Public Good Problem Re-Examined

Loren K. Fryxell

University of Oxford and Global Priorities Institute

loren.fryxell@economics.ox.ac.uk

September 9, 2024

XVIII GRASS Workshop 2024

Introduction

Preface

Given that this is a short presentation, there will be a lot of points I can't really dig into here

Feel free to see the complete paper on my website

Public Good Provision Re-Examined

I write down the government's public good provision problem from first principles

And, contrary to popular wisdom, find a solution

I call it the *cost-sharing pivotal mechanism*

Both the statement of the problem and the solution are new

Motivating Example

Consider a library which costs $4,000k to produce

#	income	effect. tax rate	distr. weight	WTP	tax burden	welfare benefit	welfare cost
4k	$25k	.04	1.93	$500	$78	965	151
5k	$50k	.08	.73	$350	$313	256	229
800	$100k	.14	.28	$250	$1,097	69	304
180	$250k	.21	.08	$250	$4,114	19	316
20	$1mil	.32	.01	$0	$25,073	0	277
10k				$3,995k	$4,000k	5,201k	2,057k

The Setting

A government would like to decide whether or not to produce a public good

- the value to each individual is **privately known**
- the distribution of values in the population is **unknown**
- the cost of production is **known**

Government's Public Good Provision Problem

The **government's public good provision problem** is to design a decision procedure that

a. provides dominant-strategy incentives for each individual to truthfully report their willingness to pay for the public good
b. induces participation in dominant strategies
c. produces the welfare-maximizing quantity of the public good via a weighted benefit-cost analysis
d. finances exactly the cost of the public good
e. taxes individuals fairly

Why Government's Problem?

Two important features of public good provision when the government is the provider are

1. the population is large
2. the government has the power to tax its citizens

Avoiding Impossibility Result #1

There are two styles of impossibility results in this domain:

The first says that there exists no mechanism which satisfies

- strategy-proofness (a),
- efficiency (c), and
- ex-post budget-balance in small populations (d)

(Green and Laffont, 1979)

Since the government's problem is constituted by large populations, this result does not apply

Avoiding Impossibility Result #2

The second says there exists no mechanism which satisfies

- strategy-proofness (a),
- individual rationality (b),
- efficiency (c), and
- ex-post budget-balance in large populations (d)

(Mailath and Postlewaite, 1990)

Avoiding Impossibility Result #2

Individual rationality presupposes that, by refusing to participate, an individual can avoid consuming the good and avoid her tax payment

But,

1. a public good is by definition non-rival and non-excludable, so if the good is ultimately produced, an individual cannot avoid its benefits (whether they participated or not)
2. the government has the power to tax its citizens, so an individual cannot avoid paying taxes simply by refusing to report her willingness to pay

Avoiding Impossibility Result #2

No participation constraint has yet been proposed which captures these two facts

I propose a new criteria, called *cost-sharing universal participation*, which does

Given this notion of participation, there is a unique solution to the problem

Moreover, we will *also* be able to satisfy an additional fairness desiderata as well

The Public Good Provision Problem

Formally, the government's public good provision problem is to design a mechanism which satisfies

a. strategy-proofness
b. cost-sharing universal participation
c. implements a weighted BCA decision rule
d. asymptotic ex-post budget-balance
e. the fair pricing principle.

(c) is a recent development in mechanism design. (b) and (e) are new to this paper.

The **cost-sharing pivotal mechanism** is the unique solution to this problem

Cost-Sharing Pivotal Mechanism vs Clarke Mechanism

The canonical mechanism in public good provision is the *Clarke mechanism*

However, the Clarke mechanism violates cost-sharing universal participation and the fair pricing principle

In fact, it violates an even more basic fairness principle I call no-extortion: if nothing is produced, no one should pay

In the library example, the Clarke mechanism chooses not to construct the library yet, disconcertingly, **charges the 20 richest individuals $20,073 each**

The Net Value Approach

The literature on strategy-proof public good provision hasn't studied much beyond the Clarke mechanism. Why?

It turns out this is due to a methodological simplification in mechanism design

The standard approach to modeling production costs in mechanism design is to

- assign cost shares c_i to each individual,
- each individual reports $\tilde{v}_i = v_i - c_i$
- proceed as if $\tilde{v}_i$ were their intrinsic value and there are no production costs

I call this the *net value approach*

The Net Value Approach

This has become the standard approach

- many papers simply study the case without production costs
- and point to the fact that we may always insert production costs into the model in this way

Indeed, Green and Laffont (1979) even contend that "there is no real alternative to this approach."

I argue it is not without loss of generality

Indeed, all of the ideas and formal results in this paper come out of a framework which keeps values for the good and costs of production separate

Model

Standard Mechanism Design Environment

A *standard mechanism design environment with transfers* is defined by

$$(I,\ Y,\ \Theta,\ \{v_i\}_{i\in I})$$

where

- I is a set of n individuals
- Y is a set of social alternatives
- $\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a type space
- $v_i : Y \times \Theta_i \to \mathbb{R}$ is individual i's willingness to pay for each alternative y given her type θ_i

i's preferences over outcomes $Y \times \mathbb{R}^n$ are quasilinear and represented by $u_i(y, t, \theta) = v_i(y, \theta_i) - t_i$

Public Good Provision Environment

To study the government's public good provision problem, I add four elements to the standard environment:

1. observable characteristics
2. production costs
3. tax burdens
4. distributional weights

Public Good Provision Environment

A *public good provision environment with transfers* is defined by

$$\mathcal{E} = (I,\ Y,\ \Theta,\ Z,\ \{v_i\}_{i\in I},\ c, \{c_i\}_{i\in I},\ \{\lambda_i\}_{i\in I})$$

where

- $Z = Z_1 \times \ldots \times Z_n$ is a space of observable characteristics for each individual i
- $c : Y \to \mathbb{R}_+$ is the production cost for each alternative y
- $c_i : Y \times Z \to \mathbb{R}_+$ is i's tax burden for alternative y given observable characteristics z (where $\sum_{i\in I} c_i(y,z) = c(y)$ for all y, z)
- $\lambda_i : Z \to \mathbb{R}_+$ is i's distributional weight

Mechanisms

In these environments, a mechanism can be represented by a pair $f = (\alpha, \tau)$, where

- $\alpha : \Theta \to Y$ is a decision rule
- $\tau : \Theta \to \mathbb{R}^n$ is a transfer rule

Each individual i reports their type θ_i to the mechanism

The mechanism implements social alternative $\alpha(\theta)$ at cost $c(\alpha(\theta))$ and collects transfers $\tau_i(\theta)$ from each i

Tax Burdens

The use of exogenous cost shares c_i is standard

I add that cost shares $c_i(\cdot)$ can be an arbitrary function of observable characteristics and interpret them as the individual's *tax burden*

Under the tax burden interpretation, each individual indirectly pays $c_i(\alpha(\theta))$ through the tax system and directly pays $\tau_i(\theta) - c_i(\alpha(\theta))$ to the mechanism

Strategy-Proofness

Definition. A mechanism is *strategy-proof* if it is a dominant strategy for each individual i to report her type $\theta_i \in \Theta_i$ truthfully.

Weighted BCA Decision Rule

Definition. A decision rule $\alpha : \Theta \to Y$ is an *unweighted BCA* (i.e., *efficient*) if, for all θ,

$$\alpha(\theta) \in \arg\max_{y \in Y} \sum_{i \in I} v_i(y, \theta_i) - c(y).$$

Definition. A decision rule $\alpha : \Theta \to Y$ is a *weighted BCA* if, for all θ,

$$\alpha(\theta) \in \arg\max_{y \in Y} \sum_{i \in I} \lambda_i(z)(v_i(y, \theta_i) - c_i(y, z)).$$

Unweighted vs Weighted BCA Decision Rules

It is standard in mechanism design to use an unweighted BCA (efficient) decision rule

However, the predominant view within welfare economics is that a weighted BCA is strongly preferred to an unweighted BCA (e.g., Blackorby and Donaldson (1990), Adler (2012), Adler (2016), Boadway (2016), Fleurbaey and Abi-Rafeh (2016), Fleurbaey et al. (2013), Bressler and Heal (2022))

The U.S. government recently (less than a year ago!) updated its official guidelines on BCA to include the use of distributional weights, so this is now official policy

Unweighted vs Weighted BCA Decision Rules

Recently, several papers in mechanism design have begun to optimize with respect to a weighted BCA objective (e.g., Dworczak, Kominers and Akbarpour (2021), Pai and Strack (2022), Barreto, Ghersengorin and Augias (2022), Akbarpour, Dworczak and Kominers (2023), Reuter and Groh (2023), Akbarpour et al. (2024))

This paper imposes a weighted BCA decision rule (implementing the first best) as an axiom

Theorem 1

Recall that cost shares $c_i(y, z)$ and distributional weights $\lambda_i(z)$ depend only on observable characteristics. Why?

Theorem 1. For any strategy-proof mechanism that implements a weighted BCA decision rule,

1. only willingness to pay can be elicited, and
2. distributional weights and cost shares must be independent of willingness to pay

Groves Mechanisms

Recall the classic result in mechanism design:

> *A mechanism is strategy-proof and implements an unweighted BCA (efficient) decision rule if and only if it is a Groves mechanism.*

Of these mechanisms, the most well-known is the pivotal mechanism (also known as the VCG mechanism)

Generalized Groves Mechanisms

It turns out that we can generalize this result quite easily to the case of a weighted BCA decision rule

> *A mechanism is strategy-proof and implements a weighted BCA decision rule if and only if it is a generalized Groves mechanism.*

And of course we can do the same to find the generalized pivotal mechanism

Generalized Pivotal Mechanism

Definition. A mechanism $f = (\alpha, \tau)$ is a *generalized pivotal mechanism* if the decision rule is a weighted BCA and the transfer rule satisfies, for all i and θ,

$$\tau_i(\theta) = \frac{1}{\lambda_i}\Big(\sum_{j\neq i}\lambda_j v_j(\alpha(0,\theta_{-i}),\theta_j) - \sum_{k\in I}\lambda_k c_k(\alpha(0,\theta_{-i}))\Big) - \frac{1}{\lambda_i}\Big(\sum_{j\neq i}\lambda_j v_j(\alpha(\theta),\theta_j) - \sum_{k\in I}\lambda_k c_k(\alpha(\theta))\Big).$$

The Cost-Sharing Pivotal Mechanism

Definition. A mechanism $f = (\alpha, \tau)$ is a *cost-sharing pivotal mechanism (CSP)* if the decision rule is weighted BCA and the transfer rule satisfies, for all i and θ,

$$\tau_i(\theta) = \frac{1}{\lambda_i}\Big(\sum_{j\neq i}\lambda_j v_j(\alpha(0,\theta_{-i}),\theta_j) - \sum_{k\in I}\lambda_k c_k(\alpha(0,\theta_{-i}))\Big)$$
$$- \frac{1}{\lambda_i}\Big(\sum_{j\neq i}\lambda_j v_j(\alpha(\theta),\theta_j) - \sum_{k\in I}\lambda_k c_k(\alpha(\theta))\Big) + c_i(\alpha(0,\theta_{-i})).$$

Solution to Government's Public Good Provision

The cost-sharing pivotal mechanism is the unique solution to the government's public good provision problem

That is, the cost sharing pivotal mechanism satisfies

a. strategy-proofness
b. cost-sharing universal participation
c. implements a weighted BCA decision rule
d. asymptotic ex-post budget-balance
e. the fair pricing principle.

And, up to small perturbations which vanish in the limit, is the *unique* mechanism which does so

Participation

Individual Rationality

Definition. A mechanism $f = (\alpha, \tau)$ is *individually-rational* if for all i, θ,

$$v_i(\alpha(\theta), \theta_i) - \tau_i(\theta) \geq 0.$$

An individual should prefer to participate in the mechanism rather than to consume nothing and pay nothing

Universal Participation

Definition. A mechanism $f = (\alpha, \tau)$ satisfies *universal participation* if for all θ and i,

$$v_i(\alpha(\theta), \theta_i) - \tau_i(\theta) \geq v_i(\alpha(0, \theta_{-i}), \theta_i).$$

An individual should prefer to participate in the mechanism rather than to consume what would have been produced without her and pay nothing

Cost-Sharing Universal Participation

Definition. A mechanism $f = (\alpha, \tau)$ satisfies *cost-sharing universal participation* if for all θ and i,

$$v_i(\alpha(\theta), \theta_i) - \tau_i(\theta) \geq v_i(\alpha(0, \theta_{-i}), \theta_i) - c_i(\alpha(0, \theta_{-i})).$$

An individual should prefer to participate in the mechanism rather than to consume what would have been produced without her and to pay her tax burden for the good

Theorem 2

The cost-sharing pivotal mechanism satisfies cost-sharing universal participation

In fact, we can *characterize* the cost-sharing pivotal mechanism in terms of cost-sharing universal participation

Theorem 2. A mechanism maximizes ex-post revenue among all mechanisms which are strategy-proof, implement a weighted BCA decision rule, and satisfy cost-sharing universal participation if and only if it is a cost-sharing pivotal mechanism.

Fairness

Individual Rationality (as Fairness)

Definition. A mechanism $f = (\alpha, \tau)$ is *individually-rational* if for all i, θ,

$$\tau_i(\theta) \leq v_i(\alpha(\theta), \theta_i).$$

It is fair for an individual to pay up to her total monetary value for what is produced and no more

Sense of Community

In my view, public good environments are characterized by a sense of community

- each individual understands that everyone needs to chip in—but not unreasonably so—for the greater good
- even if that means paying more than one's value

Indeed this is already the case. I pay taxes which go to funding a park that I may have zero value for

The Fair Pricing Principle

Definition. A mechanism $f = (\alpha, \tau)$ satisfies *the fair pricing principle* if for all i, θ,

$$\tau_i(\theta) \leq \max\left\{v_i(\alpha(\theta), \theta_i),\ c_i(\alpha(\theta))\right\}.$$

It is fair for an individual to pay up to her total monetary value for what is produced or her tax burden for the good, whichever is larger

Theorem 3

Recall:

Theorem 2. A mechanism maximizes ex-post revenue among all mechanisms which are strategy-proof, implement a weighted BCA decision rule, and satisfy cost-sharing universal participation if and only if it is a cost-sharing pivotal mechanism.

Perhaps surprisingly, the same is *also* true for the fair pricing principle (though a few more assumptions are necessary)

Theorem 3. A mechanism maximizes ex-post revenue among all mechanisms which are strategy-proof, implement a weighted BCA decision rule, and satisfy the fair pricing principle if and only if it is a cost-sharing pivotal mechanism.

Budget-Balance

Model

The cost-sharing pivotal mechanism is *asymptotically ex-post budget-balanced*

We will draw types and observable characteristics (θ_i, z_i) jointly from **any** distribution (with finite second moment), **unknown** to the designer

We want to say something about what happens when the population size n gets large, regardless of the underlying distribution

Asymptotic Ex-Post Budget-Balance

Definition. A (sequence of) mechanisms (α^n, τ^n) is *asymptotically ex-post budget-balanced* if

1. the probability of ex-post budget-balance goes to one as n goes to infinity, i.e.,

$$\mathbb{P}\Big(\sum_{i=1}^{n} \tau_i^n(\theta^n, z^n) = c^n(\alpha^n(\theta^n, z^n))\Big) \to 1 \quad \text{as} \quad n \to \infty, \quad \text{and}$$

2. the expected distance from ex-post budget-balance per capita goes to zero as n goes to infinity, i.e.,

$$\frac{1}{n}\mathbb{E}\Big(\Big|\sum_{i=1}^{n} \tau_i^n(\theta^n, z^n) - c^n(\alpha^n(\theta^n, z^n))\Big|\Big) \to 0 \quad \text{as} \quad n \to \infty.$$

Theorem 4

Theorem 4. The cost-sharing pivotal mechanism is asymptotically ex-post budget-balanced (no matter how the cost of the public good varies with n, no matter how individual cost shares vary with n).

The Clarke Mechanism

The Clarke Mechanism

The Clarke mechanism is simply the pivotal mechanism with net-values plugged in

The pivotal mechanism is simple and intuitive (and beautiful, even)

However, if one writes down the Clarke mechanism in terms of its underlying fundamentals (i.e., intrinsic values rather than net values)

Then it's clear that this "pivotal mechanism with net values" may not be the most natural generalization of the pivotal mechanism to environments with production costs

Cost-Sharing Pivotal Mechanism vs Clarke Mechanism

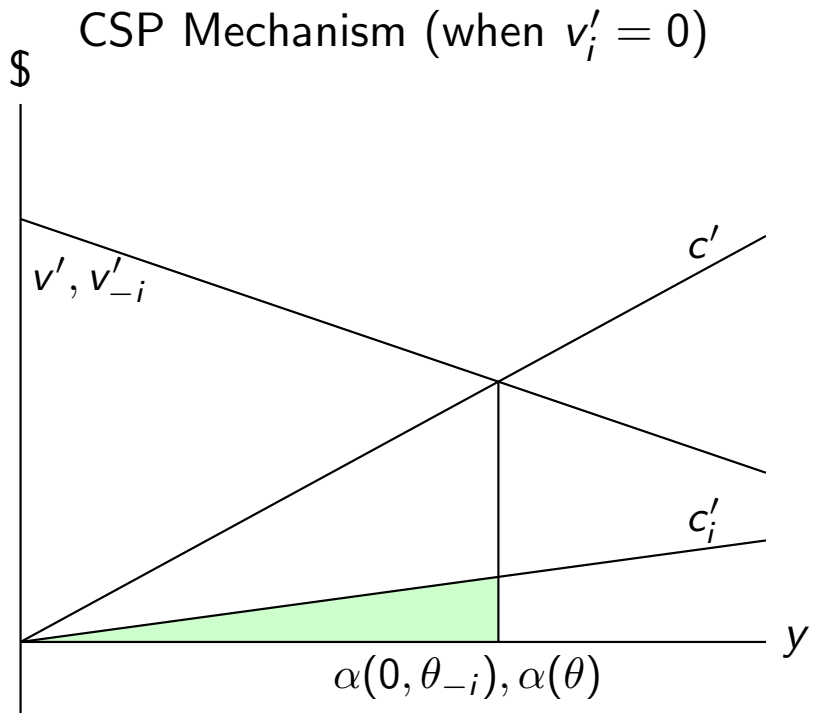

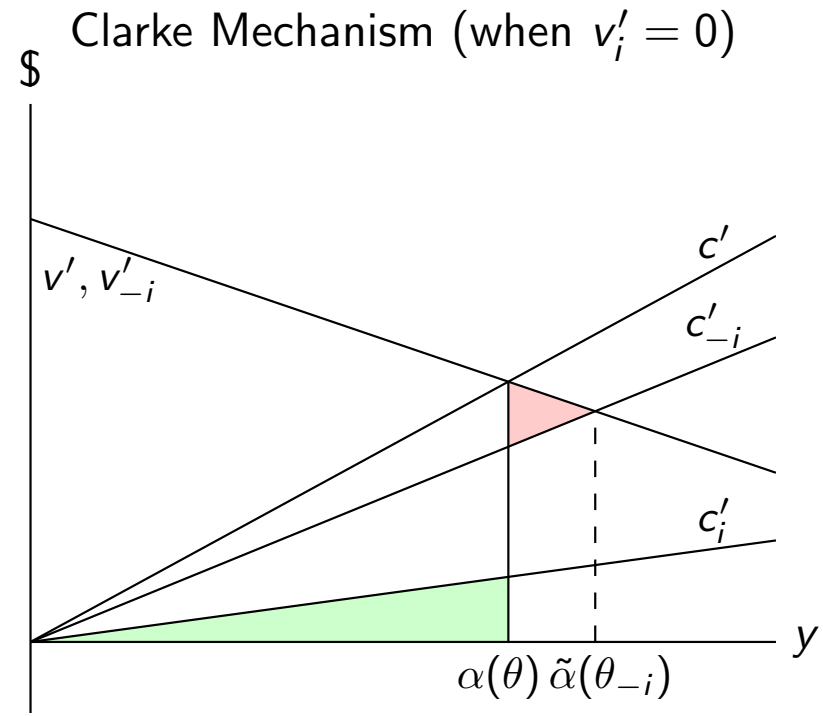

The Clarke Mechanism violates No-Extortion

Setup

- $I = \{i, j\}$ and $Y = \{\text{no park}, \text{park}\}$
- cost of the park is 4
- i's fair cost share is 2, j's fair cost share is 2
- i's value for the park is 0, j's value for the park is 3
- i's net-value for the park is $0 - 2 = -2$, j's net-value for the park is $3 - 2 = 1$

Run pivotal mechanism with net values and costless park

- efficient decision is no park (and j's total welfare is 0)
- efficient decision ignoring i's preferences is park (and j's total welfare is 1)
- so i's pivotal transfer is 1, violating no-extortion

What does the cost-sharing pivotal mechanism do?

In the cost-sharing pivotal mechanism, since the good is not produced, no one pays anything

Conclusion

Key Takeaways of the Paper

1. The classic impossibility results of public good provision do not apply to perhaps its most important case: when the government is the provider of the public good.

Key Takeaways of the Paper

2. Contrary to conventional wisdom, a natural solution exists to the government's public good provision problem. I call it the *cost-sharing generalized pivotal mechanism*. It is new and, modulo small perturbations which vanish in the limit, it is unique.

Key Takeaways of the Paper

3. Standard participation constraints are not appropriate for the government's public good provision problem. I propose a new participation constraint called *cost-sharing universal participation*.

Key Takeaways of the Paper

4. Fairness principles ought to be satisfied in public good provision, constraining what the government can fairly ask individuals to contribute to the public good. I propose a new fairness principle called the *the fair pricing principle*.

Key Takeaways of the Paper

5. An unweighted BCA (efficient) decision rule is not recommended by welfare economists nor official government policy. When it comes to public good provision, a weighted BCA decision rule should be used instead.

Key Takeaways of the Paper

6. The net value approach is not without loss of generality. Costs of producing the public good should be modeled explicitly by default, rather than concealed within the individuals' values.

Key Takeaways of the Paper

7. The canonical public good provision mechanism—the Clarke mechanism—violates asymptotic ex-post budget-balance, cost-sharing universal participation, the fair pricing principle, and an even more basic principle I call *no-extortion*: if nothing is produced, no one should pay.

Thank you for listening!

Questions, comments, or concerns?