_i^j H(r_i^j(A) \cdot k) + (\rho_1^0 - \pi_1^0) \cdot (H(k - r_1^0(A) \cdot k) - H(r_1^0(A) \cdot k)) \\
& = \sum_{i \neq 1} \sum_j \rho_i^j H(r_i^j(A) \cdot k) + \rho_1^0 H(r_1^0(A) \cdot k) + (1 - \rho_1^0) H(k - r_1^0(A) \cdot k) \\
& \quad + (\rho_1^0 - \pi_1^0) \cdot (H(k - r_1^0(A) \cdot k) - H(r_1^0(A) \cdot k)) \\
& = \sum_{i \neq 1} \sum_j \pi_i^j H(r_i^j(A) \cdot k) + \pi_1^0 H(r_1^0(A) \cdot k) + (1 - \pi_1^0) H(k - r_1^0(A) \cdot k) \\
& = \sum_i \sum_j \pi_i^j H(r_i^j(A) \cdot k).
\end{aligned}$$

We get that B is better with respect to our metric than A for π , a contradiction.

Finally, from the proof of Proposition 4 we can reuse parts of the construction showing that value monotonicity does not hold in the general case. Let us recall that the construction there ensures that one of the two committees,

$A = \{A_1, A_2, \dots, A_8\}$ or $B = \{B_1, B_2, \dots, B_8\}$, needs to be selected. Additionally we can have two attributes, X_1 and X_2 , each with three possible values. These two attributes determine whether A or B is going to be selected. We select A and B so that:

$$\begin{aligned} \langle r_1^j(A) \rangle &= (0, 1, 0) & \langle r_1^j(B) \rangle &= (0, 0, 1) \\ \langle r_2^j(A) \rangle &= (1/8, 7/8, 0) & \langle r_2^j(B) \rangle &= (0, 0, 1). \end{aligned}$$

We set $\pi_1 = (0, 0.1, 0.9)$. Now, consider $\pi_2 = (0, 1, 0)$. For the two attributes the values of committees A and B are equal to:

$$\text{committee } A: 0.1 \cdot H(8) + H(7) \approx 2.86, \quad \text{committee } B: 0.9 \cdot H(8) \approx 2.44.$$

Consequently, A will be selected by the multi-attribute d'Hondt method. Now, consider what happens when we change π_2 to $\rho_2 = (1, 0, 0)$. For the two attributes the values of committees A and B are now equal to:

$$\text{committee } A: 0.1 \cdot H(8) + H(1) \approx 1.27, \quad \text{committee } B: 0.9 \cdot H(8) \approx 2.44.$$

Yet, B has lower representation of x_2^0 than A , and ρ was obtained from π , by increasing the fraction of π_2^0 . This shows that value monotonicity is not satisfied in the general case and completes the proof. \square

References

- [1] S. Amer-Yahia, S.B. Roy, A. Chawla, G. Das, C. Yu, Group recommendation: semantics and efficiency, *Proc. VLDB Endow.* 2 (1) (2009) 754–765.
- [2] M. Balinski, P. Young, Criteria for proportional representation, *Oper. Res.* 27 (1) (1979) 80–95.
- [3] M. Balinski, P. Young, *Fair Representation: Meeting the Ideal of One Man One Vote*, second edition, Brookings Institution Press, 2001.
- [4] M.L. Balinski, G. Demange, Algorithm for proportional matrices in reals and integers, *Math. Program., Ser. A* 45 (1–3) (1989) 193–210.
- [5] M.L. Balinski, G. Demange, An axiomatic approach to proportionality between matrices, *Math. Oper. Res.* 14 (1989) 700–719.
- [6] N. Betzler, A. Slinko, J. Uhlmann, On the computation of fully proportional representation, *J. Artif. Intell. Res.* 47 (2013) 475–519.
- [7] S.J. Brams, Computer-assisted constrained approval voting, *Interfaces* 20 (5) (1990) 67–80.
- [8] S.J. Brams, D.M. Kilgour, W.S. Zwicker, The paradox of multiple elections, *Soc. Choice Welf.* 15 (2) (1998) 211–236.
- [9] I. Caragiannis, C. Kaklamanis, N. Karanikolas, A. Procaccia, Socially desirable approximations for Dodgson's voting rule, *ACM Trans. Algorithms* 10 (2014) 6.
- [10] I. Caragiannis, D. Kurokawa, H. Moulin, A.D. Procaccia, N. Shah, J. Wang, The unreasonable fairness of maximum Nash welfare, in: *Proceedings of the 2016 ACM Conference on Economics and Computation, EC-16*, 2016, pp. 305–322.
- [11] A. Casella, Storable votes, *Games Econ. Behav.* 51 (2) (2005) 391–419.
- [12] A. Casella, *Storable Votes: Protecting the Minority Voice*, Oxford University Press, Oxford, UK, 2011.
- [13] M. Cesati, Perfect code is W[1]-complete, *Inf. Process. Lett.* 81 (3) (2002) 163–168.
- [14] B. Chamberlin, P. Courant, Representative deliberations and representative decisions: proportional representation and the Borda rule, *Am. Polit. Sci. Rev.* 77 (3) (1983) 718–733.
- [15] Vincent Conitzer, Rupert Freeman, Markus Brill, Yuqian Li, Rules for choosing societal tradeoffs, in: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, February 12–17, 2016, Phoenix, Arizona, USA, 2016, pp. 460–467.
- [16] D. Cornaz, L. Galand, O. Spanjaard, Bounded single-peaked width and proportional representation, in: *Proceedings of the 20th European Conference on Artificial Intelligence, ECAI-2012*, 2012, pp. 270–275.
- [17] Z. Cvetkovski, *Convexity, Jensen's Inequality*, Springer, Berlin, Heidelberg, 2012, pp. 69–77.
- [18] M. Cygan, F. Fomin, L. Kowalik, D. Lokshantov, D. Marx, M. Pilipczuk, M. Pilipczuk, S. Saurabh, *Parameterized Algorithms*, Springer, 2015.
- [19] G. Demange, On Allocating Seats to Parties and Districts: Apportionments, Technical Report halshs-00879779, HAL, 2013.
- [20] N. Ding, F. Lin, On computing optimal strategies in open list proportional representation: the two parties case, in: *Proceedings of the 28th Conference on Artificial Intelligence, AAAI-2014*, 2014, pp. 1419–1425.
- [21] R. Downey, M. Fellows, *Parameterized Complexity*, Springer-Verlag, 1999.
- [22] E. Elkind, P. Faliszewski, P. Skowron, A. Slinko, Properties of multiwinner voting rules, *Soc. Choice Welf.* 48 (3) (2017) 599–632.
- [23] E. Elkind, P. Faliszewski, A. Slinko, Distance rationalization of voting rules, *Soc. Choice Welf.* 45 (2) (2015) 345–377.
- [24] U. Endriss, Judgment aggregation, in: F. Brandt, V. Conitzer, U. Endriss, J. Lang, A.D. Procaccia (Eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 2016, Chapter 17.
- [25] J. Flum, M. Grohe, *Parameterized Complexity Theory*, Springer-Verlag, 2006.
- [26] A. Frank, E. Tardos, An application of simultaneous Diophantine approximation in combinatorial optimization, *Combinatorica* 7 (1) (1987) 49–65.
- [27] I. Garcia, S. Pajares, L. Sebastia, E. Onaindia, Preference elicitation techniques for group recommender systems, *Inf. Sci.* 189 (2012) 155–175.
- [28] M. Garey, D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman and Company, 1979.
- [29] R. Kannan, Minkowski's convex body theorem and integer programming, *Math. Oper. Res.* 12 (3) (1987) 415–440.
- [30] Y. Koriyama, J.F. Laslier, A. Macé, R. Treibich, Optimal apportionment, *J. Polit. Econ.* 121 (3) (2013) 584–608.
- [31] J. Lang, P. Skowron, Multi-attribute proportional representation, in: *Proceedings of the 30th Conference on Artificial Intelligence, AAAI-2016*, 2016, pp. 530–536.
- [32] J. Lang, L. Xia, Voting over multiattribute domains, in: F. Brandt, V. Conitzer, U. Endriss, J. Lang, A. Procaccia (Eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 2015, Chapter 9.
- [33] I. Lari, F. Ricca, A. Scozzari, Bidimensional allocation of seats via zero-one matrices with given line sums, *Ann. Oper. Res.* 215 (1) (2014) 165–181.
- [34] H.W. Lenstra, Integer programming with a fixed number of variables, *Math. Oper. Res.* 8 (4) (1983) 538–548.
- [35] T. Lu, C. Boutilier, Budgeted social choice: from consensus to personalized decision making, in: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI-2011*, 2011, pp. 280–286.
- [36] T. Lu, C. Boutilier, Multiwinner social choice with incomplete preferences, in: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence, IJCAI-2013*, 2013, pp. 263–270.
- [37] B.L. Monroe, Fully proportional representation, *Am. Polit. Sci. Rev.* 89 (1995) 925–940.

- [38] K. Nehring, M. Pivato, The Median Rule in Judgment Aggregation, Technical report, 2016.
- [39] G. Nemhauser, L. Wolsey, M. Fisher, An analysis of approximations for maximizing submodular set functions, *Math. Program.* 14 (1) (December 1978) 265–294.
- [40] R. Niedermeier, *Invitation to Fixed-Parameter Algorithms*, Oxford University Press, 2006.
- [41] R. Potthoff, Use of linear programming for constrained approval voting, *Interfaces* 20 (5) (1990) 79–80.
- [42] A. Procaccia, J. Rosenschein, A. Zohar, On the complexity of achieving proportional representation, *Soc. Choice Welf.* 30 (3) (2008) 353–362.
- [43] F. Pukelsheim, F. Ricca, B. Simeone, A. Scozzari, P. Serafini, Network flow methods for electoral systems, *Networks* 59 (1) (2012) 73–88.
- [44] P. Serafini, B. Simeone, Parametric maximum flow methods for minimax approximation of target quotas in biproportional apportionment, *Networks* 59 (2) (2012) 191–208.
- [45] P. Skowron, What do we elect committees for? A voting committee model for multi-winner rules, in: *Proceedings of the 24th International Joint Conference on Artificial Intelligence, IJCAI-2015*, 2015, pp. 1141–1148.
- [46] P. Skowron, P. Faliszewski, J. Lang, Finding a collective set of items: from proportional multirepresentation to group recommendation, *Artif. Intell.* 241 (2016) 191–216.
- [47] P. Skowron, P. Faliszewski, A. Slinko, Achieving fully proportional representation: approximability result, *Artif. Intell.* 222 (2015) 67–103.
- [48] A. Straszak, M. Libura, J. Sikorski, D. Wagner, Computer-assisted constrained approval voting, *Group Decis. Negot.* 2 (4) (1993) 375–385.