GRAPE
Group for Research in APplied Economics

GRAPE Working Paper #70

Equity-efficiency trade-off in quasi-linear environments

Piotr Dworczak

FAME | GRAPE, 2022

Foundation of Admirers and Mavens of Economics
Group for Research in Applied Economics

Equity-efficiency trade-off in quasi-linear environments

Piotr Dworczak
Northwestern University
& FAME|GRAPE

Abstract
I study a simple equity-efficiency problem: A designer allocates a fixed amount of money to a population of agents differing in privately-observed marginal values for money. She can only screen agents by asking them to burn utility (through some socially wasteful activity). I show that giving a lump-sum payment is outperformed by a mechanism with utility burning when the agent with the lowest money-denominated cost of engaging in the wasteful activity has an expected value for money that exceeds the average value by more than a factor of two.

Keywords:
equity-efficiency trade-off, costly screening, allocation of money

JEL Classification
C78, D47, D61, D63, D82

Corresponding author
Piotr Dworczak, piotr.dworczak@northwestern.edu

Acknowledgements
I would like to thank Mohammad Akbarpour, Pawel Doligalski, Scott Duke Kominers, Philipp Strack, and Joanna Tyrowicz for helpful comments and discussions. I gratefully acknowledge the support received under the ERC Starting grant IMD-101040122

Published by: FAME | GRAPE
ISSN: 2544-2473

Foundation of Admirers and Mavens of Economics
Koszykowa 59/7
00-660 Warszawa
Poland

W | grape.org.pl
E | grape@grape.org.pl
TT | GRAPE_ORG
FB | GRAPE.ORG
PH | +48 799 012 202

1 Introduction

Governments often redistribute by allocating direct cash transfers. A natural concern that arises in these contexts is whether financial aid is received by those most in need. While some basic information about potential recipients may be available to public agencies, many of the relevant characteristics—the detailed financial situation, family circumstances, labor market opportunities—remain unobserved. When these characteristics cannot be easily verified, governments can attempt to improve targeting by requiring applicants to engage in "ordeals"—such as queuing or filling out forms—that may help screen out those who are not in need. For example, Alatas et al. (2016) show that targeting can be improved by imposing the ordeal of traveling to a registration site in the design of the Indonesia's Conditional Cash Transfer program.[1] However, ordeals—by definition—can be quite burdensome, and they decrease the utility of the recipients without offering any direct social benefit. As a result, governments may choose to forgo any screening. For example, the US government distributed financial aid during the Covid-19 pandemic by mailing checks, imposing only extremely crude eligibility criteria based on income.[2] These contrasting examples motivate the basic question: When is it optimal to improve the targeting of monetary transfers via costly screening?

To answer this question, I consider the following redistribution problem. A designer allocates a fixed amount of money to a population of agents differing in privately-observed marginal values for money (dispersion in values for money means that the designer has redistributive preferences). Absent additional tools, the designer cannot achieve any screening in an incentive-compatible mechanism—she can only offer a lump-sum payment. To introduce a trade-off, I allow the designer to pay a higher amount of money to agents who publicly "burn" some utility by completing an ordeal—an activity that is costly for the person engaging in it and does not directly benefit anyone. Any such activity entails a pure social waste since, by definition, its direct consequence is a decrease in utility. However, if only some agents find it optimal to burn utility, the designer can potentially achieve a better allocation of money. This leads to an equity-efficiency trade-off.

I show that offering cash for completing the ordeal outperforms the lump-sum transfer if agents with the lowest money-denominated cost of engaging in the ordeal have an expected value for money that exceeds the average value for money by more than a factor of **two**. This restriction on the joint distribution of costs and values for money can be reduced to a testable

[1] The idea that ordeals can be useful in screening is very well understood in economics; see, among many others, Hartline and Roughgarden (2008), Condorelli (2012), and Chakravarty and Kaplan (2013) for theoretical contributions, and Rose (2021), Zeckhauser (2021) for recent practical examples.

[2] See, for example, **?** and **?**.

condition under parametric assumptions. I also demonstrate that the conclusion is robust to allowing the designer to pick an optimal mechanism from the set of all incentive-compatible, individually rational, and budget-balanced mechanisms.

The equity-efficiency trade-off lies at the core of public economics. Classical papers, such as Diamond and Mirrlees (1971) or Atkinson and Stiglitz (1976), provided frameworks to evaluate the trade-off in relatively complex environments, where closed-form solutions are generally not available. Weitzman (1977) studied the trade-off in the simpler context of allocating a single good, and Condorelli (2013) showed how to resolve the trade-off optimally using mechanism-design tools in quasi-linear environments. Adopting the approach of Saez and Stantcheva (2016), a growing literature on inequality-aware market design models the redistributive preferences of the designer using dispersion in marginal values for money. **(alias?)** (2021) (henceforth DKA) studied a two-sided market for a homogenous good, and showed that inefficient rationing is part of an optimal market design when the expected value for money for traders with the lowest rate of substitution exceeds the average value for money by a factor of **two** or more. They call this condition "high inequality." A number of recent papers obtained an analogous condition (featuring the factor two) in various models of allocating resources to agents with quasi-linear utilities: Kang (2020) and Pai and Strack (2022)—when the allocation of the good generates externalities; Kang and Zheng (2020)—when a good and a bad is allocated, and **(alias?)** (2022)—in a setting with heterogeneous qualities of the good.

The note demonstrates that the high-inequality condition is not specific to resource allocation problems; rather, it reflects a fundamental trade-off between efficiency and redistribution in quasi-linear environments. When physical goods are allocated, efficiency requires that different agents receive a different allocation. In contrast, in the problem of allocating money, the unique efficient allocation is a lump-sum payment. Any screening must necessarily reduce efficiency, and is hence motivated entirely by redistributive motives. This clean separation of efficiency and equity allows me to provide a simple graphical intuition for why "two" is the relevant threshold in the high-inequality condition, thereby explaining its appearance in papers on optimal redistribution in resource allocation problems.

2 Problem

A designer has a budget $B > 0$ of money that she wants to allocate to a unit mass of agents. Each agent is characterized by a "marginal value for money" v. The parameter v is interpreted as a social welfare weight—it is the *social* value of giving an agent a dollar, as in Saez and Stantcheva (2016). I assume that the designer knows the distribution of values for money in the population but does not observe individual realizations. The idea is that the designer has redistributive preferences but does not directly observe characteristics that determine these preferences; for example, a designer may derive a higher value from giving a dollar to an agent in a difficult financial situation but she cannot easily verify who is really in need.[3]

If the allocation of money cannot be made contingent on any additional information available to the designer, the only feasible mechanism is to give a lump-sum payment to all agents. The total value generated for a utilitarian designer is $\mathbb{E}[v] \cdot B$, where $\mathbb{E}[v]$ is the average value for money in the population. The parameter $\mathbb{E}[v]$ is sometimes referred to as the "marginal value of public funds" in the public finance literature, and it can be normalized to 1. Under this normalization, the values v are measured in units of the opportunity cost of public funds.

Suppose, however, that the designer can ask agents to burn utility. Specifically, there is some activity—an "ordeal"—that is costly for the agents, has no intrinsic social value, and is observed by the designer. The designer can choose a difficulty (arduousness) y of the ordeal that I normalize to $[0, 1]$. (The main results are unaffected if $y \in \{0, y_0\}$, that is, when the designer cannot adjust the difficulty of the ordeal.) Each agent has a privately-observed cost c of completing the ordeal. Moreover, agents' utilities are quasi-linear: An agent with cost c and social value for money v who completes an ordeal with difficulty y and receives a monetary transfer t obtains utility

$$-cy + vt.$$

Note that the cost c in the above expression is measured in social-utility units (since the coefficient on the monetary transfer is v). From the agent's perspective, all that matters for the choice of y and t is the ratio between v and c, which I denote $k \equiv c/v$. The parameter k is the cost of the ordeal for the agent expressed in monetary units.

Assuming the designer knows the joint distribution of c and v, her objective is to maximize

$$\mathbb{E}[-cy(k) + vt(k)]$$

over $y(k)$ and $t(k)$ subject to the budget constraint (the sum of transfers is at most B) and incentive-compatibility, stating that an agent with type k chooses $(y(k),\, t(k))$ from the menu $\{(y(k'),\, t(k')\}_{k'}$ offered by the designer. The argument k in $y(k)$ and $t(k)$ indicates that agents differing in the relative cost k may choose different combinations of y and t. I

[3] If the designer has access to observable information about agents, e.g., she can verify whether an agent's income is above or below a threshold, we can interpret the problem as being solved conditional on some realization of the observable information, with the distribution of v being the conditional distribution—see **(alias?)** (2022) for a formalization of this idea.

formalize the mechanism-design problem in Section 4.[4]

Any positive level of y is a pure social waste. In particular, if the designer has no redistributive preferences (all v's are equal to 1), then the optimal mechanism is a lump-sum transfer. With redistributive preferences, however, there may be a trade-off between efficiency and equity if choosing positive levels of y for some agents allows the designer to allocate money to those with higher values v.

Throughout, I assume that k has a distribution F with a continuously differentiable density f on $[\underline{k}, \bar{k}]$, with $\underline{k} = 0$, $f(0) > 0$ and $f'(0) < \infty$. The economically meaningful assumption is that the lower bound of the support of the distribution of costs $\underline{k}$ is 0; I discuss the consequences of relaxing these assumptions in Section 5.

3 Main result

Consider the simplest possible screening mechanism: The designer offers an additional payment t_0 (on top of the lump-sum transfer) for agents who are willing to engage in some ordeal $y_0 > 0$. Agents with relative cost $k \leq t_0/y_0$ choose this option; the remaining agents decline (and only receive the lump-sum transfer).

Given the joint distribution of (v, c), define the conditional expectation function $g(k) = \mathbb{E}\left[v \middle| \frac{c}{v} = k\right]$. To avoid issues associated with conditional expectations being defined only for almost all k, I assume that g is continuous. From now on, I use the short-hand notation $\mathbb{E}\left[v \middle| k\right]$ interpreted as the expected value for money conditional on the relative cost k. Intuitively, this object will play a key role in the analysis because the mechanism screens agents based on k, while the designer ultimately wants to allocate money based on v. Thus, the function $\mathbb{E}\left[v \middle| k\right]$ determines the targeting effectiveness of the mechanism.

By a simple calculation (using the law of iterated expectations), welfare associated with the simple mechanism (assuming B is large enough) is given by

$$\int_0^{t_0/y_0} \mathbb{E}\left[v \middle| k\right](t_0 - ky_0)f(k)dk + (B - t_0 F(t_0/y_0)).$$

The first term of the welfare function captures the fact that an agent with type $k \leq t_0/y_0$ receives a payment t_0 and burns utility ky_0, thus enjoying net utility equivalent to receiving a monetary payment $t_0 - ky_0$. The designer values that utility at $\mathbb{E}\left[v \middle| k\right](t_0 - ky_0)$, since she values a dollar given to an agent with type k at $\mathbb{E}\left[v \middle| k\right]$. The second term of the welfare

[4] It is well known that in two-dimensional linear models like the current one the designer cannot do better by trying to screen c and v separately. It is without loss of optimality for the designer to condition the allocation on the ratio k of c to v—see Jehiel and Moldovanu (2001), Che et al. (2013), or DKA for proofs of analogous claims.

function captures the fact that a total amount $t_0 F(t_0/y_0)$ of money has been paid out in additional compensation, leaving less funds in the budget for lump-sum transfers.

Let us further choose t_0 so that only a small fraction of agents choose to engage in the ordeal. Specifically, let $t_0 = \epsilon y_0$ for some small $\epsilon > 0$, so that only types $k \leq \epsilon$ accept. Then, the ordeal mechanism, which I will denote by $M(\epsilon)$, outperforms the lump-sum payment mechanism if and only if

$$\int_0^\epsilon \mathbb{E}[v|k](\epsilon - k) f(k) dk > \epsilon F(\epsilon). \tag{1}$$

The condition says that the welfare gain of additional utility enjoyed by types in $[0, \epsilon]$ must exceed the opportunity cost of the required expenditure $\epsilon F(\epsilon)$. (The opportunity cost is equal to the expenditure because I normalized the value of public funds to 1.) Observe that the condition will hold for small enough ϵ if the ratio of the left hand side to the right hand side converges to a number strictly larger than 1 in the limit as ϵ goes to zero. We have

$$\lim_{\epsilon \to 0} \frac{\int_0^\epsilon \mathbb{E}[v|k](\epsilon - k) f(k) dk}{\epsilon F(\epsilon)} = \lim_{\epsilon \to 0} \frac{\int_0^\epsilon \mathbb{E}[v|k] f(k) dk}{\epsilon f(\epsilon) + F(\epsilon)} = \frac{\mathbb{E}[v|\underline{k}]}{2}, \tag{2}$$

where I used L'Hôpital's rule twice, and relied on the regularity assumptions on the distribution. We thus obtain the following result.

Proposition 1. *If the expected value for money conditional on the lowest relative cost exceeds the value of public funds by more than a factor of two, i.e., if*

$$\mathbb{E}[v|\underline{k}] > 2, \tag{$\star$}$$

then the ordeal mechanism $M(\epsilon)$ strictly outperforms giving a lump-sum transfer for some positive $\epsilon > 0$.

For intuition, notice that requiring agents to burn utility in exchange for a larger monetary transfer achieves redistribution of money to agents with the lowest relative cost of engaging in the ordeal. Achieving this redistribution is costly: As equation (2) reveals, for each dollar of public funds spent, only 1/2 of the dollar is received by agents in form of a net utility increase; the other 1/2 gets "burned" in the process of screening. Thus, the social value of targeting the monetary transfer must exceed the value of public funds by more than a factor of two for the ordeal mechanism to be socially valuable on the net.

The interpretation of condition ($\star$) is easiest when $\mathbb{E}[v|k]$ is decreasing in k. This is a natural case since $k \equiv c/v$, and thus v and k are typically inversely related. Then, the designer derives the highest expected value from giving money to the agent with the lowest cost $\underline{k}$. The value $\mathbb{E}[v|\underline{k}]$ depends both on the strength of the designer's redistributive preferences (the dispersion in v's), as well as on the targeting effectiveness of the ordeal. Condition ($\star$) states that the targeting effectiveness of the ordeal must be sufficiently high so that the designer is willing to trade off efficiency for equity.

The intuition for why "2" appears in condition ($\star$) is related to the one offered in DKA but more robust and simpler since it does not rely on the context of a goods-allocation problem. I argue that the "2" is in fact a direct consequence of quasi-linearity of preferences in money. To see that, instead of applying L'Hôpital's rule as in (2), apply the mean value theorem for integrals to the left hand side of (1):

$$\int_0^{\epsilon} \mathbb{E}[v|k](\epsilon - k)f(k)dk = \mathbb{E}[v|\delta_\epsilon]f(\delta_\epsilon)\int_0^{\epsilon}(\epsilon - k)dk \tag{3}$$

for some $\delta_\epsilon \in [0, \epsilon]$. For small ϵ, the first two terms can be approximated by $\mathbb{E}[v|\underline{k}]f(\underline{k})$. Intuitively, when ϵ is small, we can ignore the differences in welfare and probability weights applied to the utilities of different agents with costs in $[0, \epsilon]$, and instead apply the same weight $\mathbb{E}[v|\underline{k}]f(\underline{k})$ to all of them. Due to quasi-linearity of preferences, the surplus of agents who accept the ordeal can be represented as an isosceles triangle (see Figure 3.1). If the weight is constant, the total surplus is calculated as the area of this triangle. The area is exactly **half** of the area of a square with side ϵ that approximates the associated opportunity cost of public funds (see the light-grey square in Figure 3.1, with $\epsilon F(\epsilon) \approx f(\underline{k})\epsilon^2$ that holds approximately for small ϵ). Thus, to compensate for the surplus lost due to costly screening, the designer must value the area of the triangle more than twice as much as the area of the square, in social utility units. This gives us condition ($\star$).

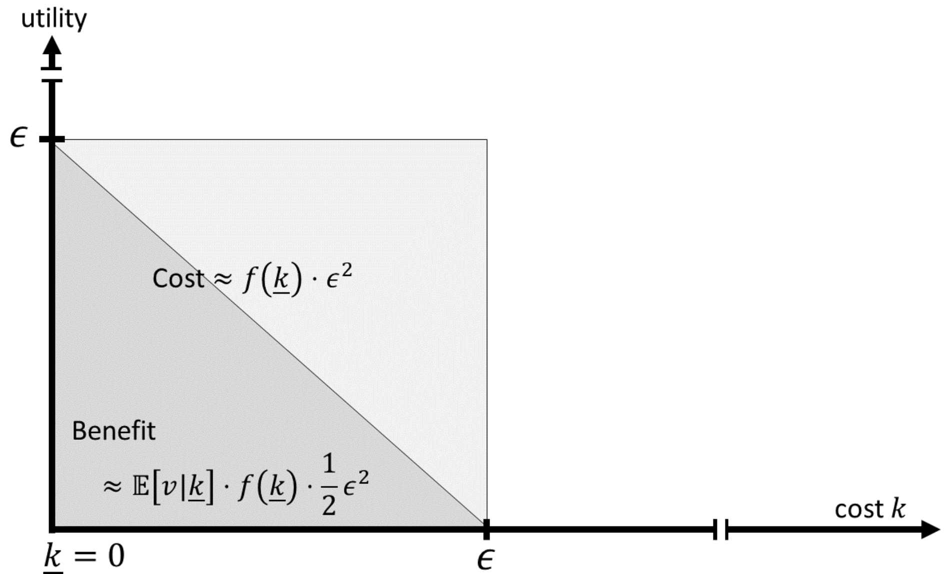

Figure 3.1: The surplus triangle and the equity-effiency trade-off

3.1 Parametric example

I present a parametric example to illustrate condition ($\star$). The parametrization is stylized but shows one possible way towards an empirical test of the condition.[5] The main obstacle is that marginal values for money are not observable. I will assume that they are equal to the slope of a (common) concave utility function for wealth at agents' respective wealth levels. This approach implicitly assumes that the program that the designer is running is small enough that it does not significantly alter the wealth distribution—then, quasi-linearity in money can be obtained as a linear approximation of the concave-utility model around agents' wealth levels.

Suppose that each agent has a utility function for wealth w,

$$U(w) = \frac{w^{1-\theta}}{1-\theta},$$

that takes the CRRA form, with coefficient of relative risk aversion $\theta > 0$. Agents differ in their wealth levels w, and thus differ in the marginal values $U'(w)$. Note that condition ($\star$) depends on the shape of the right tail of the distribution of values for money, and hence will depend on the left tail of the distribution of wealth. For this reason, I choose a family of distributions of wealth indexed by how thick the left tail is:

$$G(w) = w^{\beta},$$

for $\beta > 0$ and $w \in [0, 1]$.[6] Under this parametrization, the bottom 50% of agents with wealth below w hold $\left(\frac{1}{2}\right)^{\frac{1+\beta}{\beta}}$ of the total wealth held by agents with wealth below w, for any w. Thus, lower β corresponds to a thicker left tail. Because

$$\mathbb{P}(U'(w) \leq s) = 1 - \left(\frac{1}{s}\right)^{\frac{\beta}{\theta}},$$

marginal utilities have a Pareto distribution with tail parameter $\alpha \equiv \beta/\theta$. Assuming that $\alpha > 1$ (so that the Pareto distribution has a finite mean) and using the fact that the mean of the Pareto distribution is $\alpha/(\alpha-1)$, we can define

$$v \equiv \frac{\alpha - 1}{\alpha} \cdot U'(w),$$

[5] See Allen and Rehbeck (2021) for a discussion of the empirical testability of the high-inequality condition from DKA.

[6] It will be clear that this assumption is only relevant for small enough w; the distribution can be different for high w without affecting the results.

so that the average marginal value for money is normalized to 1. Finally, suppose that the costs c are uniformly distributed on $[0, 1]$. If c and v are assumed independent, then by direct calculation,[7]

$$\mathbb{E}\left[v|k\right] = \frac{(\alpha-1)^2}{\alpha(\alpha-2)}\frac{1-k^{\alpha-2}}{1-k^{\alpha-1}}.$$

In particular,

$$\mathbb{E}\left[v|\underline{k}\right] = \begin{cases} \frac{(\alpha-1)^2}{\alpha(\alpha-2)} & \alpha > 2, \\ \infty & \alpha \leq 2. \end{cases}$$

Therefore, condition ($\star$) holds if and only if

$$\frac{\beta}{\theta} < 1 + \sqrt{2}.$$

Intuitively, the designer has stronger redistributive preferences when there are more poor agents (β is lower) or when agents are more risk averse (θ is higher), so that they have a particularly high marginal value for money at low wealth levels. By Proposition 1, when the tail of the Pareto distribution of values for money is thick enough, it becomes optimal to sacrifice efficiency to achieve better redistribution. While the empirical estimates of θ vary widely depending on the method and context, most studies obtain that θ is weakly greater than 1. A simple empirical property of the left tail of the distribution of wealth is then sufficient for condition ($\star$): The bottom 50% of agents with wealth below some low threshold w should hold no more than $\left(\frac{1}{2}\right)^{\frac{2+\sqrt{2}}{1+\sqrt{2}}} \approx 37.5\%$ of wealth in that group.

4 Optimal mechanism

In this section, I connect the result about the simple ordeal mechanism to the question of optimal design. The main message is that the conclusions derived from the analysis of the simple mechanism carry over to the optimal mechanism.

It is well known (see footnote 4) that the optimal mechanism only screens agents based on their relative costs k, and hence the optimization problem for the designer can be written as finding the best direct mechanism of the form:

[7] For $\alpha = 2$, the term $(1 - k^{\alpha-2})/(\alpha - 2)$ is replaced by $-\log(k)$.

$$\max_{y(k)\in[0,1],\, t(k)\geq 0} \int_{\underline{k}}^{\bar{k}} \mathbb{E}[v|\, k](-ky(k)+t(k))dF(k), \tag{OBJ}$$

$$-ky(k)+t(k) \geq -ky(k')+t(k'),\, \forall k,\, k', \tag{IC}$$

$$-ky(k)+t(k) \geq 0,\, \forall k, \tag{IR}$$

$$\int_{\underline{k}}^{\bar{k}} t(k)dF(k) = B. \tag{B}$$

By standard arguments (see Appendix A), we can then derive the following result.

Proposition 2. *The optimal mechanism uses an ordeal (y is strictly positive for a positive-measure set of agents) if and only if*

$$\mathbb{E}[V(k)|\, k \leq k'] > 0 \textit{ for some } k' > 0, \tag{4}$$

where

$$V(k) = \left(\mathbb{E}\left[v|\frac{c}{v} \leq k\right] - 1\right)\frac{F(k)}{f(k)} - k. \tag{5}$$

Condition ($\star$) *implies that* $V(k) > 0$ *for small enough* k*; hence, condition* ($\star$) *implies condition* (4)*. Conversely, when* $V(k)$ *is quasi-concave, condition* (4) *implies that condition* ($\star$) *must hold as a weak inequality.*

When B *is large enough, the optimal mechanism offers a single payment* $k^\star$ *for completing the ordeal* $y = 1$*, and allocates the remaining budget as a lump-sum payment.*[8]

Thus, at least when the budget B is high enough, the simple mechanism $M(\epsilon)$ considered in Section 3 is actually optimal for some choice of ϵ (not necessarily small). Additionally, condition ($\star$) is not only sufficient for the optimality of some redistribution but also (almost) necessary under regularity conditions.

The function $V(k)$ expresses the trade-off between redistribution and efficiency. The first term in brackets shows how money gets transferred from an average agent to an agent with cost below k, which is typically a desirable effect for a designer with redistributive preferences. The inverse hazard rate $F(k)/f(k)$ appears because it measures information rents in one-dimensional screening problem in which lower types receive higher utility. The second term captures the inefficiency—it is equal to the cost k due to the normalization of the value of public funds to 1. Condition (4) states that we can find a threshold type k' such

[8] It suffices that $B \geq \bar{k}$ but a more permissive bound is derived in the proof in Appendix A. When B is low, there is no lump-sum payment, and it may be the case that the optimal mechanism allows agents to choose a "less difficult" ordeal for a smaller payment.

that the positive redistributive effect exceeds the negative inefficiency effect *on average* for types below k'. The averaging is a consequence of incentive-compatibility, since if type k' finds it optimal to engage in the ordeal, so do all types below k'.

The proof in Appendix A shows that the optimal design problem consists in maximizing $\int V(k)y(k)dF(k)$ over decreasing functions $y(k)$ with values in $[0, 1]$. Following Myerson (1981), we can replace $V(k)$ with the ironed value $\overline{V}(k)$ that is decreasing in k. It follows that the optimal $y(k)$ is a threshold rule: It is 1 for all $k \leq k^\star$ for some $k^\star$. The threshold $k^\star$ is strictly positive precisely when condition (4) holds. Under the assumption that $\underline{k} = 0$, we have $V(0) = 0$. If $V'(0) > 0$, then $V(k) > 0$ for small enough k, and thus (4) holds. By a simple calculation, $V'(0) > 0$ is equivalent to ($\star$). Additionally, if $V(k)$ is quasi-concave, then there can exist a k at which $V(k) > 0$ (which is clearly required for (4) to hold) only if $V(k)$ is increasing in the neighborhood of 0 which implies the weak inequality in ($\star$).

5 Discussion

On the assumption $\underline{k} = 0$. Throughout, I have assumed that there are some agents for whom the marginal cost of the ordeal is arbitrarily small. Suppose instead that the distribution of k is bounded away from zero by $\underline{k} > 0$. Proposition (1) is then false. Equation (2) reveals that the costs and benefits of an ϵ distortion away from a lump-sum payment are both of order ϵ^2 when $\underline{k} = 0$. However, the inefficiency is of order ϵ when $\underline{k} > 0$, and thus it is never optimal to deviate from a lump-sum payment by offering a small additional payment for a "small" ordeal. The first part of Proposition (2) still holds: A "large" ordeal is optimal if $\mathbb{E}[V(k)|\, k \leq k'] > 0$ for some k', where the phrase "large" is justified since the expectation will be negative for k' close to 0. Indeed, equation (5) reveals that $V(k)$ starts out strictly negative when $\underline{k} > 0$. A positive derivative of $V(k)$ at $\underline{k}$ (which is equivalent to condition ($\star$)) is thus no longer sufficient; however, it is still necessary if $V(k)$ is quasi-concave.

The economic meaning of the condition $\underline{k} = 0$ is that there exist perturbations of the efficient mechanism that result in a small (arbitrarily close to zero) per-agent loss in efficiency. Thus, while the condition may seem restrictive in the current model, it actually arises naturally in many other equity-efficiency problems. For example, in resource-allocation problems, a small amount of rationing induces a small per-agent loss in efficiency *regardless* of the support of agents' values for the resource. This is because rationing results in allocating the resource to agents who do not have the highest willingness to pay; the efficiency loss can be kept arbitrarily small if rationing applies to a small interval of agent values, regardless of their absolute magnitude. This explains why papers focusing on redistribution in resource-allocation problems, such as DKA and Kang and Zheng (2020), obtained a high-inequality condition analogous to ($\star$) without imposing any restrictions on the support of the distribution of values.

On the distribution of costs. I presented the intuition for condition ($\star$) under relatively strong assumptions about the distributions of relative costs k. These assumptions ensured that the distribution is well behaved around the lowest type, making local analysis possible. For contrast, consider a binary-type model. That is, suppose that the relative cost is either $\underline{k}$ or $\bar{k} > \underline{k}$. Then, there are only two candidate optimal mechanisms: a lump-sum payment, or an ordeal mechanism in which only the low-cost type completes the ordeal for an additional payment that is just low enough that the high-cost type declines. Simple algebra yields that the ordeal mechanism outperforms the lump-sum transfer if

$$\mathbb{E}[v|\,\underline{k}] > \frac{\bar{k}}{\bar{k} - \underline{k}}.$$

The ratio on the right hand side of the inequality is always above 1, meaning that the expected value for money for agents receiving the additional payment must exceed the value of public funds for an ordeal to be optimal. The magnitude of the required improvement depends on how binding the incentive-compatibility constraint is. If the high type has a relatively low cost, achieving targeting requires the low type to burn a lot of utility, and thus $\mathbb{E}[v|\,\underline{k}]$ must be very high to justify using an ordeal mechanism; however, as the cost of the high type increases, the required level of $\mathbb{E}[v|\,\underline{k}]$ declines.

The binary-cost model reveals that the more realistic continuous-cost model is needed to obtain the factor of two in condition ($\star$). This is intuitive in light of Figure 3.1—a continuous distribution combined with quasi-linearity is responsible for creating the surplus triangle.

On policy implications. While the model is purposefully simplistic and not intended to produce detailed policy implications, its analysis delivers some insights about when and which ordeals should be used. As demonstrated in Section 3.1, under additional parametric assumptions, condition ($\star$) for optimality of ordeals can be tested. A non-parametric test could also be developed under some assumption about how the marginal values of money are determined. Thus, the note provides a crude but concrete quantitative test of whether targeting monetary transfers through an ordeal is desirable.

Condition ($\star$) also provides some high-level intuition for which ordeals to use. In order for the condition to hold, y must be less costly to generate for agents with high values for money (i.e., poorer agents). For example, when the ordeal is queueing, agents should not be allowed to pay others to stand in line on their behalf. Similarly, agent's ability to produce

y should not depend on characteristics that correlate positively with wealth. For example, if the ordeal is to fill out a complicated form or apply online, more educated agents may be able to complete the task more quickly, and hence find it less costly even if they have a higher opportunity cost of time. Somewhat paradoxically then, certain types of red tape that waste petitioners' time and energy without a clear purpose may actually come closer to being optimal.

References

Alatas, V., R. Purnamasari, M. Wai-Poi, A. Banerjee, B. A. Olken, and R. Hanna (2016): "Self-targeting: Evidence from a field experiment in Indonesia," Journal of Political Economy, 124, 371–427.

Allen, R. and J. Rehbeck (2021): "Obstacles to Redistribution Through Markets and One Solution," Available at arxiv: https://ssrn.com/abstract=4180522.

Atkinson, A. B. and J. E. Stiglitz (1976): "The design of tax structure: direct versus indirect taxation," Journal of Public Economics, 6, 55–75.

Chakravarty, S. and T. R. Kaplan (2013): "Optimal allocation without transfer payments," Games and Economic Behavior, 77, 1 – 20.

Che, Y.-K., W. Dessein, and N. Kartik (2013): "Pandering to Persuade," American Economic Review, 103, 47–79.

Condorelli, D. (2012): "What money can't buy: Efficient mechanism design with costly signals," Games and Economic Behavior, 75, 613 – 624.

——— (2013): "Market and non-market mechanisms for the optimal allocation of scarce resources," Games and Economic Behavior, 82, 582–591.

Diamond, P. A. and J. A. Mirrlees (1971): "Optimal taxation and public production I: Production efficiency," American Economic Review, 61, 8–27.

Doval, L. and V. Skreta (2018): "Constrained Information Design: Toolkit," Working Paper.

Dworczak, P. ⓡ S. D. Kominers ⓡ M. Akbarpour (2021): "Redistribution through markets," Econometrica, 89, 1665–1698.

Hartline, J. D. and T. Roughgarden (2008): "Optimal mechanism design and money burning," in Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, 75–84.

Jehiel, P. and B. Moldovanu (2001): "Efficient Design with Interdependent Valuations," Econometrica, 69, 1237–1259.

Kang, M. and C. Z. Zheng (2020): "Pareto Optimality of Allocating the Bad," Working paper.

Kang, Z. Y. (2020): "Markets for goods with externalities," Working Paper.

M. Akbarpour ⓡ Dworczak, P. ⓡ S. D. Kominers (2022): "Redistributive Allocation Mechanisms," Working Paper.

Myerson, R. B. (1981): "Optimal auction design," Mathematics of Operations Research, 6, 58–73.

Pai, M. and P. Strack (2022): "Taxing Externalities Without Hurting the Poor," Available at SSRN: https://ssrn.com/abstract=4180522.

Rose, J. L. (2021): "Rationing with time: time-cost ordeals' burdens and distributive effects," Economics & Philosophy, 37, 50–63.

Saez, E. and S. Stantcheva (2016): "Generalized social marginal welfare weights for optimal tax theory," American Economic Review, 106, 24–45.

Weitzman, M. L. (1977): "Is the price system or rationing more effective in getting a commodity to those who need it most?" Bell Journal of Economics, 8, 517–524.

Zeckhauser, R. (2021): "Strategic sorting: the role of ordeals in health care," Economics & Philosophy, 37, 64–81.

A Proof of Proposition 2

The proof relies on the standard ironing technique.[9] Let u denote the utility of the highest type $\bar{k}$ in the mechanism. The envelope formula yields that in an incentive-compatible individually-rational mechanism, $u \geq 0$ and

$$-ky(k) + t(k) = u + \int_k^{\bar{k}} y(s)ds.$$

The above condition, combined with the requirement that $y(k)$ is non-increasing, is necessary and sufficient for (IC) and (IR). Using integration by parts, I can rewrite the budget constraint (B) as

$$\int_{\underline{k}}^{\bar{k}} \left(k + \frac{F(k)}{f(k)}\right) y(k)dF(k) + u = B. \tag{6}$$

[9] See Myerson (1981), and in the context of optimal redistribution Condorelli (2013) or **(alias?)** (2022), among many others.

Let α be the Lagrange multiplier on the budget constraint (B). Using integration by parts again, I can rewrite the optimal design problem as

$$\max_{y(k)\in[0,1],\, u\geq 0} \int_{\underline{k}}^{\bar{k}} \left[\left(\mathbb{E}\left[v|\frac{c}{v}\leq k\right]-\alpha\right)\frac{F(k)}{f(k)}-\alpha k\right] y(k)dF(k)+(1-\alpha)u \tag{OBJ'}$$

$$y(k) \text{ is non-increasing,} \tag{M}$$

and α must be such that a solution $(y^\star, u^\star)$ to the above problem satisfies the budget constraint (6). Existence of solution requires that $\alpha \geq 1$. I conjecture that $\alpha = 1$, and later discuss how to modify the analysis when the budget constraint (6) does not hold with that conjecture. Under the conjecture, the objective (OBJ') is equal to $\int_{\underline{k}}^{\bar{k}} V(k)y(k)dF(k)$, where V is defined by equation (5).

Next, I define the ironed value function. Let

$$\Psi(t) = -\int_t^1 V(F^{-1}(x))dx,$$

and let $\mathrm{co}(\Psi)$ denote the concave closure of Ψ. Then, I can define

$$\overline{V}(k) = (\mathrm{co}\Psi)'(F(k)),$$

as the ironed value function, and the value of the problem of maximizing $\int_{\underline{k}}^{\bar{k}} \overline{V}(k)y(k)dF(k)$ is the same as the original one. Thus, the optimal solution is

$$y^\star(k) = \mathbf{1}_{\{\overline{V}(k)\geq 0\}}.$$

Note that this solution is feasible (non-increasing) since the ironed value function is non-increasing. Let $k^\star$ be the largest k such that $\overline{V}(k) = 0$. Then, $y^\star(k) = \mathbf{1}_{\{k\leq k^\star\}}$, and the optimal mechanism is to offer a payment $k^\star$ for the ordeal $y = 1$.

Assuming that the budget constraint is satisfied, a lump-sum payment mechanism is optimal if and only if $k^\star = 0$, that is, if and only if $\overline{V}(k) \leq 0$ for all k (and otherwise, a simple ordeal mechanism is optimal). This condition is equivalent to $\mathrm{co}(\Psi)(t)$ being a decreasing function, which in turn (given that it is a concave closure of Ψ) is equivalent to $\Psi(0) \geq \Psi(t)$ for all t. Thus, a lump-sum payment mechanism is optimal if and only if

$$\int_{\underline{k}}^{k} V(k)dF(k) \leq 0, \tag{7}$$

for all k. Dividing both sides by $F(k)$ allows me to rewrite condition (7) as

$$\mathbb{E}[V(k)|\, k \le k'] \le 0,\ \forall k'.$$

Note that $V(0) = 0$, and

$$V'(0) = \mathbb{E}[v|\, \underline{k}] - 2,$$

so condition ($\star$) implies that $V(k)$ is strictly positive for small k. Thus, if ($\star$) holds, then (7) cannot hold, and using an ordeal mechanism is optimal. When $V(k)$ is quasi-concave, it can only be positive if it is increasing at $k = 0$. Thus, condition (7) can be violated only if $V'(0) \ge 0$, which requires that condition ($\star$) holds as a weak inequality.

Finally, I check whether the budget constraint holds. The budget constraint holds whenever there exists $u^\star \ge 0$ such that $k^\star F(k^\star) + u^\star = B$, that is, whenever $B \ge k^\star F(k^\star)$. If the optimal mechanism is not to use the ordeal, then $k^\star = 0$, so the condition holds. The condition also holds when B is large enough. When $B < k^\star F(k^\star)$—which can only be true if an ordeal is used—the above solution is not valid as there is no way to satisfy the budget constraint. In that case, we must have $\alpha > 1$, and hence it is uniquely optimal to set $u^\star = 0$. However, since $B > 0$, regardless of the exact optimal value of α, budget balance requires $y(k)$ to be strictly positive for a positive measure of k—an ordeal must still be used. Thus, while the form of the optimal mechanism may differ from the one described above, the conditions for *using* the ordeal are unaltered.[10]

[10]Satisfying the budget constraint with equality in the case $\alpha > 1$ might require implementing two different payments for two distinct levels of the ordeal, for reasons analogous to the ones described by Doval and Skreta (2018).