Discussion Paper Series – CRC TR 224

Discussion Paper No. 228
Project B 03

Mechanism Design for Unequal Societies

Marco Reuter [1]
Carl-Christian Groh [2]

November 2020

[1] University of Mannheim, e-mail: marco.reuter@gess.uni-mannheim.de
[2] University of Mannheim, e-mail: carl-christian.groh@gess.uni-mannheim.de

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Mechanism Design for Unequal Societies*

Marco Reuter† — University of Mannheim

Carl-Christian Groh‡ — University of Mannheim

First Version: May 20, 2020

Last Update: October 26, 2020

Abstract

We study optimal mechanisms for a utilitarian designer who seeks to assign multiple units of an indivisible good to a group of agents with unit demand. The agents have heterogeneous marginal utilities of money, which implies that utility is not perfectly transferable between them. Heterogeneous marginal utilities of money may naturally arise in environments where agents have different wealth endowments. We show that the ex post efficient allocation rule is not optimal in our setting. Firstly, a high willingness to pay may stem from a low marginal utility of money. Moreover, the transfer rule does not only facilitate implementation of the desired social choice function in our setting, but also directly affects social welfare. In the optimal mechanism, rationing may occur, which entails a conflict between ex ante and ex post efficiency. In an extension, we show that it is still not utilitarian optimal to allocate the good solely based on willingness to pay even when redistribution is not possible. Finally, we highlight how our mechanism can be implemented as an auction with minimum bids and bidding subsidies.

Keywords: optimal mechanism design, redistribution, inequality, auctions

JEL Classification: D44, D47, D61, D63, D82

---

*We thank Thomas Tröger, Volker Nocke, Hans-Peter Grüner and Vitali Gretschko for insightful comments. This work was supported by the University of Mannheim's Graduate School of Economic and Social Sciences and the SFB 884 Political Economy of Reforms funded by the German Research Foundation (DFG). Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B03) is gratefully acknowledged.

†marco.reuter@gess.uni-mannheim.de

‡carl-christian.groh@gess.uni-mannheim.de

1 Introduction

Consider the following canonical mechanism design problem with a twist: The designer owns one unit of an indivisible good and two people are vying for the allocation of this good. One of those agents is an ordinary worker with little wealth, while the other agent is a billionaire. Naturally, the ordinary worker has a much higher marginal utility of money than the billionaire. What allocation rule should a utilitarian designer choose? The canonical allocation rule, which is ex post efficient, stipulates that the good should be allocated to the agent that states the highest willingness to pay. Any other assignment rule would allow the agents to find a mutually beneficial trade among each other, where the agent with the higher willingness to pay purchases the good from the agent with the lower willingness to pay. Even when the marginal utilities of money are heterogeneous, this argument still holds as it is independent of the scaling of the utility functions. When all agents have the same marginal utility of money, the utilitarian optimal allocation rule is equivalent to the ex post efficient allocation rule. In that setting, the sole purpose of the transfer rule is to implement the desired social choice function because transfers are neutral in terms of social welfare. When the marginal utilities of money are heterogeneous, this is no longer true. This implies that the equivalence of the ex post efficient allocation rule and the utilitarian optimal allocation rule breaks down when agents have different marginal utilities of money.

Suppose that the billionaire, who gets very little utility from an additional unit of money, has a very low appreciation for the good as such. As a result, the billionaire announces a willingness to pay on the lower end of the support of his potential willingnesses to pay. Assume that, for this type realization, the virtual valuation of the billionaire is negative. As defined in Myerson [1981], the virtual valuation is the maximal amount of revenue that can be raised from an agent in exchange for the allocation of the good. Suppose further that the willingness to pay of the billionaire is still just above the willingness to pay that the ordinary worker announces. In the standard mechanism, the good will thus be allocated to the billionaire.

However, this allocation choice is not utilitarian optimal for two reasons: Firstly, the willingness to pay of the billionaire is substantially affected by her low marginal utility of money. In particular, a high willingness to pay does not necessarily imply a high consumption utility. For the criterion of ex post efficiency, only the ratio between the valuation for the good and the valuation for money, i.e. the willingness to pay, is relevant, and not the absolute levels of these valuations. However, the low marginal utility of money of the billionaire implies that a given consumption utility is mapped into a much higher willingness to pay than for the ordinary worker. This distortion must be taken into account. Secondly, the standard allocation rule fails to accurately capture the opportunity cost of this particular allocation decision. Because the virtual valuation of the billionaire is negative at her type announcement, allocating the good to the billionaire after this announcement will reduce the expected revenue the mechanism will receive from the billionaire. Under a budget balance condition, any additional dollar that the billionaire receives from the mechanism ex ante is a dollar that the ordinary worker loses ex ante. When all agents have the same marginal utility of money, this transfer is neutral in terms of social welfare. However, in the above situation, money is transferred to an agent who benefits from an additional dollar very little instead of an agent who benefits greatly from an additional dollar. This implicit movement of money caused by allocating the good to the billionaire when the billionaire's virtual valuation is negative is thus clearly not utilitarian optimal when the difference in the marginal utilities of money is sufficiently large.

Economic inequality is pervasive in reality and affects the economic incentives of agents. In particular, an individual's wealth affects the marginal utility of money of the agent. As indicated above, optimal mechanisms need to take this into account. We formalize the above ideas in the following framework: The designer initially owns m units of an indivisible good which can be allocated to N agents with unit demand for this good. We assume that the good is scarce, i.e. $m < N$. Following Dworczak et al. [2019], an agent's utility consists of two parts. An agent receives utility v^K when being allocated the good. Moreover, agents attain utility from the money that they receive from the mechanism. The marginal utility of money for an agent, which we call v^M, is constant for each agent but varies across agents. This reflects the idea that individuals may be heterogeneous in characteristics such as wealth that directly impact their marginal utility of money. Both v^K and v^M are private information. As in Dworczak et al. [2019], we abstract from wealth effects on an individual level. This approach ensures the analytical tractability of our model while still capturing the key working channels. In this framework, we characterize the utilitarian optimal mechanism which obeys individual rationality, Bayesian incentive compatibility, and ensures that the budget of the designer is balanced ex ante. The fact that we only require the budget to be balanced ex ante and not in every possible state of the world is without loss of generality, given the insights of Börgers and Norman [2009]. Further, our results enable us to analyse the effect of economic inequality on the optimal mechanism.

We derive the optimal mechanism based on the following ideas: First, we note that a sufficient statistic for individual behavior is the willingness to pay, namely $r = v^K/v^M$. We show that the designer has nothing to gain in terms of social welfare by eliciting both v^K and v^M, compared to a simpler mechanism that elicits only r. This result is based on an analogous result in Dworczak et al. [2019]. Secondly, we note that, as outlined in Myerson [1981], the revenue that can be raised from any agent in exchange for the allocation of the good is equal to her virtual valuation. On the one hand, if the virtual valuation of an agent is negative, allocating the good to this agent reduces the money that the designer can allocate ex ante. On the other hand, allocating this unit will raise the consumption utility of this agent. If any budget is allocated ex ante, it will go to the agent with the highest expected marginal utility of money. These considerations represent a trade-off between efficient allocation and redistributive concerns. A designer interested only in Pareto efficiency does not face such a trade-off, since Pareto efficiency does not rank distributions of money, as long as all the money is distributed. By contrast, a utilitarian designer would rather allocate a unit of money to an agent with a high marginal utility of money, ceteris paribus.

This trade-off is captured by the key statistic of our model: The *inequality adjusted valuation* of an individual. The inequality adjusted valuation has three components: The first component captures the designer's desire to efficiently allocate the good as such. The second component is the virtual valuation of the agent at her type realization, multiplied by the difference between the agent's expected marginal utility of money and the largest expected marginal utility of money of all agents. Allocating the good to an agent changes the expected transfer the designer pays to this agent by the virtual valuation. This change will impact the money the designer allocates ex ante to the agent with the highest marginal utility of money. The effect that this implicit movement of money has on social welfare is captured by the second term. The third component deals with the ex interim uncertainty regarding the marginal utility of money for a given agent. Whenever a realization of r implies that the marginal utility of money is likely to be greater than the ex ante expected marginal utility of money, chances are that the valuation for the good as such will be higher. This implies a positive effect on the benefits of allocation, which is captured by this term.

To the best of our knowledge, we are the first to derive a utilitarian optimal allocation rule in the above setting. The optimal allocation rule $x_i(r_i, r_{-i})$ has a bang-bang property and is defined as follows: The m individuals with the highest inequality adjusted valuations will receive the good, provided their respective inequality adjusted valuations are positive. Note that the good is not necessarily allocated to the agents with the highest willingnesses to pay. The three components of the inequality adjusted valuation capture the aforementioned trade-off which the allocation rule needs to solve optimally. By allocating the object to the agents with the highest positive inequality adjusted valuations, the designer maximizes utilitarian social welfare. Moreover, allocating the good to agents with negative inequality adjusted valuations is never optimal. Agents with negative inequality adjusted valuations necessarily have negative virtual valuations. Therefore, allocating the good to these agents reduces the revenue that can be raised from them. This reduction in revenue has opportunity costs - namely allocating this money ex ante to the agent with the highest marginal utility of money. When the inequality adjusted valuation is negative, these opportunity costs outweigh the direct benefits of allocating the good to the agent.

It turns out that this allocation rule entails rationing for certain situations, where not all units of the good are allocated. This is notable, since it implies that utilitarian social welfare maximization comes at the cost of ex post efficiency. Formally, rationing occurs whenever there are strictly fewer agents with positive inequality adjusted valuations than goods to be allocated. As seen above, allocating the good to agents with negative inequality adjusted valuations cannot be utilitarian optimal. As in Myerson [1981], rationing occurs because of its effect on the revenue that can be raised from agents. We examine when rationing will occur and analyse how the probability of rationing is impacted by the degree of wealth inequality. In doing so, we shed light on a new working channel through which economic inequality may affect the degree of allocative inefficiency.

In an extension, we investigate to what extent our results are driven by the redistributive motive of the designer. Formally, we introduce the additional constraint that no agent may receive transfers from the mechanism in expectation. This implies that the designer has no possibility for redistribution. In this setting, optimal allocation is based solely on the inferred *consumption utilities* of the agents. The ex post efficient allocation rule remains suboptimal. In the special case where the agents' marginal utilities of money are deterministic, the designer is able to perfectly elicit the consumption utility of each agent. In this case, the inequality adjusted valuation is exactly equal to the agent's consumption utility. The good is allocated to the agents with the highest inequality adjusted valuations, which are not necessarily equal to the agent's willingnesses to pay.

Finally, we describe how our mechanism can be implemented as an auction with minimum bids and bidding subsidies. We show that agents with high conditional marginal utilities of money receive bidding subsidies when competing with other agents. Thus, agents who are poor receive bidding subsidies to allow them to compete with wealthier agents. We interpret our results as a rationale for income dependant fee structures which can be found in the assignment of kindergarden places in Germany or in the assignment of college spots in the US, where students from low income households are eligible to receive discounts.

The rest of our paper proceeds as follows: We lay out the related literature in section 2. In section 3, we outline our framework. Section 4 is devoted to the characterization of the optimal mechanism. Section 5 presents some examples. Section 6 concludes.

2 Related Literature

Our work broadly relates to three strands of literature. Firstly, our research has strong connections to the contributions that characterize optimal mechanisms in non-quasilinear settings. Secondly, our work relates to the contributions from various fields which investigate the role of heterogenous marginal utilities of money. Thirdly, some of the key ideas and results of our paper resemble insights from different research areas of the public finance literature.

Previous extensions of the standard quasilinear framework incorporate one of the following two features: Differences in the marginal utilities of money between agents and wealth effects for any given agent. Our paper explicitly models settings where inequality induces heterogeneity in the utility of money in between agents. While we abstract from wealth effects, our work connects to the contributions in mechanism design that consider generalizations of the quasilinear preference framework. One of the earliest such contributions was Maskin and Riley [1984], who pin down the optimal auction in a setting with risk-averse buyers. Esö and Futo [1999] pin down the optimal auction when the seller is risk-averse. Baisa and Burkett [2019] design an auction for a setting where bidders have interdependent values and non-quasilinear preferences. Dughmi and Peres [2018] show that any allocation rule which is dominant-strategy-implementable when agents have quasilinear preferences is also dominant-strategy-implementable when players have concave utility of money in combination with a modified payment rule. Kesselheim and Kodric [2018] study the price of anarchy in settings with risk-averse agents. Several other papers incorporate generalized features of real-life decision situations other than risk-aversion into the mechanism design framework. Eisenhuth [2019] characterizes the revenue-maximizing auction when agents are loss averse and the reference point is endogenous to the choice of the mechanism. Pai and Vohra [2014] and Kotowski [2020] analyse, among others, allocation problems where buyers face heterogenous budget constraints.

Another strand of this literature focuses on characterizing the set of mechanisms that retain certain desiderata in non-quaslinear settings. Saitoh and Serizawa [2008] study the set of mechanisms that satisfy the VCG desiderata in a setting where m goods are to be allocated to people with non-quasilinear preferences. Hashimoto and Saitoh [2010] provide a generalization of the Clarke-mechanism that retains certain desiderata in non quasilinear settings. Kos and Messner [2013] pin down general necessary conditions for incentive-compatibility in a single-agent setting. Morimoto and Serizawa [2015] characterize the set of allocation rules that satisfy pareto efficiency, individual rationality, strategy proofness and a condition stating that people who are not allocated the good receive no transfers under weak assumptions on preferences. Kazumura et al. [2020] study necessary and sufficient conditions for a mechanism to be dominant-strategy incentive compatible in non-quasilinear preference domains.

Our modeling technique, in particular the utility function with two dimensional types, is based on the setup in Dworczak et al. [2019]. While the methodology of our paper closely resembles theirs, the focus is different. Dworczak et al. [2019] pin down optimal trading mechanisms in markets with a distinct buyer and seller side where the designer chooses the mechanism. We characterize the optimal mechanism for the allocation of a number of indivisible goods to a finite number of agents when the designer owns all the goods ex ante. Moreover, Dworczak et al. [2019] model a continuum of agents on both sides of the market while we focus on a finite number of agents. We purposefully work with a finite number of agents, since this is reflected in a number of real world allocation problems we attempt to model. Many of these applications are "local" allocation problems, for example the allocation of kindergarden places to students or the allocation of organs to potential recipients. In terms of outcomes, the results of Dworczak et al. [2019] mirror ours in the sense that rationing, i.e. when a designer does not allocate some goods even though some agents still demand it, can be a necessary tool to attain the social optimum in both papers. In the mechanism design framework, this comes at the cost of ex-post optimality.

The two papers most closely related to our work are Huesmann [2017] and Akbarpour et al. [2020]. Akbarpour et al. [2020] consider a setting with a unit mass of agents that have different willingnesses to pay for quality. The planner, who has an exogenous preference to raise revenue, initially owns an exogenously given distribution of qualities to assign. The problem laid out in Akbarpour et al. [2020] does not include any case by case feasibility constraint on the allocation rule due to the large markets assumption. By contrast, we consider local markets where feasibility constraints have to hold for every possible type realization. Akbarpour et al. [2020] assume that the designer has an exogenously given desire to raise revenue and cannot grant positive transfers to the agents. In contrast, we assume that the designer faces a balanced budget constraint and can grant transfers of any sign. The budget constraint endogenously pins down the strength of the designer's preference to raise revenue in our work. Thus, we determine the strength of the designer's desire to raise revenue which is left unspecified in their model when the designer faces a balanced budget constraint. Further, depending on the strength of the designer's revenue motive, which is exogenously given in Akbarpour et al. [2020], random allocation may be optimal. In our framework, random allocation is never optimal.

Huesmann [2017] examines the problem of assigning a number of indivisible goods to a number of agents. There are two types of individuals: rich and poor people with high/low ex ante wealth levels. All agents have the same underlying utility function and have concave utility-for-money. Agents only differ in their wealth levels and an agent's wealth is private information. Thus, wealth levels have to be elicited through the mechanism whereas consumption benefits do not. There are several subtle, but important differences between the contribution of Huesmann [2017] and our work. Most importantly, Huesmann [2017] assumes that an agent's preferences are fully pinned down by the agent's wealth. Our framework accommodates this as a special case while maintaining a general formulation that is able to include all sorts of characteristics that may influence an agent's marginal utility of money. Furthermore, she assumes that the utility an agent receives when consuming the good is identical across agents. Our framework allows for the idea that a higher willingness to pay may be either a reflection of higher wealth or a stronger preference for the good. On the other hand, her framework incorporates the possibility of wealth effects, which our model does not. Following the framework of Dworczak et al. [2019], we focus on the idea that agents may have different marginal utilities of money but abstract from wealth effects for any given agent. Finally, Huesmann [2017] models a situation with a continuum of agents, whereas we model a finite number of agents to understand the local allocation problems we have in mind.

Most fundamentally, we derive an optimal mechanism for a vastly different environment than the one presented in Huesmann [2017]. In doing so, we uncover, among others, the following results not present in Huesmann [2017]: We connect our utilitarian optimal allocation rule to the virtual valuation as defined in Myerson [1981] and show, in addition, how stochasticity of the marginal utility of money enters the optimal allocation rule. Further, we highlight how our mechanism can be implemented by a simple auction with bidding subsidies and minimum bids. Finally, one should note the following important difference between our results and the results of Huesmann [2017]. Optimal mechanisms in Huesmann [2017] require transfers of a given agent to only depend on the reported type of this agent and on nothing else. More specifically, they cannot be directly linked to the ex post allocation decision. If payments are not directly linked to the ex post allocation decision, situations can arise where agents have to a pay a transfer without receiving the good. Such payment schemes are typically not observed in the practical applications we have in mind.

The notion that wealth affects the marginal utility of money is incorporated in some previous contributions from neighbouring fields in economics. Esteban and Ray [2006] study a lobbying framework where different lobby groups have different wealth levels and the costs of lobbying fall in wealth. Thus, wealthier groups will have stronger incentives to lobby and resources are diverted from sectors where they would be most productive to sectors that have strong lobbies because they are wealthy. Condorelli [2013] studies the optimal allocation of goods under generalized objectives of the planner - in particular, when the goal is not necessarily welfare or revenue maximization. Kang and Zheng [2019] characterize the set of interim-pareto-optimal mechanisms in a setting where one "good" and one "bad" are to be allocated. In their framework, agents have identical values for the "good" and the "bad", but have different (constant) marginal values of money.

Moreover, our contribution is linked to some integral strands of the public finance literature. The idea of assigning different agents heterogenous welfare weights based on their economic standing was already voiced by Diamond and Mirrlees [1971] and Atkinson and Stiglitz [1976]. This is complemented by recent contributions such as Saez and Stantcheva [2016]. Here, welfare weights are based partially on variables such as wealth that enter an agent's utility but are also substantially generalized to incorporate a society's wishes for redistribution, poverty alleviation, or equality of opportunity. Our paper is also related to the contributions which analyse the usefulness of quantity controls and, in particular, rationing, to achieve the social optimum in the presence of economic inequality. Weitzman [1977] analyses when a simple rationing scheme in which all consumers get the same amount of a good is preferable to a market price mechanism in a stylized framework. Not surprisingly, the advantage of the price based system is rising in the heterogeneity of taste for the product and falling in the level of inequality. Using an envelope theorem argument, Guesnerie and Roberts [1984] make the point that a small quota can improve social welfare when there is a wedge between the social marginal costs of a commodity and the consumer price for this commodity. Lee and Saez [2012] show that the desirability of a minimum wage is closely related to the type of rationing it induces on the labor market. They show that a necessary condition for a binding minimum wage to be welfare-improving is that the minimum wage must induce workers that provide the least surplus to society to lose their jobs first.

The idea of using the public provision of goods as a redistributive tool is the main contribution of Besley and Coate [1991] and Gahvari and Mattos [2007]. Besley and Coate [1991] study a market for an indivisible and rivalrous good such as healthcare, education, etc. that can be provided either by the state or the private sector. Individuals demand at most one unit of this good but care about the quality of this good. The government can choose the quality of this good and the provision of this good is financed via a lump-sum tax. A state with utilitarian objectives will then provide an intermediate quality of the good at no costs - this will induce high-income individuals to buy this good at higher quality from the private sector while low-income individuals will buy the state-offered good. This is a redistributive act in the face of lump-sum taxation, as poor individuals will benefit from this scheme while rich individuals will not. Notably, this scheme is associated with a deadweight loss due to the in-kind nature of the redistributive program. Gahvari and Mattos [2007] build on the framework of Besley and Coate [1991] and show that there is a way to avoid this deadweight loss when the state has the availability to provide cash transfers, conditional on consuming the publicly provided good.

3 Framework

We study a finite but arbitrary number of agents $i \in \{1, 2, .., N\}$ with unit demand for an indivisible good. Initially $m < N$ units of this good are owned by the mechanism designer and to be allocated among the agents. Our model specification largely follows Dworczak et al. [2019] and assumes that the agents' behavior is described by the utility function $u_i = v_i^K x_i^K + v_i^M x_i^M$, where v_i^K represents the valuation for the good and x_i^K is a binary variable that describes whether or not the agent has received the good. What sets this specification apart from most of the literature is that the marginal utility of money may vary across agents. More precisely, the utility derived from money consists of two parts: it equals the marginal utility of money of the agent, namely v_i^M, multiplied by the amount of money received or paid by the agent in the mechanism, namely x_i^M. Both v_i^K and v_i^M are assumed to be private information. We assume that the mechanism designer is utilitarian and wants to maximize the ex ante welfare given by

$$\sum_{i=1}^{N} \mathbb{E}[v_i^K x_i^K + v_i^M x_i^M] \tag{1}$$

subject to incentive compatibility, individual rationality and budget balance constraints. Everything else equal, moving money between the agents thus impacts social welfare. We denote the allocation rule by x_i and the transfer rule by t_i. In line with the standard definitions of the literature we say that a mechanism is (Bayesian) incentive compatible if and only if for all agents i and possible types (v_i^K, v_i^M)

$$\begin{aligned} &\mathbb{E}_{-i}[v_i^K x_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M) + v_i^M t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \\ \geq\ &\mathbb{E}_{-i}[v_i^K x_i(\hat{v}_i^K, \hat{v}_i^M, v_{-i}^K, v_{-i}^M) + v_i^M t_i(\hat{v}_i^K, \hat{v}_i^M, v_{-i}^K, v_{-i}^M)] \end{aligned} \tag{2}$$

holds for all other possible type reports $(\hat{v}_i^K, \hat{v}_i^M)$. We say that participation in a mechanism is individually rational if and only if for all agents and possible types (v_i^K, v_i^M)

$$\mathbb{E}_{-i}[v_i^K x_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M) + v_i^M t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \geq \underline{U}_i \tag{3}$$

where $\underline{U}_i$ denotes the utility attached to each agent's outside option. Because utility functions are linear in both components, it can be justified that the outside option is type independent. The utility of the outside option is the utility the agent receives when she does not own the good and owns her initial level of wealth. Thus, one can interpret the transfer as the money the agent receives from or pays to the mechanism. Together with the fact that v^K represents the utility gain achieved when receiving the good, the outside option then can be normalized to 0. In the following, we thus set $\underline{U}_i = 0$.

We require our mechanism to only be budget balanced ex ante and not ex post. We say that a mechanism is budget balanced ex ante if and only if

$$\sum_{i=1}^{N} \mathbb{E}[t_i(v_i^K, v_i^M, v_{-i}^K, v_{-i}^M)] \leq 0 \tag{4}$$

Due to Börgers and Norman [2009], this assumption is without loss of generality, as for every ex ante budget balanced mechanism it is possible to find an ex post budget balanced mechanism with the same allocation rule and interim transfers for all types of all agents. Note that we do not require transfers to be negative.

Before deriving the optimal mechanism, we establish some preliminary results following the approach of Dworczak et al. [2019]. A sufficient statistic for agent behavior is the rate of substitution between the good and money of an agent

$$r_i = \frac{v_i^K}{v_i^M} \tag{5}$$

Note that von Neumann-Morgenstern utility functions are only unique up to affine transformations. Therefore, an agent's rate of substitution will be sufficient to describe her behavior. While r is sufficient to describe the agents' behavior, the mechanism designer himself is still interested in the actual values of v_i^K and v_i^M. As derived in Dworczak et al. [2019], the rate of substitution r is informative about the other parameters. Note that:

$$\mathbb{E}[v_i^K x^K + v_i^M x^M] = \mathbb{E}_{r_i}[\underbrace{\mathbb{E}[v_i^M|r_i]}_{\lambda_i(r_i)}(r_i x^K + x^M)] \tag{6}$$

The mechanism designer can infer the expected valuation of money of an agent, given her rate of substitution r_i. We assume that for every agent i the rate of substitution r_i is independently and continuously distributed on an interval $[\underline{r}_i, \bar{r}_i]$. The cdf of r_i will be denoted by $G_i(r_i)$. Further, assume that v^K and v^M are distributed such that $\lambda_i(r_i) = \mathbb{E}[v_i^M|r_i]$ is weakly decreasing in r_i. This assumption reflects the idea that a high willingness to pay is most likely to be supported, ceteris paribus, by a relatively low expected valuation for money. Further note that the factor $\lambda_i(r_i)$ represents a Pareto weight as commonly used in public finance, as discussed by Chang et al. [2018], for example.

Given that the parameter r is sufficient for understanding agent behavior, it is natural to ask whether or not the designer should elicit both v^K *and* v^M or if it sufficient to restrict the mechanism to the rate of substitution r. One might argue that the additional information that is available when eliciting v^K and v^M should allow the designer to achieve a more efficient allocation of the good and distribution of money compared to a situation where only the rate of substitution r is revealed. This line of reasoning falls short as we recall a key observation: due to the fact that utility functions are invariant to affine transformations, the behavior of two agents with different values of v^K and v^M that yield the same rate of substitution r is indistinguishable from one another. Therefore, any attempt at treating these two agents differently can not be successful. This intuition is formalized in the following proposition due to Dworczak et al. [2019]

Proposition 1 (Dworczak et al. [2019], Theorem 8) *If a mechanism is feasible (respectively, optimal) in the two dimensional model, then there exists a payoff-equivalent mechanism eliciting only r_i that is feasible (respectively, optimal) in the one dimensional model with G_i being the distribution of r_i under the joint distribution F_i for v_i^K and v_i^M, where λ_i is given by:*

$$\lambda_i(r_i) = \mathbb{E}_i[v_i^M | r_i]$$

Conversely, if a mechanism is feasible (respectively, optimal) in the one dimensional model, there exists a joint distribution F_i for v_i^K, v_i^M such that this mechanism is feasible (respectively, optimal) in the two dimensional model where r_i is distributed according to G_i and $\lambda_i(r_i)$ is defined as above.

Proof. See Dworczak et al. [2019]. ■

In light of this result we restrict ourselves to mechanisms that elicit only the rate of substitution r_i. Due to the revelation principle, we are also free to restrict our attention to direct mechanisms subject to incentive compatibility constraints. As derived in Dworczak et al. [2019], characterizing incentive compatibility is straightforward and follows the familiar formulation of the literature. Let $X_i(r_i) = \mathbb{E}_{-i}[x_i(r_i, r_{-i})]$ be the expected allocation probability of agent i, given type report r_i, and let $T_i(r_i) = \mathbb{E}_{-i}[t_i(r_i, r_{-i})]$ be the expected transfer of agent i, given type report r_i. Incentive compatibility is characterized by the following proposition:

Proposition 2 (Incentive Compatibility) *A mechanism $\{x_i(r_i, r_{-i}), t_i(r_i, t_{-i})\}_{i=1}^N$ is incentive compatible if and only if*

1. *$X_i(r_i)$ is non-decreasing in r_i (Monotonicity)*
2. *$r_i X_i(r_i) + T_i(r_i) = U_i(\underline{r}_i) + \int_{\underline{r}_i}^{r_i} X_i(s) ds$ (Integrability)*

Proof. The result follows directly from rescaling the agents' utility functions and applying the standard results from the literature. ■

By integrability, the expected transfer of an agent is given by

$$\mathbb{E}[T_i(r)] = U_i(\underline{r}_i) - \int_{\underline{r}_i}^{\bar{r}_i} X_i(r) \underbrace{J_i(r)}_{r - \frac{1-G_i(r)}{g_i(r)}} dG_i(r) \tag{7}$$

where $J_i(r)$ denotes the virtual valuation of an agent as defined in Myerson [1981]. Note that the integrability condition implies that participation of an agent in the mechanism is individually rational if and only if it is individually rational for the lowest type of any agent.

4 Optimal Mechanisms

4.1 Derivation of the optimal mechanism

Following Dworczak et al. [2019], we can use the integrability condition to rewrite the ex ante utilitarian welfare as follows:

$$\sum_i \mathbb{E}[\lambda_i(r_i)(rX_i(r_i) + T_i(r_i))] = \sum_i \left(\mathbb{E}\left[\lambda_i(r_i) \left(U_i(\underline{r}_i) + \int_{\underline{r}_i}^{r} X_i(s)ds \right) \right] \right) \tag{8}$$

$$= \sum_i \left(\underbrace{\mathbb{E}[v_i^M]}_{\Lambda_i} U_i(\underline{r}_i) + \int_{\underline{r}_i}^{\bar{r}_i} \int_s^{\bar{r}_i} X_i(s)\lambda_i(r_i)dG_i(r_i)ds \right) \tag{9}$$

$$= \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int_{\underline{r}_i}^{\bar{r}_i} X_i(s) \underbrace{\frac{\int_s^{\bar{r}_i} \lambda_i(r_i)dG_i(r_i)}{g_i(s)}}_{\Pi_i(s)} dG_i(s) \right) \tag{10}$$

where the expression is simplified using a change of the order of integration. We remark that we defined

$$\Pi_i(s) := \frac{\int_s^{\bar{r}_i} \lambda_i(r_i)dG_i(r_i)}{g_i(s)} \tag{11}$$

Note that for $\lambda_i(r_i) = 1$, we have $\Pi_i(r_i) = \frac{1-G_i(r_i)}{g_i(r_i)}$, the standard inverse hazard rate formulation. Therefore, it seems instructive to think about the function $\Pi_i(r_i)$ as an inequality adjusted inverse hazard rate.

To define our optimization problem, we plug in the expected transfers into the welfare objective and into the budget constraint. Then, the optimization problem boils down to choosing the optimal allocation rule and the optimal utility levels for the lowest type of each agent. Thus, our problem can be stated as:

$$\max_{\{x_i(r_i,r_{-i}),U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i,r_{-i}) dG(r_i,r_{-i})\right)$$

$$\text{s.t.} \sum_i \left(U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i,r_{-i}) dG(r_i,r_{-i})\right) \leq 0 \qquad \text{(Budget)}$$

$$0 \leq x_i(r_i,r_{-i}) \leq 1 \qquad \text{(Prob)}$$

$$\sum_i x_i(r_i,r_{-i}) \leq m \qquad \text{(Feas)}$$

$$X_i(r_i) \text{ non-decreasing} \qquad \text{(Mono)}$$

$$U_i(\underline{r}_i) \geq 0 \qquad \text{(IR)}$$

We define $\Lambda^* = \max\{\Lambda_i\}$. The key statistic for our allocation rule, which we call the *inequality adjusted valuation*, is:

Definition 1 (Inequality adjusted valuation) *We define the expression*

$$\varphi_i(r) := \Pi_i(r) + \Lambda^* J_i(r) \qquad (12)$$

to be the inequality adjusted valuation of agent i.

It is instructive to consider this expression in the standard case, i.e. when $\lambda_i(r) = 1$. We can easily verify that in this situation the inequality adjusted valuation $\varphi_i(r) = \frac{1-G_i(r)}{g_i(r)} + r - \frac{1-G_i(r)}{g_i(r)} = r$ equals the valuation as the inequality adjusted inverse hazard rate equals the inverse hazard rate. For the purposes of our analysis, we will assume the following about the agents' inequality adjusted valuations:

Assumption 1 *The inequality adjusted valuation $\varphi_i(r) = \Pi_i(r) + \Lambda^* J_i(r)$ is non decreasing for all agents.*

Assumption 1 will simplify the derivation of the optimal allocation rule considerably. If we find this assumption to be violated, we would have to consider an ironing procedure as in Myerson [1981]. We are now ready to state the core result of our paper.

Proposition 3 (Optimal Mechanism) *The optimal mechanism assigns the good to the m agents with the highest inequality adjusted valuations $\varphi_i(r_i) = \Pi_i(r_i) + \Lambda^* J_i(r_i)$, given that they are positive.*

Proof. We provide the proof in the appendix. ■

We outline the core idea and intuition of the proof while the detailed proof can be found in the appendix. Notice that there is a single budget constraint that forces the budget to be balanced in expectation and not on a case by case basis. This constraint has to be binding in the optimum as it is always possible to allocate any left-over funds ex ante to an agent. The key step of the proof is to find the Lagrange multiplier associated with this constraint, which represents the marginal increase in welfare after a marginal increase in the available budget. Given any Lagrange parameter μ, note that the Lagrangian of the maximization problem (neglecting the other constraints for now) reads

$$\max_{\{x_i(r_i,r_{-i}),U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left((\Lambda_i - \mu)U_i(\underline{r}_i) + \int (\Pi_i(r_i) + \mu J_i(r_i))x_i(r_i, r_{-i})dG(r_i, r_{-i}) \right) \tag{13}$$

The solution to this maximization problem is a combination of an allocation rule and a Lagrange parameter. For any given μ, the optimal allocation rule is to assign the good to the m agents with the highest $\Pi_i(r_i) + \mu J_i(r_i)$. This allocation rule induces a certain Lagrange multiplier. In the optimal solution, the Lagrange multiplier will be a fixed point of this mapping. We will argue that our allocation rule corresponds to a Lagrange multiplier value equal to Λ^*, constituting a fixed point. This is proven by showing that a marginal increase in budget leads to a marginal increase in welfare of Λ^*, given our allocation rule.

It is instructive to separately consider situations in which the virtual valuation $J_i(r_i)$ of an agent is positive and when it is negative. Note that the inequality adjusted valuation of an agent can only be negative when the virtual valuation of an agent is negative. When $J_i(r_i)$ is positive, allocating the good to the agent will raise the total revenue that can be extracted from the agent. When $J_i(r_i)$ is negative, allocating the good to agent i decreases the revenue that can be extracted from the agent.

We start off with the case of negative values of $J_i(r_i)$. We note that there is a key trade-off within the optimal mechanism: a budgetary surplus can either be distributed to an agent ex ante through the utility level of the lowest type $U_i(\underline{r}_i)$ or it can be used to subsidize the allocation of the good yielding $\Pi_i(r_i)$ at the cost of $J_i(r_i)$. If any agent receives ex ante transfers, this money must be allocated to the agent with the highest value of Λ_i, namely Λ^*. Otherwise, reallocating money ex ante from an agent i with a lower value of Λ_i to an agent j with a higher value of Λ_j would yield an immediate increase in welfare. Thus, the increase in welfare from using a marginal unit of money to fund ex ante transfers is Λ^*. This implies that the shadow value of an additional unit of budget is at least Λ^*. When subsidizing the allocation of the good to agent i, the additional social welfare obtained per unit of money spent is $\Pi_i(r_i)/(-J_i(r_i))$. We conclude that it is optimal to use an additional unit of budget to fund ex ante transfers rather than using it to subsidize allocation if and only if:

$$\Lambda^* > \frac{\Pi_i(r_i)}{-J_i(r_i)} \tag{14}$$

$$\iff 0 > \Pi_i(r_i) + \Lambda^* J_i(r_i) = \varphi_i(r_i) \tag{15}$$

This argument yields our result that the subsidized allocation of a good to an agent is only desirable if her *inequality adjusted valuation* is positive. Otherwise, the budget balance condition implies that the opportunity costs of allocating this additional unit in terms of the money that can be spent ex ante outweighs its benefits.

Now, we consider situations where the virtual valuation $J_i(r_i)$ is positive. Allocating the good to agent i in any given situation directly yields $\Pi_i(r_i)$ units of welfare from the incentive compatible allocation of the good. Further, it raises $J_i(r_i)$ units of revenue that can be distributed ex ante or used to subsidize allocation when $J_i(r_i)$ is negative. When the marginal value of an additional unit of budget is Λ^*, the inequality adjusted valuation summarizes the total benefit minus opportunity cost of allocating the good to a certain individual.

As a second step, we now show that this allocation rule corresponds to a Lagrange parameter exactly equal to Λ^* by making a case distinction. Suppose firstly that there are less agents with a positive inequality adjusted valuation than goods to be allocated for a given type profile. Recall that agents with negative inequality adjusted valuations necessarily have negative virtual valuations. By the above arguments, any additional unit of budget would thus be optimally used to fund an ex ante transfer to the agent with the highest marginal utility of money, yielding Λ^* of additional social welfare. Suppose alternatively that there are more agents with a positive inequality adjusted valuation than goods to be allocated, in which case the feasibility constraint binds. Consider a situation in which an agent has a positive inequality adjusted valuation, but she is currently not receiving the good. Recall that our allocation rule demands that the good is allocated to the agents with the m highest inequality adjusted valuations $\varphi_i(r_i) = \Pi_i(r_i) + \Lambda^* J_i(r_i)$ given that $\varphi_i(r_i) \geq 0$. This implies that in any situation where an agent i has a positive inequality adjusted valuation but currently does not receive the good, there are at least m other agents with higher inequality adjusted valuations. If the designer wants to assign the good to agent i, he would necessarily have to take the good away from another agent j who is receiving the good, as all goods are already assigned, but this change in allocation is only desirable if

$$\varphi_i(r_i) \geq \varphi_j(r_j) \tag{16}$$

which can not hold true as agent i's inequality adjusted valuation is lower than that of agent j. Overall, this implies that there are no remaining situations in which subsidizing the allocation of the good is desirable and therefore a marginal unit of extra budget would have to be assigned to the agent with the largest value of Λ_i. This implies that the Lagrange multiplier in our optimization problem is equal to Λ^*.

Recall that the above results rely on assumption 1. We now provide a condition under which assumption 1 holds.

Remark 1 *Assume that the marginal utility of money is non-stochastic for all agents.*

1. *Assumption 1 holds if the environment is regular as defined in Myerson [1981], i.e. when the virtual valuations are non-decreasing for all agents.*

2. *Assumption 1 holds if and only if the cdfs of v_i^K, call these F_i, and the corresponding pdfs f_i satisfy the following for all agents:*

$$\frac{\partial}{\partial v^K}\left[\frac{1-F_i(v^k)}{f_i(v^k)}\right] \leq 1 + \frac{\Lambda_i}{\Lambda^* - \Lambda_i}$$

Naturally, it is of interest to investigate when our allocation rule simplifies to the well known ex post efficient allocation rule which allocates the good to the m agents with the highest valuations r_i, given that they are positive. As pointed out before, our model simplifies to the standard framework when $\lambda_i(r) = 1$ for all agents and thus yields the ex post efficient allocation rule in that case. More interestingly, our allocation rule also simplifies to the standard allocation in the well studied i.i.d. environment as is shown in the next corollary:

Corollary 1 *Suppose that either of the following conditions is met:*

1. $\lambda_i(r_i) = 1$ *for all* i *and* r_i

2. *The pair* (v_i^K, v_i^M) *is i.i.d. for all agents* i *and* $\varphi_i(r_i)$ *is strictly increasing*

Then, the optimal assignment rule is equal to the standard assignment rule, i.e. the good is assigned to the m *agents with the highest valuations* r_i*, given that they are positive.*

Proof. The proof of 1. is immediate after plugging in $\lambda_i(r_i)$. To argue why condition 2 implies the standard allocation rule, we note that when (v_i^K, v_i^M) is i.i.d. for all agents i, we have $\varphi_i(r) = \varphi_j(r)$ for all i, j. By assumption, $\varphi_i(r_i)$ is strictly increasing. This implies that $\varphi(r_i)$ is a strictly increasing transformation of r_i and therefore $r_i \geq r_j$ if and only if $\varphi(r_i) \geq \varphi(r_j)$. Further, due to the i.i.d assumption on (v_i^K, v_i^M), it holds that $\Lambda^* = \Lambda_i$ for all i. Mathematically, this implies that $\varphi_i(r_i) \geq 0$ if and only if $r_i \geq 0$. ■

Case 1 of the corollary simply highlights that the standard quasi linear model is nested as a special case of the model. Case 2 is more intricate. When (v_i^K, v_i^M) is i.i.d. among all agents, the planner does not ex ante consider any agent more or less likely to be rich or poor. This reflects the idea that even though he does not know the exact valuations of money v^M, he considers each agent to be in the same situation ex ante. We recall from the discussion of proposition 3 that redistribution in the optimal mechanism is facilitated through ex ante redistribution to the agent with the highest expected marginal utility of money. However, if every agent is considered equally rich or poor ex ante, then the designer finds himself unwilling and unable to engage in any kind of redistribution and therefore follows the standard allocation rule.

To build further intuition for our results, we now present a decomposition of the inequality adjusted valuation into three components that highlight the trade-offs that the designer faces. Consider the following decomposition of the inequality adjusted valuation:

$$\varphi_i(r_i) = \Pi_i(r_i) + \Lambda^* J_i(r_i) \tag{17}$$

$$= \frac{\int_{r_i}^{\bar{r}_i} \lambda_i(s) dG_i(s)}{g_i(r_i)} + \Lambda_i r_i + (\Lambda^* - \Lambda_i) r_i - \Lambda^* \frac{1 - G_i(r_i)}{g_i(r_i)} \tag{18}$$

$$= \underbrace{\Lambda_i r_i}_{\text{Efficient allocation}} + \underbrace{(\Lambda^* - \Lambda_i) J_i(r_i)}_{\text{Ex ante transfers}} + \underbrace{\frac{\int_{r_i}^{\bar{r}_i} \lambda_i(s) - \Lambda_i dG_i(s)}{g_i(r_i)}}_{\text{Ex interim uncertainty}} \tag{19}$$

The first component highlights the desire to allocate the goods efficiently. Agents with a larger rate of substitution r_i should receive the good. The second component highlights that raising revenue through the virtual valuation $J_i(r_i)$ is beneficial for the designer. Any revenue can be redistributed ex ante to the poorest agent increasing welfare by Λ^* at the cost of the ex ante marginal utility of money of the agent from which the revenue was generated. The third component captures the fact that the marginal utility of money of an agent is not deterministic. Whenever the expected valuation of money given the rate of substitution, namely $\lambda_i(r_i)$, is larger than the ex ante expected valuation of money Λ_i, the inequality adjusted valuation increases. However, incentive compatibility demands that this comparison does not enter the inequality adjusted valuation solely on a case by case basis, but through a more complicated expression.

4.2 Rationing

We note that rationing may occur if more than $N - m$ agents have an inequality adjusted valuation that is negative. There is a particularly interesting connection between the ex ante expected valuation of money of each agent, namely Λ_i, and whether or not an agent may be subject to rationing. We say that an agent is subject to rationing when some units of the good are not allocated but this agent still has demand for the good.

Proposition 4 (Rationing) *Let i^* denote the index of the agent with $\Lambda_{i^*} = \Lambda^*$. Then*

1. *Agent i^* is never subject to rationing*
2. *All agents $i \neq i^*$ may be subject to rationing*

Proof. Consider the inequality adjusted valuation $\varphi_i(r_i)$ at the lowest possible realization $\underline{r}_i$:

$$\varphi_i(\underline{r}_i) = \Pi_i(\underline{r}_i) + \Lambda^* J_i(\underline{r}_i) \tag{20}$$

$$= \frac{\int_{\underline{r}_i}^{\bar{r}_i} \lambda(s) dG_i(s)}{g_i(\underline{r}_i)} + \Lambda^* \left(\underline{r}_i - \frac{1 - G_i(\underline{r}_i)}{g_i(\underline{r}_i)} \right) \tag{21}$$

$$= \Lambda^* \underline{r}_i + \frac{\Lambda_i - \Lambda^*}{g_i(\underline{r}_i)} \tag{22}$$

We note that for agent i^*, this expression is weakly positive if and only if $\underline{r}_i \geq 0$. Therefore, the only reason to not allocate the good to agent i^* would be that she has a negative valuation for the good. For the other agents, the expression is weakly positive if and only if

$$\Lambda^* \underline{r}_i + \frac{\Lambda_i - \Lambda^*}{g_i(\underline{r}_i)} \geq 0 \tag{23}$$

which will generally subject them to rationing unless $\underline{r}_i$ is sufficiently large. ■

Corollary 2 *Let $\underline{r}_i = 0$ for all i. Then, all agents except agent i^* are subject to rationing*

Proof. Follows directly from the proof of proposition 4 ■

Rationing is a key source of allocative inefficiency in our model and hence plays an important role for social welfare. Thus, it is instructive to understand how wealth inequality affects the incidence of rationing in our framework. To understand the quantitative magnitude of rationing in a given setting, we consider the probability that rationing occurs, i.e. the fraction of possible type realizations for which rationing would occur. We assume that the valuations for money are fixed for any agent, but can vary across agents. This allows us to obtain the following results:

Proposition 5 (Inequality and the probability of rationing) *Assume that the marginal utility of money of all agents is non stochastic. Then, it holds that:*

1. $\frac{\partial Pr(\varphi_i(r_i)<0)}{\partial \Lambda^*} \geq 0$ *holds for all agents* $i \neq i^*$*. Thus, when* Λ^* *increases, the probability with which rationing occurs weakly increases.*

2. *When* $\frac{\partial Pr(\varphi_i(r_i)<0)}{\partial \Lambda_i} < 0$*, a decrease of* $\Lambda_i \neq \Lambda^*$ *will imply an increase of the probability with which rationing will occur. Note that* $\frac{\partial Pr(\varphi_i(r_i)<0)}{\partial \Lambda_i} < 0$ *holds true if the inverse hazard rates for the distribution of* v^K *are monotonically decreasing, i.e. when* $\frac{\partial}{\partial \nu^K}\left[\frac{1-F_i(v^K)}{f_i(v^K}\right] \leq 0$*. Any such distribution is regular in the sense of Myerson [1981].*

Proof. See appendix. ■

Note that a higher probability of rationing reflects a greater extent of allocative inefficiency. Modelling the effects of an increase in wealth inequality offers several degrees of freedom. In general, the effect of such a change on the probability with which rationing occurs depends on how an increase in inequality is modelled. Proposition 5 sheds some light on these issues and allows us to make the following definitive statements: When the wealth of the poorest members of society decreases, which is reflected by an increase in Λ^*, the probability with which rationing occurs increases, *ceteris paribus*. Result 2 yields insights into the effects of an increase in wealth inequality along the lines of the development of real wages of men in the USA over the years 1990-2010. In these years, real wages of men have stagnated at the 10^{th} percentile and 50^{th} percentile, while they have gone up by around 1% (annualized) at the 90^{th} percentile - see Donovan and Bradley [2019]. Thus, over this period, the real wages of the 90^{th} percentile have risen by 22%. Within our model, this can be viewed as a decrease of Λ_i for the wealthier members of the distribution, while all other Λ_i's are left unchanged. Result 2 shows that the probability with which rationing occurs will increase as a result of these developments.

Quantifying the impact of a mean-preserving spread of the Λ's on the probability of rationing is challenging in general settings. Such a mean-preserving spread would imply an increase of Λ_j for at least one agent j and a decrease of Λ_i for at least another agent i with $\Lambda_i < \Lambda_j$. Result 2 highlights that such an exogenous change will yield opposing effects on the probability with which rationing occurs under the stated conditions. Determining which effect dominates algebraically requires specific distributional assumptions. Consider, for instance, that there are two goods to be allocated and v^K is drawn from a uniform distribution on $[0,1]$ for all agents. Consider a marginal increase of Λ_j and a marginal decrease of Λ_i for agents where $\Lambda_i < \Lambda_j$ by the same magnitude. Such a mean-preserving spread will reduce the probability of rationing in these circumstances. This is because the negative effect of the rise in Λ_j on the probability of rationing dominates the effect running through the change in Λ_i. This result shows that, in general, the effect of an increase in wealth inequality on the probability of rationing depends on the modelling choices.

4.3 Connection to auction theory

In this section we describe how our mechanism can be implemented as an auction. First, we give a brief description of the auction rules: consider a sealed bid auction in which every bidder submits a bid r_i. However, the winner of the auction is not necessarily the bidder with the highest bid. To implement our allocation rule, the auctioneer transforms these bids r_i by applying the transformation $\varphi_i(r_i)$. The winners of the auction are the m bidders that have submitted bids r_i, which have resulted in the highest values of $\varphi_i(r_i)$, given that $\varphi_i(r_i) \geq 0$. Payments of the auction are given by the integrability condition, i.e. the payment rule of the auction satisfies

$$r_i X_i(r_i) + T_i(r_i) = U_i(\underline{r}_i) + \int_{\underline{r}_i}^{r_i} X_i(s)ds \tag{24}$$

We note that by the revenue equivalence principle it does not matter whether we consider an auction format where bidders only pay when they win or whether they always pay, given a particular bid r_i. The allocation rule determines the expected transfers that are associated with a particular bid fully. Incentive compatibility of the mechanism implies that in equilibrium, every bidder submits a bid equal to her willingness to pay. However, the bidders that submit the highest willingness to pay do not necessarily win the auction.

We now highlight how this particular set of auction rules includes bidding subsidies. As described above agent i receives the good over agent j if and only if:

$$\varphi_i(r_i) \geq \varphi_i(r_j) \tag{25}$$

$$\iff r_i \geq r_j + \frac{1}{\Lambda^*}\left[\frac{1-G_i(r_i)}{g_i(r_i)}(\Lambda^* - \mathbb{E}[\lambda_i(s)|r_i \leq s \leq \bar{r}_i]) - \frac{1-G_j(r_j)}{g_j(r_j)}(\Lambda^* - \mathbb{E}[\lambda_j(s)|r_j \leq s \leq \bar{r}_j])\right] \tag{26}$$

Agent i thus receives a bidding subsidy, enabling her to win the auction despite her bid r_i not being higher than agent j's bid r_j in the following circumstances: Ceteris paribus agent i receives a bidding subsidy when competing with agent j when (i) her expected marginal utility of money, given that her willingness to pay for the good is at least r_i, is high and/or (ii) agent j's expected marginal utility of money, given that her willingness to pay for the good is at least r_j, is low. This can be interpreted as agent i being poor, while agent j is comparatively rich. The designer recognizes this situation and introduces the bidding subsidy. This bidding subsidy is optimal because it allows the poorer agent to compete with the richer agent, which offers two advantages: Firstly, it increases the revenue that the mechanism will elicit. Secondly, it ensures that the good is more likely to be allocated to the agent who derives the higher utility from consuming the good as such.

Other tools that are commonly used in auction theory are minimum bids or reserve prices. For now, consider the single unit case. Here, it is possible to find a general minimum bid that is not specific to the individual bidders. This is desirable for practical applications. We recall that the designer will not assign the good to an agent with a negative inequality adjusted valuation. Due to the bidding subsidy that is implied by the inequality adjusted valuation, any agent that is to win the auction must have a positive inequality adjusted valuation. Given that the poorest agent's inequality adjusted valuation is always weakly positive, the good must always be allocated to somebody. The only minimum bid compatible with this feature of the allocation rule is a minimum bid of 0.

Now consider the case with $m > 1$ identical units of the good. We denote by $\underline{r}_i$ the bids that solve $\varphi_i(\underline{r}_i) = 0$. For the multi-unit case, the auction that implements our mechanism sets a minimum bid of 0 for one unit of the good and bidder-specific minimum bids for the rest of the goods. One unit of the good is always allocated in our mechanism - thus, the minimum bid for this good must be zero. Moreover, consider the situation in which one unit of the good was already allocated to the poorest agent and two agents i, j have submitted bids r_i, r_j which are higher than the other agents bids (including bidding subsidies) and such that $r_i < \underline{r}_i < r_j < \underline{r}_j$. We note that a non bidder specific minimum bid may not exceed $\underline{r}_i$

to ensure the implementation of our allocation rule. If the minimum bid would exceed $\underline{r}_i$, then agent i would never receive the good in this auction when r_i is in between $\underline{r}_i$ and the minimum bid. But this would not correspond to our allocation rule in the above situation. By a similar argument, the minimum bid cannot be below $\underline{r}_i$, as it would imply the possible allocation of the good to an agent with a negative inequality adjusted valuation. Now suppose that the minimum bid exactly equals $\underline{r}_i$. Note that $\varphi_j(r_j)$ is continuous as a result of our assumptions which implies that for $r_j < \underline{r}_j$ sufficiently close to $\underline{r}_j$ we have $\varphi_i(r_i) < \varphi_j(r_j)$ which implies that agent j is awarded the good in the auction. However, this is not in in accordance with the allocation rule that was to be implemented, as $r_j < \underline{r}_j$. Therefore, the general m good scenario requires bidder specific minimum bids for $m-1$ goods where each agent i faces a minimum bid of $\underline{r}_i$.

4.4 A framework without redistributive concerns

At this stage, one may wonder to what extent our results are driven by the redistributive motive the designer possesses as a direct implication of the budget constraint. To investigate this, we now impose additional restrictions on the transfer rule which make redistribution from one agent to another impossible. This exercise places the spotlight on the pure allocative inefficiencies caused by allocating objects solely based on the stated willingness to pay.

For this section, we assume that the expected transfer of any agent must be non-negative, i.e. that:

$$\mathbb{E}[T_i(r)] \leq 0 \quad \forall i \tag{27}$$

These constraints eliminate the planner's ability to provide ex ante (positive) transfers to any agent. As a conclusion of our previous results, the above constraint is equivalent to:

$$\mathbb{E}[T_i(r)] = U_i(\underline{r}_i) - \int_{\underline{r}_i}^{\bar{r}_i} X_i(r) J_i(r) dG_i(r) \leq 0 \tag{28}$$

In such a setting, the planner thus solves the following problem:

$$\max_{\{x_i(r_i,r_{-i}),U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right)$$

$$\text{s.t. } U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \leq 0 \ \forall i \qquad \text{(Transfers)}$$

$$0 \leq x_i(r_i, r_{-i}) \leq 1 \qquad \text{(Prob)}$$

$$\sum_i x_i(r_i, r_{-i}) \leq m \qquad \text{(Feas)}$$

$$X_i(r_i) \text{ non-decreasing} \qquad \text{(Mono)}$$

$$U_i(\underline{r}_i) \geq 0 \qquad \text{(IR)}$$

In the optimal solution to the above problem, all transfer constraints must bind. Suppose, for a contradiction, that the transfer constraint is slack for some agent i in the optimal solution. Then, the designer could increase $U_i(\underline{r}_i)$ according to this constraint. This change would not violate any other constraint and would raise social welfare. Thus, our starting mechanism could not have been optimal.

Remark 2 *Consider the functions $\gamma_i(r_i)$ which are defined as follows*

$$\gamma_i(r_i) = \Pi_i(r_i) + \Lambda_i J_i(r_i)$$

Given this definition, note that $\gamma_i(r_i) \geq 0 \ \forall r_i$ holds for all agents.

Proof. See appendix. ■

Having established this, we state the optimal mechanism.

Proposition 6 *When ex ante transfers must be weakly negative, the optimal mechanism assigns the good to the m agents with the highest $\gamma_i(r_i)$. There is no rationing.*

Proof. See appendix. ■

Note that the expression $\gamma_i(r_i)$ closely resembles the inequality adjusted valuation $\varphi_i(r_i) = \Pi_i(r_i) + \Lambda^* J_i(r_i)$, which was decomposed as follows:

$$\varphi_i(r_i) = \underbrace{\Lambda_i r_i}_{\text{Efficient allocation}} + \underbrace{(\Lambda^* - \Lambda_i) J_i(r_i)}_{\text{Ex ante transfers}} + \underbrace{\frac{\int_{r_i}^{\bar{r}_i} \lambda_i(s) - \Lambda_i dG_i(s)}{g_i(r_i)}}_{\text{Ex interim uncertainty}} \tag{29}$$

Having noted this, consider the following decomposition of $\gamma_i(r)$:

$$\gamma_i(r_i) = \Pi_i(r_i) + \Lambda_i J_i(r_i) \tag{30}$$

$$= \frac{\int_{r_i}^{\bar{r}_i} \lambda_i(s) dG_i(s)}{g_i(r_i)} + \Lambda_i r_i - \Lambda_i \frac{1 - G_i(r_i)}{g_i(r_i)} \tag{31}$$

$$= \underbrace{\Lambda_i r_i}_{\text{Efficient allocation}} + \underbrace{\frac{\int_{r_i}^{\bar{r}_i} \lambda_i(s) - \Lambda_i dG_i(s)}{g_i(r_i)}}_{\text{Ex interim uncertainty}} \tag{32}$$

The only difference between $\varphi_i(r_i)$ and $\gamma_i(r_i)$ is that the second term of $\varphi_i(r_i)$ is absent in $\gamma_i(r_i)$. This holds because the second term of $\varphi_i(r_i)$ captured the designer's desire to raise revenue for redistribution, which is impossible in this setting and will thus not be present in the optimal solution.

Proposition 6, together with the decomposition of $\gamma_i(r_i)$, reinforces to the core result of our paper. In terms of utilitarian social welfare maximization, it is simply not optimal to allocate the goods to the agents with the highest willingnesses to pay. A high willingness to pay does not necessarily imply a high consumption utility, in particular for agents with a high expected marginal utility of money. In general, it is utilitarian optimal to allocate the good to the agent with the highest consumption utility as such. When all agents have deterministic marginal utilities of money, $\gamma_i(r_i)$ is exactly equal to v_i^K. When the marginal utility of money is stochastic, the inferences regarding the consumption utility as such need to be carefully weighted based on the likely distribution of the marginal utility of money.

As we have shown, redistributive concerns will influence the utilitarian optimal allocation rule when redistribution is possible. However, redistributive concerns are not the sole reason for the difference in between the utilitarian optimal allocation rule and the ex post efficient allocation rule in our setting. The main cause of this difference is the fact that the mapping from consumption utility into willingness to pay depends on the marginal utility of money of an agent.

Not surprisingly, there will be no rationing in this setting. Formally, this follows from remark 2, which states that $\gamma_i(r_i)$ will never be negative. This implies that the feasibility constraint is never slack, which means that all goods will be allocated. This is intuitive, given that rationing was a byproduct of the revenue raising motive of the designer that was present in the original setting, which was now shut down.

5 Numerical illustrations and policy implications

To further emphasize the key points of our paper, we provide some numerical illustrations. The first subsection is devoted to an example where the agents have heterogeneous, but non-stochastic marginal utilities of money. In this example, we firstly visualize the difference between the ex post efficient allocation rule and the utilitarian optimal allocation rule. We also plot the expected transfer of agents in the two mechanisms. Afterwards, we move on to show how the expected transfer rules respond to mean preserving spreads of the marginal utilities of money. We further investigate how such a mean preserving spread affects the probability with which rationing occurs to re-emphasize the link between inequality and allocative efficiency present in our model. Afterwards, we consider an example where the marginal utility of money is stochastic to provide further insights.

5.1 Deterministic marginal utility of money

In this subsection, we consider the following simplified example. There are $N = 5$ agents and $m = 2$ goods to be allocated. The valuation for money v_i^M of every agent is non-stochastic such that $v_i^M = \Lambda_i$ holds for all agents. The valuation for the good v_i^K is drawn from a uniform distribution with support $[0, 1]$ for all agents. In the numerical solution, we discretize this example by assuming that v_i^K is drawn from a finite grid. Assume that agents are ordered such that $\Lambda_1 \leq \Lambda_2 \leq \Lambda_3 \leq \Lambda_4 \leq \Lambda_5$. Assume further that that marginal utilities of money are spread evenly such that:

$$\Lambda_i = \mu + 0.5(i - 3)\sigma \;\; \forall i = \{1, .., 5\}$$

Increases in the parameter σ thus constitute a mean preserving spread of the marginal utilities of money. Assume, for the first part of this example, that $\mu = 0.01$ and $\sigma = 0.005$. For these parameters, we now plot the interim allocation probabilities both for the ex post efficient allocation rule and the utilitarian optimal allocation rule.

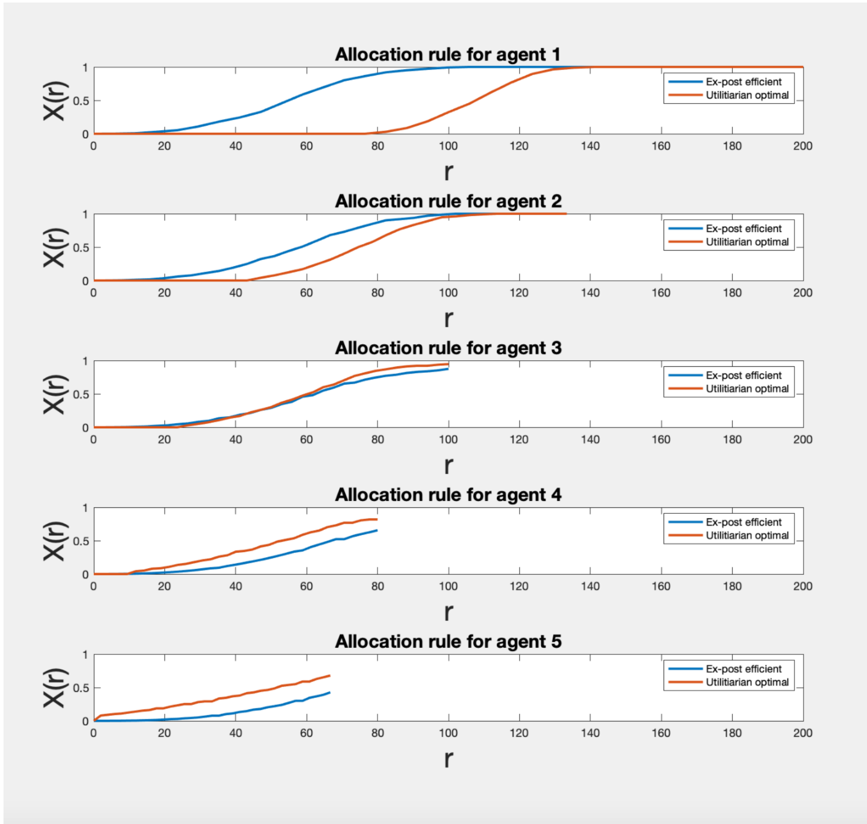

Figure 1: Interim allocation probabilities

Note the following key features: Agents 1,2,3, and 4 may be subject to rationing. Agents 1 and 2 with the relatively low marginal utilities of money have uniformly lower interim probabilities of receiving the good in the utilitarian optimal rule as long as these probabilities are interior. The converse holds true for Agent 5, who has the highest marginal utility of money. Agents 3 and 4 both have relatively similar interim allocation probabilities in the two different rules for low values of r. However, as r increases, these agents receive comparatively higher interim allocation probabilities in the utilitarian optimal rule. The intuition for these results mirrors the components of the inequality-adjusted valuations. Firstly, comparatively low values of Λ_i as such will reduce the interim allocation probability for a given r_i. When Λ_i is lower, the distribution of implied v_i^K's, conditional on r_i, gets shifted down. Allocating the good to an agent for a given r_i will thus have a lower positive effect on social welfare when Λ_i is lower, *ceteris paribus*. This working channel pushes the interim allocation probability of the agents $i = 1, 2$ in the utilitarian optimal allocation rule below the one in the ex post efficient allocation rule. Secondly, for the agents $i = 1, 2, 3, 4$ with $\Lambda_i < \Lambda^*$, negative virtual valuations make allocation less beneficial and vice versa. The virtual valuation is rising in

r_i and crosses 0 in the interior of the support of r_i. Thus, higher values of r_i will imply a higher interim probability of allocation via this channel. This working channel implies that the interim allocation probability for agents 1 and 2 approaches 1 for high values of r in the utilitarian optimal rule. For agents 3 and 4, this effect drives the spread in the two allocation rules that is created for high values of r. Agent 5 has a uniformly higher interim allocation probability under the utilitarian optimal allocation rule. This is because the second working channel is shut down for this agent and the first working channel implies a uniformly positive effect on the interim allocation probability.

Now we visualize the expected transfers of the agents under the two allocation rules. As is to be expected, the transfers of agents 1-4 are lower in the utilitarian optimal mechanism than under the ex-post efficient rule and vice versa for agent 5.

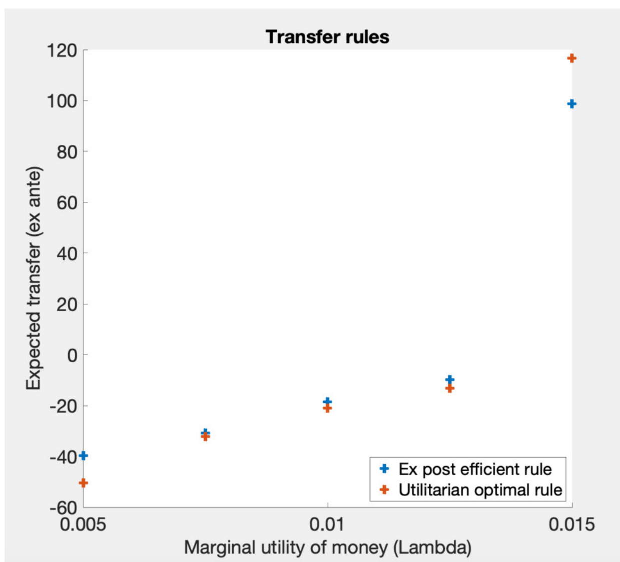

Figure 2: Ex ante transfers

Next, we visualize the expected transfers for all agents and the probability with which rationing will occur under different degrees of inequality in the society. To model different levels of inequality in the society, we will vary the parameter σ. In particular, we will consider possible realizations of σ in the interval $[0.001, 0.009]$, where a rise in σ represents a mean-preserving spread of the Λ_i's. The effects of such a spread on the transfer rules and the probability with which rationing occurs are graphed now:

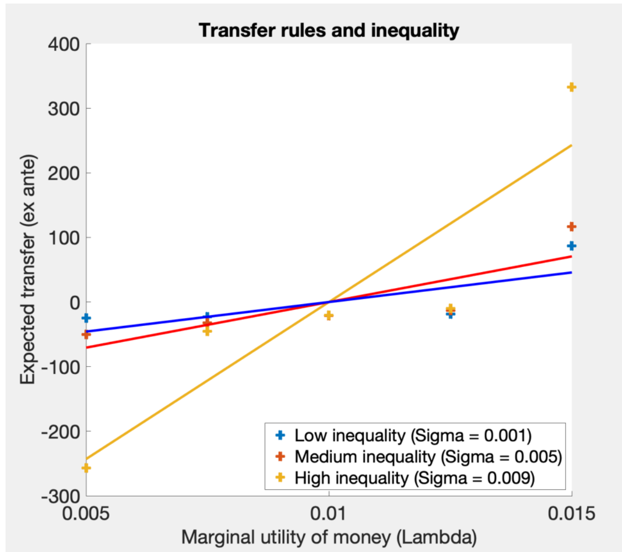

Figure 3: Transfer rules and inequality

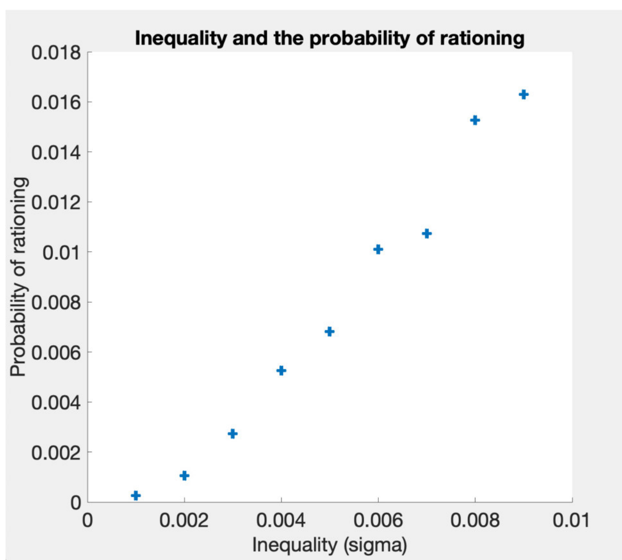

Figure 4: Inequality and the probability of rationing

A rise in σ raises the slope with which expected transfers rise as we move from agents with low Λ_i's to agents with higher Λ_i's. Whenever the marginal utility of an agent goes up, the utilitarian optimal mechanism will elicit less money from this agent and vice versa. A rise in σ will thus reduce the expected transfers for agents with low values of Λ_i and raise the expected transfers for agents with comparatively high values of Λ_i.

A rise in σ has a strictly positive and monotonous effect on the probability with which rationing occurs. Recall that rationing occurs for a given realization of types whenever there are less positive inequality adjusted valuations than goods to be allocated. Consider a rise in σ. For both agents $i = 1, 2$, this induces a fall in Λ_i and vice versa for the two agents $i = 4, 5$. In particular, Λ^* will rise. By proposition 5, this will raise $Pr(\varphi_i(r) < 0)$ for all agents $i \leq 4$. By proposition 5, the fact that Λ_i falls for agents $i = 1, 2$ implies that $Pr(\varphi_i(r) < 0)$ increases for these two agents, given that the uniform distribution is regular in the sense of Myerson [1981]. While the increase of Λ_4 implies an effect that goes in the opposite direction as compared to the effects listed previously, the latter dominate in this scenario.

5.2 Stochastic marginal utility of money

In the following section, we visualize some of our key insights for a setting where the marginal utilities of money are stochastic. Assume that there are two agents $i = 1, 2$ with $v_i^K \sim U[0, 1]$ and $v_1^M \sim Pareto(k = 3, x_{\min} = 1.5)$ while $v_2^M \sim Pareto(k = 3, x_{\min} = 1)$ with pdfs f_1 and f_2 respectively. The parameter k of the Pareto distribution describes the behavior of the distribution in the tails, where higher values of k imply thinner tails. Note that the support of the Pareto distribution is $[x_{\min}, \infty)$. Therefore, agent 2 has support on lower values of v^M than agent 1. This can naturally arise in a setting where agent 2 is ex ante more wealthy than agent 1, but there is still uncertainty about the marginal utility of money due to idiosyncrasies in consumption behaviour. Alternatively, this setup can reflect the notion that agent 2 is seen as ex ante more likely to be the wealthier of the two agents. Imagine that the designer is able to verify that agent 2 lives in a more expensive neighbourhood than agent 1 and therefore it is likely that agent 2 is wealthier than agent 1. To calculate the optimal mechanism for this example, we need to determine certain expressions of interest. For the sake of brevity, we omit the details of the calculations and provide the results only. Recall that our mechanism assigns the good to the agent with the largest inequality adjusted valuation, provided that it is positive. First, we determine the cdf $G_i(r_i)$ and pdf $g_i(r_i)$ for each agent through straightforward computations:

$$G_i(r_i) = \frac{1}{1-k}(rx_{\min})^k - \frac{k}{1-k}rx_{\min} \tag{33}$$

$$g_i(r_i) = \frac{k}{1-k}r^{k-1}x_{\min}^k - \frac{k}{1-k}x_{\min} \tag{34}$$

Note that these have support on $[0, 2/3]$ for agent 1 and $[0, 1]$ for agent 2. The next step is to determine the expected valuation of money, given the rate of substitution, namely $\lambda_i(r_i) = \mathbb{E}[v_i^M | r_i]$. This is:

$$\lambda_i(r_i) = \frac{(k-1)(x_{min}^{2-k} - r^{k-2})}{(k-2)(x_{min}^{1-k} - r^{k-1})} \tag{35}$$

Using this, we determine the inequality adjusted hazard rate. Furthermore, recall that our regularity condition requires that $\Pi(r) + \Lambda^* J(r)$ is non decreasing. That this regularity condition holds true in this particular example can be easily verified.

This example demonstrates how the allocation rule is modified in the presence of inequality. We recall that our optimal assignment rule assigns the good to agent i if and only if $\Pi_i(r_i) + \Lambda^* J_i(r_i) \geq \Pi_k(r_k) + \Lambda^* J_k(r_k)$ and $\Pi_i(r_i) + \Lambda^* J_i(r_i) \geq 0$. For our particular example, the utilitarian optimal allocation rule is illustrated in figure 5. The red line illustrates the ex post efficient allocation rule where an agent is assigned the object if and only if her rate of substitution r exceeds that of the other agent. The blue line represents the utilitarian optimal allocation rule. We recall that $v_1^M \sim Pareto(k = 3, x_{\min} = 1.5)$ and $v_2^M \sim Pareto(k = 3, x_{\min} = 1)$. Note that our example featured a setting where agent 1 was perceived to be poorer than agent 2 in the sense of having higher marginal utilities of money, on average. Our mechanism takes this inequality between the two agents into account and as a result agent 1 receives the good more often than under the standard assignment rule. This result is driven by two forces:

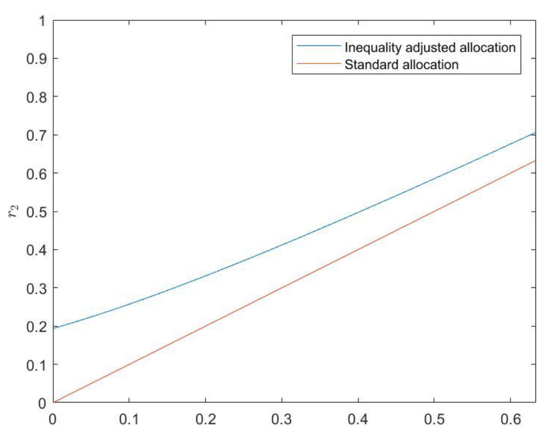

Figure 5: Utilitarian optimal allocation rule vs. ex post efficient allocation rule

First, the mechanism realizes that a low willingness to pay $r = v^K/v^M$ does not necessarily imply that the valuation of the good, v^K, is low. When the poor agent reports low values of r this is often driven by a high marginal utility of money, not by a low consumption utility.

Second, the designer has a preference to redistribute money from the rich agent to the poor agent. When the rich agent has a negative virtual valuation, the designer is not able to generate revenue for redistribution. This working channel does not exist for the poorer agent, since money would always be redistributed to her.

5.3 Policy implications

Our results provide a rationale for why a mechanism in a local allocation situation should use information about the wealth of the agents that participate. This explains, for example, why parents have to reveal their income when applying for kindergarden places for their children in German cities such as Munich, Hamburg, Cologne, or Bonn. For an overview of this, see Geis-Thöne [2018]. Our numerical results indicate that ignoring wealth inequality does not only imply that the utilitarian optimum will not be reached, but also favors wealthier people. Firstly, the optimal transfers wealthier people pay are lower in the standard setup where all agents have the same marginal utility of money. Secondly, our numerical results indicate that wealthier people receive the good more often in the standard ex-post efficient allocation rule than in the utilitarian optimal mechanism. Our results provide a rationale for some allocative tools that are already applied in practice. Our result that wealthier people should pay higher transfers in the optimal mechanism is mirrored in the fact that the transfer schemes for kindergarden place allocations are progressive in parental wealth in the aforementioned cities.

6 Discussion and Conclusion

We have derived the utilitarian optimal mechanism for an assignment problem in which the designer initially owns m units of an indivisible good which are to be distributed among a finite number N of agents, where $N > m$. In contrast to the usual assumption of the literature, our model works with heterogeneous marginal utilities of money. This implies that utility is not perfectly transferable between agents. We have formalized this notion by adapting the model of Dworczak et al. [2019] to our framework. While they analyze a two sided market with a clear distinction of buyers and sellers and a designer who sets the market rules, we consider a designer who owns the goods and wants to assign them.

The utilitarian optimal mechanism revolves around a key statistic which we call the *inequality adjusted valuation*. The inequality adjusted valuation condenses three critical considerations. First, the designer would like to allocate the good as efficiently as possible. Second, the designer is aware that transfers are not welfare-neutral and must thus consider how the allocation rule is linked to the transfer rule via the incentive compatibility condition. Third, there is ex interim uncertainty about each agent's marginal utility of money as the reported rate of substitution r_i is not perfectly informative about the marginal utility of money of a certain agent.

We showed that the mechanism which optimally balances these three considerations allocates the good to an agent if and only if (i) her inequality adjusted valuation is among the m highest inequality adjusted valuations and (ii) her inequality adjusted valuation is positive. Further, we have shown that the framework with perfectly transferable utility is included in our framework as a special case. In that special case, the inequality adjusted valuation of an agent simplifies to the standard valuation. Therefore, the ex post efficient allocation rule corresponds to the utilitarian optimal allocation rule. Further, we have shown that the ex post efficient allocation rule also arises when the inequality adjusted valuations satisfy a regularity condition and the preference parameters are distributed independently and identically. Under these assumptions, the inequality adjusted valuation is a strictly increasing transformation of the willingness to pay, which implies that the ex post efficient allocation rule is utilitarian optimal. In general, the optimal mechanism may feature rationing. Rationing is a byproduct of the revenue motive of the designer and may occur for all agents, except for the poorest agent. Whether the other agents are subject to rationing depends on the specific distributional assumptions, mirroring Myerson [1981], and the question of whether or not the virtual valuation of an agent is negative.

We have illustrated our results with two examples and have shown that agents who are perceived to be poorer have a higher chance of receiving the good in the utilitarian optimal mechanism than in the ex post efficient allocation rule. Their high marginal utility of money results in a comparatively low willingness to pay for the good, which the inequality adjusted valuation takes into account. Further, we investigated how wealth inequality, which we incorporate into our model through the heterogeneous marginal utilities of money, affects the probability of rationing. While rationing is useful to facilitate redistribution, it creates ex post inefficiencies that are undesirable. This highlights a novel channel through which wealth inequality contributes to allocative inefficiency.

References

Mohammad Akbarpour, Piotr Dworczak, and Scott Duke Kominers. Redistributive allocation mechanisms. Technical report, 2020.

Anthony Barnes Atkinson and Joseph E Stiglitz. The design of tax structure: direct versus indirect taxation. *Journal of public Economics*, 6(1-2):55–75, 1976.

Brian Baisa. Efficient multiunit auctions for normal goods. *Theoretical Economics*, 15(1): 361–413, 2020.

Brian Baisa and Justin Burkett. Efficient ex post implementable auctions and english auctions for bidders with non-quasilinear preferences. *Journal of Mathematical Economics*, 82:227–246, 2019.

Timothy Besley and Stephen Coate. Public provision of private goods and the redistribution of income. *The American Economic Review*, 81(4):979–984, 1991.

Tilman Börgers and Peter Norman. A note on budget balance under interim participation constraints: the case of independent types. *Economic Theory*, 39(3):477–489, 2009.

Bo Hyun Chang, Yongsung Chang, and Sun-Bin Kim. Pareto weights in practice: A quantitative analysis across 32 oecd countries. *Review of Economic Dynamics*, 28:181–204, 2018.

Daniele Condorelli. Market and non-market mechanisms for the optimal allocation of scarce resources. *Games and Economic Behavior*, 82:582–591, 2013.

Peter A Diamond and James A Mirrlees. Optimal taxation and public production i: Production efficiency. *The American Economic Review*, 61(1):8–27, 1971.

Sarah Donovan and David Bradley. Real wage trends, 1979 to 2018. *Congressional research service*, 2019.

Moritz Drexl and Andreas Kleiner. Optimal private good allocation: The case for a balanced budget. *Games and Economic Behavior*, 94:169–181, 2015.

Shaddin Dughmi and Yuval Peres. Mechanisms for risk averse agents, without loss. *arXiv preprint arXiv:1206.2957*, 2018.

Piotr Dworczak, Scott Duke Kominers, and Mohammad Akbarpour. Redistribution through markets. *Becker Friedman Institute for Research in Economics Working Paper*, (2018-16), 2019.

Roland Eisenhuth. Reference-dependent mechanism design. *Economic Theory Bulletin*, 7 (1):77–103, 2019.

Péter Esö and Gabor Futo. Auction design with a risk averse seller. *Economics Letters*, 65 (1):71–74, 1999.

Joan Esteban and Debraj Ray. Inequality, lobbying, and resource allocation. *American Economic Review*, 96(1):257–279, 2006.

Firouz Gahvari and Enlinson Mattos. Conditional cash transfers, public provision of private goods, and income redistribution. *American Economic Review*, 97(1):491–502, 2007.

Wido Geis-Thöne. Familien müssen für die gleiche betreuung in der kita unterschiedlich viel zahlen: Ein vergleich der gebührenordnungen der größten städte in deutschland. IW-Report 50/2018, Köln, 2018. URL `http://hdl.handle.net/10419/190947`.

Roger Guesnerie and Kevin Roberts. Effective policy tools and quantity controls. *Econometrica: Journal of the Econometric Society*, pages 59–86, 1984.

Kazuhiko Hashimoto and Hiroki Saitoh. Domain expansion of the pivotal mechanism. *Social Choice and Welfare*, 34(3):455–470, 2010.

Katharina Huesmann. Public assignment of scarce resources under income effects. *Working paper*, 2017.

Mingshi Kang and Charles Z Zheng. Necessity of auctions for redistributive optimality. Technical report, Working Paper, 2019.

Tomoya Kazumura, Debasis Mishra, and Shigehiro Serizawa. Mechanism design without quasilinearity. *Theoretical Economics*, 15(2):511–544, 2020.

Thomas Kesselheim and Bojana Kodric. Price of anarchy for mechanisms with risk-averse agents. *arXiv preprint arXiv:1804.09468*, 2018.

Nenad Kos and Matthias Messner. Incentive compatibility in non-quasilinear environments. *Economics Letters*, 121(1):12–14, 2013.

Maciej H Kotowski. First-price auctions with budget constraints. *Theoretical Economics*, 15 (1):199–237, 2020.

David Lee and Emmanuel Saez. Optimal minimum wage policy in competitive labor markets. *Journal of Public Economics*, 96(9-10):739–749, 2012.

Eric Maskin and John Riley. Optimal auctions with risk averse buyers. *Econometrica: Journal of the Econometric Society*, pages 1473–1518, 1984.

Shuhei Morimoto and Shigehiro Serizawa. Strategy-proofness and efficiency with non-quasilinear preferences: A characterization of minimum price walrasian rule. *Theoretical Economics*, 10(2):445–487, 2015.

Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

Mallesh M Pai and Rakesh Vohra. Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, 150:383–425, 2014.

Emmanuel Saez and Stefanie Stantcheva. Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1):24–45, 2016.

Hiroki Saitoh and Shigehiro Serizawa. Vickrey allocation rule with income effect. *Economic Theory*, 35(2):391–401, 2008.

Martin L Weitzman. Is the price system or rationing more effective in getting a commodity to those who need it most? *The Bell Journal of Economics*, pages 517–524, 1977.

Appendices

A Proofs

A.1 Proof of proposition 3:

The optimal mechanism needs to solve:

$$\max_{\{x_i(r_i,r_{-i}),U_i(\underline{r}_i\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i,r_{-i}) dG(r_i,r_{-i})\right)$$

$$\text{s.t.} \sum_i \left(U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i,r_{-i}) dG(r_i,r_{-i})\right) \leq 0 \qquad \text{(Budget)}$$

$$0 \leq x_i(r_i,r_{-i}) \leq 1 \qquad \text{(Prob)}$$

$$\sum_i x_i(r_i,r_{-i}) \leq m \qquad \text{(Feas)}$$

$$X_i(r_i) \text{ non-decreasing} \qquad \text{(Mono)}$$

$$U_i(\underline{r}_i) \geq \underline{U}_i \qquad \text{(IR)}$$

We ignore the monotonicity constraint for now and later show that this will be fulfilled in the optimal solution we present. Then the corresponding Lagrangian with the Kuhn-Tucker constraints is:

$$L = \sum_i \left[\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i,r_{-i}) dG(r_i,r_{-i}) + \mu\left(\int J_i(r_i) x_i(r_i,r_{-i}) dG(r_i,r_{-i}) - U_i(\underline{r}_i)\right)\right] +$$

$$\sum_i \mu_{1,i}(r_i,r_{-i})\big(x_i(r_i,r_{-i})\big) + \sum_i \mu_{2,i}(r_i,r_{-i})\big(1 - x_i(r_i,r_{-i})\big) + \rho(r_i,r_{-i})\left(m - \sum_i x_i(r_i,r_{-i})\right) +$$

$$\sum_i \delta_i \left(U_i(\underline{r}_i) - \underline{U}_i\right)$$

The first-order-conditions on $x_i(r)$ and $U_i(\underline{r}_i)$ then are:

$$\frac{\partial L}{\partial x_i(r)} = \Pi_i(r_i) g(r) + \mu J_i(r_i) g(r) + \mu_{1,i}(r) - \mu_{2,i}(r) - \rho(r) = 0$$

$$\frac{\partial L}{\partial U_i(\underline{r}_i)} = \Lambda_i - \mu + \delta_i = 0$$

Define the j-th order statistic for a collection of N random variables $\{x_i\}_{i=1}^N$ as $x_{j:N}$. Note that the IR constraints will bind for all agents except agent i^*. To see this, consider an agent $i \neq i^*$ and assume that $U_i(\underline{r}_i) > \underline{U}_i$.

Then, consider a reduction of $U_i(\underline{r}_i)$. This only affects the budget constraint and allows for an increase of $U_{i^*}(\underline{r}_{i^*})$. This increase implies a rise of the objective by $\Lambda^* > \Lambda_i$, which is thus optimal. This also implies that the budget constraint, which will bind in the optimum, can be rewritten as:

$$U_{i^*}(\underline{r}_{i^*}) = \int J_i(r_i)x_i(r_i, r_{-i})dG_i(r_i, r_{-i})$$

Claim: The following is a solution to the above Kuhn-Tucker problem, with appropriate Lagrange multipliers on the other constraints:

$$\mu = \Lambda^* = \max_i \Lambda_i$$

$$x_i(r) = \begin{cases} 1 & \Pi_i(r_i) + \Lambda^* J_i(r_i) > 0 \ \& \ \Pi_i(r_i) + \Lambda^* J_i(r_i) \geq [\Pi_i(r_i) + \Lambda^* J_i(r_i)]_{N-m+1:N} \\ 0 & \text{else} \end{cases}$$

We first prove why the above allocation rule is optimal for a general value of μ. Note that the budget constraint must always be binding in the optimal solution (otherwise it is always possible to raise ex ante transfers, which raises social welfare, without violating any other constraints).

Thus, it is without loss of generality to take the first line of the above Lagrangian as the objective function and maximize this subject to the remaining constraints. Note that this is linear in $x_i(r)$, which implies that the optimal allocation rule has a bang-bang property.

When $\Pi_i(r) + \mu J_i(r) < 0$, it is optimal to set $x_i(r) = 0$. Setting $x_i(r) < 0$ is not possible. Setting $x_i(r) > 0$ is not optimal, since moving to $x_i(r) = 0$ would be better. This change would not violate any constraints if the original point was a solution and would raise the objective function.

Suppose $\Pi_i(r) + \mu J_i(r) \geq 0$ and is among the m highest inequality adjusted valuations for a given r. Setting $x_i(r)$ equal to 1 is optimal, since it respects all constraints and there is no other allocation that would yield a higher utility.

Reducing the allocation probability $x_i(r)$ for any of these agents would, as such, reduce the objective function. In compensation, you could only increase $x_i(r)$ for an agent who has not received any of the good yet, but such an agent must have, by construction of the allocation rule, an inequality adjusted valuation below the one of the agent listed above. Thus, this allocation rule is optimal.

Setting $x_i(r) = 0$ when $\Pi_i(r) + \mu J_i(r) \geq 0$, but is not among the m highest inequality adjusted valuations is also optimal. Suppose $x_i(r) > 0$ for such a situation. Then, there must exist another agent with a higher inequality-adjusted valuation for which the probability constraints are not yet binding who should receive the good instead.

We now prove that $\mu = \Lambda^* = \max_i \Lambda_i$. We prove this by making a fixed point claim. Suppose that the Lagrange parameter on the feasibility constraint is $\Lambda = \Lambda^*$.

To understand the shadow value of an additional unit of budget, note that an additional unit of budget can either be used to raise ex-ante transfers or to change the allocation rule $x_i(r)$ for some type combination r. When tweaking the allocation rule, an additional unit of budget can be used to raise the allocation probability for an agent with a negative virtual valuation or to reduce the allocation probability for an agent with a positive virtual valuation.

Consider first type realizations r where the feasibility constraint binds, i.e. where $\rho(r) > 0$. This means there are more agents with positive inequality adjusted valuations than goods to be allocated. Note that the allocation rule stipulates that the individuals with the highest inequality adjusted valuations already receive the good for this realization of the vector r.

The only agents with a negative virtual valuation where more of the good can still be allocated must not have received the good yet. Consider such an agent i and examine the impacts of an increase in $x_i(r)$. This will change $U_{i^*}(\underline{r}_{i^*})$ by $J_i(r_i)g(r)$. This changes the objective function marginally by Λ^*. The total effect on the objective function is thus:

$$\big[\Pi_i(r_i) + \Lambda^* J_i(r_i)\big] g(r)$$

Given that the feasibility constraint binds, such a rise in $x_i(r)$ must be compensated by a fall of $x_j(r)$ (of the same magnitude) for an agent j who has already received the good. Similar arguments can be made for the reduction of $x_j(r)$, which will have the following total effect on the objective function:

$$\big[\Pi_j(r_j) + \Lambda^* J_j(r_j)\big] g(r)$$

The total effect of this change on the objective function is thus:

$$\Big[\big(\Pi_i(r_i) + \Lambda^* J_i(r_i)\big) - \big(\Pi_j(r_j) + \Lambda^* J_j(r_j)\big)\Big] g(r) < 0$$

This is negative because $\varphi_i(r_i) < \varphi_j(r_j)$ holds by the allocation rule, given that i has not yet received the good but j has received the good. By analogous arguments, there exist no optimal reductions of $x_j(r)$ for agents with positive virtual valuations.

Consider realizations of r such that the feasibility constraint is slack, i.e. that $\rho(r) = 0$. This implies that $[\Pi_i(r_i) + \Lambda^* J_i(r_i)]_{N-m+1:N} < 0$. For all people that have already received the good, no more of the good can be allocated - see the solution for $x_i(r)$.

All people that have not yet received the good must have $[\Pi_i(r_i) + \Lambda^* J_i(r_i)] < 0$. The key point is that it is not optimal to raise these values of $x_i(r)$. We show this as follows.

For the agents i for which nothing is allocated sofar, it holds that:

$$\Pi_i(r_i) + \Lambda^* J_i(r_i) < 0 \iff \Pi_i(r_i) < -\Lambda^* J_i(r_i)$$

For these agents, $J_i(r_i) < 0$ must hold. Thus, the above implies that:

$$\Pi_i(r_i) < -\Lambda^* J_i(r_i)$$

$$\iff$$

$$-\Pi_i(r_i) > \Lambda^* J_i(r_i) \iff \frac{-\Pi_i(r_i)}{J_i(r_i)} < \Lambda^*$$

Suppose the BC is made marginally slacker. The designer can either raise $U_i(\underline{r}_i)$ by 1 for agent i with $\Lambda^* = \max_j \Lambda_j$. Alternatively, the designer can raise such an $x_i(r)$ by $-1/J_i(r_i)g(r)$. Such a rise in $x_i(r)$ will imply a rise of the original objective by:

$$-\Pi_i(r_i)g(r)/J_i(r_i)g(r)$$

But this is worse than just raising $U_i(\underline{r}_i)$, as shown above. This completes the proof. The designer would always use additional funds in the budget to raise $U_i(\underline{r}_i)$, which implies that the Lagrange multiplier of the budget constraint corresponds to Λ^*.

Thus, no changes in the allocation rule are optimal.

Finally, it remains to show that the monotonicity constraint will hold in the optimal solution. This requires that $X_i(r_i)$ is non-decreasing. To see this, note that:

$$X_i(r_i) = \mathbb{1}[\varphi_i(r_i) > 0]Pr(\varphi_i(r_i) \geq [\varphi_j]_{N-m+1:N})$$

Note our key assumption that $\varphi_i(r_i)$ is increasing in r_i. For values of r_i where the inequality adjusted valuation is negative, monotonicity holds. Now consider values of r_i where $\varphi_i(r_i) > 0$. Here, since $\varphi_i(r_i)$ is rising in r_i, this implies that the above probability cannot be falling.

A.2 Proof of Remark 1:

Assume that the marginal utility of money is non-stochastic for all agents, such that $\Lambda_i = \lambda_i(r_i) = v_i^M$

Part 1: Since v^M is deterministic, we have that:

$$\varphi_i(r_i) = \Lambda_i r_i + (\Lambda^* - \Lambda_i)J_i(r_i)$$

When the environment is regular, i.e. when the virtual valuation $J_i(r)$ is rising in r, then the inequality-adjusted valuation will also be increasing in r.

Part 2:

We are interested in the distribution $G_i(r)$ for the random variable $r = p_i(v^K)$, with $p_i(x) = x/\Lambda_i$. Naturally, the derivative of $p_i^{-1}(r)$ is Λ_i. Thus, when the pdf of v^K is f_i, one can show that:

$$G_i(r) = F_i(\Lambda_i r) \implies g_i(r) = \Lambda_i f_i(\Lambda_i r)$$

We thus have:

$$\frac{1-G_i(s)}{g_i(s)} = \frac{1-F_i(\Lambda_i s)}{\Lambda_i f_i(\Lambda_i s)}$$

Recall that:

$$J_i(r) = r - \frac{1-G_i(r)}{g_i(r)} \implies \frac{\partial J_i(r)}{\partial r} = 1 - \frac{\partial}{\partial r}\left[\frac{1-G_i(r)}{g_i(r)}\right]$$

The necessary and sufficient condition for the inequality adjusted valuation to weakly rise in r_i when v^M is deterministic is thus:

$$\frac{\partial \varphi_i}{\partial r} = \Lambda_i + (\Lambda^* - \Lambda_i)\frac{\partial J_i(r)}{\partial r} \geq 0 \iff \frac{\Lambda_i}{\Lambda^* - \Lambda_i} \geq -\frac{\partial J_i(r)}{\partial r} = -1 + \frac{\partial}{\partial r}\left[\frac{1-G_i(r)}{g_i(r)}\right]$$

$$\iff$$

$$1 + \frac{\Lambda_i}{\Lambda^* - \Lambda_i} \geq \frac{\partial}{\partial r}\left[\frac{1-G_i(r)}{g_i(r)}\right]$$

Finally, we can note the following by the chain rule:

$$\frac{\partial}{\partial r}\left[\frac{1-G_i(r)}{g_i(r)}\right] = \frac{\partial}{\partial r}\left[\frac{1-F_i(\Lambda_i r)}{\Lambda_i f_i(\Lambda_i r)}\right] = (\Lambda_i/\Lambda_i)\frac{\partial}{\partial v^K}\left[\frac{1-F_i(v^K)}{f_i(v^K)}\right]$$

This translates into the following necessary and sufficient condition:

$$\frac{\partial}{\partial v^K}\left[\frac{1-F_i(v^K)}{f_i(v^K)}\right] \leq 1 + \frac{\Lambda_i}{\Lambda^* - \Lambda_i}$$

A.3 Proof of Proposition 5

Preliminaries

We suppose that the marginal utility of money is fixed for all agents, such that $v_i^M = \Lambda_i$. In this case, the inequality-adjusted valuation is:

$$\varphi_i(r_i) = \Lambda_i r_i + (\Lambda^* - \Lambda_i)J_i(r_i)$$

Note that:

$$J_i(r_i) = r_i - \frac{1-G_i(r_i)}{g_i(r_i)} = r_i - \frac{1}{\Lambda_i}\frac{1-F_i(\Lambda_i r_i)}{f_i(\Lambda_i r_i)}$$

The first component of the inequality adjusted valuation is only 0 at $r_i = 0$, at which we have:

$$J_i(0) = 0 - \frac{1}{\Lambda_i}\frac{1-F_i(0)}{f_i(0)} = -\frac{1}{\Lambda_i}\frac{1}{f_i(0)} < 0$$

Given that we work with distributions without mass points, we can note the following: Defining $\hat{r}$ as the valuation which solves $\varphi(\hat{r}) = 0$, it must hold that $\hat{r} > 0$. Also note continuity of $J_i(r)$ in r.

Proof of point 1:

We need to show that:

$$\frac{\partial Pr(\varphi_i(r) > 0)}{\partial \Lambda^*} \geq 0$$

Consider any agent i and note the following:

$$\frac{\partial \varphi_i(r)}{\partial \Lambda^*} = J_i(r) \;\; \forall r \in [\underline{r}_i, \bar{r}_i]$$

To understand the effect of a rise in Λ^* on the probability with which agent i is rationed , note firstly that the stochastic variable is r_i.

Firstly, consider realizations of r_i where $\varphi_i(r_i) < 0$ ex ante. Because $\varphi_i(r_i) < 0$, it must hold that $J_i(r_i) < 0$ at these realizations of r_i. For these realizations, a rise in Λ^* will thus reduce $\varphi_i(r_i)$, keeping this negative for all these realizations of r_i.

Secondly, consider realizations of r_i where $\varphi_i(r_i) \geq 0$ and $J_i(r_i) \geq 0$ holds true. For these realizations of r_i, the rise in Λ^* will imply a weak increase of $\varphi_i(r_i)$, such that $\varphi_i(r_i) \geq 0$ will still hold after the rise in Λ^*.

Thirdly and finally, consider realizations of r_i where $\varphi_i(r_i) \geq 0$ and $J_i(r_i) < 0$ holds true. For these realizations of r_i, the rise in Λ^* will imply a (weak) decrease of φ_i, which can potentially push those into the negative region, even though they were positive ex ante. This working channel has a weakly positive effect on the probability that this agent is rationed.

It remains to show that the rise in $Pr(\varphi_i < 0)$ implies a rise in the probability with which rationing occurs.

To do this, we introduce some terminology. Consider agent i, for which we know that $Pr(\varphi_i < 0)$ has increased due to some exogenous change.

Profile of Agents A profile of agents is a certain arrangement of the signs of the inequality adjusted valuations for all agents $i \in I$, where either $\varphi_i(r) < 0$ or $\varphi_i(r) \geq 0$ holds.

Partial profile of agents A partial profile of agents is a certain arrangement of the signs of the inequality adjusted valuations for all agents $j \neq i$, where either $\varphi_j(r) < 0$ or $\varphi_j(r) \geq 0$ holds.

We focus on the probability with which rationing occurs. This probability is a sum over profiles of agents where at least $N - m + 1$ agents have negative inequality adjusted valuations.

In this probability, we can loop over partial profiles of agents.

Consider first partial profiles of agents where weakly less than $N - m - 1$ agents j have $\varphi_j < 0$. These partial profiles will not enter the probability of rationing, since rationing would not occur for such a partial profile, no matter the realization of φ_i. An increase of $Pr(\varphi_i < 0)$ will thus have no effect for these type profiles.

Consider now partial profiles where weakly more than $N - m + 1$ agents have $\varphi_j < 0$. For these partial profiles, rationing will occur, no matter the realizations of φ_i. Thus, for every such partial profile, there are two profiles that enter the probability with which rationing occurs. Raising $Pr(\varphi_i > 0)$ will then have zero effect on the probability of rationing via these partial profiles, since the effects exactly cancel out.

Consider finally partial profiles where exactly $N - m$ agents $j \neq i$ have negative inequality adjusted valuations. For these partial profiles, $\varphi_i < 0$ must hold for rationing to occur. Hence, only one profile with such a partial profile will enter the probability of rationing, namely the one where $\varphi_i < 0$, with the corresponding probabilities. When $Pr(\varphi_i < 0)$ increases, you will unambiguously raise the probability of rationing for these profiles.

<u>Point 2:</u>

It was shown previously that an increase of $Pr(\varphi_i(r) < 0)$ will imply an increase in the probability with which rationing occurs. This implies the first sentence of this point.

We thus only have to show:

- $\frac{\partial}{\partial \nu^K}\left[\frac{1-F_i(v^K)}{f_i(v^K}\right] \leq 0 \implies \frac{\partial Pr(\varphi_i(r_i)<0)}{\partial \Lambda_i} < 0.$

Consider any agent i. The derivative of φ_i with respect to Λ_i is:

$$\frac{\partial \varphi_i(r_i)}{\partial \Lambda_i} = r_i + (\Lambda^* - \Lambda_i)\frac{\partial J_i(r_i)}{\partial \Lambda_i} + (-1)J_i(r_i) = r_i + (\Lambda^* - \Lambda_i)\frac{\partial J_i(r_i)}{\partial \Lambda_i} + (-1)\left[r_i - \frac{1 - G_i(r_i)}{g_i(r_i)}\right] =$$

$$(\Lambda^* - \Lambda_i)\frac{\partial J_i(r_i)}{\partial \Lambda_i} + \frac{1 - G_i(r_i)}{g_i(r_i)}$$

When the virtual valuation is weakly rising in Λ_i, the inequality adjusted valuation will uniformly rise in Λ_i.

A uniform increase in $\varphi_i(r_i)$ as a result of a change in in Λ_i implies that the probability with which rationing occur will fall.

Recall that:

$$J_i(r_i) = r_i - \frac{1 - G_i(r_i)}{g_i(r_i)} = r_i - \frac{1}{\Lambda_i}\frac{1 - F_i(\Lambda_i r_i)}{f_i(\Lambda_i r_i)}$$

The derivative is thus:

$$\frac{\partial J_i(r)}{\partial \Lambda_i} = +\frac{1}{\Lambda_i^2}\frac{1-F_i(\Lambda_i r_i)}{f_i(\Lambda_i r_i)} - \frac{1}{\Lambda_i}\frac{\partial(\Lambda_i r_i)}{\partial \Lambda_i}\left[\frac{\partial}{\partial v^K}\frac{1-F_i(v^K)}{f_i(v^K)}\right] =$$

$$\frac{1}{\Lambda_i^2}\frac{1-F_i(\Lambda_i r_i)}{f_i(\Lambda_i r_i)} - \frac{1}{\Lambda_i} r_i\left[\frac{\partial}{\partial v^K}\frac{1-F_i(v^K)}{f_i(v^K)}\right]$$

Sufficient for this derivative to be positive is hence that

$$\frac{\partial}{\partial v^K}\frac{1-F_i(v_i^K)}{f_i(v_i^K)} \leq 0$$

A.4 Proof of remark 2:

We need to show that $\gamma_i(r_i) \geq 0$ for all i, where:

$$\gamma_i(r) = \Pi_i(r) + \Lambda_i J_i(r_i) = \frac{\int_r^{\bar{r}_i} \lambda_i(s) dG_i(s)}{g_i(r)} + \Lambda_i\left(r - \frac{1-G_i(r)}{g_i(r)}\right) = \Lambda_i r_i + \frac{\int_r^{\bar{r}_i} \lambda_i(s) - \Lambda_i dG_i(s)}{g_i(r)}$$

The first component is positive. To see that the second component is positive, note that $\lambda_i(s)$ is rising in s by assumption. This implies that:

$$\int_r^{\bar{r}_i} (\lambda_i(s) - \Lambda_i)\, g_i(s) ds = [1-G_i(r)] \int_r^{\bar{r}_i} \lambda_i(s) \frac{g_i(s)}{1-G_i(r)} ds - [1-G_i(r)]\Lambda_i =$$

$$[1-G_i(r)]\Big(\mathbb{E}[\lambda_i(s)|s \geq r] - \mathbb{E}[\lambda_i(s)]\Big) \geq 0$$

A.5 Proof of proposition 6

$$\max_{\{x_i(r_i,r_{-i}),U_i(\underline{r}_i)\}_{i=1}^N} \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i})\right)$$

$$\text{s.t. } U_i(\underline{r}_i) - \int J_i(r_i) x_i(r_i, r_{-i}) dG_i(r_i, r_{-i}) \leq 0 \;\; \forall i \qquad \text{(Transfers)}$$

$$0 \leq x_i(r_i, r_{-i}) \leq 1 \qquad \text{(Prob)}$$

$$\sum_i x_i(r_i, r_{-i}) \leq m \qquad \text{(Feas)}$$

$$X_i(r_i) \text{ non-decreasing} \qquad \text{(Mono)}$$

$$U_i(\underline{r}_i) \geq 0 \qquad \text{(IR)}$$

One can easily show that the transfer constraints must bind for all agents.

Thus, we can set up the following Lagrangian (ignoring monotonicity constraint for now):

$$L = \sum_i \left(\Lambda_i U_i(\underline{r}_i) + \int \Pi_i(r_i) x_i(r_i, r_{-i}) dG(r_i, r_{-i}) \right) + \sum_i \mu_i \left(\int J_i(r_i) x_i(r_i, r_{-i}) dG_i(r_i, r_{-i}) - U_i(\underline{r}_i) \right)$$
$$+ \ldots$$

For a given set of Lagrange parameters $\{\mu_i\}_{i \in I}$, the optimal allocation rule is to allocate the good to the agents with the highest values of $\Pi_i(r) + \mu_i J_i(r)$. It remains to pin down the Lagrange parameters.

Consider an arbitrary agent i. We now show that $\mu_i = \Lambda_i$ for all agents. Suppose this holds.

Consider a marginal increase of the expected transfer that can be allocated to agent i. What can the designer do with this additional dollar? Firstly, the designer could change the allocation rule and raise $x_i(r)$ for a realization r for which agent i has not yet received the good.

Suppose the feasibility constraint binds for this realization of r. Then, you could raise $x_i(r)$ only when you reduce $x_j(r)$ for another agent j who has previously received the good. This has a direct negative effect on the social welfare equal to $-\Pi_j(r_j)$. In addition, via the binding transfer constraint of agent j, $U_j(\underline{r}_j)$ will change by $-J_j(r_j)$, which impacts social welfare by Λ_j. Conversely, the total rise of the objective function via the utility of agent i is $\Pi_i(r_i) + \Lambda_i J_i(r_i)$. However, this change will then not be optimal by the above allocation rule.

The feasibility constraint can never be slack under the above Lagrange parameters. This holds because γ_i can never be negative - see remark 2. Assuming that the feasibility constraint is slack quickly leads to a contradiction.

This completes the proof. Whenever the transfer constraint becomes marginally slacker, the only way the designer can use this additional dollar is to raise $U_i(\underline{r}_i)$, which yields additional social welfare equal to Λ_i.

Monotonicity will also hold in the above setting, given that γ_i is rising in r_i and the allocation probability is rising in γ_i.

A.6 Results for uniform distribution and a mean-preserving spread

We have to show:

Assume v^K is uniformly distributed on $[0, 1]$ and that there are two goods to be allocated. When $\Lambda_i < \Lambda_j$, it holds that:

$$Pr(\varphi_i(r_i) < 0) \frac{\partial Pr(\varphi_j(r_j) < 0)}{\partial \Lambda_j} - Pr(\varphi_j(r_j) < 0) \frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i} < 0$$

Assume $m = 2$, v_i^M is fixed for all agents, and $v^K \sim U[0,1]$. Then, rationing will occur whenever all the agents except the poorest agent have a negative inequality adjusted valuation. Hence, the probability that rationing occurs is:

$$\prod_{k \neq i^*} Pr(\varphi_k(r) < 0)$$

Suppose you raise Λ_j marginally and reduce Λ_i by the same magnitude. The effect of this on the probability with which rationing occurs will be:

$$\prod_{k \neq i,j,i^*} Pr(\varphi_k(r) < 0)\bigg(\frac{\partial Pr(\varphi_j(r_j) < 0)}{\partial r_j} Pr(\varphi_i(r_i) < 0) - \frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial r_i} Pr(\varphi_j(r_j) < 0)\bigg)$$

For any person $k \neq i^*$, $\varphi_k(r)$ becomes the following in the above setting:

$$\varphi_k(r) = \big(1 - \Lambda^*/\Lambda_k\big) + r_i\big(2\Lambda^* - \Lambda_i\big)$$

Let's compute the probability that this is below 0:

$$Pr(\varphi_k < 0) = Pr\Bigg(\big(1 - \Lambda^*/\Lambda_k\big) + r_k\big(2\Lambda^* - \Lambda_k\big) < 0\Bigg) =$$

$$Pr\Bigg(r_k < \frac{\Lambda^*/\Lambda_k - 1}{2\Lambda^* - \Lambda_k}\Bigg) = G_k\Bigg(\frac{\Lambda^*/\Lambda_k - 1}{2\Lambda^* - \Lambda_k}\Bigg) = \frac{\Lambda^* - \Lambda_k}{2\Lambda^* - \Lambda_k}$$

Taking the derivative with respect to Λ_i yields that:

$$\frac{\partial}{\partial \Lambda_k}\Bigg[Pr\Bigg(r_k < \frac{\Lambda^*/\Lambda_k - 1}{2\Lambda^* - \Lambda_k}\Bigg)\Bigg] = \frac{(2\Lambda^* - \Lambda_k)(-1) - (\Lambda^* - \Lambda_k)(-1)}{(2\Lambda^* - \Lambda_k)^2} = -\frac{(2\Lambda^* - \Lambda_k) - (\Lambda^* - \Lambda_k)}{(2\Lambda^* - \Lambda_k)^2} =$$

$$-\frac{\Lambda^*}{(2\Lambda^* - \Lambda_k)^2} < 0$$

Now consider two agents i, j with $\Lambda_i < \Lambda_j$.

$$Pr(\varphi_i(r_i) < 0)\frac{\partial Pr(\varphi_j(r_j) < 0)}{\partial \Lambda_j} - Pr(\varphi_j(r_j) < 0)\frac{\partial Pr(\varphi_i(r_i) < 0)}{\partial \Lambda_i}$$

This is smaller than 0 iff:

$$Pr(\varphi_i < 0)\frac{\partial Pr(\varphi_j < 0)}{\partial \Lambda_j} < Pr(\varphi_j < 0)\frac{\partial Pr(\varphi_i < 0)}{\partial \Lambda_i} =$$

$$\frac{\Lambda^* - \Lambda_i}{2\Lambda^* - \Lambda_i}\Bigg(-\frac{\Lambda^*}{(2\Lambda^* - \Lambda_j)^2}\Bigg) < \frac{\Lambda^* - \Lambda_j}{2\Lambda^* - \Lambda_j}\Bigg(-\frac{\Lambda^*}{(2\Lambda^* - \Lambda_i)^2}\Bigg) \iff$$

$$\frac{\Lambda^*(\Lambda^* - \Lambda_i)}{(2\Lambda^* - \Lambda_i)(2\Lambda^* - \Lambda_j)^2} > \frac{\Lambda^*(\Lambda^* - \Lambda_j)}{(2\Lambda^* - \Lambda_j)(2\Lambda^* - \Lambda_i)^2} \iff$$

$$\frac{(\Lambda^* - \Lambda_i)}{(2\Lambda^* - \Lambda_j)} > \frac{(\Lambda^* - \Lambda_j)}{(2\Lambda^* - \Lambda_i)} \iff (\Lambda^* - \Lambda_i)(2\Lambda^* - \Lambda_i) > (\Lambda^* - \Lambda_j)(2\Lambda^* - \Lambda_j)$$

Note that:

$$\frac{\partial}{\partial \Lambda_k}\Big[(\Lambda^* - \Lambda_k)(2\Lambda^* - \Lambda_k)\Big] < 0$$

This implies that the above inequality always holds true,since $\Lambda_j > \Lambda_i$.