Resource Allocation with Positive Externalities*

Dhruva Bhaskar[†] and Evan Sadler[‡]

November 7, 2019

Abstract

In many allocation problems, transfers are unavailable, but incentives are partially aligned due to positive externalities. We study how a designer can exploit this alignment to allocate a resource between n players. We identify a natural mechanism that partitions types into intervals and allocates among players in the highest interval. While interim allocations are identical for all types in the same interval, the exact allocation depends on the *lowest* reported type. This novel feature is a crucial source of incentives. In a class of distributions, our mechanism is optimal, and it is approximately efficient when n is large.

JEL Classification: D82, D44, D72

1 Introduction

Allocation decisions often affect people beyond the direct recipients of resources. Campaign funds used to advertise for one congressional candidate may yield spillovers for other candidates in the same party. Even without these direct spillovers, other party members likely value having more ideological allies in congress. If a firm's CEO sets aside funds for a project within a particular division, it may well impact employees and managers in other divisions. The project could improve firm-wide profits, allowing larger bonuses for all, or it

*We are deeply grateful to Joyee Deb, Elliot Lipnowski, Alessandro Lizzeri, David Pearce, Debraj Ray, Ariel Rubinstein, and Maher Said for their feedback and suggestions. We also thank Alp Atakan, Ala Avoyan, Abhijit Banerjee, Elchanan Ben-Porath, Kfir Eliaz, Alex Frankel, Drew Fudenberg, Ben Golub, Sergiu Hart, Jason Hartline, Johannes Hörner, Navin Kartik, Aditya Kuvalekar, Shengwu Li, Laurent Mathevet, Andrew McClellan, Mallesh Pai, Mikhail Panov, Motty Perry, João Ramos, Michael Richter, Arunava Sen, Ran Spiegler, Ennio Stacchetti, Philipp Strack, Balasz Szentes, Juuso Toikka, Jerome Williams, Alex Wolitzky, Chris Woolnough, and Cherry Yao for helpful conversations, and we thank seminar participants at NYU and MIT for their comments. Bhaskar's work on the project was supported by a grant from the European Research Council (ERC 681460, Informative Prices).

†Queen Mary University of London – d.bhaskar@qmul.ac.uk

‡Columbia University – es3668@columbia.edu

could provide tools that other divisions use to enhance their own work. Likewise, infrastructure spending in one city—say on roads or telecommunication networks—benefits visitors and many businesses based elsewhere. Donors, managers, and governments need to elicit accurate information about the value of different projects, but a reliance on transfers to align incentives seems problematic in these settings. Positive externalities offer another lever.

This paper studies a resource allocation problem with positive externalities and no transfers. The designer has a unit of a resource to allocate between n players.[1] Each player i has a private marginal value $t_i \in [0,1]$ for the resource that is drawn independently from a smooth distribution F. Player i values every other player j's allocation at αt_j for some $\alpha \in (0,1)$. While player i benefits from receiving more of the resource, she prefers j to have it if $t_j > \frac{1}{\alpha} t_i$. Our interest is the extent to which the designer can exploit this limited preference alignment to extract information from the players.

Interdependent preferences and a lack of transfers distinguish our setting from more classical allocation problems, and these features necessitate a different analytical approach. Our approach is non-constructive. We identify a natural allocation rule that partitions the type space into intervals and allocates the resource among players in the highest reported interval. We subsequently show that an incentive compatible mechanism exists that has two key properties: the mechanism fully allocates the resource, and if one player reports type 0, the mechanism always allocates to the highest type. These properties allow us to bound the welfare achieved in our mechanism. For a family of type distributions that includes the uniform, we show that our mechanism is optimal. More generally, we show that the designer can get arbitrarily close to efficiency when n is large enough, no matter how small α is.

A crucial feature of our mechanism is that the allocation among players in the highest reported interval depends on the *lowest* reported interval. Conditional on m players reporting in the highest interval, each receives an interim expected allocation of $\frac{1}{m}$. However, the exact allocation becomes more efficient when the lowest reported interval is lower. If the lowest type is relatively high, we allocate to the lowest type in the highest interval. If the lowest type is low, we allocate to the highest type. This feature helps provide incentives for players to report lower types: misreporting upward not only takes away from higher types in expectation, it also results in a less efficient allocation *within* the highest reported interval.

After describing the model in section 2, we introduce *bin mechanisms* in section 3. A bin mechanism partitions the type space $[0,1]$ into a set of intervals, or bins, and delivers interim allocations that depend only on the reported intervals, not on finer details on the type profile. We highlight two types of allocation rules that give rise to bin mechanisms. Simple bin mechanisms always divide the resource evenly among those in the highest reported bin, while better bin mechanisms implement an allocation that depends on the lowest reported bin. Simple bin mechanisms provide a useful benchmark because they are implementable without commitment power.[2] The difference in welfare between the two mechanisms sheds light on the value of commitment in our setting.

[1] Equivalently, the designer allocates probability units of an indivisible good.

[2] Formally, this means the allocation can be implemented as the equilibrium of a cheap talk game between the players and the designer.

Section 4 provides a complete characterization of optimal mechanisms when types are uniformly distributed. In particular, we show that any incentive compatible mechanism is optimal if it fully allocates the resource and always allocates to the highest type when one player reports type 0. A better bin mechanism with a countably infinite number of bins satisfies both properties and is therefore optimal. Section 5 turns to general type distributions, providing an upper bound on the welfare from any incentive compatible mechanism and a lower bound on the welfare from our better bin mechanism. The bounds coincide for a one-dimensional family of type distributions. Regardless of the type distribution, the lower bound implies that our mechanism approaches efficiency as n becomes large, no matter how small the externality α. We defer a discussion of related work to section 6.

2 Allocating a Budget

A designer must allocate one unit of a resource between n players $i \in \{1, 2, ..., n\}$. A feasible allocation is a vector $\mathbf{b} = (b_1, b_2, ..., b_n) \in \mathbb{R}^n$ with $0 \le b_i$ for each i and $\sum_{i=1}^n b_i \le 1$. Write $\mathcal{B}$ for the space of feasible allocations. Each player i has a type $t_i \in [0, 1]$ drawn independently from the distribution F. We assume F is continuous with full support and has density uniformly bounded away from zero. If the allocation is $\mathbf{b}$, and types are $\mathbf{t}$, then player i earns utility

$$u_i(\mathbf{b}, \mathbf{t}) = b_i t_i + \alpha \sum_{j \neq i} b_j t_j,$$

where $\alpha \in (0, 1)$. That is, each player earns the value of her own allocation plus a fraction α of the value of the other players' allocations. The utility functions and the type distribution are common knowledge to the players and the designer, while type realizations are private information to each player.

An *allocation rule* is a function $\mathbf{b}(\mathbf{t}) : [0, 1]^n \to \mathcal{B}$ mapping vectors of types to feasible budget allocations. An allocation rule gives total expected welfare of

$$W = \int_{[0,1]^n} \sum_{i=1}^{n} b_i(\mathbf{t}) t_i \mathrm{d}F(\mathbf{t}). \tag{1}$$

The designer seeks a feasible and incentive compatible mechanism to maximize W. Without loss of generality, we consider symmetric, direct mechanisms. Players report types to the designer who implements an allocation satisfying

$$b_i(t_1, t_2, ..., t_n) = b_{\pi(i)}(t_{\pi(1)}, t_{\pi(2)}, ..., t_{\pi(n)}) \equiv b(\mathbf{t})$$

for any permutation of the indices π. Given an allocation rule b, the interim expected utility that player i with true type t earns from reporting type t' is

$$U(t, t') = \int_{[0,1]^{n-1}} \left[t b(t', \mathbf{t}_{-i}) + \alpha \sum_{j \neq i} t_j b(t_j, t', \mathbf{t}_{-ij}) \right] \mathrm{d}F(\mathbf{t}_{-i}). \tag{2}$$

A mechanism is *(Bayesian) incentive compatible* if truthful reporting constitutes a Bayes' Nash Equilibrium of the induced game. Incentive compatibility requires that

$$U(t,t) \geq U(t,t')$$

for all $t' \in [0,1]$. An optimal mechanism maximizes expected welfare subject to the incentive compatibility constraint.

We can interpret the allocation $b(\mathbf{t})$ as either a share of a divisible good, such as a monetary budget, or as a probabilistic allocation of an indivisible good, such as a project. In the former case, we implicitly assume that each player has a constant marginal value for money, and there is limited liability—the designer cannot demand payments. We eliminate transfers and introduce a particular form of interdependent preferences to an otherwise classical mechanism design setting.

The source of interdependence is a positive payoff externality that partially aligns players' preferences. If a player with type t *knew* that another player had type $t' > \frac{t}{\alpha}$, she would want to allocate the resource to the other player. The parameter α measures the extent of preference alignment. If $\alpha = 0$, a player has no incentive to give up a claim on the resource, and the only incentive compatible mechanisms pool all types. If $\alpha = 1$, players share the designer's welfare maximization objective, and we can attain the efficient allocation. In between, the efficient allocation is clearly not incentive compatible, but partial preference alignment allows the designer to extract some useful information.

2.1 Preliminary Analysis

As a first step, we rewrite the designer's problem using reduced form allocation rules. Define $Q(t)$ as a type t player's interim expected allocation and $M(t)$ as a type t player's interim expected benefit from the other players' allocations:

$$Q(t) = \int_{[0,1]^{n-1}} b(t, \mathbf{t}_{-1}) \mathrm{d}F(\mathbf{t}_{-1}), \qquad M(t) = \alpha \int_{[0,1]^{n-1}} \sum_{j=2}^{n} t_j b(t_j, t, \mathbf{t}_{-j2}) \mathrm{d}F(\mathbf{t}_{-2}).$$

A type t player's interim expected payoff is then $U(t) = tQ(t) + M(t)$. We can rewrite the welfare (1) as

$$W = n \int_0^1 tQ(t) \mathrm{d}F(t). \tag{3}$$

Standard arguments yield the following lemma.

Lemma 1. *If the allocation rule b is incentive compatible, then $Q(t)$ is weakly increasing, and $U(t)$ is increasing and convex with $U'(t) = Q(t)$ almost everywhere.*

Consequently, the allocation rule b is incentive compatible if and only if $Q(t)$ is weakly increasing and

$$tQ(t) + M(t) = M(0) + \int_0^t Q(x) \mathrm{d}x. \tag{4}$$

Using (3) and Lemma 1, we can write the designer's problem as

$$\max_{b} \quad n\int_0^1 tQ(t)\mathrm{d}F(t) \tag{5}$$
$$s.t. \qquad tQ(t)+M(t)=M(0)+\int_0^t Q(x)\mathrm{d}x, \quad t\in[0,1]$$
$$Q(t)\geq Q(t'), \quad \forall\, t>t'$$
$$\sum_{i=1}^{n} b(t_i,\mathbf{t}_{-i})\leq 1, \quad b(\mathbf{t})\geq 0, \quad \mathbf{t}\in[0,1]^n.$$

This formulation highlights two key difficulties that distinguish our setting from typical problems. First, the value $M(t)$ plays the role of a "transfer," providing incentives for truthful reporting, but the designer can only control $M(t)$ through the allocation itself. As a consequence, we cannot write feasibility constraints using the reduced form allocation as in Border (1991). Each choice of Q imposes its own feasibility constraints on M and vice versa. Nevertheless, we make progress through a guess-and-verify approach.

3 Bin Mechanisms

To illustrate one possible approach, suppose $n=2$, and consider a mechansim based on a single threshold $\tau\in(0,1)$. If one player's type is above τ and the other is below, we allocate the resource to the player with the higher type. If both types fall on the same side of τ, we divide the resource evenly. One can check that this mechanism is incentive compatible precisely when $\tau=\alpha\mathbb{E}_F[X]$. The benefit to a player from reporting above τ is a higher allocation. The cost is a reduction in the externality from the other player's allocation. For a type below τ, misreporting upwards entails taking the resource from a significantly higher type in expectation, which results in a lower expected payoff.

Bin mechanisms generalize this idea, partitioning $[0,1]$ into a collection of intervals and delivering interim allocations that depend only on the interval in which a player's type lies. Formally, a mechanism is a bin mechanism if there exists a collection of thresholds $\{\tau_k\}_{k=0}^K$, where $1=\tau_0>\tau_1>\tau_2>...>\tau_K=0$ and a collection of constants $\{Q_k,M_k\}_{k=0}^K$ such that $Q(t)=Q_k$ and $M(t)=M_k$ whenever $t\in I_k\equiv(\tau_k,\tau_{k-1}]$. We refer to the intervals I_k as bins. The number of bins K may be finite or infinite. When K is finite, we assume the Kth bin is the interval $[\tau_K,\tau_{K-1}]$, and when K is infinite, the final bin I_∞ is simply the point $t=0$.

Figure 1 illustrates a simple allocation rule for a mechanism with 3 bins when $n=2$. The horizontal axis shows player 1's type, and the vertical axis shows player 2's type. If types are realized in an off-diagonal box, the entire budget goes to the player with the highest type: this is the efficient outcome. For the boxes on the diagonal, the budget is divided evenly between the players, which is inefficient.

Our analysis centers on bin mechanisms with $K=\infty$. We consider two types of allocation rules that give rise to bin mechanisms. In a *simple bin mechanism*, we take the players in the highest reported bin and allocate the budget evenly between them. In a *better bin*

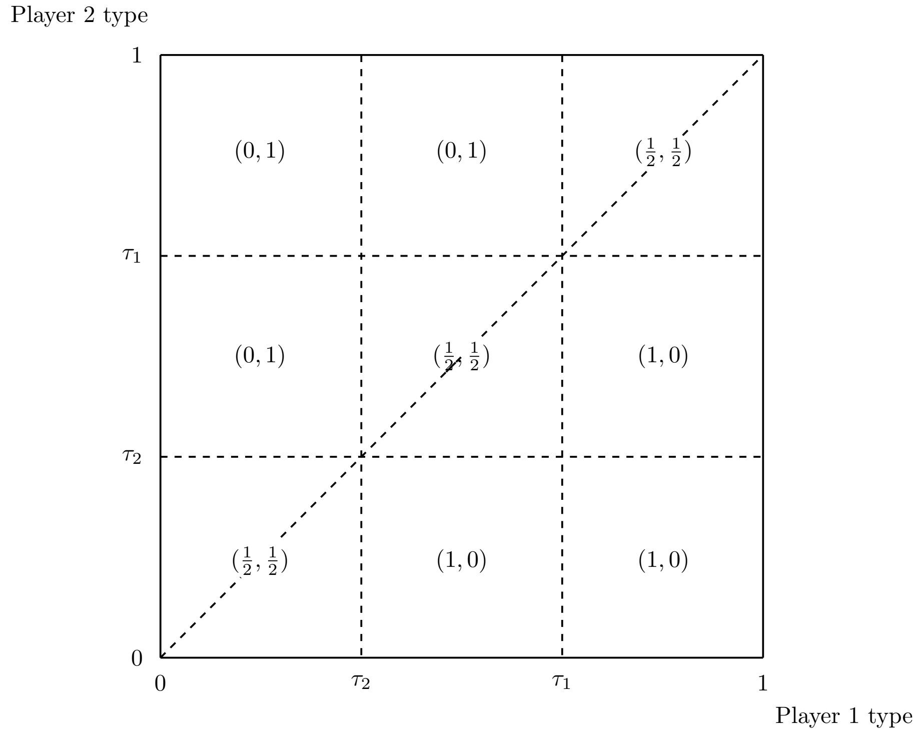

Figure 1: The allocation rule for a mechanism with 3 bins.

mechanism, we again allocate only to players in the highest reported bin, but the exact allocation amongst them depends on the lowest reported type—the lower is this type, the more efficient the allocation.

Formal definitions require some notation. Given a collection of thresholds $\{\tau_k\}_{k=0}^{\infty}$ and a vector of types $\mathbf{t}$, let $\overline{k}(\mathbf{t})$ denote the smallest k such that $t_i \in I_k$ for some $i = 1, 2, ..., n$. Similarly, let $\underline{k}(\mathbf{t})$ denote the largest k such that $\exists\, t_i \in I_k$. Let $m(\mathbf{t})$ denote the number of i such that $t_i \in I_{\overline{k}(\mathbf{t})}$—that is, the number of players in the highest bin. Assuming that $t_1 \in I_{\overline{k}(\mathbf{t})}$, let $r(\mathbf{t}) \in \{1, 2, ..., m(\mathbf{t})\}$ denote the rank of t_1 among types in the highest bin. Finally, let $F_{k,m}$ denote the cdf of the *lowest* type, conditional on k being the highest occupied bin and on m players in that bin. We slightly absue notation by writing $F_{k,m}(I)$ for the measure of an interval under this distribution. To economize on notation, we suppress the dependence of $\overline{k}$, $\underline{k}$, m, and r on the type vector $\mathbf{t}$.

Definition 1. *In a* ***simple bin mechanism (SBM)****, the allocation rule satisfies*

$$b(\mathbf{t}) = \begin{cases} \frac{1}{m} & \text{if } t_1 \in I_{\overline{k}} \\ 0 & \text{otherwise.} \end{cases}$$

In a ***better bin mechanism (BBM)****, we have* $b(\mathbf{t}) = 0$ *if* $t_1 \notin I_{\overline{k}}$*. If* $t_1 \in I_{\overline{k}}$ *the allocation* *rule satisfies* $b(\mathbf{t}) = \frac{1}{m}$ *if* $m = n$*, and otherwise*

$$b(\mathbf{t}) = \begin{cases} \frac{1}{mF_{\overline{k},m}(I_{\underline{k}})} & \text{if } F_{\overline{k},m}(\tau_{\underline{k}}) < \frac{r-1}{m}, \text{ and } \frac{r}{m} < F_{\overline{k},m}(\tau_{\underline{k}-1}) \\ \min\left\{1, \frac{r-mF_{\overline{k},m}(\tau_{\underline{k}})}{mF_{\overline{k},m}(I_{\underline{k}})}\right\} & \text{if } \frac{r-1}{m} \leq F_{\overline{k},m}(\tau_{\underline{k}}) < \frac{r}{m} \\ \min\left\{1, \frac{mF_{\overline{k},m}(\tau_{\underline{k}-1})-(r-1)}{mF_{\overline{k},m}(I_{\underline{k}})}\right\} & \text{if } \frac{r-1}{m} < F_{\overline{k},m}(\tau_{\underline{k}-1}) \leq \frac{r}{m} \\ 0 & \text{otherwise.} \end{cases}$$

Conditional on having m reports in the highest bin, a better bin mechanism divides the conditional distribution of the lowest type into m quantiles. If the lowest type falls in quantile k, then the allocation goes to the kth highest type within the highest bin. The interim expected allocation of each player in the highest bin is the same—each receives the budget with probability $\frac{1}{m}$—but we use the lowest type report as a randomization device. Our first result shows that there exist thresholds such that the corresponding SBM is incentive compatible, and there exists another set of thresholds such that the corresponding BBM is incentive compatible.

Proposition 1. *There exists an incentive compatible SBM and an incentive compatible BBM.*

Proof. Let $\mathcal{T}$ denote the space of weakly decreasing infinite sequences $\{\tau_k\}_{k=0}^{\infty}$ with $\tau_0 = 1$ and $\lim_{k\to\infty} \tau_k = 0 \equiv \tau_\infty$. Pick either allocation rule, and consider player reports of the form (k, t), where k is the reported bin, and t is a type within that bin—note if $\tau_k = \tau_{k+1}$, there is only one possible value of t. In both mechanisms, the interim allocations only depend on the bin a player reports, so we can characterize best responses via the first coordinate of the report. Given truthful reports of the other players, note that a player's reported k in a best response is monotonic in her own type.

Assuming other players report truthfully, let $BR(\boldsymbol{\tau})$ denote the thresholds in a best response—a player whose type is inbetween $BR(\boldsymbol{\tau})_k$ and $BR(\boldsymbol{\tau})_{k-1}$ would report in bin k. The map $BR(\boldsymbol{\tau})$ is monotone in the natural order on $\mathcal{T}$—i.e. $\boldsymbol{\tau} \geq \boldsymbol{\tau'}$ if $\tau_k \geq \tau'_k$ for each k. Given a particular $\overline{\tau} \in \mathcal{T}$, let $\mathcal{T}(\overline{\tau})$ denote the set of $\tau \in \mathcal{T}$ such that $\tau \leq \overline{\tau}$. The proof of existence rests on the following lemma.

Lemma 2. *There exists* $\overline{\tau} \in \mathcal{T}$ *such that* $BR(\tau) \in \mathcal{T}(\overline{\tau})$ *for any* $\tau \in \mathcal{T}(\overline{\tau})$*.*

The space $\mathcal{T}(\overline{\tau})$ is a complete lattice, so given the lemma, we can apply Tarski's fixed point theorem to find a fixed point τ^* of the map BR, and this fixed point corresponds to an incentive compatible mechanism. Note we cannot have $\tau_k^* = 0$ for any $k < \infty$ because a positive mass of types near zero can profitably deviate by reporting bin $k' > k$. Hence, the fixed point uses infinitely many bins.

For the SBM, it is straightforward to show that taking $\overline{\tau}_{k+1} = \alpha\overline{\tau}_k$ provides a suitable upper bound. The proof that such a bound exists for the BBM is in the Appendix.

□

Simple bin mechanisms are a useful benchmark because they are implementable without commitment. The designer can ask in which bin each player's type lies, and without further information, dividing the budget among players in the highest bin is ex-post optimal. In fact, simple bin mechanisms are the best the designer can do without commitment.[3] When $n = 2$, simple bin mechanisms and better bin mechanisms are identical. Examining the case in which $n = 3$ helps clarify how they differ. If only one player is in the highest bin, this player gets the entire budget. When all three are in the same bin, they share it equally. If exactly two players are in the highest bin, then each receives a share of $\frac{1}{2}$ in expectation. However, the bin of the lowest player affects the exact allocation. If the third player has a high type, the resource goes to the *lowest* type among those in the highest bin. If the third player has a low type, then the resource goes to the *highest* type.

We introduce this more complicated allocation rule because it helps provide incentives to report low types. When a player deviates upwards, she reduces her expected payoff from the externality in two ways. First, she may receive a positive allocation that could have gone to a player with a higher type. Second, even if she does not receive a positive allocation, she may negatively distort the allocation between players in the highest bin. By adding this second source of incentives, a better bin mechanism can achieve higher welfare when $n > 2$.

We note two properties of better bin mechanisms that are important for our analysis going forward. First, a BBM never burns money: we fully allocate the budget. Second, whenever a player reports type 0, the entire budget goes to the player with the highest type. To see why, note that having infinitely many bins means there must be a concentration around zero. If we fix the highest bin, then all sufficiently low bins are contained in the first quantile of the lowest-type distribution. Hence, a sufficiently low report makes the budget go to the highest type. A report of zero uniquely results in allocation to the highest type for any realized highest bin. The next section shows that when types are uniformly distributed, these two properties fully characterize the set of optimal mechanisms.

4 Uniform Types

If F is uniform on $[0, 1]$, we obtain a tight characterization of optimal mechanisms.

Theorem 1. *Suppose F is uniform on $[0, 1]$. The better bin mechanism achieves the optimal welfare of*

$$W^* = \frac{1 + \alpha(n-1)}{2 + \alpha(n-1)}. \tag{6}$$

Moreover, any incentive compatible mechanism is optimal if it both

(a) Fully allocates the budget, and

(b) Allocates to the highest type whenever a player reports type zero.

[3] See Li et al. (2016) for a detailed analysis of a similar allocation problem with cheap talk.

Proof. Integrating the incentive compatibility constraint (4) yields

$$\begin{aligned}\int_0^1 tQ(t)+M(t)\mathrm{d}t &= M(0)+\int_0^1\int_0^t Q(x)\mathrm{d}x\mathrm{d}t \\ &= M(0)+\int_0^1\int_x^1 Q(x)\mathrm{d}t\mathrm{d}x \\ &= M(0)+\int_0^1 (1-t)Q(t)\mathrm{d}t.\end{aligned}$$

From the definition of $M(t)$, we have

$$\int_0^1 M(t)\mathrm{d}t = \alpha(n-1)\int_0^1 tQ(t)\mathrm{d}t.$$

Some slight rearranging now yields

$$\int_0^1 (2+\alpha(n-1))\, tQ(t)\mathrm{d}t = M(0)+\int_0^1 Q(t)\mathrm{d}t,$$

implying

$$W = \frac{n}{2+\alpha(n-1)}\left(M(0)+\int_0^1 Q(t)\mathrm{d}t\right).$$

The value of $M(0)$ is bounded above by $\alpha\mathbb{E}_{F^{n-1}}[X] = \frac{\alpha(n-1)}{n}$—i.e. the externality from always allocating to the highest type among the other players. The value of $\int_0^1 Q(t)\mathrm{d}t$ is bounded above by $\frac{1}{n}$ since we are allocating one unit and the mechanism is symmetric. The better bin mechanism achieves both bounds, so we conclude that it is optimal. □

The better bin mechanism achieves the optimum by efficiently providing incentives. In a simple bin mechanism, the allocation within a bin plays no role for incentive provision. In contrast, the better bin mechanism leverages the within-bin allocation to provide incentives for players with low types. Such a player is unlikely to receive any of the resource, so her direct allocation has a limited effect on incentives. In the better bin mechanism, her report still matters even when she receives nothing, and this compensates the agent for the small chance at a higher allocation she would get by reporting a higher type. The lower the type, the greater the compensation. An important feature is that the amount of compensation is continuous at zero. Were this not the case, a small but positive mass of low types would have an incentive to deviate and report type zero.

We make two observations from Theorem 1. First, while the BBM is an optimal mechanism, it is far from the only one. This is easiest to see in the case with $n=2$. For each $w\in(0,1)$, consider an alternative allocation rule in which each player receives a weight w^k when reporting in bin k, and player i recieves an allocation of

$$\frac{w^{k_i}}{w^{k_{-i}}+w^{k_i}}.$$

One can repeat the argument in Proposition 1 to show that thresholds exist that make this allocation incentive compatible. Since $w < 1$, an agent's weight goes to zero as she reports higher k, so we clearly have $M(0) = \mathbb{E}_F[X] = \frac{1}{2}$, the maximum possible. Each $w \in (0, 1)$ corresponds to a distinct optimal mechanism, so there exists a continuum of optimal mechanisms. It seems likely there is nothing particularly special about the bin structure, and one may be able to find many alternative classes of optimal mechanisms for this allocation problem.

Second, the value of commitment increases with n. In a simple bin mechanism, as $n \to \infty$ the highest threshold τ_1 converges to

$$\overline{\tau}_1 = \frac{\alpha}{2 - \alpha}.$$

To see why, note that for n large there is almost certainly at least one player in this interval, and this is precisely the point at which

$$\frac{1 + \tau_1}{2} = \frac{\tau_1}{\alpha}.$$

That is, the expected type of a player in the highest bin is $\frac{1}{\alpha}$ times the lower threshold of the bin. Consequently, as $n \to \infty$, the optimal welfare converges to

$$\frac{1 + \overline{\tau}_1}{2} = \frac{1}{2 - \alpha}.$$

In contrast, we immediately see from (6) that welfare in a better bin mechanism converges to 1. This implies that the highest threshold must converge to 1. How is this possible? Reporting in a higher bin affects the allocation of the budget within the bin. In particular, when we add a player to the highest bin, we change the partition of the lower bins that determines the allocation. For instance, if a player of type τ_1 reports in bin 2, we might allocate to the second highest type in bin 1, but if she misreports in bin 1, we might allocate to the third highest type in bin 1 instead. This shift can lead to a relatively large loss of efficiency if our player expects to share the budget with many others in the highest bin, thus providing incentives to truthfully report.

5 Welfare Bounds for General Distributions

This section provides welfare bounds that apply to any smooth type distribution with full support. We establish an upper bound based on the incentive compatibility constraint, and a lower bound that a better bin mechanism achieves. For a class of distributions, the bounds coincide, showing that better bin mechanisms are optimal. Regardless of the distribution, the lower bound converges to 1 as $n \to \infty$, showing that our earlier point about commitment power holds much more generally.

To state our result, we define the inverse hazard rate

$$h(t) = \frac{1 - F(t)}{f(t)},$$

and the function

$$g(\lambda, t) = \lambda \mathbb{E}_F \left[X \mid X \geq t \right] + \frac{1-\lambda}{1+\alpha(n-1)} \left(\alpha n \mathbb{E}_{F^{n-1}}[X] + \mathbb{E}_F \left[h(X) \mid X \geq t \right] \right). \tag{7}$$

Theorem 2. *In any feasible and incentive compatible mechanism, we have*

$$W \leq \min_{\lambda \in [0,1]} \max_{t \in [0,1]} g(\lambda, t).$$

Any incentive compatible mechanism that

(a) Fully allocates the budget, and

(b) Allocates to the highest type whenever a player reports type zero,

achieves welfare of at least

$$W \geq \max_{\lambda \in [0,1]} \min_{t \in [0,1]} g(\lambda, t).$$

Proof. Integrating the incentive compatibility constraint (4) gives

$$\begin{aligned}
\frac{1+\alpha(n-1)}{n} W = \int_0^1 tQ(t) + M(t) \mathrm{d}F(t) &= M(0) + \int_0^1 \int_0^t Q(x) \mathrm{d}x \mathrm{d}F(t) \\
&= M(0) + \int_0^1 \int_x^1 Q(x) \mathrm{d}F(t) \mathrm{d}x \\
&= M(0) + \int_0^1 \left(1 - F(t)\right) Q(t) \mathrm{d}t \\
&= M(0) + \int_0^1 h(t) Q(t) \mathrm{d}F(t).
\end{aligned}$$

Define $c_n = 1 + \alpha(n-1)$. For any $\lambda \in [0,1]$, we have

$$\begin{aligned}
\frac{c_n}{n} W &= \lambda c_n \int_0^1 tQ(t) \mathrm{d}F(t) + (1-\lambda) \left(M(0) + \int_0^1 h(t) Q(t) \mathrm{d}F(t) \right) \\
&\leq \lambda c_n \int_0^1 tQ(t) \mathrm{d}F(t) + (1-\lambda) \left(\alpha \mathbb{E}_{F^{n-1}}[X] + \int_0^1 h(t) Q(t) \mathrm{d}F(t) \right).
\end{aligned}$$

Write $q(t)$ for the derivative of $Q(t)$, in the sense of distribution, and define $\overline{\mu} = \mathbb{E}_{F^{n-1}}[X]$. Moreover, define the functions

$$G(t) = \int_0^t s \mathrm{d}F(s), \quad \text{and} \quad H(t) = \int_0^t 1 - F(s) \mathrm{d}s = \int_0^t h(s) \mathrm{d}F(s).$$

Note that $G(1) = H(1) = \mathbb{E}_F[X]$.

Integrating by parts gives

$$\begin{aligned}\frac{c_n}{n}W &\le \lambda c_n \left(Q(1)G(1) - \int_0^1 q(t)G(t)\mathrm{d}t\right) \\ &\qquad + (1-\lambda)\left(\alpha\overline{\mu} + Q(1)H(1) - \int_0^1 q(t)H(t)\mathrm{d}t\right) \\ &= (1-\lambda)\alpha\overline{\mu} + \int_0^1 \left(\lambda c_n \left(G(1)-G(t)\right) + (1-\lambda)\left(H(1)-H(t)\right)\right) q(t)\mathrm{d}t.\end{aligned}$$

Feasibility implies that

$$\frac{1}{n} \ge \int_0^1 Q(t)\mathrm{d}F(t) = Q(1) - \int_0^1 q(t)F(t)\mathrm{d}t = \int_0^1 q(t)\left(1-F(t)\right)\mathrm{d}t,$$

and incentive compatibility requires that $q(t) \ge 0$. Hence, for each $\lambda \in [0,1]$, the solution to

$$\begin{aligned}\max_{q(t)\ge 0} \quad & (1-\lambda)\alpha\overline{\mu} + \int_0^1 \left(\lambda c_n \left(G(1)-G(t)\right) + (1-\lambda)\left(H(1)-H(t)\right)\right) q(t)\mathrm{d}t \\ s.t. \quad & \int_0^1 q(t)\left(1-F(t)\right)\mathrm{d}t \le \frac{1}{n},\end{aligned} \tag{8}$$

is an upper bound on $\frac{c_n}{n}W$. This program has a straightforward solution. Viewing q as a distribution, if we put a unit of mass at t, it contributes the amount

$$\lambda c_n \left(G(1)-G(t)\right) + (1-\lambda)\left(H(1)-H(t)\right)$$

to our objective at the "cost" $1-F(t)$. Hence, the optimal q places a point mass of $\frac{1}{n(1-F(t))}$ at a t that maximizes

$$\begin{aligned}&\lambda c_n \frac{G(1)-G(t)}{1-F(t)} + (1-\lambda)\frac{H(1)-H(t)}{1-F(t)} \\ &\qquad = \lambda c_n \mathbb{E}_F\left[X \mid X \ge t\right] + (1-\lambda)\mathbb{E}_F\left[h(X) \mid X \ge t\right].\end{aligned}$$

Hence, the maximum value of the program (8) is

$$\max_{t\in[0,1]} (1-\lambda)\alpha\overline{\mu} + \frac{\lambda c_n}{n}\mathbb{E}_F\left[X \mid X \ge t\right] + \frac{1-\lambda}{n}\mathbb{E}_F\left[h(X) \mid X \ge t\right] = \max_{t\in[0,1]} \frac{c_n}{n} g(\lambda, t).$$

Since this is an upper bound on $\frac{c_n}{n}W$ for every $\lambda \in [0,1]$, the first claim in the theorem follows.

For the second claim, note if an incentive compatible mechanism satisfies properties (a) and (b), then a similar calculation gives

$$\frac{c_n}{n}W = \lambda c_n \int_0^1 tQ(t)f(t)\mathrm{d}t + (1-\lambda)\left(\alpha\overline{\mu} + \int_0^1 h(t)Q(t)\mathrm{d}F(t)\right)$$

for any $\lambda \in [0,1]$. We conclude welfare is bounded below by $\min_{t\in[0,1]} g(\lambda, t)$, and the result follows.

$\square$

Theorem 2 offers a simple way to obtain bounds on the optimal welfare. Note that despite the suggestive form of the bounds, we cannot apply standard minimax theorems to show that they are equal because the max and the min swap both order and variables. Nevertheless, there is an easy special case in which the bounds are tight.

Corollary 1. *If $h(t)$ is linear—that is, if $F(t) = 1 - (1-t)^\gamma$ for some $\gamma > 0$—then a better bin mechanism is optimal.*

Proof. If $h(t)$ is linear, we can choose $\lambda \in [0, 1]$ so that $g(\lambda, t)$ is constant as a function of t. Theorem 2 now implies the result. □

Finally, we note that welfare from a better bin mechanism converges to 1 as $n \to \infty$.

Corollary 2. *As $n \to \infty$, the welfare from a better bin mechanism converges to* 1.

Proof. Take $\lambda = 0$ in the lower bound, and note that $g(0,t) > \frac{\alpha n \mathbb{E}_{F^{n-1}}[X]}{1+\alpha(n-1)} \to 1$ as $n \to \infty$. The result follows. □

6 Related Work

Our problem involves allocating a good with interdependent values and no transfers, bridging two literatures in mechanism design. Work on mechanism design with interdepenent values largely focuses on settings with transfers (e.g. Jehiel et al., 1996, 1999; Jehiel and Moldovanu, 2001). McLean and Postlewaite (2015) show that a modification of the VCG mechanism can achieve efficient allocations at low cost when players are "informationally small." Kucuksenel (2012) looks at an allocation problem with altruistic preferences and characterizes the set of interim efficient allocations with transfers. The availability of transfers leads to very different mechanisms derived via different techniques.

A recent and growing literature studies cardinal mechanism design without transfers (e.g. Ben-Porath et al., 2014; Mylovanov and Zapechelnyuk, 2017; Goldlücke and Tröger, 2018). A key idea in this work is that when there are multiple goods to allocate, one can link decisions to satisfy incentives—if a player receives one good, it affects her chances of receiving other goods.[4] Borgers and Postl (2009) and Miralles (2012) study two related problems. Borgers and Postl study two players who choose among three alternatives. The players' ordinal rankings of the alternatives are commonly known and diametrically opposed. The value of the middle alternative to each player is private information and is drawn independently from a known distribution. Here, the probability that a player obtains her most preferred alternative acts as a type of money, allowing the designer to extract information. Miralles (2012) studies the allocation of two goods between multiple players without transfers. The optimal mechanism treats the probability of receiving one good as a numeraire, which is used to create "transfers." One might view our model as allocating two goods—giving resource to one player versus giving it to others—but a crucial difference is that players cannot observe

[4] Hylland and Zeckhauser (1979) introduced the idea of using units of goods as "money."

their value from other players receiving the good. This significantly complicates interim expected payoffs and prevents us from applying standard techniques. (e.g. Myerson, 1981).

A number of papers explore the tradeoff that players face when they are biased towards certain actions but strong signals can overcome this bias. This work focuses on settings without commitment (e.g. Li et al., 2001). For instance, Li et al. (2016) study a cheap talk model with two senders. Each sender is responsible for a project and has private information regarding its value. While the senders share common interests with the receiver, each has own-project bias. The model is similar to ours when $n = 2$, but there is no commitment, and the analysis is limited to uniform types. The authors show that all equilibria are partitional, and the best equilibrium has infinitely many partition elements—this corresponds to our simple bin mechanism. Our results show that, for two senders, the best equilibrium outcome of the cheap talk game is also the best that a designer with commitment power can achieve.

7 Final Remarks

Externalities from resource allocation provide a tool to elicit information without transfers—people may willingly reveal low values if the benefit from allocating to others is sufficiently high. An intuitive bin mechanism proves optimal under some conditions and approximately efficient when there are many players. Two features of the mechanism stand out as noteworthy. First, it uses coarse information, providing interim allocations that depend only on a partition of the type space into intervals. In this respect, the allocation appears similar to equilibria of cheap talk games. However, the exact allocation depends on the *lowest* reported type, becoming more efficient when this type is lower. This dependence is a crucial source of incentives for reporting low types, and it requires commitment power to implement.

One insight of our analysis is that the value of commitment power depends on group size. In our setting, a small group of players limits the value of commitment. One can bound the welfare difference between a simple bin mechanism and a better bin mechanism by comparing the corresponding values of $M(0)$. When n is small, this difference is small—when $n = 2$, commitment power has no value. A second insight is that large groups require very little altruism to align incentives. As long as α is not zero, we can achieve an approximately efficient allocation when n is sufficiently large.

There are natural directions to extend our analysis. First, one might seek a more general characterization of optimal mechanisms. As noted in section 4, there is nothing special about the bin structure we use. We provide one way to construct a mechanism that achieves two key bounds, but other welfare equivalent mechanisms exist. Moreover, outside the class of distributions with linear inverse hazard rates, it is not at all clear that our mechanism is optimal. In fact, one can construct examples with point masses in which money burning yields an improvement. Another extension would consider asymmetric type distributions or asymmetric externalities. While it seems likely that one could derive similar bounds based on reduced form allocation rules, it is not immediately clear how to adjust our mechanism. The general problem may require new techniques to solve.

References

Ben-Porath, Elchanan, Eddie Dekel, and Barton Lipman (2014), "Optimal Allocation with Costly Verification." *American Economic Review*, 104, 3779–3813.

Border, Kim (1991), "Implementation of Reduced Form Auctions: A Geometric Approach." *Econometrica*, 59, 1175–1187.

Borgers, Tilma and Peter Postl (2009), "Efficient Compromising." *Journal of Economic Theory*, 144, 2057–2076.

Goldlücke, Susanne and Thomas Tröger (2018), "Assigning an Unpleasant Task Without Payment." Working Paper.

Hylland, Aanund and Richard Zeckhauser (1979), "The Efficient Allocation of Individuals to Positions." *Journal of Political Economy*, 67, 293–314.

Jehiel, Philippe and Benny Moldovanu (2001), "Efficient Design with Interdependent Valuations." *Econometrica*, 69, 1237–1259.

Jehiel, Philippe, Benny Moldovanu, and Ennio Stacchetti (1996), "How (Not) to Sell Nuclear Weapons." *American Economic Review*, 86, 814–829.

Jehiel, Philippe, Benny Moldovanu, and Ennio Stacchetti (1999), "Multidimensional Mechanism Design for Auctions with Externalities." *Journal of Economic Theory*, 85, 258–293.

Kucuksenel, Serkan (2012), "Behavioral Mechanism Design." *Journal of Public Economic Theory*, 14, 767–789.

Li, Hao, Sherwin Rosen, and Wing Suen (2001), "Conflicts and Common Interests in Committees." *American Economic Review*, 91, 1478–1497.

Li, Zhuozheng, Heikki Rantakari, and Huanxing Yang (2016), "Competitive Cheap Talk." *Games and Economic Behavior*, 96, 65–89.

McLean, Rich and Andrew Postlewaite (2015), "Implementation with Interdependent Valuations." *Theoretical Economics*, 10, 923–952.

Miralles, Antonio (2012), "Cardinal Bayesian Allocation Mechanisms Without Transfers." *Journal of Economic Theory*, 147, 179–206.

Myerson, Roger (1981), "Optimal Auction Design." *Mathematics of Operations Research*, 6, 58–73.

Mylovanov, Tymofiy and Andriy Zapechelnyuk (2017), "Optimal Allocation with Ex Post Verification and Limited Penalties." *American Economic Review*, 107, 2666–2694.

Proof of Lemma 2

Note it is sufficient to show that there exists some $\overline{\tau}$ such that $BR(\overline{\tau}) \leq \overline{\tau}$—since BR is monotone, this means $BR(\tau) \leq \overline{\tau}$ for all $\tau \in \mathcal{T}(\overline{\tau})$. Fix a player i, and let $\overline{k}$ be the highest ranked bin of other $n-1$ players. Let $\overline{m}$ denote the number of players other than i in bin $\overline{k}$. Let r denote the bin in which player i reports. Let $\overline{t}$ denote the type of the player to whom the budget is allocated, conditional on this player being different from i. The payoff to i with type t reporting in bin k of a BBM is

$$\begin{aligned} t\mathbb{P}\left(\overline{k} > k\right) + \mathbb{P}\left(\overline{k} = k\right)\left(t\mathbb{E}\left[\frac{1}{\overline{m}+1} \,|\, \overline{k} = k\right] + \alpha\mathbb{E}\left[\frac{\overline{m}}{\overline{m}+1} \,|\, \overline{k} = k\right]\mathbb{E}[\overline{t} \,|\, r = \overline{k} = k]\right) \\ + \alpha\mathbb{P}\left(\overline{k} < k\right)\mathbb{E}[\overline{t} \,|\, r = k,\, \overline{k} < k] \end{aligned}$$

Taking the difference between this and the same condition for bin $k+1$, we find the value of the threshold $BR(\tau)_k$ is the t that solves

$$\begin{aligned} &t\left(\mathbb{P}(\overline{k} = k+1)\mathbb{E}\left[\frac{\overline{m}}{\overline{m}+1} \,|\, \overline{k} = k+1\right] + \mathbb{P}(\overline{k} = k)\mathbb{E}\left[\frac{1}{\overline{m}+1} \,|\, \overline{k} = k\right]\right) \\ &\quad = \alpha\mathbb{P}(\overline{k} = k+1)\mathbb{E}\left[\frac{\overline{m}}{\overline{m}+1} \,|\, \overline{k} = k+1\right]\mathbb{E}[\overline{t} \,|\, r = \overline{k} = k+1] \\ &\quad\quad + \alpha\mathbb{P}(\overline{k} = k)\mathbb{E}\left[\frac{1}{\overline{m}+1} \,|\, \overline{k} = k\right]\mathbb{E}[\overline{t} \,|\, r = \overline{k} = k] \\ &\quad\quad + \alpha\mathbb{P}(\overline{k} \leq k)\left(\mathbb{E}[\overline{t} \,|\, r = k+1,\, \overline{k} \leq k] - \mathbb{E}[\overline{t} \,|\, r = k,\, \overline{k} \leq k]\right), \end{aligned}$$

We need to bound three terms on the right hand side. Dividing through by the coefficient on the left hand side, we see that the first two terms are a weighted average of

$$\alpha\mathbb{E}_F[X \,|\, X \in (\tau_{k+1}, \tau_k)] \quad \text{and} \quad \alpha\mathbb{E}_F[X \,|\, X \in (\tau_k, \tau_{k-1})].$$

The first term is bounded by $\alpha\tau_k$, and the second is bounded by $\alpha\tau_{k-1}$. Hence, the weighted average is bounded by $\alpha\tau_{k-1}$, regardless of the sequence $\{\tau_k\}_{k\in\mathbb{N}}$

To bound the last term, we pick a specific sequence $\{\tau_k\}_{k\in\mathbb{N}}$. Fix some small $\epsilon > 0$, and choose the thresholds such that $\mathbb{P}_F(X \leq \tau_k) = (1-\epsilon)^k$. Note this implies that

$$\mathbb{P}(\overline{k} > k) = (1-\epsilon)^{(n-1)k} \quad \text{and} \quad \mathbb{P}(\overline{k} = k) = (1-\epsilon)^{(n-1)(k-1)}\left(1 - (1-\epsilon)^{n-1}\right).$$

For a binomial random variable X with $n-1$ trials and success probability p, we have $\mathbb{E}\left[\frac{1}{X+1}\right] = \frac{1-(1-p)^n}{np}$. From this we can deduce that

$$\mathbb{E}\left[\frac{1}{\overline{m}+1} \,|\, \overline{k} = k\right] = \frac{1-(1-\epsilon)^n}{n\epsilon}, \quad \mathbb{E}\left[\frac{\overline{m}}{\overline{m}+1} \,|\, \overline{k} = k\right] = 1 - \frac{1-(1-\epsilon)^n}{n\epsilon}.$$

We can therefore write

$$\begin{aligned}
&\mathbb{P}(\overline{k} = k+1)\mathbb{E}\left[\frac{\overline{m}}{\overline{m}+1} \,|\, \overline{k} = k+1\right] + \mathbb{P}(\overline{k} = k)\mathbb{E}\left[\frac{1}{\overline{m}+1} \,|\, \overline{k} = k\right] \\
&\qquad = (1-\epsilon)^{(n-1)k}\left(1-(1-\epsilon)^{n-1}\right)\left(1 - \frac{1-(1-\epsilon)^n}{n\epsilon}\right) \\
&\qquad\quad + (1-\epsilon)^{(n-1)(k-1)}\left(1-(1-\epsilon)^{n-1}\right)\left(\frac{1-(1-\epsilon)^n}{n\epsilon}\right) \\
&\qquad = (1-\epsilon)^{(n-1)(k-1)}\left(1-(1-\epsilon)^{n-1}\right)\left(1 + \left(1-(1-\epsilon)^{n-1}\right)\frac{1-(1-\epsilon)^n}{n\epsilon}\right) \\
&\qquad \geq (1-\epsilon)^{(n-1)(k-1)}\left(1-(1-\epsilon)^{n-1}\right).
\end{aligned}$$

For the other side, we can rewrite the term

$$\mathbb{P}(\overline{k} \leq k)\left(\mathbb{E}[\overline{t} \,|\, r = k+1,\, \overline{k} \leq k] - \mathbb{E}[\overline{t} \,|\, r = k,\, \overline{k} \leq k]\right)$$

as

$$\sum_{j \leq k} \mathbb{P}(\overline{k} = j)\left(\mathbb{E}[\overline{t} \,|\, r = k+1,\, \overline{k} = j] - \mathbb{E}[\overline{t} \,|\, r = k,\, \overline{k} = j]\right).$$

For $j < k$, any difference between $\mathbb{E}[\overline{t} \,|\, r = k+1,\, \overline{k} \leq k]$ and $\mathbb{E}[\overline{t} \,|\, r = k,\, \overline{k} \leq k]$ derives from the event that a) multiple other players report in bin $\overline{k}$, and b) bin k is the lowest bin among the other agents. We can bound the difference by bounding the probabilities of a) and b). The probability that at least two other agents report in bin j, conditional on j being the highest reported bin, is

$$1 - (1-\epsilon)^{n-2} \leq (n-2)\epsilon.$$

The probability that bin k is the lowest bin among the other agents is

$$\left(1-(1-\epsilon)^k\right)^{n-1}.$$

The last piece we need to construct our bound comes from the fact that the change in report from $k+1$ to k has no effect on efficiency if k is much lower than j—i.e. if k is low enough so that a report in bin k would still result in the highest type getting the allocation. In particular, there is a constant $c > 0$ such that the change in report has no effect if

$$(1-\epsilon)^{k-j} < c.$$

Let $\mathcal{C}_k$ denote the set of $j < k$ such that $(1-\epsilon)^{k-j} < c$. Moreover, since the density of the type distribution is uniformly bounded away from zero, there exists another constant c' such that

$$\tau_{j-1} - \tau_j \leq c'(\tau_{k-1} - \tau_k)(1-\epsilon)^{k-j}$$

for every j. Since the loss of efficiency within bin j is bounded by $\tau_{j-1} - \tau_j$, we can bound the sum over bins $j < k$ by

$$\sum_{j \in \mathcal{C}_k} (\tau_{j-1} - \tau_j)(1-\epsilon)^{(n-1)(j-1)} \left(1 - (1-\epsilon)^{n-1}\right) \left(1 - (1-\epsilon)^k\right)^{n-1} (n-2)\epsilon$$
$$\leq c' \sum_{j \in \mathcal{C}_k} (\tau_{k-1} - \tau_k)(1-\epsilon)^{(n-1)(j-1)+k-j} \left(1 - (1-\epsilon)^{n-1}\right) \left(1 - (1-\epsilon)^k\right)^{n-1} (n-2)\epsilon.$$

For $j = k$, we can bound the efficiency loss by

$$\mathbb{P}(\overline{k} = k)(\tau_{k-1} - \tau_k) = (1-\epsilon)^{(n-1)(k-1)} \left(1 - (1-\epsilon)^{n-1}\right) (\tau_{k-1} - \tau_k).$$

Using our lower bound on the left hand side coefficient, we can bound the contribution of the third term by α times

$$\begin{aligned}(\tau_{k-1} - \tau_k) &\left(1 + \sum_{j \in \mathcal{C}_k} (1-\epsilon)^{(n-2)(j-k)} \left(1 - (1-\epsilon)^k\right)^{n-1} (n-2)\epsilon\right) \\ &\leq (\tau_{k-1} - \tau_k)\left(1 + (c')^{-(n-2)}(n-2)\epsilon\right) \\ &= (\tau_{k-1} - \tau_k)\left(1 + c''\epsilon\right).\end{aligned}$$

Consequently, we have for every threshold that

$$BR(\overline{\tau})_k \leq \alpha \tau_{k-1} + \alpha(1 + c''\epsilon)(\tau_{k-1} - \tau_k).$$

The ratio $\frac{\tau_{k-1}}{\tau_k}$ is uniformly bounded and converges to 1 as ϵ becomes small, so we can choose ϵ small enough that this is less than τ_k for all k.