

Funding Public Projects: A Case for the Nash Product Rule

FLORIAN BRANDL, Princeton University

FELIX BRANDT, Technische Universität München

DOMINIK PETERS, Carnegie Mellon University

CHRISTIAN STRICKER, Technische Universität München

WARUT SUKSOMPONG, University of Oxford

We study a mechanism design problem where a community of agents wishes to fund public projects via voluntary monetary contributions by the community members. This serves as a model for participatory budgeting without an exogenously available budget, as well as donor coordination when interpreting charities as public projects and donations as contributions. Our aim is to identify a mutually beneficial distribution of the individual contributions. In the preference aggregation problem that we study, agents report linear utility functions over projects together with the amount of their contributions, and the mechanism determines a socially optimal distribution of the money. We identify a specific mechanism—the Nash product rule—which picks the distribution that maximizes the product of the agents’ utilities. This rule is Pareto efficient, and we prove that it satisfies attractive incentive properties: the Nash rule spends an agent’s contribution only on projects the agent finds acceptable, and it provides strong participation incentives. We also discuss issues of strategyproofness and monotonicity.

1 INTRODUCTION

Many cities run *participatory budgeting* programs, in which they set aside a fraction of their budget and allow their residents to vote for project proposals—such as school renovations, park construction, or accessibility additions—usually using an online voting interface. This process allows city governments to discover which projects the public cares most about, and thus can lead to a more efficient use of the budget [Cabannes, 2004].

In a city, there is an authority which provides an exogenous budget to be spent. For other communities without such a budget, a similar system could help determine the best public projects to implement, but these projects will need to be funded via voluntary monetary contributions from community members. Examples of such communities include residents of an apartment complex (who want to coordinate spending on gardening in a courtyard, or on cleaning services), homeowners on a city street (to coordinate tree care, snow removal, or security patrols), or student clubs in a university (to coordinate funding for events and meet-ups). A simple solution for these use cases is to set up an account for each proposed project and invite community members to contribute to projects as they wish. But this can lead to Pareto dominated outcomes. For example, in an apartment complex consisting of several buildings, residents might each give to projects specific to their own building (because for each individual, those projects provide the greatest benefit), when all residents would agree that it would be better to pool the money for projects benefiting the entire complex. To discover this kind of improvement, a voting-based system seems promising: if agents report their preferences to a mechanism, it can aggregate the preferences to identify a socially optimal funding scheme. Our aim is to develop the theory of such systems, and to show that a mechanism based on maximizing the Nash product, that is, the product of the agents utilities, provides remarkably strong incentives for community members to contribute financially.

The public projects funded by these voting-based funding systems could, in particular, be charities. Many companies run programs to encourage their employees to give charitably. For example, Microsoft runs an annual charity matching program, and in 2019, employees donated a total of \$181 million to charities.¹ Donations to charities have the characteristics of public projects, since

¹<https://www.microsoft.com/en-us/corporate-responsibility/philanthropies/employee-engagement>

each donor prefers more money being allocated to charities she perceives as impactful. A voting-based system like the one we will propose can help the employees of a company to coordinate their giving, and thereby identify a more efficient money distribution. The system can also be used in higher-stakes applications involving major philanthropic foundations. Notably, the Open Philanthropy Project, which grants more than \$100 million a year to various organizations, has called on academics to develop mechanisms to combine different staff members’ views on the most effective giving opportunities, and also to help coordinate the giving of different philanthropic organizations [Muehlhauser, 2017]. Interest in donor coordination mechanisms has also been expressed in the effective altruism community [Peters, 2019].

We imagine the typical user experience in a funding system as follows. Users are presented with a list of proposed public projects. Among these, each user designates an arbitrary number of projects to be *acceptable*, meaning that the user will be satisfied with a distribution as long as the mechanism directs her contribution to these projects. Optionally, the user can specify fine-grained preferences among acceptable projects; in this paper we will allow cardinal utilities which are linear and weakly increasing in the amount spent on a project and are additive between projects. (Unlike some other works on participatory budgeting, we focus on divisible projects that can use an arbitrary amount of funding; this model is particularly suitable for charities.) Finally, the user is asked to commit a monetary amount to contribute to the funding system, and this money is transferred to the clearinghouse. The mechanism then aggregates the submitted preference information, decides how much money each project should receive, and then disburses the collected contributions to the projects accordingly.

As described, this system with pre-collected contributions requires the participants to trust that the system will in fact pay out the collected money (and not steal it). Our discussion is also applicable when this trust cannot be established (or when decentralization is desired); in such cases, we can interpret the mechanism output as recommendations to the agents about how they should distribute their contribution. However, this comes with the risk that participants will not follow these recommendations, and thereby jeopardize the efficiency and fairness properties of the mechanism’s overall distribution.

While many preference aggregation mechanisms are imaginable, we will focus on a specific mechanism called the *Nash product rule* and argue that it is remarkably well-adapted to the applications we have outlined. This mechanism performs a type of social welfare maximization. Specifically, it selects a distribution δ of the overall collected contributions which maximizes the product of voters’ utilities, with voters weighted by the size of their contribution. A voter i ’s utility in distribution δ is defined as $u_i(\delta) := \sum_{x \in A_i} u_i(x)\delta(x)$, where A_i is the set of projects acceptable to i , $u_i(x)$ is the utility per unit of money that i assigns to project x , and $\delta(x)$ is the amount spent on x . (The formal definition implicitly sets the utility of unacceptable projects to 0.) Then the optimization objective will be $\prod_{i \in N} u_i(\delta)^{C_i}$, where C_i is the size of i ’s contribution. The idea of maximizing the product of voter utilities originates in the Nash bargaining solution, which is why we refer to this mechanism as the Nash product rule. The Nash product has recently become popular among researchers in various fields, including in the allocation of indivisible private items [Caragiannis et al., 2016], committee elections [Lackner and Skowron, 2018], and participatory budgeting [Fain et al., 2016, 2018]. In all these areas, the Nash rule satisfies strong fairness and proportionality properties appropriate to the setting.

Since we are considering a model with voluntary contributions, certain incentive properties that have not been studied in other models become crucial for us. For starters, given the interpretation of voter preferences, it is important that the contribution of user i is only spent on projects that are acceptable to i —otherwise, the user can justifiably complain that the money is being misused. In the case where we interpret the mechanism output only as a recommendation, this property

becomes essential, since a recommendation to send money to an unacceptable project is likely to be ignored. We say that a distribution δ is *implementable* if we can decompose it as $\delta = \delta_1 + \dots + \delta_n$ such that δ_i spends exactly C_i , and spends it only on projects acceptable to i . In Theorem 4.1, we prove that the distribution selected by the Nash product is always implementable.

Our main result concerns participation incentives offered by the Nash product; we wish to assure agents that it is beneficial for them to contribute their money to the mechanism. Intuitively, we want to guarantee that if an agent i contributes additional money to the mechanism, then (a) this money will be distributed only to acceptable projects and (b) the money that was already in the system is distributed in a way that gives i at least as much utility as before the increased contribution. Of course, (a) will be true for any implementable mechanism; it is condition (b) that makes pooling participation difficult to achieve. Consider for instance the mechanism that maximizes leximin welfare [see, e.g., Moulin, 2003] among all implementable distributions. This mechanism will fail (b), because it will take additional money as an opportunity to increase the welfare of low-utility agents, who may have interests different from i . We introduce an axiom called *pooling participation*, requiring that i 's utility, upon contributing, is at least the utility i would have derived from others' money had i not participated, plus the worst-case utility i gets from spending her contribution on an acceptable project.² We establish in Theorem 4.2 that the Nash product rule satisfies pooling participation. Our proof of this fact is rather involved, because we need to reason about the trajectory of the maximizer of the Nash product as agent i 's contribution continuously increases. Using the Taylor expansion of the Nash objective and a carefully chosen perturbation, we are able to estimate the derivative of i 's utility as i 's contribution changes, which after an integration yields the result.

Since the Nash product rule maximizes a monotonic function of voter utilities, its outcome is guaranteed to be Pareto efficient. This feature is important in order to gain user acceptance, and to avoid sub-optimal outcomes such as those discussed earlier. While Pareto efficiency is often considered a minimal condition, it can be difficult to achieve by mechanisms that need to satisfy additional incentive properties. Implementation and pooling participation can very easily be satisfied by non-efficient mechanisms;³ the remarkable feature of the Nash product is that it provides these incentive properties while ensuring Pareto efficiency. The Nash product is the only rule known to us that is both Pareto efficient and satisfies pooling participation. While we can construct other rules that are Pareto efficient and implementable, these are artificial, and Nash is the only natural such rule that we are aware of.

While the Nash product rule is incentive-compatible in the sense of implementability and pooling participation, it still allows for other strategic behavior. In particular, agents may have an incentive to misrepresent their utility functions. Because the Nash product penalizes distributions in which some agents obtain very low utility, it can be beneficial for agents to pretend to like popular projects less, or even to mark them as unacceptable. This will make the Nash product rule worry that those agents will be underserved, and thus increase the funding of other projects acceptable to them. Unfortunately, by a result due to Hylland [1980, Thm. 2], every Pareto efficient mechanism will be vulnerable to misrepresentation of preferences. In Section 5, we sharpen this result by showing that every (symmetric) mechanism that is Pareto efficient and implementable can be manipulated by the same strategy effective in the case of the Nash product rule, namely by understating one's utilities. We additionally note that the Nash product rule fails a natural monotonicity property, and

²We also consider a version with best-case spending, but show that this version is incompatible with Pareto efficiency.

³For example, by the uncoordinated "mechanism" that tells each user to direct her entire contribution to the project to which she assigned the highest utility, without doing any aggregation of preferences.

we leave the intriguing open question whether there exists a mechanism that is monotonic and matches the other axiomatic virtues of the Nash rule.

In Section 6, we discuss strengthened versions of implementability and pooling participation that are appropriate for a utility model in which agents do not distinguish only between acceptable and unacceptable projects, and when agents have maximally attractive outside options. While these two stronger properties may be desirable in the context of donor coordination, we show that each of them is incompatible with Pareto efficiency.

Overall, our discussion suggests that the Nash rule is a prime candidate for deployment in participatory budgeting settings with voluntary contributions. It combines Pareto efficiency with strong incentive properties, and as we will mention in Section 2, it also satisfies important fairness and proportionality properties. Finally, the rule is simple to define, can be computed to arbitrary precision using convex programming, and because it is implementable, its distribution decisions can be easily understood by users.

2 RELATED WORK

Participatory budgeting has mostly been studied under the assumption that the budget is provided by an outside source (such as the city government); the recent literature on aggregation mechanisms was surveyed by Aziz and Shah [2020]. In the most common model, projects come with a fixed cost, and they can either be fully funded or not at all. For this model, researchers have studied issues of preference elicitation [Benade et al., 2017, Goel et al., 2019], monotonicity properties [Talmon and Faliszewski, 2019], and proportionality concepts [Aziz et al., 2018, Fain et al., 2018]. In the case where each project has the same cost (the *unit cost* case), this model captures committee elections, which have received much recent attention. The literature on committee elections compares different voting rules, studies their winner determination problems, and considers axioms; for an overview, see Faliszewski et al. [2017].

In a related model, the projects can receive an arbitrary amount of funding (like in our paper) but the budget is still exogenous and of fixed size. In the computer science literature, the study of this setting was initiated by Fain et al. [2016], who argued that allocations in *Lindahl equilibrium* [Foley, 1970] are particularly desirable. The Lindahl equilibrium is a market equilibrium in an artificial market for public goods. In these markets, each agent faces personalized prices (usually interpreted as taxes) for the public goods, and in equilibrium each agent demands the same bundle of public goods. Under standard assumption, Foley [1970] showed that a Lindahl equilibrium exists (by reduction to the Arrow–Debreu private goods case), and is Pareto efficient. He also showed that equilibrium allocations are in the *core*: no coalition of agents can afford (using only a fraction of the budget proportional to their size) an allocation that each coalition member prefers to the equilibrium. For the case of additive linear utilities, Fain et al. [2016] proved that the Nash product rule yields an allocation in Lindahl equilibrium, and hence is in the core.⁴ The core can be interpreted as guaranteeing agents proportional representation: if a fraction of $\alpha\%$ of voters assign positive utility only to some set A' of projects, then the Nash product rule will spend at least $\alpha\%$ of the budget on projects in A' .

Our model is also closely related to *probabilistic social choice* [see, e.g., Brandt, 2017]: a division of a fixed budget among projects is the same as a probability distribution over alternatives. Bogomolnaia et al. [2005] have studied the special case of dichotomous preferences in probabilistic social choice and introduced several rules (including the Nash product rule). Their results can be interpreted as results for participatory budgeting with projects that can receive an arbitrary amount of funding and

⁴This mirrors the canonical result that the Nash product yields an equilibrium in Fisher markets for private goods under additive valuations [Eisenberg and Gale, 1959].

where each voter assigns only utilities of either 0 or 1 to projects. Later, Duddy [2015] introduced another rule for this setting, and Aziz et al. [2019] undertook a more thorough axiomatic analysis. In particular, Aziz et al. [2019] showed that the Nash product rule guarantees *average fair share*: for any group of $\alpha\%$ of the voters which is *cohesive* (there is a project that they all approve), for an average group member, the Nash product spends at least $\alpha\%$ of the budget on approved projects. They also proved that Nash satisfies *strict participation*. This property, which was introduced by Brandl et al. [2015], makes sense for the fixed-budget setting, but in our variable-budget interpretation it is weak.⁵ Our main result that Nash satisfies pooling participation implies Aziz et al.’s [2019] result.

The key difference between all of the above literatures and our model is that in our model, the individual contributions to the endowment are *owned* by the agents. This suggests the definitions of the axioms of implementability and pooling participation, which—to the best of our knowledge—have not been considered in previous work. On the other hand, as we already mentioned, the Nash product rule has featured prominently in many streams of research, and has been shown to have desirable fairness and efficiency properties in many settings [e.g., Aziz et al., 2019, Caragiannis et al., 2016, Conitzer et al., 2017, Fain et al., 2016, 2018].

A major application of our model is for allocating donations to charities. Conitzer and Sandholm [2004, 2011] have also considered this application from a mechanism design perspective. They let agents incentivize other agents to donate more by devising “matching offers”, where a donation is made conditional on how much other agents donate, and to which charities. Rather than using matching offers, in our model, donations are incentivized by (implicitly) letting agents vote over how others donate their money, and weighting votes by the size of the voter’s contribution.

3 MODEL AND AXIOMS

Let A be a finite set of m projects (e.g., charities or joint activities) and N a finite set of n agents. For each $i \in N$, agent i ’s *contribution* is $C_i \in \mathbb{R}_{>0}$, giving rise to a total *endowment* of $C = \sum_{i=1}^n C_i$.

A *distribution* δ is a function that describes how some amount V (e.g., an individual contribution or the entire endowment) is distributed among the projects, so $\delta : A \rightarrow \mathbb{R}_{\geq 0}$ such that $\sum_{x \in A} \delta(x) = V$. For convenience, we write distributions as linear combinations of projects, so that $a + 2b$ denotes for distribution δ with $\delta(a) = 1$ and $\delta(b) = 2$. The set of all distributions of value V is denoted by $\Delta(V)$. Each distribution of the endowment $\delta \in \Delta(C)$ can be divided into n individual distributions $\delta_i \in \Delta(C_i)$ such that $\delta = \sum_{i \in N} \delta_i$. Clearly, the division of a distribution into individual distributions is not unique.

Apart from her contribution, each agent possesses a utility function $u_i : A \rightarrow \mathbb{R}$, where $u_i(a)$ describes how much utility an agent derives if one unit of money goes to project a . So agent i ’s utility for a distribution $\delta \in \Delta(C)$ is

$$u_i(\delta) = \sum_{x \in A} \delta(x) \cdot u_i(x).$$

A project is said to be *acceptable* by an agent if it gives her positive utility, and *unacceptable* otherwise. We assume that the utility assigned to unacceptable projects is 0.⁶ We normalize utility functions such that the utility assigned to least-preferred acceptable projects is 1. (In the special

⁵For dichotomous preferences, strict participation implies that if before contributing, $\beta\%$ of others’ money was spent on i ’s approved projects, then strictly more than $\beta\%$ of i ’s additional contribution will be spent on i ’s approved projects, while others’ money is not spent in a worse way for i . On the other hand, our pooling participation ensures that if agent i contributes money, *all* of it will be spent on i ’s approved projects, while again others’ money is not spent in a worse way for i .

⁶If we allowed agents to differentiate between unacceptable projects, they could influence other agents’ distribution to those projects, while being guaranteed to never have their own contributions spent on unacceptable projects. This would lead to “bossiness”, so we prohibit this.

case that an agent is completely indifferent, we label all projects as acceptable and so set all utilities to 1). A utility function u_i is *dichotomous* if we have $u_i(x) \in \{0, 1\}$ for all $x \in A$, so that agent i only distinguishes between acceptable and unacceptable projects without further discriminating between the acceptable ones. The *type* of an agent i is $\theta_i = (u_i, C_i)$ and the set of all types Θ . A *profile* $\theta = (\theta_i)_{i \in N} \in \Theta^N$ assigns a type to every agent. We denote by θ_{-i} the profile including all agents but i . A *mechanism* f is a function mapping a profile θ to a distribution $\delta \in \Delta(C)$.

We now discuss the main properties of distribution mechanisms that we are interested in. Recall that the purpose of such mechanisms is to identify a high-quality distribution and to exploit any synergies in the agents' preferences. Thus, we aim for distributions that allocate a large portion of the endowment to projects that many agents value highly. On the other hand, we have to ensure that the distribution can be implemented even if agents are self-interested. In addition, it should be beneficial for agents to participate in the process of pooling their resources, since the presence of additional agents can only increase the potential for mutual benefits.

A minimal requirement for the first objective is Pareto efficiency (or Pareto optimality). Indeed, if a mechanism produces a distribution so that we could reshuffle money between projects and make every agent better off, it has not made full use of the potential for mutual gains.

Definition 3.1 (Efficiency). Let $\theta \in \Theta^n$ be a type profile. A distribution $\delta' \in \Delta(C)$ *Pareto dominates* another distribution $\delta \in \Delta(C)$ if $u_i(\delta') \geq u_i(\delta)$ for all $i \in N$ and $u_i(\delta') > u_i(\delta)$ for some $i \in N$. A distribution $\delta \in \Delta(C)$ is *efficient* if no distribution dominates it.

We say that a mechanism f is efficient if $f(\theta)$ is efficient for all $\theta \in \Theta^n$.

Depending on the concrete application, the mechanism might operate in a decentralized setting, and may not be able to directly control the use of the agents' contributions (for example, when a donor coordination service does not actually collect money from its participants). In such cases, the mechanism's output δ is better understood as a recommendation to the agents about how they should use their resources. We would then need to decompose δ into personalized recommendations $\delta_i \in \Delta(C_i)$, so that if every agent spends her money according to δ_i , we recover δ . In order for the distribution to be implementable, it should ask every agent to spend her contributions exclusively on acceptable projects.

Definition 3.2 (Implementability). Let $\theta \in \Theta^n$ be a type profile. A distribution $\delta \in \Delta(C)$ is *implementable* if it can be divided into individual distributions $(\delta_i)_{i \in N}$ with $\delta_i \in \Delta(C_i)$ for all $i \in N$ and $\delta = \sum_{i \in N} \delta_i$ such that for all $i \in N$, we have $\delta_i(x) > 0$ only if $u_i(x) > 0$.

We say that a mechanism f is implementable if $f(\theta)$ is implementable for all $\theta \in \Theta^n$. In Section 6, we discuss a strengthening of implementability which requires that $\delta_i(x) > 0$ only if i has assigned maximum utility to x , i.e., only if $u_i(x) \geq u_i(y)$ for all $y \in A$. However, this requirement turns out to be too strong; it clashes with efficiency.

We cannot force agents to participate in the mechanism. Even when not participating, agents still benefit from the projects funded by the participating agents; this fact reduces participation incentives (a kind of free-riding). However, there is a downside to not participating: the mechanism only considers the utility functions of those who contribute to the mechanism. Thus, agents who do not participate forego the opportunity to influence the spending of other agents.

To motivate the definition of our participation axiom, focus on agent i , and recall that we have normalized i 's reported utilities so that the utility of each acceptable project is at least 1. Agent i considers two possible actions:

- (1) contribute 0 to the mechanism and instead spend the amount C_i on some outside option that gives i a utility of 1 per unit of money, for a total utility of $u_i(f(\theta_{-i})) + C_i$;
- (2) contribute C_i to the mechanism, for a total utility of $u_i(f(\theta))$.

Pooling participation requires that i obtains at least as much utility when choosing action (2) than when choosing (1).

Definition 3.3 (Pooling participation). A mechanism f satisfies *pooling participation* if for each $\theta \in \Theta^n$ and $i \in N$, $u_i(f(\theta)) \geq u_i(f(\theta_{-i})) + C_i$.

While pooling participation is phrased as a choice between contributing 0 or $C_i > 0$, it will also imply a similar property saying that it is beneficial for agents to increase their contribution from C_i to $C_i + \Delta$; this implication holds for all reasonable rules that treat agents with the same utility function as a single agent (like the Nash product rule does).

In Section 6, we again discuss a strengthening of pooling participation, where the quality of the outside option is higher, so that i obtains utility $u_i(f(\theta_{-i})) + C_i \cdot \max_{y \in A} u_i(y)$ when choosing action (1). Again, this stronger version is incompatible with efficiency.⁷

4 THE NASH PRODUCT RULE

The *Nash product*, which refers to the product of agent utilities, is often seen as a compromise between utilitarian and egalitarian welfare [Moulin, 1988]. Maximizing the Nash product has been found to yield fair and proportional outcomes in many preference aggregation settings, and it also turns out to be attractive in our context. Formally, it is defined as follows.

$$NASH(\theta) = \arg \max_{\delta \in \Delta(C)} \prod_{i \in N} (u_i(\delta))^{C_i} = \arg \max_{\delta \in \Delta(C)} \sum_{i \in N} C_i \log(u_i(\delta)). \quad (\text{Nash Product Rule})$$

The distribution $NASH(\theta)$ maximizes a sum of functions, namely $C_i \log(\cdot)$, that are strictly increasing in the agents' utilities. Thus, it is efficient [see, e.g., Moulin, 1988].

A simple example of $NASH$ for a type profile θ with two agents and two projects is shown in Table 1. We have $\delta = NASH(\theta) = \arg \max_{\delta \in \Delta(2)} \delta(a) \cdot (\delta(a) + 3 \delta(b)) = 1.5 a + 0.5 b$. Pooling participation is satisfied in this example because $u_1(\delta) = 1.5 > 1 = u_1(b) + 1 = u_1(NASH(\theta_{-1})) + C_1$ and $u_2(\delta) = 3 > 2 = u_2(a) + 1 = u_2(NASH(\theta_{-2})) + C_2$. On the other hand, simply maximizing utilitarian welfare—that is, the sum of individual utilities—in this example would result in $\delta' = 2 b$, which violates both implementability, as project b is unacceptable for agent 1, and pooling participation because agent 1 would prefer an outside option to participating in the mechanism.

In general, $NASH$ can be efficiently approximated using convex programming [see, e.g., Bogomolnaia et al., 2005]. However, as $NASH$ can return distributions with irrational values (see θ' in Figure 1 for such an example), its exact computation is not possible in general.

	$u_i(a)$	$u_i(b)$	C_i
Agent 1	1	0	1
Agent 2	1	3	1

Table 1. Type profile $\theta = (\theta_1, \theta_2)$ with contributions $C_1 = C_2 = 1$ and $NASH(\theta) = 1.5 a + 0.5 b$.

It is not obvious whether the distribution returned by $NASH$ is always implementable or if $NASH$ satisfies pooling participation. We will deal with these questions in turn and find that the answer in both cases is yes. The proof for pooling participation is rather involved.

⁷Note that implementability and pooling participation are logically independent properties, even when utilities are dichotomous.

The Nash distribution is the solution to an optimization problem, and thus satisfies the first-order conditions of optimality. By manipulating these conditions, we can show that *NASH* is implementable.⁸

THEOREM 4.1. *NASH satisfies implementability.*

PROOF. We have to show that there is a decomposition of $NASH(\theta)$ into $\delta_i \in \Delta(C_i)$, $i \in N$, such that $\sum_{i \in N} \delta_i(x) = \delta(x)$ for all x .

We consider the Karush–Kuhn–Tucker conditions and write the Lagrangian as

$$\mathcal{L}(\delta, \lambda, \mu_1, \dots, \mu_m) = \sum_{i \in N} C_i \log(u_i(\delta)) + \lambda \left(C - \sum_{x \in A} \delta(x) \right) + \sum_{x \in A} \mu_x \delta(x),$$

where $\lambda \in \mathbb{R}$ is the Lagrange multiplier for the constraint $\sum_{x \in A} \delta(x) = C$ and $\mu_x \geq 0$ is the multiplier for the constraint $\delta(x) \geq 0$.

Suppose δ is an optimal solution. By complementary slackness, we must have $\mu_x = 0$ whenever $\delta(x) > 0$. Also, we must have $\partial \mathcal{L} / \partial \delta(x) = 0$, that is, $\sum_{i \in N} C_i u_i(x) / u_i(\delta) - \lambda + \mu_x = 0$. By case distinction based on whether $\delta(x) > 0$, it follows that $\lambda \delta(x) = \sum_{i \in N} C_i \delta(x) u_i(x) / u_i(\delta)$ for all $x \in A$. Hence,

$$\lambda \cdot C = \sum_{x \in A} \lambda \delta(x) = \sum_{x \in A} \sum_{i \in N} C_i \delta(x) u_i(x) / u_i(\delta) = \sum_{i \in N} C_i u_i(\delta) / u_i(\delta) = \sum_{i \in N} C_i = C.$$

So $\lambda = 1$, and hence $\sum_{i \in N} C_i u_i(x) / u_i(\delta) = 1$ for all $x \in A$ such that $\delta(x) > 0$.

Now, for each $i \in N$, define an individual distribution $\delta_i \in \Delta(C_i)$ with $\delta_i(x) = C_i \delta(x) u_i(x) / u_i(\delta)$ for all $x \in A$. Clearly, $\text{supp}(\delta_i) \subseteq \{a \in A : u_i(a) > 0\}$ and $\delta_i \in \Delta(C_i)$, since $\sum_{x \in A} \delta(x) u_i(x) = u_i(\delta)$. To see that $\delta = \sum_{i \in N} \delta_i$, note that for $x \in A$ with $\delta(x) = 0$ we have $\delta_i(x) = 0$ for all $i \in N$, and for $x \in A$ with $\delta(x) > 0$, we have

$$\sum_{i \in N} \delta_i(x) = \sum_{i \in N} C_i \delta(x) u_i(x) / u_i(\delta) = \delta(x) \sum_{i \in N} C_i u_i(x) / u_i(\delta) = \delta(x). \quad \square$$

By inspecting the proof, we see that the distribution δ_i of agent i even satisfies a stronger notion of implementability: the fraction of her contribution that she gives to project x is proportional to the utility $\delta(x) u_i(x)$ she derives from x in the Nash distribution δ [see also Guerdjikova and Nehring, 2014]. So if for example, half of agent i 's utility $u_i(\delta)$ is due to the amount $\delta(x)$ spent on x , then she transfers half of her contribution to x . Thus, given δ , agents can easily compute their individual distributions without the need for a central clearinghouse to tell the agents their individual distributions. Agents do not even need to know the other agents' utility functions or contributions.

The proof that *NASH* satisfies pooling participation is technically involved and requires a number of lemmas, whose proofs we defer to the appendix. At a high level, we estimate the rate of change of an agent's utility as her contribution increases, and integrate this quantity as she goes from not participating to participating in the mechanism to obtain the desired result. The estimation entails expressing the logarithm of the utilities as a Taylor expansion and analyzing the relationship between the change in an agent's contribution and the change in these utilities at the distribution returned by *NASH*.

THEOREM 4.2. *NASH satisfies pooling participation.*

⁸This proof is similar to a result by Guerdjikova and Nehring [2014] who consider *NASH* with dichotomous preferences, and establish an equivalent property in this restricted setting.

Recall that we normalized utilities so that the utility assigned to least-preferred acceptable projects is 1 and so that the utility assigned to unacceptable projects is 0. Formally, in order to prove Theorem 4.2, we must show that for all $\theta \in \Theta^n$ and $i \in N$, $u_i(\text{NASH}(\theta)) \geq u_i(\text{NASH}(\theta_{-i})) + C_i$. Consider the function $g: \Theta^n \rightarrow \Delta(1)$ with $g(\theta) = \text{NASH}(\theta)/C_\theta$ for all $\theta \in \Theta^n$, where C_θ denotes the sum of the contributions in θ . We will show that

$$u_i(g(\theta)) \geq \frac{1}{C_\theta} ((C_\theta - C_i)u_i(g(\theta_{-i})) + C_i), \quad (1)$$

which is equivalent to the inequality above for NASH .

Denote by $\mathcal{P}_\theta \subseteq \mathbb{R}^n$ the polytope of feasible utility profiles scaled by $1/C_\theta$, i.e., $\mathcal{P}_\theta = \{u(\delta) : \delta \in \Delta(1)\}$. Note that \mathcal{P}_θ is convex. For $U \in \mathcal{P}_\theta$, let $F_\theta(U) = \sum_{i \in N} C_i \log U_i$.

LEMMA 4.3. *For all $\theta \in \Theta^n$, F_θ has a unique maximizer $U \in \mathcal{P}_\theta$. Moreover, $C_i \leq U_i \leq C_\theta u_i^{\max}$ for all $i \in N$.*

Since by Lemma 4.3, F_θ has a unique maximizer for all $\theta \in \Theta^n$, we can define the function $\mathcal{U}: \Theta^n \rightarrow \mathbb{R}_{>0}^n$ that returns this unique maximizer. Observe that $\mathcal{U}(\theta) \in \mathcal{P}_\theta$ for all $\theta \in \Theta^n$.

LEMMA 4.4. *\mathcal{U} is continuous in C_N on $\mathbb{R}_{>0}^n$ and \mathcal{U}_i is weakly increasing in C_i for all $i \in N$.*

LEMMA 4.5. *For every $\theta \in \Theta^n$ and $U \in \mathcal{P}_\theta$, there is $\varepsilon > 0$ such that for all $dU \in \mathbb{R}^n$ with $|dU| \leq \varepsilon$ and $U + dU \in \mathcal{P}_\theta$, we have $U + tdU \in \mathcal{P}_\theta$ for all $t \in [0, 2]$.*

The next three lemmas will be useful for analyzing error terms obtained in the main analysis.

LEMMA 4.6. *Let $\theta \in \Theta^n$, $U = \mathcal{U}(\theta)$, and $dU \in \mathbb{R}^n$ such that $U + dU \in \mathcal{P}_\theta$. Then,*

$$\sum_{i \in N} C_i \frac{dU_i}{U_i} \leq 0.$$

If also $U - dU \in \mathcal{P}_\theta$, then equality holds.

LEMMA 4.7. *Let $\theta \in \Theta^n$, $x \in \mathbb{R}^n$, and $\alpha, \beta > 0$ such that $\sum_{i \in N} C_i x_i = 0$ and $-\alpha \leq x_i \leq \beta$ for all $i \in N$. Then,*

$$\sum_{i \in N} C_i x_i^2 \leq \alpha \beta \sum_{i \in N} C_i.$$

LEMMA 4.8. *For all $\mu \in (0, 2)$ there is $\varepsilon^* \in (0, 1)$ with the following property: For any $\Phi: [0, 2] \rightarrow \mathbb{R}$ such that $\Phi(1) = \max_{t \in [0, 2]} \Phi(t)$ and*

$$at - (1 + \varepsilon)\beta t^2 \leq \Phi(t) \leq at - (1 - \varepsilon)\beta t^2,$$

for some $\alpha, \beta \geq 0$ and $\varepsilon \in (0, \varepsilon^)$ and all $t \in [0, 2]$, it holds that $\alpha \geq \mu\Phi(1)$.*

PROOF OF THEOREM 4.2. We want to prove (1) with $i = 1$ as the focal agent. The first step, and also the bulk of the proof, is to derive a lower bound on the derivative of $\mathcal{U}_1(\theta) = u_1(g(\theta))$ as a function of C_1 whenever $\mathcal{U}_1(\theta) < 1$. Then, integrating this derivative and using monotonicity of \mathcal{U}_1 (see Lemma 4.4) gives (1).

Step 1. Let $\theta \in \Theta^n$ such that $\mathcal{U}_1(\theta) < 1$ and $U = \mathcal{U}(\theta)$. Moreover, let $\mu \in (0, 2)$ and let ε^* be such that the conclusion of Lemma 4.8 holds; let $\varepsilon \in (0, \varepsilon^*)$. Considering the Taylor expansion of the logarithm, $\varepsilon' > 0$ exists such that for all $i \in N$ and $|r| < \varepsilon'$,

$$\left| \log(U_i + r) - \log U_i - \frac{r}{U_i} + \frac{1}{2} \left(\frac{r}{U_i} \right)^2 \right| \leq \frac{\varepsilon}{4} \left(\frac{r}{U_i} \right)^2. \quad (2)$$

Now let $C'_N \in \mathbb{R}_{>0}^n$ such that $C'_1 = C_1 + dC_1$ with $0 < dC_1 < \min\{\varepsilon', \frac{\varepsilon}{(2+\varepsilon)} C_1\}$ and $C'_i = C_i$ for all $i \in N \setminus \{1\}$, and let $\theta' = (u, C')$. Consider the function $\phi: \mathbb{R}^n \rightarrow \mathbb{R}$ defined on dU with $|dU| < \varepsilon^*$, such that

$$\phi(dU) := F_{\theta'}(U + dU) - F_{\theta}(U) - dC_1 \log U_1 = \sum_{i \in N} C_i \frac{dU_i}{U_i} + dC_1 \frac{dU_1}{U_1} - \psi(dU),$$

for some $\psi: \mathbb{R}^n \rightarrow \mathbb{R}$ with

$$(1 - \varepsilon) \frac{1}{2} \sum_{i \in N} C_i \left(\frac{dU_i}{U_i} \right)^2 \leq \psi(dU) \leq (1 + \varepsilon) \frac{1}{2} \sum_{i \in N} C_i \left(\frac{dU_i}{U_i} \right)^2.$$

The existence of ψ is guaranteed by (2) and the bound on dC_1 .

Now let $U' = \mathcal{U}(\theta')$ and $dU' = U' - U$. Note that, since the only term in $\phi(dU)$ that depends on dU is $F_{\theta'}(U + dU)$, dU' maximizes ϕ among all $dU \in \mathbb{R}^n$ with $U + dU \in \mathcal{P}_{\theta}$. By Lemma 4.5, there is $\varepsilon'' > 0$ such that, for all $dU \in \mathbb{R}^n$ with $|dU| \leq \varepsilon''$ and $U + dU \in \mathcal{P}_{\theta}$, we have $U + rdU \in \mathcal{P}_{\theta}$ for all $r \in [0, 2]$. Since \mathcal{U} is continuous in C_N by Lemma 4.4, $|dU'|$ will be small if dC_1 is small and we can choose dC_1 to be even smaller if necessary so that $2|dU'| \leq \min\{\varepsilon', \varepsilon''\}$. Then, the function $\Phi: [0, 2] \rightarrow \mathbb{R}$ with $\Phi(r) = \phi(rdU')$ is well-defined and satisfies the prerequisites of Lemma 4.8 with

$$\alpha = \sum_{i \in N} C_i \frac{dU'_i}{U_i} + dC_1 \frac{dU'_1}{U_1} \quad \text{and} \quad \beta = \frac{1}{2} \sum_{i \in N} C_i \left(\frac{dU'_i}{U_i} \right)^2.$$

Hence, it follows from Lemma 4.8 that

$$\sum_{i \in N} C_i \frac{dU'_i}{U_i} + dC_1 \frac{dU'_1}{U_1} \geq \mu \Phi(1).$$

Since U maximizes F_{θ} , by Lemma 4.6, $\sum_{i \in N} C_i \frac{dU'_i}{U_i} \leq 0$. It follows that

$$dC_1 \frac{dU'_1}{U_1} \geq \mu \Phi(1). \quad (3)$$

Next, let $\delta = g(\theta)$. Let $H_1 = \sum_{a \in A: u_1(a) > 0} \delta(a)$ be the fraction spent on agent 1's acceptable projects, i.e., those that agent 1 assigns positive utility. Recall that $U_1 < 1$, and so $H_1 < 1$. Since *NASH* satisfies the positive share property, i.e., it gives each agent positive utility, we have $H_1 > 0$. From $0 < H_1 < 1$, we get that $\delta(a) < 1$ for all $a \in A$. Thus, for $|t| > 0$ small enough, take the distribution δ^t with $\delta^t(a) = (1+t)\delta(a)$ for all $a \in A$ with $u_1(x) > 0$ and $(1 - \frac{H_1}{1-H_1}t)\delta(a)$ for all $a \in A$ with $u_1(a) = 0$. One can check that $\delta^t \in \Delta(1)$:

$$\begin{aligned} \sum_{a \in A} \delta^t(a) &= \sum_{\substack{a \in A \\ u_1(a) > 0}} (1+t)\delta(a) + \sum_{\substack{a \in A \\ u_1(a) = 0}} (1 - \frac{H_1}{1-H_1}t)\delta(a) \\ &= (1+t) \sum_{\substack{a \in A \\ u_1(a) > 0}} \delta(a) + (1 - \frac{H_1}{1-H_1}t) \sum_{\substack{a \in A \\ u_1(a) = 0}} \delta(a) \\ &= (1+t)H_1 + (1 - \frac{H_1}{1-H_1}t)(1 - H_1) = 1. \end{aligned}$$

Let $dU^t = u(\delta^t) - U$. For $|t|$ small enough, we have that $U + dU^t \in \mathcal{P}_{\theta}$ and $U - dU^t \in \mathcal{P}_{\theta}$. Indeed, $U + dU^t = u(\delta^t)$, and for the second statement we can perturb δ infinitesimally in the opposite direction. This is a valid perturbation because $\delta(a) < 1$ for all $a \in A$, and for $a \in A$ such that

$\delta(a) = 0$ we have $\delta^t(a) = \delta(a)$. Thus, by Lemma 4.6, we have

$$\sum_{i \in N} C_i \frac{dU_i^t}{U_i} = 0.$$

So for sufficiently small $|t|$, we have

$$\phi(dU^t) = dC_1 \frac{dU_1^t}{U_1} - \psi(dU^t) \geq dC_1 \frac{dU_1^t}{U_1} - (1 + \varepsilon) \frac{1}{2} \sum_{i \in N} C_i \left(\frac{dU_i^t}{U_i} \right)^2.$$

Since $dU_1^t = u_1(\delta^t) - U_1 = (1 + t)U_1 - U_1$, we have that $\frac{dU_1^t}{U_1} = t$. Similarly, it follows that $-\frac{H_1}{1-H_1} t \leq \frac{dU_i^t}{U_i} \leq t$ for all $i \in N$.

Now, by definition of H_1 , we have $U_1 \geq H_1$. Thus $1 - U_1 \leq 1 - H_1$. Hence $-\frac{U_1}{1-U_1} \leq -\frac{H_1}{1-H_1}$. Thus, applying Lemma 4.7 with $\alpha = \frac{U_1}{(1-U_1)} t$, $\beta = t$, and $x_i = \frac{dU_i^t}{U_i}$, it follows that

$$\phi(dU^t) \geq dC_1 t - (1 + \varepsilon) \frac{1}{2} \frac{U_1 C_\theta}{1 - U_1} t^2.$$

Now let $t := \frac{1-U_1}{U_1 C_\theta} dC_1$. If dC_1 is small enough, then t is also small enough and, recalling that dU maximizes ϕ among all $dU \in \mathbb{R}^n$ with $U + dU \in \mathcal{P}_\theta$, we get

$$\Phi(1) = \phi(dU') \geq \phi(dU^t) \geq \frac{1}{2}(1 - \varepsilon) \frac{1 - U_1}{U_1 C_\theta} (dC_1)^2.$$

Thus, by (3), we get

$$dC_1 \frac{dU_1'}{U_1} \geq \frac{\mu}{2}(1 - \varepsilon) \frac{1 - U_1}{U_1 C_\theta} (dC_1)^2,$$

from which it follows from $dC_1 > 0$ that

$$dU_1' \geq \frac{\mu}{2}(1 - \varepsilon) \frac{1 - U_1}{C_\theta} dC_1.$$

Since $\mu \in (0, 2)$ was arbitrary and $\varepsilon > 0$ can be chosen arbitrarily small, it follows that

$$dU_1' \geq \frac{1 - U_1}{C_\theta} dC_1.$$

Step 2. Now let $\theta \in \Theta^N$ be arbitrary. We show (1). Let $C_N^s \in \mathbb{R}^N$ such that $C_1^s = s$ and $C_i^s = C_i$ for all $i \in N \setminus \{1\}$, $\theta^s = (u, C_N^s)$, and $\tilde{\mathcal{U}}(s) = \mathcal{U}(\theta^s)$. Let $C_1' = \sup\{s \in [0, C_1] : \tilde{\mathcal{U}}_1(s) < 1\}$.

Case 1. Consider the case $C_1' = 0$. By Lemma 4.4, $\tilde{\mathcal{U}}_1$ is monotonic in s , and so

$$\tilde{\mathcal{U}}_1(C_1) \geq \tilde{\mathcal{U}}_1(0) = \frac{1}{C_\theta} \left((C_\theta - C_1) \tilde{\mathcal{U}}_1(0) + C_1 \tilde{\mathcal{U}}_1(0) \right) \geq \frac{1}{C_\theta} \left((C_\theta - C_1) \tilde{\mathcal{U}}_1(0) + C_1 \right)$$

which proves (1).

Case 2. If $C_1' > 0$, let $\varepsilon \in (0, C_1')$ be arbitrary. By Step 1, the lower right derivative of $\tilde{\mathcal{U}}_1$ at $s \in (\varepsilon, C_1')$ is at least $\frac{1 - \tilde{\mathcal{U}}_1(s)}{C_\theta - C_1 + s}$. Integrating this estimate from ε to C_1' yields

$$-\int_\varepsilon^{C_1'} \frac{\frac{\partial \tilde{\mathcal{U}}_1(s)}{\partial s}}{1 - \tilde{\mathcal{U}}_1(s)} ds \leq -\int_\varepsilon^{C_1'} \frac{1}{C_\theta - C_1 + s} ds$$

from which we get

$$\log(1 - \tilde{\mathcal{U}}_1(C_1')) - \log(1 - \tilde{\mathcal{U}}_1(\varepsilon)) \leq -(\log(C_\theta - C_1 + C_1') - \log(C_\theta - C_1 + \varepsilon)).$$

Exponentiation yields $\frac{1-\tilde{u}_1(C'_1)}{1-\tilde{u}_1(\epsilon)} \leq \frac{C_\theta - C_1 + \epsilon}{C_\theta - C_1 + C'_1}$. Since ϵ was arbitrary and \tilde{u}_1 is monotonic, we get $\frac{1-\tilde{u}_1(C'_1)}{1-\tilde{u}_1(0)} \leq \frac{C_\theta - C_1}{C_\theta - C_1 + C'_1}$. Rewriting this equation gives us

$$\tilde{u}_1(C'_1) \geq \frac{1}{C_\theta - C_1 + C'_1} \left((C_\theta - C_1)\tilde{u}_1(0) + C'_1 \right)$$

Finally, we note that if $C'_1 < C_1$, then $\tilde{u}_1(C'_1) = 1$. Together with monotonicity, this gives us

$$\begin{aligned} \tilde{u}_1(C_1) &= \frac{1}{C_\theta} \left((C_\theta - C_1 + C'_1)\tilde{u}_1(C'_1) + (C_1 - C'_1)\tilde{u}_1(C'_1) \right) \\ &\geq \frac{1}{C_\theta} \left((C_\theta - C_1)\tilde{u}_1(0) + C'_1 + (C_1 - C'_1)\tilde{u}_1(C'_1) \right) = \frac{1}{C_\theta} \left((C_\theta - C_1)\tilde{u}_1(0) + C_1 \right) \end{aligned}$$

which is (1). \square

NASH is not the only efficient rule that satisfies implementability. Whether a distribution is efficient only depends on its support [see, e.g., Aziz et al., 2014]. So we can, for example, take the support of the Nash distribution and let every agent assign her contribution to one of her most preferred projects within the support. Since the Nash distribution is implementable, this support contains at least one acceptable project for every agent. The allocation rule we obtain in this way is thus implementable, although it is artificial.

Efficiency and pooling participation seem to be harder to satisfy simultaneously. We are not aware of any rule other than *NASH* that satisfies these properties.

5 STRATEGYPROOFNESS AND MONOTONICITY

Even when agents do participate in the mechanism, they may be strategic and try to misrepresent their preferences in a way that induces the mechanism to choose a distribution that gives them higher utility. Mechanisms that are immune to strategic misrepresentation are called *strategyproof*.

Definition 5.1 (Strategyproofness). A mechanism is *strategyproof* if for all $\theta, \theta' \in \Theta^n$ with $\theta = \theta'$ except $u_i \neq u'_i$, it holds that $u_i(f(\theta)) \geq u_i(f(\theta'))$.

In other words, strategyproofness disincentivizes an agent from misreporting her utility function u_i . The other part of her type is the contribution C_i , which she might also misreport (more precisely, underreport), but this worry is already captured by pooling participation

In many mechanism design settings, only degenerate mechanisms that ignore most of the information, such as dictatorships, are strategyproof. Our setting is no exception: Hylland's [1980] Theorem 2 implies that only mechanisms that allocate the entire endowment to the most preferred projects of a single fixed agent are efficient and strategyproof, when $m \geq 3$. If we look at *NASH* specifically, it can be manipulated by agents who *understate* their utility for projects, i.e., by reporting a lower utility for some project, or by labeling a project as unacceptable when it really is acceptable. As we now show, this misrepresentation strategy can be used against every efficient and implementable rule. For technical reasons, we additionally add a mild symmetry axiom.

Definition 5.2. A mechanism is *symmetric* if it assigns the same amount of funding to projects x and y whenever they are symmetric, in the sense that swapping x and y leaves the type profile as a whole unchanged, possibly after also renaming the agents.

PROPOSITION 5.3. *Every symmetric mechanism which is efficient and implementable can be manipulated by understating utilities, when $m \geq 5, n \geq 5$.*

PROOF. We prove the incompatibility for $m = 5$ and $n = 5$, and this proof (and also the following proofs) can be adapted to larger values by adding agents who are indifferent between all projects or by adding projects which no one finds acceptable.

Assume that f is a symmetric mechanism which is efficient and implementable. Now consider a profile θ with uniform contributions ($C_i = 1$ for all agents $i \in N$) and the sets of acceptable projects for each agent

$$(\{a\}, \{abc\}, \{abd\}, \{ace\}, \{de\}).$$

We assume that each agent is indifferent between acceptable projects, so $u_i(x) = 1$ whenever x is acceptable to i . Let $\delta = f(\theta)$ be the distribution returned by the mechanism. By implementability for 5, we must have $\delta(d) + \delta(e) \geq 1$. Since the profile is symmetric under the permutation $\sigma = (bc)(de)$, we can assume without loss of generality that $\delta(d) > 0$. It follows that $u_4(\delta) < C$ because a positive amount is spent on project d , which is not one of agent 4's acceptable projects.

Now, suppose that the fourth agent changes her utility for a from 1 to 0, i.e., marks a as unacceptable. Then we get the profile θ' with the following sets of acceptable projects for each agent, with agents again being indifferent between all acceptable projects:

$$(\{a\}, \{abc\}, \{abd\}, \{ce\}, \{de\}).$$

Let $\delta' = f(\theta')$ be the distribution now returned by the mechanism. By efficiency, we must have $\delta'(b) = 0$ since otherwise we can redistribute resources from b to a to get a Pareto improvement. Next, suppose that both $\delta'(c)$ and $\delta'(d)$ are positive, say $\delta'(c) \geq \epsilon$ and $\delta'(d) \geq \epsilon$ for some $\epsilon > 0$. Then δ' is Pareto dominated by the distribution obtained from δ' by moving ϵ from c to a and ϵ from d to e . This contradicts efficiency of f , so $\delta'(c) = 0$ or $\delta'(d) = 0$. Since projects c and d are symmetric in θ' , we must have $\delta'(c) = \delta'(d) = 0$. Hence, δ' distributes the entire endowment between projects a and e , and so $u_4(\delta') = C$.

Thus, agent 4 has successfully manipulated f by understating her utilities. \square

Each use of the implementability axiom in this proof can easily be replaced by pooling participation. Hence, every efficient and symmetric mechanism that satisfies pooling participation can also be manipulated by understating utilities.

Proposition 5.3 is related to a result of Bogomolnaia et al. [2005, Prop. 6] which shows that no mechanism can satisfy efficiency, strategyproofness, and a condition called “fair welfare share” for dichotomous utility functions. Their proof works for cases with at least 17 agents and 5 projects.

We have focused on incentive issues for the voters, but the projects may also be strategic. The proposers or managers of specific projects wish to acquire funding, and charities hope for large donations. From this point of view, it is important that a mechanism satisfies *monotonicity*: a project should not receive less funding when it becomes more popular. A failure of monotonicity may result in perverse incentives for a project manager, such as deliberately making her own project seem unattractive to agents in order to acquire more funding.

Definition 5.4 (Monotonicity). A mechanism f is *monotonic* if for every agent $i \in N$, project $x \in A$, and type profiles $\theta, \theta' \in \Theta^n$ with $\theta' = \theta$ except $u'_i(a) > u_i(a)$, it holds that $f(\theta')(x) \geq f(\theta)(x)$.

So all else being equal, a project must not earn less funding if some agent states a higher utility for it. We show that monotonicity is weaker than strategyproofness.

PROPOSITION 5.5. *If a mechanism is strategyproof, it is also monotonic.*

PROOF. Let $\theta, \theta' \in \Theta^n$ be two profiles with $\theta' = \theta$, except that $u'_i(x) > u_i(x)$ for some agent i and project x ; let $\delta = f(\theta)$ and $\delta' = f(\theta')$. We examine the restrictions strategyproofness imposes

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	u_i
δ_1	1	·	·	·	3
δ_2	1	·	·	·	3
δ_3	1	·	·	·	3
δ_4	·	·	·	2	6
δ_5	·	·	·	2	6
δ_6	·	·	·	2	6
Σ	3	·	·	6	27
	θ with $A_1 = \{a\}$				

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	u_i
δ_1	0.21	·	·	0.79	7.46
δ_2	0.6	0.3	·	·	2.31
δ_3	0.6	·	0.3	·	2.31
δ_4	·	0.21	0.21	1.59	7.46
δ_5	·	0.23	·	1.77	6.69
δ_6	·	·	0.23	1.77	6.69
Σ	1.54	0.77	0.77	5.92	32.92
	θ' with $A'_1 = \{a, d\}$				

Fig. 1. Type profiles with dichotomous utility functions used in the proof of Proposition 5.6. The last row of each table shows the rounded distribution returned by *NASH*. The other rows show a rounded division of this distribution into individual distributions (zeros omitted) and the agents' utilities. Approved projects are highlighted in grey.

when i changes her utility function from u_i to u'_i in the profile θ and from u'_i to u_i in θ' :

$$\sum_{a \neq x} \delta'(a)u_i(a) + \delta'(x)u_i(x) \leq \sum_{a \neq x} \delta(a)u_i(a) + \delta(x)u_i(x) \quad \text{and}$$

$$\sum_{a \neq x} \delta(a)u_i(a) + \delta(x)u'_i(x) \leq \sum_{a \neq x} \delta'(a)u_i(a) + \delta'(x)u'_i(x).$$

Adding the two inequalities and cancelling identical terms gives

$$\delta(x)(u'_i(x) - u_i(x)) \leq \delta'(x)(u'_i(x) - u_i(x)).$$

Since $u'_i(x) - u_i(x) > 0$ we get $\delta'(x) \geq \delta(x)$. \square

Monotonicity also seems desirable independently of any incentive concerns. For instance, voters will trust a voting system more if they know that reporting higher utilities for a project will tend to lead to more funding for that project. Somewhat surprisingly, *NASH* fails monotonicity, as the following example shows.

PROPOSITION 5.6. *NASH violates monotonicity.*

PROOF. Consider the type profile θ with dichotomous utility functions represented by the sets of approved projects of each agent ($\{a\}$, $\{a, b\}$, $\{a, c\}$, $\{b, c, d\}$, $\{b, d\}$, $\{c, d\}$) and contributions $C_i = 1$ for $1 \leq i \leq 3$ and $C_i = 2$ for $4 \leq i \leq 6$. Figure 1 illustrates this situation. We have $NASH(\theta) = 3a + 6d$. If agent 1 additionally approves project d , i.e., $\theta' = \theta$ besides $u'_1(d) = 1$, we get $NASH(\theta') = 2\kappa a + \kappa b + \kappa c + (9 - 4\kappa)d$ with $\kappa = \frac{(7 - \sqrt{22})}{3}$. Hence, the contribution for project d decreased from 6 to $9 - 4\kappa \approx 5.92055\dots$, violating monotonicity. \square

Since *NASH* is the only rule we know to be efficient and satisfying pooling participation, we do not know if a rule exists that satisfies these two properties and is also monotonic.

6 LIMITS OF EFFICIENT MECHANISMS

In this section, we discuss the limits that we run into if we try to strengthen our notions of implementability and pooling participation as described in Section 3. Specifically, we show that these strengthenings are incompatible with efficiency.

First, we consider a *strong implementability*, where we require that $\text{supp}(\delta_i) \subseteq \arg \max\{u_i(a) : a \in A\}$ instead of $\text{supp}(\delta_i) \subseteq \{a \in A : u_i(a) > 0\}$. In other words, each agent is only asked to spend her contributions on her favorite projects.

PROPOSITION 6.1. *No efficient mechanism satisfies strong implementability when $m \geq 3, n \geq 2$.*

PROOF. To see that strong implementability is in conflict with efficiency, consider the example in Table 2 below. Here, both agents 1 and 2 have a pet project a and b , respectively, which the other agent dislikes; there is also a compromise project x , which is close to optimal for both. It is best for an agent to spend her entire contribution on her pet project independently of what the other agent is doing. So the only allocation we can implement in the above sense is $a + b$, which gives utility $1 + \varepsilon$ for both. But this fails to make use of the mutual interest in x : if they spent the total endowment of 2 on x , they could achieve utility 2 each. \square

	$u_i(a)$	$u_i(b)$	$u_i(x)$	C_i
Agent 1	$1 + \varepsilon$	0	1	1
Agent 2	0	$1 + \varepsilon$	1	1

Table 2. Type profile with $0 < \varepsilon < 1$ showing the incompatibility of strong implementability and efficiency.

The strengthening of pooling participation we discussed in Section 3 requires that $u_i(f(\theta)) \geq u_i(f(\theta_{-i})) + C_i u_i^{\max}$, where $u_i^{\max} = \max_{y \in A} u_i(y)$. Again, this notion of *strong pooling participation* cannot be satisfied in conjunction with efficiency.

PROPOSITION 6.2. *No efficient mechanism satisfies strong pooling participation when $m, n \geq 4$.*

	$u_i(a)$	$u_i(b)$	$u_i(c)$	$u_i(x)$	C_i
Agent 1	$2 - \varepsilon$	0	0	1	1
Agent 2	0	$2 - \varepsilon$	0	1	1
Agent 3	0	0	0	1	1
Agent 4	0	0	$2 - \varepsilon$	1	1

Table 3. Type profile with $0 < \varepsilon < 0.5$ used in the proof of Proposition 6.2.

PROOF. Assume for contradiction that there exists a mechanism f satisfying strong pooling participation and efficiency. We first determine a sequence of agents to be used in the proof. For $\theta = (\theta_1, \theta_2, \theta_3, \theta_4)$ as in Table 3, the distribution $\delta = f(\theta)$ should only allocate resources to at most one of a, b , and c . Otherwise, if there is any subset $\{y, z\} \subset \{a, b, c\}$, $y \neq z$ with $\delta(y) > 0$ and $\delta(z) > 0$, the distribution

$$(\delta(y) - \kappa) y + (\delta(z) - \kappa) z + (\delta(x) + 2\kappa) x$$

with $\kappa = \min(\delta(y), \delta(z))$ is strictly preferred by all four agents. Thus, without loss of generality, we can assume that $\delta(a) \geq 0$ and $\delta(b) = \delta(c) = 0$.

Starting with agent 1, we let the other agents join one after another and, using strong pooling participation, derive lower bounds on the amount of money allocated to project x . It will turn out

that after agent 4 has joined, the mechanism would have to allocate more than the total endowment of 4 to x in order to satisfy strong pooling participation, which is a contradiction.

Let $\theta' = (\theta_1, \theta_2)$ and $\delta' = f(\theta')$. As above, efficiency implies that either $\delta'(a) = 0$ or $\delta'(b) = 0$. Otherwise, if $\delta'(a) > 0$ and $\delta'(b) > 0$, the distribution

$$(\delta'(a) - \kappa') a + (\delta'(b) - \kappa') b + (\delta'(x) + 2\kappa') x$$

with $\kappa' = \min(\delta'(a), \delta'(b))$ is strictly preferred by both agents with a utility improvement of $2\kappa'\varepsilon > 0$. We assume that $\delta'(b) = 0$. Treating the case $\delta'(a) = 0$ requires no more than switching the order of agents 1 and 2.

By strong pooling participation, agent 2 must get at least the same utility as if both agents acted in an uncoordinated manner: $u_2(\delta') \geq u_2(a) + (2 - \varepsilon) = 2 - \varepsilon$ and with $\delta'(b) = 0$, we have $\delta'(x) \geq \frac{2-\varepsilon}{u_2(x)} = 2 - \varepsilon$.

If now agent 3 joins, we get $\theta'' = (\theta_1, \theta_2, \theta_3)$ and $\delta'' = f(\theta'')$. Strong pooling participation for agent 3 requires that $u_3(\delta'') \geq u_3(\delta') + 1 \geq (2 - \varepsilon) + 1 = 3 - \varepsilon$. However, agent 3 gets positive utility only from project x , thus we have $\delta''(x) \geq 3 - \varepsilon$.

If we finally add the agent 4 to get the full type profile $\theta = (\theta_1, \theta_2, \theta_3, \theta_4)$ with $\delta = f(\theta)$, we already know that $\delta(c) = 0$. Applying strong pooling participation for agent 4 yields $u_4(\delta) \geq u_4(\delta'') + 1 \geq (3 - \varepsilon) \cdot 1 + (2 - \varepsilon) = 5 - 2\varepsilon$. As $\delta(c) = 0$, agent 4 can only get positive utility from project x , and thus $\delta(x) \geq 5 - 2\varepsilon > 4$ for $0 < \varepsilon < 0.5$, which exceeds the total endowment of 4. \square

The reason for this incompatibility is structurally similar to that for implementability: efficiency requires spending resources on the compromise project x , but strong pooling participation can only be satisfied if the pet projects a , b , and c are funded.

7 CONCLUSION

We have studied preference aggregation mechanisms that can be used for funding public projects through voluntary individual contributions. We have identified the Nash product rule (*NASH*) as a mechanism that combines efficiency with attractive incentive properties, and are excited for the possibility of implementing such a system in the real world.

Throughout this paper, we have assumed that agents have linear utilities. In many applications, more general utilities make sense, both to capture complementarities or substitutes between projects, and to allow agents to specify decreasing returns for additional funding of a project. For the case of linear utilities, *NASH* returns a Lindahl equilibrium [Fain et al., 2016], but this equivalence breaks down for more general utilities. Since Lindahl equilibria in general satisfy a kind of implementability notion, a natural candidate for a mechanism that satisfies pooling participation for more general utilities would be a mechanism returning Lindahl equilibria. However, our proof technique does not easily extend to this case.

In the opposite direction, we can also further restrict the class of allowed utility functions. A natural choice is to only allow dichotomous utilities, making the setting a type of approval voting for divisible projects, which has been well-studied [Aziz et al., 2019, Bogomolnaia et al., 2005, Duddy, 2015]. Notably, for dichotomous utility functions, our axioms of implementability and pooling participation coincide with the strong notions we have mentioned in Section 6. Thus, the impossibilities we have proved there do not apply in the dichotomous setting, and *NASH* becomes an even stronger candidate mechanism, though some of the rules discussed by Aziz et al. [2019] become attractive alternatives (this is particularly true for Duddy's [2015] *conditional utilitarian rule*, which returns the implementable distribution with the highest utilitarian welfare).

As mentioned before, we do not know rules other than *NASH* that satisfy both efficiency and pooling participation. It would be desirable to obtain an axiomatic characterization of *NASH*

using these kinds of axioms. Two related results are already known: In the dichotomous setting, Guerdjikova and Nehring [2014] characterize *NASH* using several mild conditions (such as convexity, continuity, and reinforcement) plus a core condition that is weaker than implementability or pooling participation. Bogomolnaia et al. [2002, Prop. 6] show that *NASH* is the only rule that satisfies “fair group share” (again, a condition weaker than implementability or pooling participation) among rules that maximize a quantity of form $\sum_{i \in N} f(u_i(\delta))$ for some function f .

ACKNOWLEDGEMENTS

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-1 and BR 2312/12-1. We are grateful to Fedor Nazarov for suggesting the proof technique for Theorem 4.2.

REFERENCES

- H. Aziz, A. Bogomolnaia, and H. Moulin. 2019. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, 753–781.
- H. Aziz, F. Brandl, and F. Brandt. 2014. Universal Pareto Dominance and Welfare for Plausible Utility Functions. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*. ACM Press, 331–332.
- H. Aziz, B. E. Lee, and N. Talmon. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, 23–31.
- H. Aziz and N. Shah. 2020. Participatory Budgeting: Models and Approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*, T. Rudas and P. Gábor (Eds.). Springer.
- G. Benade, S. Nath, A. D. Procaccia, and N. Shah. 2017. Preference Elicitation For Participatory Budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 376–382.
- A. Bogomolnaia, H. Moulin, and R. Stong. 2002. Collective choice under dichotomous preferences. (2002). Mimeo.
- A. Bogomolnaia, H. Moulin, and R. Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.
- F. Brandl, F. Brandt, and J. Hofbauer. 2015. Incentives for Participation and Abstention in Probabilistic Social Choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, 1411–1419.
- F. Brandt. 2017. Rolling the Dice: Recent Results in Probabilistic Social Choice. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access, Chapter 1, 3–26.
- Y. Cabannes. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.
- I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (ACM-EC)*, 305–322.
- V. Conitzer, R. Freeman, and N. Shah. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (ACM-EC)*. 629–646.
- V. Conitzer and T. Sandholm. 2004. Expressive Negotiation over Donations to Charities. In *Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC)*. ACM Press, 51–60.
- V. Conitzer and T. Sandholm. 2011. Expressive markets for donating to charities. *Artificial Intelligence* 175, 7–8 (2011), 1251–1271.
- C. Duddy. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73 (2015), 1–5.
- E. Eisenberg and D. Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- B. Fain, A. Goel, and K. Munagala. 2016. The core of the participatory budgeting problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE) (Lecture Notes in Computer Science (LNCS))*, 384–399.
- B. Fain, K. Munagala, and N. Shah. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*. 575–592.
- P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, U. Endriss (Ed.). Chapter 2.
- D. K. Foley. 1970. Lindahl’s Solution and the Core of an Economy with Public Goods. *Econometrica* 38, 1 (1970), 66–72.
- A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. 2019. Knapsack Voting for Participatory Budgeting. *ACM Transactions on Economics and Computation* 7, 2 (2019), 8:1–8:27.
- A. Guerdjikova and K. Nehring. 2014. Weighing Experts, Weighing Sources: The Diversity Value. (2014). Working paper.

- A. Hylland. 1980. Strategyproofness of Voting Procedures with Lotteries as Outcomes and Infinite Sets of Strategies. (1980). Mimeo.
- M. Lackner and P. Skowron. 2018. Consistent Approval-Based Multi-Winner Rules. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*. 47–48.
- H. Moulin. 1988. *Axioms of Cooperative Decision Making*. Cambridge University Press.
- H. Moulin. 2003. *Fair Division and Collective Welfare*. The MIT Press.
- L. Muehlhauser. 2017. Technical and Philosophical Questions That Might Affect Our Grantmaking. Open Philanthropy Project, <https://www.openphilanthropy.org/blog/technical-and-philosophical-questions-might-affect-our-grantmaking>.
- D. Peters. 2019. Economic Design for Effective Altruism. In *The Future of Economic Design*, J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker (Eds.). Springer-Verlag.
- N. Talmon and P. Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2181–2188.

A OMITTED PROOFS

A.1 Proof of Lemma 4.3

Assume for contradiction that there are two distinct $U', U'' \in \mathcal{P}_\theta$ which maximize F_θ . As a positive linear combination of strictly concave functions, F_θ is a strictly concave function. Hence, for $U = \frac{1}{2}(U' + U'') \in \mathcal{P}_\theta$, by strict concavity of F_θ , we have

$$F_\theta(U) > \frac{1}{2}(F_\theta(U') + F_\theta(U'')) = F_\theta(U'),$$

which contradicts the assumption that U' maximizes F_θ over \mathcal{P}_θ .

Now let $i \in N$. Clearly, U_i is upper bounded by $C_\theta u_i^{\max}$. For the lower bound on U_i , recall from Theorem 4.1 that *NASH* is implementable. Thus the distribution *NASH*(θ) is the sum of distributions δ_j , $j \in N$, such that $\delta_j \in \Delta(C_j)$ and $\delta_j(x) > 0$ only if $u_j(x) > 0$. In particular, $u_i(\delta_i) \geq C_i$, since the have normalized utility functions so that the lowest positive utility of each agent is 1.

A.2 Proof of Lemma 4.4

First we show that \mathcal{U} is continuous in C_N on $\mathbb{R}_{>0}^n$. Let $\theta \in \Theta^n$, $\theta = (u, C_N)$ with $C_N \in \mathbb{R}_{>0}^n$, $(C_N^k)_{k \in \mathbb{N}} \subseteq \mathbb{R}_{>0}^n$ converging to C_N , and $\theta^k = (u, C_N^k)$. Further, let $U^k = \mathcal{U}(\theta^k)$ and $U = \mathcal{U}(\theta)$. Observe that since C_N^k converges to C_N , by Lemma 4.3, $0 < \lambda \leq U_i^k \leq \Lambda$ for all i and some $\lambda, \Lambda > 0$ and large enough k . Hence, by passing to a subsequence if necessary, we may assume that U^k converges to U^* for some $U^* \in \mathcal{P}_\theta$. Since the family of functions F_θ, F_{θ^k} , $k \in \mathbb{N}$, is uniformly equicontinuous on $[\lambda, \Lambda]^n$, it follows that $F_{\theta^k}(U^k)$ converges to $F_\theta(U^*)$. Moreover, as U^k maximizes F_{θ^k} , we have $F_{\theta^k}(U^k) \geq F_{\theta^k}(U)$, which converges to $F_\theta(U)$. Hence, U^* maximizes F_θ , which, since F_θ has a unique maximizer by Lemma 4.3, implies that $U^* = U$. Hence, U^k converges to U .

To prove that \mathcal{U}_i is monotonically increasing in C_i , let $\theta, \theta' \in \Theta^n$ and $t > 0$ such that $\theta' = (u, C')$ with $C'_1 = C_1 + t$ and $C_i = C'_i$ for all $i \in N \setminus \{1\}$. We show that $\mathcal{U}_1(\theta') \geq \mathcal{U}_1(\theta)$. Let $U = \mathcal{U}_1(\theta)$ and $U' = \mathcal{U}_1(\theta')$ and assume for contradiction that $U'_1 < U_1$. Then,

$$F_{\theta'}(U') = \sum_{i \in N} C_i \log U'_i + t \log U'_1 < \sum_{i \in N} C_i \log U_i + t \log U_1 = F_\theta(U),$$

where the inequality follows from the assumption that $U'_1 < U_1$ and the fact that U is a maximizer of F_θ . This contradicts the assumption that U' maximizes $F_{\theta'}$.

A.3 Proof of Lemma 4.5

Since \mathcal{P}_θ is a polytope, it is an intersection of a finite number of closed half-spaces H_i . Observe that the desired property holds for each H_i . Indeed, if the point U is in the interior of H_i , we can take ε to be half of the distance from U to the boundary of H_i , while if U is on the boundary of H_i , the entire ray $\{U + tU \mid t \geq 0\}$ is contained in H_i and we can take ε to be any positive real number. It follows that the desired property also holds for the intersection of the half-spaces H_i , which is \mathcal{P}_θ .

A.4 Proof of Lemma 4.6

Consider the function $\tau: [0, 1] \rightarrow \mathbb{R}$ with $\tau(t) = F_\theta(U + tU)$ and observe that τ attains its maximum at 0. Since $U_i > 0$ for all $i \in N$ by Lemma 4.3, τ is differentiable at 0. Hence, the right derivative of τ at 0 is non-positive, i.e.,

$$\left. \frac{\partial \tau}{\partial t} \right|_{t=0} = \left. \frac{\partial}{\partial t} \left(\sum_{i \in N} C_i \log(U_i + tU_i) \right) \right|_{t=0} = \sum_{i \in N} C_i \frac{dU_i}{U_i} \leq 0.$$

If additionally $U - dU \in \mathcal{P}_\theta$, the first part implies $-\sum_{i \in N} C_i \frac{dU_i}{U_i} \leq 0$, from which equality follows.

A.5 Proof of Lemma 4.7

Since $-\alpha \leq x_i \leq \beta$, we have $\left|x_i - \frac{\beta-\alpha}{2}\right| \leq \frac{\beta+\alpha}{2}$. It follows that

$$\begin{aligned} \sum_{i \in N} C_i x_i^2 &= \sum_{i \in N} C_i \left(x_i - \frac{\beta-\alpha}{2}\right)^2 - \left(\frac{\beta-\alpha}{2}\right)^2 \sum_{i \in N} C_i \\ &\leq \left(\frac{\beta+\alpha}{2}\right)^2 \sum_{i \in N} C_i - \left(\frac{\beta-\alpha}{2}\right)^2 \sum_{i \in N} C_i \\ &= \alpha\beta \sum_{i \in N} C_i, \end{aligned}$$

as claimed.

A.6 Proof of Lemma 4.8

We first prove an auxiliary lemma.

LEMMA A.1. *Let $\lambda^* \in (0, \frac{1}{2})$. Then, there are $\varepsilon^* \in (0, 1)$ and $t \in [1, 2]$ such that*

$$t - \lambda \frac{1+\varepsilon}{1-\varepsilon} t^2 > 1 - \lambda \quad \text{for all } \lambda \in [0, \lambda^*] \text{ and } \varepsilon \in (0, \varepsilon^*).$$

PROOF. The inequality in the statement can be rewritten as $\lambda < \frac{t-1}{\frac{1+\varepsilon}{1-\varepsilon} t^2 - 1}$. Choose an arbitrary $t \in (1, \frac{1}{\lambda^*} - 1)$. We have $t \in [1, 2]$ and $\lambda^* < \frac{1}{1+t}$. Since $\lim_{\varepsilon \rightarrow 0} \frac{t-1}{\frac{1+\varepsilon}{1-\varepsilon} t^2 - 1} = \frac{1}{1+t}$, we can choose $\varepsilon^* \in (0, 1)$ such that $\lambda^* < \frac{t-1}{\frac{1+\varepsilon}{1-\varepsilon} t^2 - 1}$ for all $\varepsilon \in (0, \varepsilon^*)$. It follows that $\lambda < \frac{t-1}{\frac{1+\varepsilon}{1-\varepsilon} t^2 - 1}$ for all $\lambda \in [0, \lambda^*]$ and $\varepsilon \in (0, \varepsilon^*)$, as desired. \square

We now proceed to prove Lemma 4.8. If $\mu \leq 1$, then by choosing any $\varepsilon^* \in (0, 1)$, we have $\mu\Phi(1) \leq \Phi(1) \leq \alpha$ by assumption. Assume henceforth that $\mu > 1$. For any $\varepsilon^* \in (0, 1)$ that we choose later, note that if $\alpha = 0$, then taking the given $\varepsilon \in (0, \varepsilon^*)$ yields $\Phi(1) \leq 0$, so $\alpha \geq \mu\Phi(1)$ always holds. Hence it suffices to consider $\alpha > 0$. Let $\lambda^* := 1 - \frac{1}{\mu} > 0$ and choose $\varepsilon^* > 0$ and $t^* \in [1, 2]$ such that

$$t^* - \lambda \frac{1+\varepsilon}{1-\varepsilon} (t^*)^2 > 1 - \lambda$$

for all $\lambda \in [0, \lambda^*]$ and $\varepsilon \in (0, \varepsilon^*)$, which is possible by Lemma A.1.

Let $\lambda := \frac{\alpha - \Phi(1)}{\alpha} \geq 0$. Assume for the sake of contradiction that the desired conclusion is not true, i.e., $\alpha < \mu\Phi(1)$. This is equivalent to $\lambda < \lambda^*$. Since the function $\Psi(t) := \alpha t - \Phi(t)$ satisfies $\beta(1-\varepsilon)t^2 \leq \Psi(t) \leq \beta(1+\varepsilon)t^2$, by substituting $t = t^*$ and $t = 1$, we have $\Psi(t^*) \leq \Psi(1) \frac{1+\varepsilon}{1-\varepsilon} (t^*)^2$. It follows that

$$\Phi(t^*) = \alpha t^* - \Psi(t^*) \geq \alpha \left(t^* - \frac{\Psi(1)}{\alpha} \frac{1+\varepsilon}{1-\varepsilon} (t^*)^2 \right) = \alpha \left(t^* - \lambda \frac{1+\varepsilon}{1-\varepsilon} (t^*)^2 \right) > \alpha(1-\lambda) = \Phi(1).$$

This contradicts the assumption that $\Phi(1) = \max_{t \in [0, 2]} \Phi(t)$.