Designing Participatory Budgeting Mechanisms
Grounded in Judgment Aggregation

Simon Rey, **Ulle Endriss**, **Ronald de Haan**

Institute for Logic, Language and Computation (ILLC), University of Amsterdam

{s.j.rey, u.endriss, r.dehaan}@uva.nl

Abstract

We introduce a new approach for designing rules for participatory budgeting (PB), the problem of deciding on the use of public funds based directly on the views expressed by the citizens concerned. The core idea is to embed instances of the participatory budgeting problem into judgment aggregation, a powerful general-purpose framework for modelling collective decision making. Taking advantage of the possibilities offered by judgment aggregation, we enrich the familiar setting of participatory budgeting with additional constraints, namely dependencies between projects and quotas regarding different types of projects. We analyse the rules obtained in both algorithmic and axiomatic terms.

1 Introduction

Participatory budgeting (PB) is an instrument intended to improve the democratic process by allowing citizen to directly express their views regarding the use of public funds (Cabannes, 2004). Since its first use for municipal budget, PB has now been adopted across the world (Shah, 2007). PB proceeds in two stages. First, citizens submit project proposals, some of which are shortlisted. Then they vote on which of the shortlisted projects to fund, given the limitations set by the budget available. In this paper we introduce a new approach for designing voting rules for this second stage. The central idea is to embed PB into judgment aggregation (JA), a highly expressive framework for collective decision making that has been extensively studied in the field of (computational) social choice (List and Pettit, 2002; Endriss, 2016).

Most existing formal work in social choice theory regarding PB views voting on projects as a generalisation of approval-based multiwinner voting (see, e.g., Aziz, Lee, and Talmon, 2018; Talmon and Faliszewski, 2019). While this can provide useful intuitions regarding, for instance, the type of normative desiderata we may wish to postulate for PB, it does not allow for great flexibility, e.g., when it comes to modelling expressive forms of PB that allow us to specify dependencies between projects and the like. This is why we take a complementary approach and study project selection as a special case of the more general problem of JA.

While JA is very expressive—e.g., it naturally generalises many forms of preference aggregation and voting (Dietrich and List, 2007; Lang and Slavkovik, 2013; Endriss, 2018)—computing outcomes for JA is typically computationally in-

tractable (Endriss, Grandi, and Porello, 2012). The central challenge we address in this paper thus is to find ways of implementing PB via JA in an *efficient* manner. To do so, the core idea we explore is to look for tractable fragments of JA, by further developing the approach of De Haan (2018) of modelling JA problems using Boolean circuits in decomposable negation normal form (DNNF). This allows us to model PB problems with multiple resources, dependencies between projects, and quotas of different types of projects.

Of course, an expressive framework for modelling PB scenarios and a set of algorithmically efficient PB rules alone are not sufficient. We also require a good understanding of whether the rules we design are normatively adequate. We therefore provide an *axiomatic analysis* of the rules we propose, focusing in particular on the notion of *exhaustiveness* (ruling out any under-use of the budget) and the *monotonicity axioms* proposed by Talmon and Faliszewski (2019).

Related work. According to the terminology of Aziz and Shah (2020), we focus on *combinatorial PB with binary projects and approval ballots*. For this framework, Aziz, Lee, and Talmon (2018) and Talmon and Faliszewski (2019) analysed several rules in both axiomatic and algorithmic terms, proposing greedy algorithms and dynamic programming techniques. Using a different approach, Fain, Goel, and Munagala (2016) and Freeman et al. (2019) instead studied PB solutions as market equilibria, in the spirit of the public decision making setting of Conitzer, Freeman, and Shah (2017). Particularly relevant to our work, Fain, Munagala, and Shah (2018) considered a general setting of public decision making with matroid, matching, and packing constraints, allowing for great flexibility on what can be modelled. Jain, Sornat, and Talmon (2020) studied the computational complexity of PB when projects have types and utilities are defined over the types. Lu and Boutilier (2011) considered yet another extension of PB, where the cost of a project might depend on the number of agents choosing it. Finally, De Haan (2018) was the first to discuss the idea of embedding PB into JA.

Paper outline. We recall relevant definitions from PB and JA in Section 2 and then introduce our central definition of an *embedding* of PB into JA. Section 3 is devoted to the study of efficient embeddings for basic PB and the extensions we propose, Section 4 discusses exhaustiveness, and Section 5 contains the remainder of our axiomatic analysis.

2 Frameworks

In this section we recall basic definitions regarding the frameworks of participatory budgeting (PB) and judgment aggregation (JA). We also define the main concept of this paper, namely embeddings of PB instances into JA.

2.1 Participatory Budgeting

We mainly adopt the notation of Aziz and Shah (2020). PB is about selecting a set of projects to be funded, given a (possibly multi-dimensional) budget limit. The set of (binary) *projects* is denoted by $\mathcal{P} = \{p_1, \dots, p_m\}$. Let $\mathcal{R} = \{r_1, \dots, r_d\}$ be a set of *resources* and $\boldsymbol{b} = (b_1, \dots, b_d)$ a *budget limit vector*, with $b_i \in \mathbb{R}_{\geq 0}$ indicating the limit in terms of resource r_i. The costs of the projects are defined by a *cost function* $c : \mathcal{P} \times \mathcal{R} \to \mathbb{R}_{\geq 0}$, indicating for a given project the cost in terms of the given resource. Slightly overloading notation, we use $c(p) = (c(p, r_1), \dots, c(p, r_d))$ to denote the cost vector of project p. Moreover, for any subset $P \subseteq \mathcal{P}$, let $c(P, r) = \sum_{p \in P} c(p, r)$ and $c(P) = \sum_{p \in P} c(p)$. A *problem instance* $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ for PB consists of a set of resources, a budget limit vector, a set of projects, and a cost function. $\mathcal{I}$ is the set of all such instances.

A solution of a PB problem instance, called a *budget allocation*, is a subset of projects $A \subseteq \mathcal{P}$. A budget allocation A is said to be *feasible* if $c(A) \leq \boldsymbol{b}$. For a given $I \in \mathcal{I}$, the set of all feasible budget allocations is denoted by $\mathcal{A}(I)$.

Before deciding which budget allocation to recommend, we consult the *agents* belonging to a set $\mathcal{N} = \{1, \dots, n\}$. Each agent $i \in \mathcal{N}$ submits an *approval ballot* $A_i \subseteq \mathcal{P}$, giving rise to a *profile* $\boldsymbol{A} = (A_1, \dots, A_n)$. For any given project p, its *approval score* under profile $\boldsymbol{A}$ is $\sum_{i \in \mathcal{N}} \mathbb{1}_{p \in A_i}$, the number of agents approving of p. W.l.o.g., we assume that every project has an approval score of at least 1, as projects with approval score 0 can be removed in a pre-processing step. Finally, a *PB rule* is a function $F : \mathcal{I} \times (2^{\mathcal{P}})^n \to 2^{2^{\mathcal{P}}} \setminus \{\emptyset\}$ mapping any given instance I and profile $\boldsymbol{A}$ to a nonempty set $F(I, \boldsymbol{A}) \subseteq \mathcal{A}(I)$ of feasible budget allocations.[1] Returning a set allows us to model ties.

2.2 Judgment Aggregation

The JA framework we use is known as *binary aggregation with integrity constraints* (Grandi and Endriss, 2011).[2]

Let $\mathcal{L}_{\boldsymbol{P}}$ be the set of propositional formulas over a given set $\boldsymbol{P}$ of propositional atoms, using the usual connectives $\neg, \vee, \wedge, \to$, and logical constants $\bot$ and $\top$. Propositional atoms and their negations are called literals. For any $P \subseteq \boldsymbol{P}$, we write $Lit(P) = P \cup \{\neg x \mid x \in P\}$ for the set of literals corresponding to P. We often use x_i to denote atoms and ℓ_{x_i} to denote literals corresponding to x_i, i.e., $\ell_{x_i} \in \{x_i, \neg x_i\}$. We say that ℓ_{x_i} is positive if $\ell_{x_i} = x_i$ and negative if $\ell_{x_i} = \neg x_i$. A truth assignment $\alpha : \boldsymbol{P} \to \{0, 1\}$ is a mapping indicating for each atom its truth value. For $\ell_{x_i} \in Lit(\boldsymbol{P})$,

set $\alpha(\ell_{x_i}) = \alpha(x_i)$ if ℓ_{x_i} is positive and $\alpha(\ell_{x_i}) = 1 - \alpha(x_i)$ otherwise. We write $\alpha \models \varphi$ whenever α is a model of φ.

In the context of JA, the atoms in $\boldsymbol{P}$ represent *propositions* an agent may either *accept* or *reject*. A *judgment* J is a set $J \subseteq \boldsymbol{P}$, indicating which propositions are accepted. Let $ext(J) = J \cup \{\neg x \mid x \in \boldsymbol{P} \setminus J\}$ for any given judgment J. Observe that a judgment J can be equivalently described as the truth assignment α such that $\alpha(x) = 1$ if and only if $x \in J$. In our examples, when we do not explicitly specify some propositions, it is assumed that we only consider judgments (and truth assignments) for which the unspecified propositions are rejected (mapped to 0).

An *integrity constraint* $\Gamma \in \mathcal{L}_{\boldsymbol{P}}$ is a formula used to constrain the range of admissible judgments. A judgment J *satisfies* Γ (written $J \models \Gamma$), if J, interpreted as a truth assignment, is a model of Γ. Let $\mathfrak{J}(\Gamma) = \{J \subseteq \boldsymbol{P} \mid J \models \Gamma\}$. A *problem instance* for JA is simply an integrity constraint Γ.

We again use $\mathcal{N} = \{1, \dots, n\}$ to denote the set of *agents*. Each agent $i \in \mathcal{N}$ provides us with a judgment J_i, resulting in a *judgment profile* $\boldsymbol{J} = (J_1, \dots, J_n)$. For a profile $\boldsymbol{J}$ and a literal $\ell \in Lit(\boldsymbol{P})$, we write $n_\ell^{\boldsymbol{J}} = \sum_{i \in \mathcal{N}} \mathbb{1}_{\ell \in ext(J_i)}$ for the number of supporters of ℓ. The *majoritarian outcome* for a profile, denoted by $m(\cdot)$, is the set of literals supported by a majority of agents: $m(\boldsymbol{J}) = \{\ell \in Lit(\boldsymbol{P}) \mid n_\ell^{\boldsymbol{J}} > \frac{n}{2}\}$.

A *JA rule* is a function $F : \mathcal{L}_{\boldsymbol{P}} \times (2^{\boldsymbol{P}})^n \to 2^{2^{\boldsymbol{P}}} \setminus \{\emptyset\}$ taking as input an integrity constraint Γ and a judgment profile $\boldsymbol{J}$ and returning a nonempty set $F(\Gamma, \boldsymbol{J}) \subseteq \mathfrak{J}(\Gamma)$ of admissible judgments. Observe that no assumption is made about the profile. In particular, we do not require $J_i \models \Gamma$ for any $i \in \mathcal{N}$.

Before reviewing a number of well-known concrete JA rules, let us first introduce a very general class of such rules.

Definition 1 (Additive rules). *A judgment aggregation rule F is an additive rule if there exists a function $f : (2^{\boldsymbol{P}})^n \times Lit(\boldsymbol{P}) \to \mathbb{R}$ such that, for every integrity constraint Γ and every profile $\boldsymbol{J} \in (2^{\boldsymbol{P}})^n$, we have:*

$$F(\Gamma, \boldsymbol{J}) = \underset{J \in \mathfrak{J}(\Gamma)}{\operatorname{argmax}} \sum_{\ell \in ext(J)} f(\boldsymbol{J}, \ell).$$

This class generalises both the *scoring rules* of Dietrich (2014) and the *additive majority rules* (AMRs) defined by Nehring and Pivato (2019). A scoring rule is associated with a scoring function $s : 2^{\boldsymbol{P}} \times Lit(\boldsymbol{P}) \to \mathbb{R}$ and corresponds to the additive rule with $f(\boldsymbol{J}, \ell) = \sum_{i \in \mathcal{N}} s(J_i, \ell)$. An AMR is associated with a non-decreasing gain function $g : \{0, \dots, n\} \to \mathbb{R}$ with $g(k) < g(k')$ for any $k < \frac{n}{2} \leq k'$ and is an additive rule with $f(\boldsymbol{J}, \ell) = g(n_\ell^{\boldsymbol{J}})$. Three additive rules are of particular importance for our purposes:

- The *Slater rule* (Miller and Osherson, 2009; Lang et al., 2011) selects the admissible outcome closest to the majoritarian outcome in terms of the number of propositions they agree on. It is the AMR associated with the gain function g with $g(x) = 1$ if $x \geq \frac{n}{2}$ and $g(x) = 0$ otherwise.

- The *Kemeny rule* (Pigozzi, 2006; Miller and Osherson, 2009) selects the feasible outcome that is the closest to the profile as a whole. It is both an AMR with $g(x) = x$ and a scoring rule with $s(J, \ell) = \mathbb{1}_{\ell \in ext(J)}$.

---

[1]Observe that $\mathcal{A}(I)$ is never empty as the empty set is always feasible. This is not true for the extensions discussed in Section 3.

[2]While this framework is most convenient for our purposes, the original framework of List and Pettit (2002) could be used as well, given that it is known that the former can be efficiently embedded into the latter (Endriss et al., 2016).

Figure 1: Reduction from PB to JA

- The *leximax rule* (Everaere, Konieczny, and Marquis, 2014; Nehring and Pivato, 2019) favours the propositions supported by the largest majorities. It is the AMR defined by the gain function $g(x) = |\boldsymbol{P}|^x$.

Note that the three rules presented above are all *majority-consistent*, meaning that whenever the majoritarian outcome is admissible, it is the unique judgement returned by the rules.

2.3 Embedding PB into JA

The aim of this paper is to design rules to solve PB problems. To this end, we want to embed PB into JA and then use JA rules to compute budget allocations (see also Figure 1).

For a given PB instance, we introduce one proposition for each project to obtain $\boldsymbol{P}$. So we have a direct correspondence between budget allocations $A \subseteq \mathcal{P}$ and judgments $J \subseteq \boldsymbol{P}$, and thus also between PB profiles and JA profiles. Similarly, any JA outcome can be translated back into the PB setting.

Definition 2 (Outcome translation). *Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a PB instance and let $\Gamma \in \mathcal{L}_{\boldsymbol{P}}$ be an integrity constraint expressed over the atoms $\boldsymbol{P} = \{x_p \mid p \in \mathcal{P}\}$. The outcome translation $\tau : 2^{\boldsymbol{P}} \to 2^{\mathcal{P}}$ maps any judgment $J \in 2^{\boldsymbol{P}}$ to a budget allocation $A = \tau(J) = \{p \in \mathcal{P} \mid x_p \in J\}$.*

We moreover extend the outcome translation to sets $\mathcal{J} \subseteq 2^{\boldsymbol{P}}$ of judgments by stipulating that $\tau(\mathcal{J}) = \{\tau(J) \mid J \in \mathcal{J}\}$.

An *embedding* is a function $E : \mathcal{I} \to \mathcal{L}_{\boldsymbol{P}}$ that takes a PB instance as input and returns an integrity constraint (i.e., a JA instance). Given an embedding, we can now translate any input of a PB rule into an input for a JA rule, apply the JA rule, and finally translate the result obtained into a set of budget allocations (see Figure 1). However, to be meaningful, the integrity constraint should express the budget constraint of the PB instance. This is captured by the notion of correctness.

Definition 3 (Correct embedding). *An embedding $E : \mathcal{I} \to \mathcal{L}_{\boldsymbol{P}}$ is said to be correct if, for every PB instance $I \in \mathcal{I}$, we have $\tau(\mathfrak{J}(E(I))) = \mathcal{A}(I)$.*

3 Efficient Embeddings

In this section we present specific embeddings of enriched PB instances into JA. Given that the problem of computing outcomes for the JA rules defined in Section 2.2 is known to be highly intractable (Lang and Slavkovik, 2014; Endriss and De Haan, 2015), we need to ensure that PB instances are mapped into JA instances that permit efficient outcome determination. To this end, we first present a class of Boolean functions (to encode integrity constraints) for which the outcome determination can be solved efficiently.

3.1 Tractable Languages for JA

As shown by De Haan (2018), computing outcomes under Kemeny and Slater can be done efficiently when the integrity constraint is a Boolean circuit in decomposable negation normal form (DNNF). We are going to extend this result to all additive rules. But let us first recall the definition of a DNNF circuit (Darwiche and Marquis, 2002).

Definition 4 (DNNF circuits). *A circuit in negation normal form (NNF) is a rooted directed acyclic graph whose leaves are labelled with $\top, \bot, x$ or $\neg x$, for $x \in \boldsymbol{P}$ and whose internal nodes are labelled with $\land$ or $\lor$. A DNNF circuit C is an NNF circuit that is decomposable: for every conjunction in C, no two conjuncts share a common propositional variable.*

For a given JA rule F, we define the outcome determination problem as the following decision problem:

$\textsc{Outcome}(F)$
Input: An integrity constraint Γ, a judgment profile $\boldsymbol{J}$, and a subset of literals $L \subseteq Lit(\boldsymbol{P})$.
Question: Is there a $J \in F(\Gamma, \boldsymbol{J})$ such that $L \subseteq ext(J)$?

We now show that for any additive JA rule F we can solve $\textsc{Outcome}(F)$ efficiently when Γ is given as a DNNF circuit.

Theorem 1. *Let F be an additive JA rule defined w.r.t. some polynomial-time computable function f. Then $\textsc{Outcome}(F)$ is polynomial-time solvable if the integrity constraint Γ in the input is represented as a DNNF circuit.*

Proof. We show that when Γ is a DNNF circuit, we can use the Algebraic Model Counting (AMC) problem to solve $\textsc{Outcome}(F)$. Given a propositional formula $\varphi \in \mathcal{L}_{\boldsymbol{P}}$, a commutative semi-ring $\langle A, \oplus, \otimes, e^{\oplus}, e^{\otimes} \rangle$, and a labelling function $\lambda : Lit(\boldsymbol{P}) \to A$, the AMC problem is to compute:

$$\mathrm{AMC}(\varphi) = \bigoplus_{\substack{\alpha: \boldsymbol{P} \to \{0,1\} \\ \alpha \models \varphi}} \bigotimes_{\substack{\ell \in Lit(\boldsymbol{P}) \\ \alpha(\ell)=1}} \lambda(\ell).$$

The pair $\langle \oplus, \lambda \rangle$ is called *neutral* if and only for every propositional atom $x \in \boldsymbol{P}$, $\lambda(x) \oplus \lambda(\neg x) = e^{\otimes}$. Kimmig, Van den Broeck, and De Raedt (2017) proved that when φ is a DNNF circuit, $\oplus$ is idempotent, and $\langle \oplus, \lambda \rangle$ is neutral, then the AMC problem can be solved in polynomial time.

We now show that $\textsc{Outcome}(F)$ can be solved using the AMC problem when F is an additive rule. Consider the max-plus algebra—a commutative and idempotent semiring (Akian, Bapat, and Gaubert, 2006)—defined by $A = \mathbb{R} \cup \{-\infty, \infty\}$, $e^{\oplus} = -\infty$, and $e^{\otimes} = 0$, where $\oplus$ and $\otimes$ are the usual $+$ and $\max$ over $\mathbb{R} \cup \{-\infty, \infty\}$. For $\langle \oplus, \lambda \rangle$ to be neutral, we require $\max(\lambda(x), \lambda(\neg x)) = 0$ for every $x \in \boldsymbol{P}$.

Consider an additive JA rule F associated with f. For a profile $\boldsymbol{J}$ we introduce a labelling function $\lambda_{\boldsymbol{J}}(\cdot)$ defined as follows for every literal $\ell_x \in Lit(\boldsymbol{P})$ of the atom $x \in \boldsymbol{P}$:

$$\lambda_{\boldsymbol{J}}(\ell_x) = f(\boldsymbol{J}, \ell_x) - \max\left[f(\boldsymbol{J}, x), f(\boldsymbol{J}, \neg x)\right].$$

For every such labelling function, we can show that $\mathrm{argmax}_{J \in \mathfrak{J}(\Gamma)} \sum_{\ell_x \in ext(J)} \lambda_{\boldsymbol{J}}(\ell_x) = F(\Gamma, \boldsymbol{J})$. This implies that $\textsc{Outcome}(F)$ can be solved by solving the AMC problem twice: first for $\varphi = \Gamma$ and then for $\varphi = \Gamma'$, where Γ' is

obtained from Γ by fixing the value of the atoms as in L, the subset of literals given as input to OUTCOME(F).

Finally, observe that the pair $\langle \max, \lambda \rangle$ is neutral. Thus, AMC can be solved in polynomial time when φ is a DNNF circuit. Hence, the OUTCOME(F) problem can also be solved in polynomial time when Γ is a DNNF circuit. □

This general result immediately implies tractability of outcome determination for the rules we are interested in here and will allow us to use these rules to compute budget allocations for PB instances embedded into JA.

Corollary 2. *When the integrity constraint is represented as a DNNF circuit, then the problem* OUTCOME*(F) can be solved in polynomial time when F is either the Kemeny, the Slater, or the leximax rule.*

3.2 DNNF Circuit Embeddings

We now move on to the description of embeddings returning integrity constraints represented as DNNF circuits. In doing so, we follow De Haan (2018) but use a slight generalisation of his approach, allowing us to deal with PB instances with multiple resources. The basic idea is that every $\lor$-node in the DNNF circuit will represent the choice of selecting or not a given project. At each of these nodes, we need to keep track of the amount of resources already used to determine whether a project can be selected without exceeding the budget.

For a project index j and a vector of used quantities per resources $\boldsymbol{v} \in \mathbb{R}_{\geq 0}^d$, we introduce the $\lor$-node $N(j, \boldsymbol{v})$, corresponding to the situation where we previously made a choice on projects with indices 1 to $j-1$, and where for these choices we used resources according to $\boldsymbol{v}$. These nodes $N(j, \boldsymbol{v})$ are defined as follows.

$$
\begin{cases}
\top & \text{if } j = m + 1 \\
\lor \begin{pmatrix} (x_{p_j} \land N(j+1, \boldsymbol{v} + c(p_j))) \\ (\neg x_{p_j} \land N(j+1, \boldsymbol{v})) \end{pmatrix} & \text{if } \boldsymbol{v} + c(p_j) \leq \boldsymbol{b} \\
(\neg x_{p_j} \land N(j+1, \boldsymbol{v})) \lor (x_{p_j} \land \bot) & \text{otherwise}
\end{cases}
$$

For a PB instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$, the *tractable embedding TE(I)* returns the integrity constraint defined by $N(1, \boldsymbol{0}_d)$, where $\boldsymbol{0}_d$ denotes the vector of length d whose components are all equal to 0.

Proposition 3. *The tractable embedding TE is correct, and for any given PB instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ returns an integrity constraint TE(I) represented as a DNNF circuit of size in $\mathcal{O}(m \times |\{c(A) \mid A \subseteq \mathcal{A}(I)\}|)$.*

Proof. Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a PB instance, and Γ an integrity constraint such that $\Gamma = TE(I)$.

We first show that Γ is represented as DNNF circuit. First, observe that Γ is a Boolean circuit rooted in $N(0, \boldsymbol{0}_d)$. Next, observe that every $\lor$-node is of the form $(x \land \beta_1) \lor (\neg x \land \beta_2)$, where $x \in \boldsymbol{P}$ is a propositional atom and β_1, β_2 are either $\lor$-nodes, $\bot$, or $\top$. This implies that Γ is represented as an NNF circuit. Because each project is only considered once, the propositional atom corresponding to the project cannot appear in two distinct conjuncts. Hence, Γ is a DNNF circuit.

Observe that there are at most $m \times |\{c(A) \mid A \subseteq \mathcal{A}(I)\}|$ $\lor$-nodes in Γ—one for each $N(j, \boldsymbol{v})$ for which the budget is

not exceeded—all of them having at most two child $\land$-nodes. There are moreover $2m + 2$ leaves, one per literal and two for $\bot$ and $\top$, hence the size of the DNNF circuit.

We now show that the tractable embedding is correct. Observe that a branch leading to the $\bot$-leaf is chosen if and only if a project p_j exceeding the budget limit has been chosen. Hence, finding an assignment that does not lead to a $\bot$ leaf in Γ can only be done by selecting feasible projects. The set of such assignments defines the set of outcomes satisfying Γ, so $\tau(E(I)) \subseteq \mathcal{A}(I)$. Now, consider $A \in \mathcal{A}(I)$. Since A is feasible, it is clear that there exists a branch in the DNNF circuit Γ along which the selected projects correspond exactly to those that are in A. We thus have $\tau(E(I)) = \mathcal{A}(I)$. □

At this point, it should be noted that the exponential factor in the size of the embedding, namely $|\{c(A) \mid A \subseteq \mathcal{A}(I)\}|$, is bounded from above by the sum of the budget limits for each resource. Hence, the corresponding DNNF circuit is of size in $\mathcal{O}(m \times \sum_{r \in \mathcal{R}} b_r)$, making it pseudo-polynomial in the size of the PB instance. The embeddings can thus be argued to be efficient for realistic scenarios.

In the remainder of this section we investigate to what extent this approach allows us to introduce additional distributional constraints for PB.

3.3 Dependencies between Projects

We now consider the situation where some projects can only be achieved if some others are also achieved.

Take a PB instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$. We introduce a set of implications, $Imp \subseteq \mathcal{L}_{\boldsymbol{P}}$, linking projects together. A set of implications is a set of propositional formulas of the form $\ell_{x_p} \to \ell_{x_{p'}}$ for p and p' two projects in $\mathcal{P}$ with ℓ_{x_p} and $\ell_{x_{p'}}$ being the corresponding literals. Note that this corresponds to 2-CNF formulas. Such an implication indicates that if ℓ_{x_p} is positive (resp. negative), p can be selected (resp. not selected) only if p' is selected (resp. not selected) if $\ell_{x_{p'}}$ is positive (resp. negative). A budget allocation A satisfies the set of implications Imp if and only if the previously described semantics is satisfied. Moreover, we will write $\ell_{x_p} \to^* \ell_{x_{p'}}$ if there is a chain of implication in Imp linking ℓ_{x_p} to $\ell_{x_{p'}}$.

First of all, we show that finding a feasible budget allocation when there are implications between project is a NP-complete problem.

Proposition 4. *Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a PB instance and Imp a set of implications over I. Deciding whether there exists a feasible budget allocation for I satisfying Imp is NP-complete, and NP-hardness holds even for a single resource.*

Proof. It is straightforward to show that the problem of is in NP. To show that the problem is NP-hard, we reduce from the NP-complete problem 2-CNF MINIMAL MODEL (Ben-Eliyahu and Dechter, 1996).

2-CNF MINIMAL MODEL	
Input:	A formula $\varphi \in \mathcal{L}_{\boldsymbol{P}}$ in conjunctive normal form with exactly two literals per clauses and $k \in \mathbb{N}$.
Question:	Is there a model α such that $\alpha \models \varphi$ and $|\{p \in \boldsymbol{P} \mid \alpha(p) = 1\}| \leq k$?

Take an instance $\langle \varphi, k \rangle$ of 2-CNF MINIMAL MODEL. We construct the following participatory budgeting instance I. The set of resources is $\mathcal{R} = \{r\}$ with budget limit $\boldsymbol{b}_r = k$. There is one project per propositional atom in φ, $\mathcal{P} = \{p_x \mid x \in P\}$, and $c(p) = 1$ for every $p \in \mathcal{P}$. Finally, the set of implications Imp is the set of clauses in φ.

We claim that there exists a model of φ setting no more than k variables to true if and only if there exists a feasible budget allocation for I that satisfies the set of implications Imp. The proof is omitted for space reasons. $\square$

Based on this result, we cannot hope to find an embedding into a DNNF ciruit of polynomial size. However, we can still define an interesting parameterized embedding, in the spirit of parameterized complexity (Downey and Fellows, 2013). To that end we introduce the interconnection graph $G = \langle \mathcal{P}, E \rangle$ of a set of implications Imp where there is an edge $(p_i, p_j) \in E$ between p_i and p_j if and only if there exists an implication in Imp linking the two projects. We will then look at the pathwidth of this graph (Bodlaender, 1998).

Theorem 5. *Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a participatory budgeting instance and Imp a set of implications over I. There exists a correct embedding from I and Imp to an integrity constraint expressed as a DNNF circuit Γ whose size is in $\mathcal{O}\left(m \times |\{c(A) \mid A \subseteq \mathcal{A}(I)\}| \times 2^k\right)$, where k is the pathwidth of the interconnection graph of Imp.*

Proof. We will present our embedding and sketch the proof idea. In the following, we assume that $\boldsymbol{P} = \{x_p \mid p \in \mathcal{P}\}$.

Let $G = \langle \mathcal{P}, E \rangle$ be the interconnection graph of Imp. We order the projects according to the order in which they are introduced in an optimal path decomposition of G.[3] We introduce $\vee$-nodes $N(j, \boldsymbol{v}, L)$ where j is a project index, $\boldsymbol{v} \in \mathbb{R}^d_{\geq 0}$ a vector of used quantities per resource and $L \subseteq Lit(\boldsymbol{P})$ a subset of literals. Intuitively, the set L specifies the literals that we selected and that we should remember. If $j = m + 1$, then $N(j, \boldsymbol{v}, L) = \top$. If the positive literal x_{p_j} is implied by some literal in L w.r.t. *Imp*, then $N(j, \boldsymbol{v}, L) = N(j+1, \boldsymbol{v}+c(p_j), L \cup \{x_{p_j}\})$. Similarly, if the negative literal $\neg x_{p_j}$ is implied by some literal in L w.r.t. *Imp*, then $N(j, \boldsymbol{v}, L) = N(j+1, \boldsymbol{v}, L \cup \{\neg x_{p_j}\})$. Otherwise, if $\boldsymbol{v}+c(p_j) \leq \boldsymbol{b}$, then $N(j, \boldsymbol{v}, L) = (x_{p_j} \wedge N[j+1, \boldsymbol{v}+c(p_j), L \cup \{x_{p_j}\}]) \vee (\neg x_{p_j} \wedge N[j+1, \boldsymbol{v}, L \cup \{\neg x_{p_j}\}])$, and otherwise, $N(j, \boldsymbol{v}, L) = (x_{p_j} \wedge \bot) \vee (\neg x_{p_j} \wedge N[j+1, \boldsymbol{v}, L \cup \{\neg x_{p_j}\}])$.

The tractable embedding with dependencies, written $TE_{dep}(\cdot)$, refers to the integrity constraint defined by $N(j, \boldsymbol{0}_m, \emptyset)$.

By naively taking all possible sets $L \subseteq Lit(\boldsymbol{P})$, the number of $\vee$-nodes is exponential in m. However, since we ordered the projects according to a path decomposition, in each node $N(j, \boldsymbol{v}, L)$, we can "forget" all literals in L corresponding to a project that will not be introduced after project j in the path decomposition—thereby reducing the set L for node $N(j, \boldsymbol{v}, L)$ to size at most k. Then, for each j and $\boldsymbol{v}$, there are at most 2^k $\vee$-nodes $N(j, \boldsymbol{v}, L)$. We omit the detailed proof that the resulting DNNF circuit Γ has the required

---

[3]This can be done in in $\mathcal{O}(2^k)$ (Bodlaender and Kloks, 1996).

size, and that this embedding is correct—this proof is entirely similar to the proof of Proposition 3. $\square$

We conclude with a small detour to the field of Knowledge Compilation (see, e.g., Marquis, 2015). The question that we consider is equivalent to asking whether the conjunction of a 2-CNF formula with a DNNF circuit can be efficiently encoded as a DNNF circuit of polynomial size. It turns out to be impossible unless the Polynomial Hierarchy collapses.[4]

3.4 Quotas on Types of Projects

Another very natural constraint is to consider types and quotas over the projects. The idea is that the projects belong to various types (heath, education, environment to name a few) and that some quotas over these types are to be respected by the final budget allocation (at least two heath-related projects for instance). We model this idea by defining a type system.

Formally, for a given PB instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$, a type system is a tuple $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ where $\mathcal{T} \in 2^{2^{\mathcal{P}}}$ is a set of types, each type being a subset of projects; $\mathcal{Q} = \langle Q, +, 0, \leq_Q \rangle$ is an ordered group over which the quotas are expressed; $q : \mathcal{T} \to Q^2$ is a quota function such that for any type $t \in \mathcal{T}$, $q(t) = (a, b) \in Q^2$ with $a \leq_Q b$ and $f : \mathcal{T} \times \mathcal{A}(I) \to Q$ is a type aggregator.

For $t \in \mathcal{T}$ such that $q(t) = (a, b)$, we write $q(t)^- = a$ and $q(t)^+ = b$, which indicate the lower and upper quota for type t respectively. A budget allocation A is feasible if the quotas are respected—that is, if and only if for every type $t \in \mathcal{T}$, we have $q(t)^- \leq_Q f(t, A) \leq_Q q(t)^+$.

The type aggregator $f(\cdot)$ can be defined in several different ways. We provide two type aggregators that are very natural.

- **Cardinality-type aggregator**: the quotas express lower and upper bounds on the number of projects selected for each type. We have $Q = \mathbb{N}$, $\leq_Q$ is the usual order on $\mathbb{N}$, and the type aggregator is $f^{\text{card}}(t, A) = |A \cap t|$.

- **Cost-type aggregator**: the quotas define lower bound and upper bound on the total cost of the selected projects for each type. Here $Q = \mathbb{R}^d_{\geq 0}$, $\leq_Q$ is the component-wise order defined in the preliminaries, and the type aggregator is $f^{\text{cost}}(t, A) = \sum_{p \in A \cap t} c(p)$.

We first show that deciding whether there is a feasible budget allocation with a given type system is NP-complete, for both of the type aggregators.

Proposition 6. *Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a PB instance and $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ a type system over I. Deciding whether there exists a feasible budget allocation A is NP-complete when f is either the cardinality or the cost-type aggregator, and NP-hardness holds even for a single resource.*

Proof. It is straightforward to show that the problem is in NP. To show NP-hardness, we reduce from the NP-complete problem SET SPLITTING (Garey and Johnson, 1979).

---

[4]Due to space constraints, the proof is omitted. In short, one can prove that if 2-CNF formulas can be compiled into polynomial-size DNNF circuits, then the NP-complete problem CLIQUE is in P/poly, using similar techniques as Cadoli et al. (2002). This entails the Polynomial Hierarchy collapsing (Karp and Lipton, 1980).

Let $\langle C, S \rangle$ be an instance of SET SPLITTING. We construct a participatory budgeting instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ such that $\mathcal{R} = \{r\}$, and $\boldsymbol{b}_r = |S|$. There is one project per element in S, $\mathcal{P} = \{p_s \mid s \in S\}$, and $c(p_s) = 1$ for every $s \in S$. Thus, the budget limit can never be exceeded. The corresponding set of types is $\mathcal{T} = \{\{p_s \mid s \in c\} \mid c \in C\}$ and for a given $t \in \mathcal{T}$, the quota is $q(t) = (1, |t| - 1)$. With one resource and projects whose costs are in $\{0, 1\}$, the cardinality-type aggregator and the cost-type aggregator coincide.

We claim that $\langle C, S \rangle$ is a yes-instance of SET SPLITTING if and only if there exists a feasible budget allocation in the instance I with the previous type system. For a given partition of S, (S_1, S_2), a suitable corresponding budget allocation is $A = S_1$ (or equivalently $A = S_2$). We omit the full proof of this claim. □

Once again, this implies that no efficient embedding can be defined for this extension. In the following we present a parameterized embedding for PB with types and quotas.

The embedding works for any *additive* type aggregator $f : \mathcal{T} \times \mathcal{A}(I) \to Q$, that is, any type aggregator f for which there exists a *score type function* s that takes as input a project $p \in \mathcal{P}$ and returns an element in Q such that for every type $t \in \mathcal{T}$ and every allocation $A \in \mathcal{A}(I)$, $f(t, A) = \sum_{p \in A} s(p)$. The two type aggregators described above are both additive, with $s^{\text{card}}(p) = 1$ and $s^{\text{cost}}(p) = c(p)$.

Let I be an instance and $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ a type system over I, the overlap graph of the type system is the graph $G = \langle \mathcal{T}, E \rangle$, where there is an edge $\{t, t'\}$ in E if and only $t \cap t' \neq \emptyset$.

Theorem 7. *Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a PB instance and $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ a type system where f is an additive type aggregator defined w.r.t. the score type function s. There exists a correct embedding for I and $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ that returns an integrity constraint represented as a DNNF circuit whose size is in $\mathcal{O}\left(m \times |\{c(A) \mid A \subseteq \mathcal{A}(I)\}| \times k^*\right)$ where $k^* = \max_{t \in \mathcal{T}}(|\{f(t, A) \mid A \in \mathcal{A}(I)\}|)^{k+1}$ and where k is the pathwidth of the overlap graph of $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$.*

Proof. We use a similar strategy as for Theorem 5. The general idea is that because the type aggregator is additive, we can keep track of the current value of the quotas, and then we deciding whether a project can be selected or not we can check the current quota value before making our choice.

We will define the $\vee$-nodes $N(j, \boldsymbol{v}, \boldsymbol{q})$, where j is a project index, $\boldsymbol{v} \in \mathbb{R}_{\geq 0}^d$ a vector of used resources and $\boldsymbol{q} \in Q^{|\mathcal{T}|}$ is a vector of current quota value. If there is a type $t \in \mathcal{T}$ such that $\boldsymbol{q}_t >_Q q(t)^+$, then $N(j, \boldsymbol{v}, \boldsymbol{q}) = \bot$. If there is a type t whose projects have all been considered and such that $\boldsymbol{q} <_Q q(t)^-$, then $N(j, \boldsymbol{v}, \boldsymbol{q}) = \bot$. If $j = m + 1$ we have $N(j, \boldsymbol{v}, \boldsymbol{q}) = \top$. Otherwise, we define $N(j, \boldsymbol{v}, \boldsymbol{q})$ as follows. Let $T_{p_j} = \{t \in \mathcal{T} \mid p_j \in t\}$. Define $\boldsymbol{q}'$ such that $\boldsymbol{q}'_t = \boldsymbol{q}_t$ for every $t \notin T$, and such that $\boldsymbol{q}'_t = \boldsymbol{q}_t + s(p_j)$ for every $t \in T_{p_j}$. If $\boldsymbol{v} + c(p_j) \leq \boldsymbol{b}$, then $N(j, \boldsymbol{v}, \boldsymbol{q}) =$

$(x_{p_j} \wedge N[j+1, \boldsymbol{v} + c(p_j), \boldsymbol{q}']) \vee (\neg x_{p_j} \wedge N[j+1, \boldsymbol{v}, \boldsymbol{q}])$; otherwise, $N(j, \boldsymbol{v}, \boldsymbol{q}) = (x_{p_j} \wedge \bot) \vee (\neg x_{p_j} \wedge N[j+1, \boldsymbol{v}, \boldsymbol{q}])$. The tractable embedding for types and quotas, written TE_{quo}, refers to the integrity constraint defined by $N(1, \boldsymbol{0}_m, \boldsymbol{0}_{|\mathcal{T}|})$.

Similarly to the proof of Theorem 5, we can order the projects according to the ordering of types in a suitable path decomposition of the overlap graph. By doing so, in each node $N(j, \boldsymbol{v}, \boldsymbol{q})$, we can "forget" all types in $\boldsymbol{q}$ for which we already considered all projects—thereby reducing the number of nodes $N(j, \boldsymbol{v}, \boldsymbol{q})$ for each j and $\boldsymbol{v}$. □

The $\max_{t \in \mathcal{T}}(|\{f(t, A) \mid A \in \mathcal{A}(I)\}|)^{k+1}$ factor in the size of the integrity constraint can be very high. However, for the cardinality and the cost-type aggregators we have:

$$\max_{t \in \mathcal{T}} |\{f^{\text{card}}(t, A) \mid A \in \mathcal{A}(I)\}| = \max_{t \in \mathcal{T}} q(t)^+ \leq |\mathcal{P}|,$$

$$\max_{t \in \mathcal{T}} |\{f^{\text{cost}}(t, A) \mid A \in \mathcal{A}(I)\}| = \max_{t \in \mathcal{T}} q(t)^+ \leq \prod_{r \in \mathcal{R}} \boldsymbol{b}_r.$$

Although in general the problem of finding a feasible budget allocation satisfying the type system is hard, it is not when types are not overlapping. This is a very natural case to consider and in this case the pathwidth of the overlap graph would be 0, so the embedding would be efficient.

Corollary 8. *The tractable embedding for types and quotas is efficient when the type aggregator is additive and the overlap graph is the empty graph (types do not overlap).*

4 Enforcing Exhaustiveness

Amongst the very basic requirements of a budget allocation is that of *exhaustiveness* (Aziz, Lee, and Talmon, 2018), or *inclusion maximality* (Talmon and Faliszewski, 2019). It requires that the budget be used as much as possible.

Definition 5 (Exhaustiveness). *Given a PB problem instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$, a budget allocation $A \in \mathcal{A}(I)$ is said to be exhaustive if, for every project $p \in \mathcal{P} \setminus A$, there exists at least one resource $r \in \mathcal{R}$ such that $c(A \cup \{p\}, r) > \boldsymbol{b}_r$.*

For a given instance I, we denote by $\mathcal{A}_{EX}(I)$ the set of feasible and exhaustive budget allocations. An embedding $E : \mathcal{I} \to \mathcal{L}_{\boldsymbol{P}}$ is said to be exhaustive if for every instance $I \in \mathcal{I}$, we have $\tau(\mathfrak{J}(E(I))) \subseteq \mathcal{A}_{EX}(I)$. A JA rule F is said to be exhaustive if for every correct embedding E, every instance $I \in \mathcal{I}$ and every profile $\boldsymbol{A}$ it is the case that $\tau(F(E(I), \boldsymbol{A})) \subseteq \mathcal{A}_{EX}(I)$. Similarly, a PB rule is said to be exhaustive if it only returns exhaustive budget allocations. Finally, an exhaustive embedding E is correct if $\mathcal{A}_{EX}(I) = \tau(\mathfrak{J}(E(I)))$, for every instance I.

Because the scenarios typically modelled using JA are rather different from PB, the exhaustiveness axiom is not satisfied by the main JA rules. This has to do with the semantics of rejection (of a proposition) in the context of JA.

Proposition 9. *No majority-consistent JA rule is exhaustive.*

Proof. Consider a correct but not exhaustive embedding E (for instance TE). As E is not exhaustive, there exists a PB instance I such that there is at least one admissible JA outcome $J \in \mathfrak{J}(E(I))$ with $\tau(J) \notin \mathcal{A}_{EX}(I)$. Now consider a profile $\boldsymbol{A}$ with n agents in which $\lceil n/2 \rceil + 1$ agents only

approve of the projects in $\tau(J)$; the other agents are not constrained. On the JA side, the majoritarian outcome will be J. Since the majoritarian outcome is admissible, any majority-consistent rule F must return $\{J\}$ on $E(I)$ and A, which does not correspond to an exhaustive budget allocation. □

This result is far-reaching, because most JA rules have been specifically designed to be majority-consistent.[5]

To circumvent this problem and enforce exhaustiveness, we will investigate two approaches: either encoding exhaustiveness in the integrity constraint or designing new JA rules.

4.1 Exhaustive Embeddings

We introduce the *exhaustive tractable embedding*, which is an adaptation of the tractable embedding to maintain exhaustiveness when there is exactly one resource.

Consider a PB instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ with $\mathcal{R} = \{r\}$. Similarly to the previous embeddings, we define the $\vee$-nodes of the integrity constraint as $N(j, v, c^*)$, where j is a project index, v is the budget used in terms of resource r, and c^* is the cost of the cheapest non-selected project. If $j = m + 1$, we have $N(j, v, c^*) = \top$ in case $c^* > v$ and $N(j, v, c^*) = \bot$ otherwise. If $j \leq m$, we have $N(j, v, c^*) = \big(x_{p_j} \wedge N[j+1, \boldsymbol{v} + c(p_j), c^*]\big) \vee \big(\neg x_{p_j} \wedge N[j+1, \boldsymbol{v}, \min(c^*, c(p_j))]\big)$ in case $\boldsymbol{v} + c(p_j) \leq \boldsymbol{b}$ and $N(j, v, c^*) = \big(\neg x_{p_j} \wedge N[j+1, \boldsymbol{v}, c^*]\big) \vee (x_{p_j} \wedge \bot)$ otherwise. The exhaustive tractable embedding $ETE(\cdot)$ returns the integrity constraint defined by $N(1, 0, \max_{p \in \mathcal{P}} c(p))$.

Proposition 10. *The exhaustive tractable embedding is correct and exhaustive, and returns an integrity constraint Γ represented as a DNNF circuit of size $\mathcal{O}(m^2 \times |\{c(P) \mid P \subseteq \mathcal{P}\}|)$ for any $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$.*

The proof is omitted but very similar to earlier proofs. Just note that a budget allocation is exhaustive if and only if the cheapest non-selected project does not fit in it.

This embedding is only defined for instances with a single resource and, unfortunately, the idea does not generalise. The reason is that, when there are several resources, then there could be exponentially many "cheapest projects".

Proposition 11. *Let $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ be a PB instance with $|\mathcal{R}| \geq 2$. Then I cannot be embedded in polynomial time into an exhaustive DNNF circuit, unless the polynomial hierarchy collapses to the second level.*

Proof (sketch). The idea is that, if we were to find a DNNF circuit encoding, then we would be able to check in polynomial time whether an exhaustive budget allocation exists that selects at least k projects. We would then be able to solve in polynomial time any instance of size n of the NP-complete problem of 3-DIMENSIONAL MATCHING (Karp, 1972) using a single DNNF circuit of size polynomial in n. This would imply that NP $\subseteq$ P/poly, which means that the Polynomial Hierarchy collapses (Karp and Lipton, 1980). □

Since exhaustive embeddings cannot be used efficiently when there are multiple resources, we turn to another way to enforce exhaustiveness: Asymmetric judgment aggregation.

4.2 Asymmetric Judgment Aggregation Rules

In the context of PB, when an agent does not include a project in her approval ballot, this does not imply that she does not want to see the project being funded, but rather that it is not one of her top projects. Therefore, to implement PB via JA we need to adapt the JA rules so that not selecting a project (i.e., not accepting a proposition) is not interpreted as a rejection. To this end we introduce the new family of *asymmetric* JA rules. They avoid the symmetric treatment of acceptance and rejection common in most standard JA rules.

Definition 6 (Asymmetric Additive Rules). *Let F be an additive JA rule associated with $f : (2^P)^n \times Lit(\boldsymbol{P}) \to \mathbb{R}_{\geq 0}$. Then its asymmetric counterpart F_{asy} is the rule for which, for every integrity constraint Γ and every profile $\boldsymbol{J}$, we have:*

$$F_{asy}(\Gamma, \boldsymbol{J}) = \operatorname*{argmax}_{J \in \mathfrak{J}(\Gamma)} \sum_{\substack{\ell \in ext(J) \\ \ell \text{ is positive}}} f(\boldsymbol{J}, \ell) + \epsilon.$$

Here ϵ is any positive constant that is smaller than $\frac{1}{|\boldsymbol{P}|}$ of $\min\{f(\boldsymbol{J}, \ell) \neq 0 \mid \boldsymbol{J} \in (2^P)^n, \ell \in ext(J), \ell \text{ is positive}\}$.

Importantly, this definition applies only if f is $\mathbb{R}_{\geq 0}$-valued. The use of ϵ guarantees that accepting positive literals will always be more appealing than accepting negative ones, while being small enough so as to not impact the relative values of positive literals. Note that $\epsilon = \frac{1}{|\boldsymbol{P}|+1}$ is a suitable choice for the three rules defined near the end of Section 2.2.

Proposition 12. *Let F be an additive JA rule associated with an $\mathbb{R}_{\geq 0}$-valued function f. Then the asymmetric counterpart of F satisfies exhaustiveness.*

Proof. Executing F_{asy} involves computing a score for every admissible candidate outcome J. By definition, no negative literal in J can contribute to its score, while every positive literal makes a strictly positive contribution of at least ϵ. Thus, flipping a negative literal always results in an increased score. So F_{asy} only returns admissible judgments for which flipping any negative literal would violate the integrity constraint. This corresponds to exhaustiveness. □

Observe that the asymmetric counterpart of any additive rule is additive itself (and similarly for scoring rules, albeit not for AMRs). Finally, it is interesting to note that the asymmetric variant of the leximax rule is very similar to the well-known *greedy approval rule* for PB (Aziz and Shah, 2020).[6]

5 Axiomatic Analysis

Axioms are means for encoding formal properties related to the normative adequacy of mechanisms for collective decision-making (Thomson, 2001). Exhaustiveness is an example for such an axiom. In this section we investigate to

---

[5]Were it not for our assumption that every project must have at least one supporter (which rules out certain profiles), Proposition 9 could be strengthened to say that no unanimous JA rule is exhaustive (F is *unanimous* if $F(J, \ldots, J) = \{J\}$ for all judgments J).

[6]This rule selects greedily projects based on their approval score until the budget limit is reached. It actually corresponds to the asymmetric variant of the *ranked-agenda rule* (Lang et al., 2011).

| | KEM | | SLAT | | LEXIMAX | |
	sym	asym	sym	asym	sym	asym
Exhaustiveness	✗	✓	✗	✓	✗	✓
Limit Mono.	✗	✗	✗	✗	✗	✗
Discount Mono.	✓	✓	✓	✓	✓	✓
Splitting Mono.	✗	✓	✗	✓	✗	✓
Merging Mono.	✗	✗	✗	✗	✗	✗

Table 1: Axiomatic results: "sym" denotes the usual rule and "asym" its asymmetric counterpart.

what extent other important axioms proposed in the literature on PB are satisfied by JA rules when used for the purpose of PB. The results of this section are summarised in Table 1.

The literature on axioms for PB is still sparse. We focus on the monotonicity axioms introduced by Talmon and Faliszewski (2019), generalising their definitions to allow for multiple resources and irresolute rules. Formally, for a given PB axiom $\mathfrak{X}$, we say that the JA rule F *satisfies* $\mathfrak{X}$ *w.r.t. embedding E* if, for every PB instance I, the PB rule mapping $\boldsymbol{A}$ to $\tau(F(E(I), \boldsymbol{A}))$ for any given profile $\boldsymbol{A}$ satisfies $\mathfrak{X}$.

Moreover, for a resolute rule F, the axioms of Talmon and Faliszewski (2019) are usually stated as "when one moves from an instance/profile pair $(I, \boldsymbol{A})$ to another pair $(I', \boldsymbol{A}')$, then if $F(I, \boldsymbol{A})$ satisfies a certain property, $F(I', \boldsymbol{A}')$ should satisfy a corresponding property." We generalise these axioms to the irresolute case by requiring that, if *every* budget allocation returned by our rule for $(I, \boldsymbol{A})$ satisfies the property in question, then *every* budget allocation for $(I', \boldsymbol{A}')$ should satisfy the corresponding property. Note that Baumeister, Boes, and Seeger (2020) chose a different generalisation with existential instead of universal quantifiers.

The first axiom is called *limit monotonicity*. It states that after any increase in the budget limit that is not so substantial as to make some previously unaffordable project affordable, any funded project should continue to get funded. This axiom is closely related to that of committee monotonicity for multiwinner voting rules (Elkind et al., 2017).

Definition 7 (Limit monotonicity). *A PB rule F is said to be limit-monotonic if, for any two PB instances $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ and $I' = \langle \mathcal{R}, \boldsymbol{b}', \mathcal{P}, c \rangle$ with $\boldsymbol{b} \leq \boldsymbol{b}'$ and $c(p) \leq \boldsymbol{b}$ for all projects $p \in \mathcal{P}$, it is the case that $\bigcap F(I, \boldsymbol{A}) \subseteq \bigcap F(I', \boldsymbol{A})$ for all profiles $\boldsymbol{A}$.*

The following example shows that this axiom is not satisfied by any of the JA rules of interest, even when $|\mathcal{R}| = 1$.

Example 1. Consider the following three-agent profile for PB instance I with one resource, $b_1 = 3$, and three projects:

Project	p_1	p_2	p_3
Cost	2	2	1
Agents 1 and 2	✓	✓	✗
Agent 3	✗	✗	✓

The exhaustive allocations are $\{p_1, p_3\}$ and $\{p_2, p_3\}$. Consider the instance I' identical to I, except for the budget limit of $b_1' = 4$. For I', every rule would return $\{\{p_1, p_2\}\}$ when used with a correct exhaustive embedding. Project p_3 is then a witness of the violation of limit monotonicity. △

We move on to *discount monotonicity*, an axiom stating that, if the cost of a selected project is reduced, then that project should continue to be selected.

Definition 8 (Discount monotonicity). *A PB rule F is said to be discount-monotonic if, for any two PB instances $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ and $I' = \langle \mathcal{R}, \boldsymbol{b}', \mathcal{P}, c \rangle$ with $c(p) \geq c'(p)$ and $c(p') = c'(p')$ for all $p' \in \mathcal{P} \setminus \{p\}$ for some distinguished project $p \in \mathcal{P}$, it is the case that $p \in \bigcap F(I, \boldsymbol{A})$ implies $p \in \bigcap F(I', \boldsymbol{A})$ for all profiles $\boldsymbol{A}$.*

To study how JA rules deal with discount monotonicity, we introduce a new axiom for JA. This axiom is relevant for us, since it is a sufficient condition for discount monotonicity.

Definition 9 (Constraint monotonicity). *A JA rule F is said to be constraint-monotonic if, for any two integrity constraints $\Gamma, \Gamma' \in \mathcal{L}_{\boldsymbol{P}}$ with $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$ and any profile $\boldsymbol{J}$, it is the case that $F(\Gamma', \boldsymbol{J}) \setminus F(\Gamma, \boldsymbol{J}) \subseteq \mathfrak{J}(\Gamma') \setminus \mathfrak{J}(\Gamma)$.*

Lemma 13. *Every constraint-monotonic JA rule is discount-monotonic w.r.t. any correct embedding.*

Proof. Let F be a JA rule that is constraint-monotonic. Let E be a correct embedding. Consider the instances $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ and $I' = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c' \rangle$, where a project $p \in \mathcal{P}$ became cheaper from I to I' as in Definition 8.

Let $\boldsymbol{A}$ be an arbitrary profile with $p \in \bigcap \tau(F(E(I), \boldsymbol{A}))$. We need to show that $p \in \bigcap \tau(F(E(I'), \boldsymbol{A}))$. Observe that $\mathcal{A}(I) \subseteq \mathcal{A}(I')$. Because E is correct, we also have $\mathfrak{J}(E(I)) \subseteq \mathfrak{J}(E(I'))$. Moreover, for every $A \in \mathcal{A}(I') \setminus \mathcal{A}(I)$, we have $p \in A$ as only $c'(p)$ changed in I'. Hence, for every outcome $J \in \mathfrak{J}(E(I')) \setminus \mathfrak{J}(E(I))$, we have $p \in \tau(J)$. From constraint-monotonicity, we have that for every profile $\boldsymbol{A}$, $F(E(I'), \boldsymbol{A}) \subseteq F(E(I), \boldsymbol{A}) \cup \mathfrak{J}(E(I')) \setminus \mathfrak{J}(E(I))$. Hence, for every $J \in F(E(I'), \boldsymbol{A})$, we have $p \in \tau(J)$. □

Our axiom turns out to be satisfied by many JA rules.

Proposition 14. *Every additive rule is constraint-monotonic.*

Proof. Consider any additive rule F. Suppose, that F is *not* constraint-monotonic. Then there exist two integrity constraints Γ and Γ' with $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$ and a profile $\boldsymbol{J}$ for which there exists a $J \in F(\Gamma', \boldsymbol{J}) \setminus F(\Gamma, \boldsymbol{J})$ with $J \in \mathfrak{J}(\Gamma) \setminus \mathfrak{J}(\Gamma')$. As $J \notin F(\Gamma, \boldsymbol{J})$, there exists some $J' \in \mathfrak{J}(\Gamma)$ with a higher total score than that of J. Moreover, since $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$, this same J' would outperform J also under Γ'. This implies that $J \notin F(\Gamma', \boldsymbol{J})$, which is a contradiction, so we are done. □

Corollary 15. *The Kemeny, Slater, and leximax rules as well as their asymmetric counterparts are all discount-monotonic w.r.t. any correct embedding.*

The last two axioms we consider deal with situations where projects are split into subprojects (and the dual operation of merging projects). First, *splitting monotonicity* states that, if a selected project is split into a set of projects approved by the same agents, then some of these new projects should still be selected. The axiom of *merging monotonicity* expresses a similar condition when merging projects.

Given a PB instance $I = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}, c \rangle$ and a profile $\boldsymbol{A}$, we say that $I' = \langle \mathcal{R}, \boldsymbol{b}, \mathcal{P}', c' \rangle$ and $\boldsymbol{A}'$ are *the result of splitting*

project $p \in \mathcal{P}$ into P (with $P \cap \mathcal{P} = \emptyset$), if $\mathcal{P}' = (\mathcal{P} \setminus \{p\}) \cup P$, for all $p' \in P$, $c(p) \neq \mathbf{0}_d$, $c'(P) = c(p)$, $c'(p') = c(p')$ for all $p' \in \mathcal{P}' \setminus P$, $A'_i = A_i$ for all $i \in \mathcal{N}$ with $p \notin A_i$, and $A'_i = (A_i \setminus \{p\}) \cup P$ for all other $i \in \mathcal{N}$. We also say that I and $\mathbf{A}$ are *the result of merging P into p given I' and $\mathbf{A}'$*.

Definition 10 (Splitting monotonicity). *A PB rule F is said to be splitting-monotonic if, for any two PB instances $I = \langle \mathcal{R}, \mathbf{b}, \mathcal{P}, c \rangle$ and $I' = \langle \mathcal{R}, \mathbf{b}, \mathcal{P}', c' \rangle$ with corresponding profiles $\mathbf{A}$ and $\mathbf{A}'$ such that I' and $\mathbf{A}'$ are the result of splitting project p into P given I and $\mathbf{A}$, it is the case that if $p \in \bigcap F(I', \mathbf{A})$ then $A' \cap P \neq \emptyset$ for all $A' \in F(I', \mathbf{A})$.*

Definition 11 (Merging monotonicity). *A PB rule F is said to be merging-monotonic if, for any two PB instances $I = \langle \mathcal{R}, \mathbf{b}, \mathcal{P}, c \rangle$ and $I' = \langle \mathcal{R}, \mathbf{b}, \mathcal{P}', c' \rangle$ with corresponding profiles $\mathbf{A}$ and $\mathbf{A}'$ such that I' and $\mathbf{A}'$ are the result of merging project set P into project p given I and $\mathbf{A}$, it is the case that $P \subseteq \bigcap F(I, \mathbf{A})$ implies $p \in \bigcap F(I', \mathbf{A})$.*

We first show that splitting monotonicity is satisfied by the asymmetric counterpart of any AMR.

Proposition 16. *Every asymmetric counterpart of an AMR is splitting-monotonic.*

Proof. Let F be the asymmetric counterpart of an AMR and let E be a correct exhaustive embedding. Consider a PB instance $I = \langle \mathcal{R}, \mathbf{b}, \mathcal{P}, c \rangle$ and a profile $\mathbf{A}$. Let $I' = \langle \mathcal{R}, \mathbf{b}, \mathcal{P}', c' \rangle$ and $\mathbf{A}'$ be the instance and profile resulting from splitting $p \in \bigcap \tau(F(E(I), \mathbf{A}))$ into P.

Consider any outcome $J_1 \in F(E(I), \mathbf{A})$. Note that for all $J \in \mathfrak{J}(E(I)) \cap \mathfrak{J}(E(I'))$, we have $p \notin \tau(J)$ and $\tau(J) \cap P = \emptyset$. Because $p \in \bigcap \tau(F(E(I), \mathbf{A}))$, this implies that J_1 has a higher total score than any $J_2 \in \mathfrak{J}(E(I)) \cap \mathfrak{J}(E(I'))$.

Consider now a possible outcome $J'_1 = (J_1 \setminus \{p\}) \cup \{p'\}$ for some $p' \in P$. Based on the definition of $\mathbf{A}'$, it is clear that $n_x^{\mathbf{A}} = n_x^{\mathbf{A}'}$ for every $x \in J_1 \setminus \{x_p\}$ and that $n_{x_p}^{\mathbf{A}} = n_{x_{p'}}^{\mathbf{A}'}$. Hence, because F is the asymmetric counterpart of an AMR, J_1 and J_2 have the same total score. Indeed, the total score for an AMR only depends on the approval score of each literal and because F is asymmetric we only consider positive literals. This implies that J'_1 has a higher total score than any $J_2 \in \mathfrak{J}(E(I)) \cap \mathfrak{J}(E(I'))$. Thus, $\mathfrak{J}(E(I)) \cap F(E(I'), \mathbf{A}) = \emptyset$. As for every $J' \in \mathfrak{J}(E(I')) \setminus \mathfrak{J}(E(I))$ we have $P \cap \tau(J') \neq \emptyset$, every outcome returned by F would have a nonempty intersection with P. □

While this last result is promising, it unfortunately does not extend to symmetric rules.

Example 2. Consider the following pairs of instances and three-agent profiles: I and $\mathbf{A}$ on the left and I' and $\mathbf{A}'$ on the right. Both involve just one resource, with $b_1 = 4$.

Project	p_1	p_2	p_3	p_4 ‖	p_1	$\{p_2^1, p_2^2, p_2^3, p_2^4\}$	p_3	p_4
Cost	2	2	1	1 ‖	2	0.5	1	1
Agents 1 and 2	✓	✗	✗	✗ ‖	✓	✗	✗	✗
Agent 3	✗	✓	✓	✓ ‖	✗	✓	✓	✓

Observe that I' and $\mathbf{A}'$ are the result of splitting p_2 into $\{p_2^1, p_2^2, p_2^3, p_2^4\}$, given I and $\mathbf{A}$. We leave the relevant calculations to the reader, but the Kemeny, the Slater, and the

leximax rules would all return $\{\{p_1, p_2\}\}$ for $(I, \mathbf{A})$ when used with a correct exhaustive embedding. However, they would return $\{\{p_1, p_3, p_4\}\}$ for $(I', \mathbf{A}')$. Hence, p_2 is a witness of a violation of splitting monotonicity. △

We finally investigate merging monotonicity. It turns out that none of the rules we are considering in this paper satisfy it. A simple counterexample is described in the following. Consider an instance with one resource, a budget of 4, and six projects: four of cost 1 and two of cost 2. Consider now a profile with a single agent approving of every project. Then all of our rules (Kemeny, Slater, and leximax) would select the four projects of cost 1. Now, if the four projects of cost 1 are merged into one project of cost 4, all our rules would select the two projects of cost 2. This is a violation of merging monotonicity as the newly created project is not selected. As the only agent approves of every project, the same hold for the asymmetric counterpart of the rules.

To conclude this section, we shortly discuss the overall axiomatic picture for JA rules. The most striking results are that none of our rules satisfy limit and merging monotonicity. For limit monotonicity, it should be noted that no PB rule we know of satisfies it (Talmon and Faliszewski, 2019). It seems to be too strong a requirement. For merging monotonicity, the situation is less clear-cut: Some PB rules satisfy it but none that are widely used. Other axiomatic results are in line with Talmon and Faliszewski (2019). Overall, JA rules perform similarly to other PB rules in normative terms.

6 Conclusion

We have proposed an efficient way of solving PB problems by means of JA rules. The richness of the JA framework allowed us to develop embeddings for generalised forms of PB. While the resulting problems are computationally hard in general, we nevertheless were able to present useful parameterized embeddings for them. Regarding the axiomatic properties of JA rules for PB, we observed that a naïve way of embedding PB into JA leads to rules that violate the crucial exhaustiveness requirement of PB. We then suggested two ways of enforcing exhaustiveness: Through exhaustive embeddings or by using asymmetric JA rules. We also analysed some common JA rules and their asymmetric counterparts in view of monotonicity axioms for PB and found that the asymmetric rules fare better than the original rules.

In terms of future work, it would be interesting to study more PB axioms, to allow us to better differentiate between different JA rules. Indeed, for now, the Kemeny, Slater, and leximax rule cannot be distinguished based on the axioms we studied. A particularly exciting direction would be to investigate proportionality axioms such as those introduced by Aziz, Lee, and Talmon (2018) and Haret et al. (2020).

Beyond its immediate significance to the theory and practice of PB, we believe our work also highlights some important aspects of working with different frameworks for collective decision making. The high expressive power of JA permits us to encode many problems of practical interest as well as properties and constraints. Finding efficient ways of solving decision problems embedded into JA can be hard, but once identified, these methods allow for great flexibility.

Acknowledgements

References

Akian, M.; Bapat, R.; and Gaubert, S. 2006. Max-plus algebra. In Hogben, L., ed., *Handbook of Linear Algebra*. Chapman and Hall, London. chapter 35.

Aziz, H., and Shah, N. 2020. Participatory budgeting: Models and approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*. Springer.

Aziz, H.; Lee, B. E.; and Talmon, N. 2018. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 23–31.

Baumeister, D.; Boes, L.; and Seeger, T. 2020. Irresolute approval-based budgeting. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1774–1776.

Ben-Eliyahu, R., and Dechter, R. 1996. On computing minimal models. *Annals of Mathematics and Artificial Intelligence* 18(1):3–27.

Bodlaender, H. L., and Kloks, T. 1996. Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *Journal of Algorithms* 21(2):358–402.

Bodlaender, H. L. 1998. A partial k-arboretum of graphs with bounded treewidth. *Theoretical Computer Science* 209(1-2):1–45.

Cabannes, Y. 2004. Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization* 16(1):27–46.

Cadoli, M.; Donini, F. M.; Liberatore, P.; and Schaerf, M. 2002. Preprocessing of intractable problems. *Inf. Comput.* 176(2):89–120.

Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 629–646.

Darwiche, A., and Marquis, P. 2002. A knowledge compilation map. *Journal of Artificial Intelligence Research* 17:229–264.

Dietrich, F., and List, C. 2007. Arrow's Theorem in judgment aggregation. *Social Choice and Welfare* 29(1):19–33.

Dietrich, F. 2014. Scoring rules for judgment aggregation. *Social Choice and Welfare* 42(4):873–911.

Downey, R. G., and Fellows, M. R. 2013. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer.

Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare* 48(3):599–632.

Endriss, U., and de Haan, R. 2015. Complexity of the winner determination problem in judgment aggregation: Kemeny, Slater, Tideman, Young. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Endriss, U.; Grandi, U.; de Haan, R.; and Lang, J. 2016. Succinctness of languages for judgment aggregation. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Endriss, U.; Grandi, U.; and Porello, D. 2012. Complexity of judgment aggregation. *Journal of Artificial Intelligence Research* 45:481–514.

Endriss, U. 2016. Judgment aggregation. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 17.

Endriss, U. 2018. Judgment aggregation with rationality and feasibility constraints. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 946–954.

Everaere, P.; Konieczny, S.; and Marquis, P. 2014. Counting votes for aggregating judgments. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1177–1184.

Fain, B.; Goel, A.; and Munagala, K. 2016. The core of the participatory budgeting problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*, 384–399.

Fain, B.; Munagala, K.; and Shah, N. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 575–592.

Freeman, R.; Pennock, D. M.; Peters, D.; and Wortman Vaughan, J. 2019. Truthful aggregation of budget proposals. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, 751–752.

Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability*, volume 29. W.H. Freeman.

Grandi, U., and Endriss, U. 2011. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 204–209.

de Haan, R. 2018. Hunting for tractable languages for judgment aggregation. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 194–203.

Haret, A.; Lackner, M.; Pfandler, A.; and Wallner, J. P. 2020. Proportional belief merging. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2822–2829.

Jain, P.; Sornat, K.; and Talmon, N. 2020. Participatory budgeting with project interactions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*.

Karp, R. M., and Lipton, R. J. 1980. Some connections between nonuniform and uniform complexity classes. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC)*, 302–309.

Karp, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*. Springer. 85–103.

Kimmig, A.; van den Broeck, G.; and de Raedt, L. 2017. Algebraic model counting. *Journal of Applied Logic* 22:46–62.

Lang, J., and Slavkovik, M. 2013. Judgment aggregation rules and voting rules. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 230–243.

Lang, J., and Slavkovik, M. 2014. How hard is it to compute majority-preserving judgment aggregation rules? In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, 501–506.

Lang, J.; Pigozzi, G.; Slavkovik, M.; and van der Torre, L. 2011. Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 238–246.

List, C., and Pettit, P. 2002. Aggregating sets of judgments: An impossibility result. *Economics & Philosophy* 18(1):89–110.

Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.

Marquis, P. 2015. Compile! In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 4112–4118.

Miller, M. K., and Osherson, D. 2009. Methods for distance-based judgment aggregation. *Social Choice and Welfare* 32(4):575–601.

Nehring, K., and Pivato, M. 2019. Majority rule in the absence of a majority. *Journal of Economic Theory* 213–257.

Pigozzi, G. 2006. Belief merging and the discursive dilemma: An argument-based account to paradoxes of judgment aggregation. *Synthese* 152(2):285–298.

Shah, A., ed. 2007. *Participatory budgeting*. Public Sector Governance and Accountability Series. Washington, DC: The World Bank.

Talmon, N., and Faliszewski, P. 2019. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2181–2188.

Thomson, W. 2001. On the axiomatic method and its recent applications to game theory and resource allocation. *Social Choice and Welfare* 18(2):327–386.