

Redistributive Allocation Mechanisms*

Mohammad Akbarpour [Ⓐ] Piotr Dworzak [Ⓐ] Scott Duke Kominers[†]

May 23, 2020

Abstract

Many scarce public resources are allocated below market-clearing prices (and sometimes for free). Such “non-market” mechanisms necessarily sacrifice some surplus, yet they can potentially improve equity by increasing the rents enjoyed by agents with low willingness to pay. In this paper, we develop a model of mechanism design with redistributive concerns. Agents are characterized by a privately observed willingness to pay for quality, and a publicly observed label. A market designer controls allocation and pricing of a set of objects of heterogeneous quality, and maximizes a linear combination of revenue and total surplus—with Pareto weights that depend both on observed and unobserved agent characteristics. We derive structural insights about the form of the optimal mechanism and describe how social preferences influence the use of non-market mechanisms.

Keywords: optimal mechanism design, redistribution, inequality, welfare

JEL codes: C78, D47, D61, D63, D82

*The “Ⓐ” symbol indicates that the authors’ names are in certified random order, as described by [Ray Ⓐ Robson \(2018\)](#). The authors thank Cassiano Machado Alves, Jeff Ely, Ravi Jagadeesan, Simon Loertscher, Ellen Muir, Alessandro Pavan, Stefanie Stantcheva, and Alex Teytelboym for helpful comments. All three authors gratefully acknowledge the support of the Washington Center for Equitable Growth and thank the Becker Friedman Institute both for sparking ideas that led to this work and for hosting the authors’ collaboration. Additionally, Kominers gratefully acknowledges the support of National Science Foundation grant SES-1459912, the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications (CMSA), an Alfred P. Sloan Foundation grant to the CMSA, and the Human Capital and Economic Opportunity (HCEO) Working Group sponsored by the Institute for New Economic Thinking (INET).

[†]*Akbarpour*: Graduate School of Business, Stanford University; mohamwad@stanford.edu. *Dworczak*: Department of Economics, Northwestern University; piotr.dworczak@northwestern.edu. *Kominers*: Entrepreneurial Management Unit, Harvard Business School; Department of Economics and Center of Mathematical Sciences and Applications, Harvard University; and National Bureau of Economic Research; kominers@fas.harvard.edu.

1 Introduction

Many scarce public resources—such as public housing, road access, schools seats, national park permits, and certain types of healthcare—are allocated at below-market-clearing prices, and sometimes for free. Such “non-market” allocation mechanisms naturally raise concerns among economists since they necessarily sacrifice some allocative surplus, failing to allocate resources to those with the highest willingness to pay. Yet, policymakers often justify non-market mechanisms on fairness grounds. If resources were allocated using market-clearing prices, agents with the lowest willingness to pay would often be excluded from enjoying their benefits. Because low willingness to pay for many goods and services is likely to be correlated with adverse social and economic circumstances—such as low wealth, unemployment or health problems—marketplace designers may be naturally more concerned about the welfare of such agents. But how should we think about the resulting trade-offs?

In this paper, we develop a model of optimal allocation with redistributive concerns. A policymaker/market designer controls allocation and pricing in a single marketplace, and trades off allocative efficiency against redistribution. We characterize the optimal mechanism for an arbitrary set of Pareto weights that can depend on both observable and unobservable characteristics of market participants. We show that under broad conditions, it is optimal to augment market mechanisms with lotteries that allocate some of the objects at a reduced price (potentially for free), as long as the designer attaches a sufficiently high Pareto weight to either agents with low willingness to pay (an unobservable characteristic), or to agents that belong to a “preferred” group (an observable characteristic, as in affirmative action).

Framework. Our model is as follows. The designer seeks to allocate a set of indivisible objects of heterogeneous quality to a large number of prospective buyers, i.e. *agents*. Each agent is characterized by her *type* r , which reflects her willingness to pay for quality. More precisely, an agent who receives a good with quality q in exchange for a monetary transfer t will have utility $rq - t$. Each agent also has a publicly observable *label* i , which is (potentially) informative about the agent’s type, and may indicate the relative social value placed on the agent’s utility. We think of agents having the same label as comprising a group of agents with the same observable characteristics. For example, labels could represent neighborhoods—in which case agents from low-income neighborhoods might have lower willingness to pay on average (reflecting high opportunity cost of money, or equivalently low wealth). At the same time, social preferences could reflect labels—perhaps targeting socioeconomic mobility by placing more weight on agents from lower-income neighborhoods.

The designer chooses a mechanism that allocates the goods among buyers (with some

goods potentially unassigned) to maximize a weighted sum of revenue and welfare, subject to individual rationality and incentive-compatibility constraints. Crucially, when measuring welfare, the designer considers a weighted sum of agents' utilities, where weights can depend both on labels and on willingness to pay. Thus, in particular, we allow the market designer to have redistributive preferences; for example, she can put higher weights on agents with lower willingness to pay, perhaps reflecting the idea that those agents are more likely to be poor. Because of individual-rationality and incentive-compatibility constraints, the designer faces a non-trivial trade-off between efficiency and redistribution (the conclusion of the second welfare theorem does not hold).

The weight attached to revenue in the designer's objective function is also an important preference parameter, and is interpreted as the marginal (social) opportunity cost of a unit of money. If the designer returns the revenue to agents in the form of a lump-sum transfer (or the revenue subsidizes a budget financed otherwise through taxes), its marginal value is equal to the average Pareto weight. However, giving money to agents with certain observable characteristics might not always be possible or desirable. For example, lump-sum transfers might be politically or institutionally infeasible; the designer may be concerned about entry or agents' incentives to affect their labels (e.g., becoming unemployed); or paying the lump-sum transfer might be frictional or administratively costly. In all these cases, the weight on revenue might be below the average Pareto weight attached to a given group. Finally, if money can be used more effectively for some outside cause, then the weight on revenue might exceed the average Pareto weight attached to buyers.

We think of the designer in our model as either a marketplace regulator, such as a local government or agency, or more abstractly as a reflection of social preferences. The implicit assumption underlying our analysis is that the designer can only control the marketplace in question, and does not have access to macroeconomic tools such as income taxation. As a consequence, the designer takes the inequality in the market as given, and does not take into account how her mechanism might potentially influence the Pareto weights. This is in contrast to many classical public finance models (such as the canonical framework of [Diamond and Mirrlees, 1971](#)), where inequality is determined endogenously within the model, and redistribution is primarily accomplished via taxing income of workers with privately observed ability. While we share the interest in the redistribution question with the public finance literature, our market-design approach to the problem—with quasi-linear utilities, unit demand, and emphasis on market-clearing rules—is distinct and complementary.

Preview of results. We construct an optimal mechanism for this setting in two steps: First, we solve the optimal allocation problem for each group of labeled agents, conditional

on a given allocation of goods to that group. Then, we show how that “within”-label solution gives rise to a simple greedy algorithm for optimally allocating goods “across” the groups.

For a fixed label and set of goods to allocate, the optimal “within” mechanism partitions agents into blocks according to their willingness to pay, and within each block allocates the good either randomly or assortatively (with higher-quality goods assigned to agents with higher willingness to pay). Which allocative mechanism is used is determined by a trade-off: Randomization increases the probability that agents with lower willingness to pay receive the good and lowers prices for all agents with higher willingness to pay, but also leads to allocative inefficiency. Meanwhile, assortative matching allocates the good to those with highest willingness to pay, thus increasing the allocative efficiency and raising more revenue.

Given our solution for an optimal mechanism “within” label groups, we find the optimal allocation of goods “across” groups by identifying a simple statistic characterizing which group will have the highest marginal return to a unit of quality q . Then the optimal mechanism allocates goods “across” groups greedily with respect to group value, starting with lowest-quality goods and proceeding to the highest.

The first key implication of our analysis is that under relatively broad conditions, the designer always allocates some number of lowest-quality objects at random and for free in “favored” groups, if she cannot give them money directly (i.e., lump-sum transfers to this group are not possible). More precisely, if the average Pareto weight on agents with a certain label exceeds the weight that the designer attaches to revenue, and the good is “essential”—which we model as willingness to pay for it being bounded away from 0—then the optimal mechanism offers all objects of quality below a certain threshold in a lottery at a price of 0. Such a policy has an intuitive positive effect on the welfare of agents who choose the lottery and receive the object for free. However, the result makes no assumptions on the relative welfare weights attached to these agents, and thus this direct effect in itself would not imply the result. Instead, there is an indirect effect that allows the designer to transfer the benefit to all agents: By using a lottery with a price of 0 for the agents with lowest willingness-to-pay, the designer can lower prices for agents who buy higher-quality objects at positive prices—without violating the incentive constraints. At the same time, allocating *all* objects in a lottery might be suboptimal because it both *(i)* lowers revenue (which could still have significant weight in the designer’s objective) and *(ii)* decreases allocative efficiency. This result justifies the wide-spread use of non-market mechanisms for allocation of goods such as public housing, food, or health care to agents satisfying some verifiable eligibility criteria, in cases where it would be politically infeasible (or otherwise impractical) to subsidize these groups of agents directly with monetary transfers.

More generally, the size and the composition of the blocks of random and assortative

matching depend on the underlying distribution of agents' willingness to pay, as well as the weights. When the designer has relatively strong preferences for raising revenue, we show that in typical cases there is assortative matching at the top of the distribution of willingness to pay. Still, assortative matching will fail to be optimal globally if the Pareto weights on agents' surplus are sufficiently dispersed. Assortative matching shifts more surplus to agents with high willingness to pay, so if Pareto weights are decreasing in willingness to pay, it may still be optimal for the designer to use random matching to allocate to agents with lower willingness to pay.

However, we show that the effect of dispersion in Pareto weights is non-monotone. Fixing the average Pareto weight, and assuming that the weight on revenue is weakly larger, fully assortative matching arises in the limit both when Pareto weights become constant, and when they become entirely concentrated on agents with the lowest willingness to pay. The reason is that when the Pareto weights are extremely skewed, the designer attaches nearly no weight to the surplus of agents with relatively high willingness to pay, and thus the revenue-maximizing motive dominates for these types. If only a small proportion of "poorest" agents receive a positive weight, the corresponding random-allocation region must also be small. This suggests that the scope for using a random allocation is maximized when the designer has intermediate preferences for redistribution towards agents with low willingness to pay.

Finally, we explore the implications of the structure of the optimal "within" mechanisms for how objects should be allocated across groups. Our main insight is that when faced with a group where matching is random and a group where matching is assortative, the designer should in general allocate the intermediate-quality objects to the former, and lowest- and highest-quality object to the latter. This stems from different ways in which the type of matching affects the value of the marginal unit of quality. In the random-matching groups, this value is constant, and hence only average quality matters; in the assortative-matching group, the marginal value is low when the good is allocated to a low type but can be high when it is allocated to a high type, and hence dispersion in the distribution of quality is desirable for the designer. This suggests that in settings where subsidized programs with non-market allocations for eligible agents coexist with a free market, it is not necessarily optimal to allocate only the lowest-quality objects in these programs.

1.1 Related work

We are far from being the first to study the question of how to allocate goods when the designer's preferences account for various measures of agent welfare. This question was asked at least as early as the work of [Weitzman \(1977\)](#), who showed that a fully random allocation

can be better than competitive pricing when the agents’ needs (as reflected in the designer’s objective function) are not well expressed by their willingness to pay. Yet, we would like to emphasize that while a myriad of papers have been written on the problem of maximizing revenue and efficiency in mechanism design, the question of optimal redistribution (in the absence of the second welfare theorem) has received incomparably less attention, despite its obvious practical relevance.

The two most related papers to ours are those of [Condorelli \(2013\)](#) and [Dworczak & Kominers & Akbarpour \(2019\)](#). [Condorelli \(2013\)](#) provided conditions for the optimality of market and non-market mechanisms for allocating k identical objects to n agents in an environment where the designer maximizes agents’ values that may be different from their willingness to pay. We extend [Condorelli’s](#) analysis by considering heterogeneous quality of objects and observable agent characteristics. Moreover, we allow the designer to have preferences over revenue, and accommodate cases when lump-sum transfers are not feasible. These additional elements lead to new economic insights about the structure of optimal mechanisms. For example, many of our conclusions depend crucially on how the average Pareto weight relates to the weight on revenue; the dispersion in Pareto weights as a function of willingness to pay is a key determinant of the use of randomized allocation; and heterogeneous quality makes the “across” problem and various notions of matching meaningful—all these elements are absent from [Condorelli’s](#) analysis. [Dworczak & Kominers & Akbarpour \(2019\)](#) (henceforth DKA) studied a similar question to ours but in a two-sided market with buyers and sellers, goods of homogenous quality, with mechanisms that clear the market and maintain budget balance. They show that the optimal mechanism is simple in that it only uses a combination of lump-sum payments to one side of the market and rationing (offering the good at an attractive price but with an interior probability of transacting). The current paper extends this analysis and examines which of the conclusions are robust to adding realistic features—heterogeneous quality of objects, richer preferences of the designer over revenue, additional observable information about the agents, and a restriction on the use of lump-sum transfers. Heterogeneity of object qualities makes the optimal mechanism more complicated in general, with potentially multiple “regions” of assortative versus random matching (which generalizes rationing). When revenue receives a sufficiently high weight in the designer’s objective, the optimal mechanism may leave some objects unassigned. Additional observable information creates a novel problem of how to split the objects among groups of agents sharing the same public characteristics, allowing applications to affirmative action. Finally, we show that when lump-sum transfers are not feasible, randomization in the mechanism may be optimal even under conditions that would make rationing suboptimal in the setting of DKA.

The introduction of observable characteristics to our model—while novel for the redistribution question within the market-design context—is a classical idea in the taxation literature. For example, [Akerlof \(1978\)](#) describes how “tagging” is used in the tax system for redistributive purposes. The interpretation of Pareto weights in our model is also closely analogous to how they are used in public finance; specifically [Saez and Stantcheva \(2016\)](#) introduced generalized social marginal welfare weights in the context of optimal tax theory and interpreted them as the value that society puts on providing an additional dollar of consumption to any given individual.

Although we analyze a model with monetary transfers, in the special case when the designer attaches no weight to revenue maximization, our framework becomes mathematically equivalent to a costly-signaling (“money-burning”) setting, in which assortative matching can only be obtained when agents engage in a socially wasteful activity to separate from lower types. Several papers have analyzed conditions on the distribution of willingness to pay under which screening maximizes total surplus (see, for example, [Hartline and Roughgarden, 2008](#), [Condorelli, 2012](#), [Chakravarty and Kaplan, 2013](#)).¹ Roughly, assortative matching is optimal when the inverse hazard rate is non-decreasing (this requires that the willingness to pay is unbounded from above), and a fully random allocation is optimal in the opposite case. Related results appear in the literature on matching contests, e.g., [Damiano and Li \(2007\)](#), [Hoppe et al. \(2009\)](#), and [Olszewski and Siegel \(2019\)](#), as well as two-sided matching markets, e.g., [Gomes and Pavan \(2016, 2018\)](#). We contribute to this literature by adding the Pareto weights that depend on agent’s willingness to pay, and showing how their monotonicity and dispersion interacts with the monotonicity of the inverse hazard rate.

Our paper also relates to a large literature in economics on price controls as a redistributive tool. [Viscusi et al. \(2005\)](#) study how price regulations can lead to allocative costs. [Bulow and Klemperer \(2012\)](#), meanwhile, characterized when price controls can be harmful to all market participants. Public provision of goods as a form of redistribution has also been studied before, see, e.g., [Besley and Coate \(1991\)](#); [Blackorby and Donaldson \(1988\)](#); [Gahvari and Mattos \(2007\)](#).

While certain additional steps are needed to accommodate several features of our framework, the methods we use to solve the “within” problem are not novel and can be seen as a generalization of the ironing technique developed by [Myerson \(1981\)](#). Following the intuitive approach to ironing developed by [Bulow and Roberts \(1989\)](#), [Hartline and Roughgarden \(2008\)](#) apply to it to a problem with multiple goods, and [Condorelli \(2012\)](#) to multiple

¹Similar conditions are obtained as early as [McAfee and McMillan \(1992\)](#) in a setting where bidders collude but cannot share payments among each other; then, bidding in the auction becomes equivalent to burning utility.

goods with heterogeneous quality. More recently, [Muir and Loertscher \(2020\)](#) apply similar techniques to solve a problem of a revenue-maximizing seller in the presence of resale; [Kleiner, Moldovanu, and Strack \(2020\)](#) show that these procedures can be obtained as a special case of a general property of extreme points that arise in optimization problems involving majorization constraints.

2 The model

Framework. A designer allocates a set of objects of heterogeneous quality to a set of agents differing in both their observable characteristics and their willingness to pay for quality. There is a unit mass of agents, with each agent characterized by a type $r \in \mathbb{R}_+$, and a label $i \in I$, where I is a finite set. There are $\mu_i > 0$ agents in group i . Let G_i denote the cumulative distribution function of r in group i , with a continuous density denoted by g_i , fully supported on $[\underline{r}_i, \bar{r}_i]$. We assume that i is observed by the designer, while r is private information of the agent. There is a mass $\mu \leq 1$ of objects, with each object characterized by a one-dimensional quality $q \in Q \subseteq [0, 1]$, where Q is a compact set. Let F denote the (generalized) cumulative distribution function of q , that is, $F(q)$ is the total mass of objects of quality equal to or less than q . If an agent with type r is assigned a good with quality q in exchange for a monetary transfer t , that agent’s utility is $rq - t$. We will refer to r as the *willingness to pay*.

Our framework incorporates a few strong assumptions about the environment. First, we assumed that agent’s utility is quasi-linear in money. Second, agents differ only in their “intensity” of preferences but they agree on the ranking of qualities. Third, each agent’s utility only depends on the *expected* quality of the good—agents are risk neutral.²

Because not being assigned any object is equivalent to being assigned an object of quality 0, we can without loss of generality assume that $F(1) = 1$ by putting enough mass at $q = 0$. Thus, all agents receive an object (but possibly with quality 0).

Assignments and Mechanisms. An assignment Γ is a collection of $|I|$ measurable functions $\Gamma_i : [\underline{r}_i, \bar{r}_i] \rightarrow \Delta(Q)$ with $\Gamma_i(q|r)$ interpreted as the probability that an agent with

²That said, we can incorporate risk aversion into our model by defining the agent’s utility to be $rv(q) - t$ for some concave function v . In that case, we could define a new random variable $\tilde{q} = v(q)$ with CDF \tilde{F} , called “risk-adjusted quality,” and apply our results to the model with \tilde{F} as the primitive distribution of quality.

type r in group i is assigned an object with quality q or less.³ The assignment Γ is *feasible* if

$$\Gamma_i(\cdot|r) \text{ is a CDF, } \forall i, \forall r \in [\underline{r}_i, \bar{r}_i]; \quad (2.1)$$

$$\sum_{i \in I} \mu_i \int_{\underline{r}_i}^{\bar{r}_i} \Gamma_i(q|r) dG_i(r) \geq F(q), \forall q \in Q. \quad (2.2)$$

Condition (2.2) states that the distribution of *assigned* qualities is first-order stochastically dominated by the distribution of *available* qualities. The condition reflects the availability of free disposal—a decrease in quality can be achieved by randomizing between a given quality and quality 0. Because the utility of agents only depends on the expected quality, it will be convenient to denote

$$Q^{\Gamma_i}(r) = \int_0^1 q d\Gamma_i(q|r).$$

We will write $Q_i(r)$ if the reference to the underlying assignment Γ_i is irrelevant.

To describe feasible mechanisms, we rely on the Revelation Principle. A direct mechanism $(\Gamma_i, t_i)_{i \in I}$ asks agents to report their rates of substitution r , assigns objects according to $\Gamma_i(q|r)$ in group i , and charges agents according to the transfer function $t_i(r)$. We assume that transfers are non-negative: $t_i(r) \geq 0$ for all i and r . This is equivalent (for incentive-compatible mechanisms) to the condition that the utility \underline{U}_i of the lowest type \underline{r}_i within each group i satisfies $\underline{U}_i \leq Q^{\Gamma_i}(\underline{r}_i)\underline{r}_i$. This assumption rules out direct lump-sum cash transfers to agents, and is needed for existence of optimal mechanism when no budget constraint is imposed. However, as we will explain later, lump-sum transfers can be accommodated in our analysis through an appropriate assumption on the designer's objective function (intuitively, lump-sum transfers to a group of agents are effectively allowed if the designer's payoff from generating a dollar of revenue is equal to the value of giving a dollar to all agents in that group).

Formally, a mechanism $(\Gamma_i, t_i)_{i \in I}$ is *feasible* if

- Γ is a feasible assignment, i.e., it satisfies conditions (2.1)-(2.2);
- Each agent reports her rate of substitution truthfully:

$$rQ^{\Gamma_i}(r) - t_i(r) \geq rQ^{\Gamma_i}(\hat{r}) - t_i(\hat{r}), \forall i, r, \hat{r}; \quad (2.3)$$

- Each agent receives non-negative utility from the mechanism but does not receive a

³We assume that all agents with the same type r are assigned the same lottery over objects.

positive money transfer:

$$0 \leq \underline{U}_i \leq Q^{\Gamma_i}(r_i)r_i, \forall i. \quad (2.4)$$

By standard arguments (see [Myerson, 1981](#)), we can represent the utility of an agent with type r in an incentive-compatible individually-rational mechanism as

$$U_i(r) \equiv rQ^{\Gamma_i}(r) - t_i(r) = \underline{U}_i + \int_{r_i}^r Q^{\Gamma_i}(\tau)d\tau. \quad (2.5)$$

Moreover, a mechanism is feasible if and only if Γ is a feasible assignment, $Q^{\Gamma_i}(r)$ is non-decreasing in r for all i , and $t_i(r)$ satisfies (2.5) for some $\underline{U}_i \in [0, Q^{\Gamma_i}(r_i)r_i]$.

The objective function. We assume that the designer maximizes a weighted sum of revenue and social surplus, for an arbitrary set of Pareto weights:

$$\alpha \underbrace{\sum_{i \in I} \mu_i \left(\int_{r_i}^{\bar{r}_i} t_i(r) dG_i(r) \right)}_{\text{revenue}} + \sum_{i \in I} \mu_i \underbrace{\left(\int_{r_i}^{\bar{r}_i} \lambda_i(r) U_i(r) dG_i(r) \right)}_{\text{social surplus with weights } \lambda_i},$$

where $\alpha \geq 0$ and $\{\lambda_i(r) \geq 0 : i \in I, r \in [r_i, \bar{r}_i]\}$ is the set of Pareto weights. We assume throughout that $\lambda_i(r)$ is continuous in r for each i . A simple calculation shows that this objective function can be represented by a function of the form

$$\sum_{i \in I} \mu_i \left(\int_{r_i}^{\bar{r}_i} V_i(r) Q^{\Gamma_i}(r) dG_i(r) + v_i \underline{U}_i \right), \quad (2.6)$$

by setting $v_i = \bar{\lambda}_i - \alpha$, where $\bar{\lambda}_i = \int_{r_i}^{\bar{r}_i} \lambda_i(\tau) dG_i(\tau)$ is the average Pareto weight for group i , and

$$V_i(r) = \alpha J_i(r) + \Lambda_i(r) h_i(r),$$

where $h_i(r) = (1 - G_i(r))/g_i(r)$ denotes the inverse hazard rate of G_i , $J_i(r) = r - h_i(r)$ is the virtual surplus function, and

$$\Lambda_i(r) = \mathbb{E}_{\tilde{r} \sim G_i}[\lambda(\tilde{r}) | \tilde{r} \geq r],$$

is the average Pareto weight attached to agents with types above r . Previous work has analyzed special cases of this objective function. With fully transferable utility, the objective function reduces to $J_i(r) + h_i(r) = r$ yielding the usual measure of allocative efficiency. With

$\alpha = 0$ and constant Pareto weights, the objective function reduces to $h_i(r)$, and the designer maximizes total agent surplus under “money burning.” In our general setting, $h_i(r)$, which is a measure of information rents, is weighted by the function $\Lambda_i(r)$ representing the Pareto weights.

Our objective function is quite general but has important limitations within the context of redistribution. Primarily, the approach of using exogenous Pareto weights reflects the assumption that the designer takes inequality as given. With this formulation, she cannot express preferences over the inequities created by the mechanism itself. In particular, we do not accommodate quotas that control the overall fairness of the outcome, and are popular in some contexts, such as school choice (see for example [Bodoh-Creed and Hickman, 2018](#)).

Interpretation. As discussed in the introduction, we think of α —the weight on revenue—as representing the marginal social value (or opportunity cost) of a dollar spent by the designer on some cause. For example, if a city mayor designs a public housing program, the revenue she raises can be used to subsidize the city budget, or invest in the construction of new homes.

Several special cases are of particular interest. We will refer to the problem with $\alpha = \bar{\lambda}_i$ as *internal redistribution within group i* , and to the problem with $\alpha = \bar{\lambda} := \sum_i \mu_i \bar{\lambda}_i$ simply as *internal redistribution*. The interpretation is that a dollar of revenue has the same worth to the designer as giving a dollar to a randomly selected agent within group i , or to a randomly selected agent from the entire population, respectively. This is mathematically equivalent to allowing lump-sum payments to agents in group i , or all agents, respectively.

We view the internal redistribution case as particularly natural when the set of agents is representative of the entire (local) population, so that we can think of the designer’s revenue as being the property of the agents. At the same time, this case rules out lump-sum payments to “preferred” groups i that are subject to affirmative action, understood as $\bar{\lambda}_i > \bar{\lambda}$. This can reflect political constraints, or an unmodeled inability of the designer to prevent agents from pretending that they have label i (e.g., if the label i denotes the income level, agents can misrepresent their income, or even purposefully decrease it to be eligible for a free cash payment).

More generally, the assumption of non-negative transfers has bite whenever $\bar{\lambda}_i > \alpha$. One interpretation is that lump-sum payments are allowed but there are frictions (such as administrative costs) that decrease their marginal value. In the extreme case $\alpha = 0$, our model becomes mathematically equivalent to a costly-screening (utility-burning) model in which an agent’s payment to the designer is more appropriately interpreted as a costly activity (such as standing in a queue) that is socially wasteful.

On the other hand, when $\alpha > \bar{\lambda}_i$, the designer has a higher value from spending a dollar outside of group i . This could be because there is another group j with $\bar{\lambda}_j > \bar{\lambda}_i$ to which the designer can give a lump-sum payment. Another interpretation is that the designer in this case is a private firm that maximizes profits but is potentially concerned to some extent about the welfare of her customers.

We will say that the good is *essential* for group i if $\underline{r}_i > 0$, that is, the agents' willingness to pay is bounded away from 0. This assumption should be economically interpreted as saying that the distribution G_i is mostly concentrated on values of r above \underline{r}_i rather than that there are literally no agents with values below \underline{r}_i (our results that assume essential goods continue to hold for distributions that attach a small enough mass to $r \in [0, \underline{r}_i]$).

3 Optimal mechanisms

The optimal mechanism is found in two steps:

1. The objects are allocated “across” groups: F is split into I CDFs F_i^* ;
2. The objects are allocated “within” groups: For each label i , the objects F_i^* are allocated optimally according to the expected-quality schedule Q_i^* .

3.1 The “within” problem.

In this step, we assume that F_i is the CDF of object qualities that are to be allocated to agents with label i . Formally, we will refer to the within problem for group i as maximizing (2.6) subject to feasibility with $I = \{i\}$, $\mu_i = 1$, and $F = F_i$. For a function Ψ , let $\text{co}(\Psi)$ denote the concave closure of Ψ (the point-wise smallest concave function that bounds Ψ from above) and let $\text{dco}(\Psi)$ denote the decreasing concave closure of Ψ (the point-wise smallest concave decreasing function that bounds Ψ from above).

We will say that there is *assortative matching* among types $r \in [a, b]$, if $Q_i^*(r) = F_i^{-1}(G_i(r))$ for all $r \in [a, b]$. We will say that there is *random matching* among types $r \in [a, b]$ if $Q_i^*(r) = \bar{q}$ for some \bar{q} and all $r \in [a, b]$.⁴

Remark 1. Because we have not imposed any assumptions on F (for example, we haven't ruled out degenerate distributions of quality), assortative and random matching could actually coincide (if F_i is constant in the relevant range). In particular, the two concepts do not

⁴Throughout, $H^{-1}(x)$ denotes the generalized inverse of a right-continuous non-decreasing function H on $[a, b]$: $H^{-1}(x) = \min\{y \in [a, b] : H(y) \geq x\}$, for all $x \leq \max_y H(y)$.

differ when all types in a given interval are not allocated any objects. The distinction between random and assortative matching can be guaranteed to be meaningful for each group i by assuming that $F(0) = 0$, F is continuous, and it is optimal to allocate all objects within group i which is implied by $\int_{r_i}^{r_i} (\alpha J_i(r) + \Lambda_i(r) h_i(r)) dG_i(r) \geq 0$ for all r_i .

Theorem 1. *Define*

$$\Psi_i(t) := \int_t^1 V_i(G_i^{-1}(x)) dx + \max\{0, \bar{\lambda}_i - \alpha\} r_i \mathbf{1}_{\{t=0\}}.$$

The value of the within problem for group i is given by

$$\int_0^1 dco(\Psi_i)(F_i(q)) dq,$$

An optimal solution is given by an expected-quality schedule

$$Q_i^*(r) = \Phi_i^*(G_i(r)) \mathbf{1}_{\{r \geq G_i^{-1}(x_i^*)\}},$$

where $[0, x_i^]$ is the maximal interval on which $dco(\Psi_i)$ is constant, and Φ_i^* is non-decreasing and satisfies*

$$\Phi_i^*(x) = \begin{cases} \frac{\int_a^b F_i^{-1}(y) dy}{b-a} & \text{if } x \in (a, b) \text{ and } (a, b) \text{ is a maximal interval on which } co(\Psi_i) \text{ is linear,} \\ F_i^{-1}(x) & \text{otherwise,} \end{cases}$$

*for almost all x .*⁵

Moreover, it is optimal to set $\underline{U}_i = 0$ when $\alpha \geq \bar{\lambda}_i$, and $\underline{U} = Q_i^(r_i) r_i$ when $\alpha \leq \bar{\lambda}_i$.*

As mentioned in the literature review, the proof of Theorem 1 uses relatively standard techniques known as “generalized ironing” that extend Myerson’s methods to richer environments. For completeness, and because several features of our setting (primarily the non-negativity of transfers and the continuous distribution of quantity) require these methods to be adjusted, we present a complete argument in the appendix. In the proof, we work with an arbitrary objective function of the form (2.6), not necessarily coming from maximizing a weighted sum of revenue and surplus.

The theorem describes a simple procedure to obtain a closed-form solution to the within problem:

⁵An interval (a, b) is a maximal interval on which $co(\Psi_i)$ is linear if $co(\Psi_i)$ is linear on (a, b) and no interval $(c, d) \supsetneq (a, b)$ has that property.

1. Compute the function Ψ_i that is a non-linear transformation of the original objective function. A noteworthy feature of Ψ_i is that it incorporates the constraint that transfers are non-negative: Whenever $\bar{\lambda}_i > \alpha$, this constraint must bind, and hence \underline{U}_i is set to the maximal feasible level $Q_i(\underline{r}_i)\underline{r}_i$. In the transformed objective function Ψ_i , this corresponds to an upward jump at 0. (This potential discontinuity of Ψ_i at 0 will be responsible for one of our key results in the next section.)
2. Compute the concave closure $\text{co}(\Psi_i)$ and the decreasing concave closure $\text{dco}(\Psi_i)$ of Ψ_i .
3. If on some initial interval $[0, x_i^*]$, $\text{co}(\Psi_i) < \text{dco}(\Psi_i)$, then objects of quality below the x_i^* quantile of F_i are not allocated (the designer uses the free disposal option), and hence agents with willingness to pay below $r_i^* = G_i^{-1}(x_i^*)$ are assigned quality 0. This can only happen if Ψ_i is not decreasing everywhere, which requires $V_i(r)$ to be negative for some r , which is possible when $J_i(r) < 0$.
4. The remaining object qualities are partitioned into intervals; the remaining agents are partitioned in the order of increasing willingness to pay to match the mass of objects within each interval; whenever $\text{co}(\Psi_i)$ is linear on a (maximal) interval, the matching between types and quality is (fully) random within that interval; whenever $\text{co}(\Psi_i)$ is strictly concave on an interval, the matching between types and quality is assortative.

The function Ψ_i plays a key role in determining the properties of the optimal mechanism. To gain intuition, we can use integration by parts and substitution, and obtain that for any $r > \underline{r}_i$,

$$\Psi_i(G_i(r)) = \int_r^{\bar{r}_i} \tau \lambda_i(\tau) dG_i(\tau) + (\alpha - \Lambda_i(r))r(1 - G_i(r)). \quad (3.1)$$

Thus, the value of Ψ_i at some quantile $x = G_i(r)$, is simply the payoff to the designer from selling quality 1 at a price of r .

3.2 The “across” problem.

Based on the solution to the within problem for each i separately, we can now formulate the “across” problem as

$$\max_{(F_i)_{i \in I}} \sum_{i \in I} \mu_i \int_0^1 \text{dco}(\Psi_i)(F_i(q)) dq, \quad (3.2)$$

$$\text{s.t.} \quad \sum_{i \in I} \mu_i F_i(q) = F(q), \quad \forall q \in Q. \quad (3.3)$$

Once the optimal F_i^* are found that solve the problem (3.2)–(3.3), the optimal solution within each group i is described by Theorem 1.

Our second technical result describes a solution procedure for the “across” problem. Let $s_i(x) \equiv |\text{dco}(\Psi_i)'(x)|$ denote the (absolute value of the) slope of $\text{dco}(\Psi_i)$ at quantile x .

Theorem 2. *There exists a non-decreasing non-negative function $S(q)$ such that for all i and q , the optimal solution $(F_i^*)_{i \in I}$ to the problem (3.2)–(3.3) satisfies*

$$\begin{cases} F_i^*(q) = 0 & \text{if } s_i(0) > S(q), \\ F_i^*(q) = 1 & \text{if } s_i(1) < S(q), \\ F_i^*(q) \text{ solves } s_i(F_i^*(q)) = S(q) & \text{otherwise.} \end{cases}$$

Moreover, $S(q) = \min_{i: F_i^*(q) < 1} s_i(F_i^*(q))$.

The proof is in the appendix. Intuitively, we solve the program (3.2)–(3.3) by considering a relaxed problem in which the constraint that $F_i(q)$ is a CDF is dropped, and later verifying that there exists a solution to the relaxed program that is feasible. The index $S(q)$ is the Lagrange multiplier on the resource constraint (3.3) for the relaxed problem.

The theorem describes a “greedy” procedure that allocates all objects to the $|I|$ groups. Roughly speaking, the algorithm can be seen as allocating the objects by gradually increasing the CDFs F_i^* , in the order of increasing slopes $s_i(\cdot) = |\text{dco}(\Psi_i)'(\cdot)|$. The function $S(q)$ keeps track of the running minimum over these slopes. Starting from the lowest quality, we first increase the CDF F_i^* for group i with the smallest slope s_i at 0 (in the case where there are several such groups, the proof of Theorem 2 describes how to break the ties). At any q , we increase the CDF of group(s) i with the lowest slope s_i at $F_i^*(q)$. That is, only groups i with $s_i(F_i^*(q)) = S(q)$ are allocated objects with quality q . When some $F_i^*(q)$ reaches 1, we stop increasing the CDF for that group. Of course, what allows the algorithm to be “greedy” is the fact that $\text{dco}(\Psi_i)$ is a non-increasing concave function for all i , and thus the slopes s_i are non-negative and non-decreasing.

4 Economic implications

Proposition 1. *If the average Pareto weight $\bar{\lambda}_i$ in some group i is strictly larger than the weight on revenue α , and the good is essential ($\underline{r}_i > 0$), then there exists $r_i^* > \underline{r}_i$ such that the allocation is random at a price of 0 for all types $r \leq r_i^*$.*

Our first result states that it is always optimal to allocate the objects randomly to the lowest-willingness-to-pay agents at a price of 0 if (i) the designer cares more about the

surplus of an average agent within group i than about revenue, and (ii) the good is essential. The first assumption is likely to hold in cases when label i is associated with preferential treatment or affirmative action but targeting a purely monetary transfer to group i is not feasible. For the internal redistribution problem ($\alpha = \bar{\lambda}$), we have $\alpha < \bar{\lambda}_i$ if the label i is associated with a group of agents that the designer wants to redistribute to; for example, if i could capture having income below some threshold, or being unemployed. The second assumption says that the good is essential, that is, the willingness to pay is bounded away from 0. This is likely to hold for goods like food, housing or health care.

Mathematically, the result is an immediate consequence of Theorem 1: The two assumptions imply that there is an upward jump in Ψ_i at 0, and thus the concave closure $\text{co}(\Psi_i)$ must be strictly above Ψ_i (and hence linear) in the neighborhood of 0, resulting in a random allocation.

To understand the result from an economic perspective, note that when $\bar{\lambda}_i > \alpha$, holding fixed the allocation, the designer would like to minimize the transfers that agents are paying in the mechanism. The non-negative transfers condition prevents the designer from giving a monetary transfer to agents directly, and implementing assortative matching requires prices to be increasing. Consider, instead, providing some goods for free to the lowest- r agents; this policy raises their utility if they value quality, that is, if all r are bounded away from 0. Proposition 1 predicts that the designer can always improve her overall objective this way. However, the reason why this policy is effective is *not* that the designer is concerned about the welfare of agents with the lowest willingness to pay; note that the only assumption we made is about the average Pareto weight that is consistent with, for example, placing no weight on agents with the lowest r . Rather, the correct intuition is that providing the goods for free to agents with the lowest willingness to pay allows the designer to lower prices—and hence increase utility—for all higher types. Of course, a caveat is that providing the goods for free precludes any screening of the corresponding types (reducing allocative efficiency and welfare). However, it can be shown that the reduction in allocative efficiency is always second-order compared to the benefits when the region of random matching is small (see Appendix A for a formal argument that confirms this intuition). The optimal mechanism determines the size of the random-allocation interval by trading off providing more goods for free against a decrease in allocative efficiency.

The next result describes conditions under which this trade-off is resolved towards full randomization. By Theorem 1, full randomization is optimal if and only if $\text{dco}(\Psi_i)$ is linear which is when the graph of Ψ_i lies below the line connecting $(0, \Psi_i(0))$ with $(1, \Psi_i(1))$:

$$\Psi_i(G_i(r)) \leq (1 - G_i(r)) \int_{r_i}^{\bar{r}_i} r \lambda_i(r) dG_i(r),$$

for all r . Using equation (3.1), we can interpret this condition as stating that the designer’s payoff from a price mechanism with price r (with quality normalized to 1) is smaller than the payoff from allocating the same set of objects uniformly at random at a price of 0, for any r . We focus on some more interpretable implications of this observation.

Proposition 2. *A necessary condition for full randomization to be optimal within group i is that*

$$\alpha \bar{r}_i \leq \int_{r_i}^{\bar{r}_i} r \lambda_i(r) dG_i(r). \quad (4.1)$$

A sufficient condition is (4.1) and quasi-convexity of $\alpha J_i(r) + \Lambda_i(r) h_i(r)$.

A first noteworthy aspect of Proposition 2 is that optimality of full randomization requires that the average Pareto weight $\bar{\lambda}_i$ is strictly higher than the weight on revenue α (this is a direct consequence of inequality (4.1)). In particular, if lump-sum transfers to group i are feasible for the designer, then a fully random allocation cannot be optimal. As discussed above, if α is taken to be equal to the average Pareto weight for all agents, then this implies that full randomization can be optimal only in groups associated with preferential treatment. The necessary condition comes from a hypothetical scenario in which the designer has only one (infinitesimal) unit of the object with quality 1 to allocate: It must be that the value of revenue from selling that object at a maximal price to the highest willingness-to-pay agent is smaller than the value of allocating this object uniformly at random at a price of 0.

Perhaps more surprisingly, this necessary condition becomes sufficient if $\alpha J_i(r) + \Lambda_i(r) h_i(r)$ is quasi-convex, which is a relatively permissive assumption. For example, $\alpha J_i(r) + \Lambda_i(r) h_i(r)$ could be increasing, so that Ψ_i is strictly concave on $(0, 1]$, and yet full randomization will be optimal ($\text{dco}(\Psi_i)$ will be linear) when the upward jump at 0 in Ψ_i is large enough. Economically, this means that the force identified in Proposition 1 can be strong enough (when $(\bar{\lambda}_i - \alpha) \underline{r}_i$ gets large) to induce a random allocation for all types.

The type of in-kind redistribution predicted by Proposition 2 is quite common in practice (e.g., food stamps). In this case, the label i can be seen as capturing criteria that an individual must satisfy to be eligible for such a program. The result says that for such form of redistribution to be optimal, several conditions must be met.⁶ First, the designer should not be able to give these “eligible” agents a direct lump-sum transfer, as only then it is possible that $\bar{\lambda}_i > \alpha$. Second, the weight on revenue (measuring how effectively it can be used for other purposes) should be small relative to the Pareto weights on agents that are eligible. Third, such programs are more likely to be optimal for essential goods; note that if the Pareto weights are non-increasing, the right-hand side of (4.1) is bounded above by

⁶We emphasize that our model does not capture important economic, social, and psychological reasons for which an in-kind distribution might be preferable to cash transfers.

$(1/2)(\bar{r}_i + \underline{r}_i)\bar{\lambda}_i$, and thus when $\underline{r}_i = 0$, the average Pareto weight must be at least twice as large as the weight on revenue. However, when \underline{r}_i is large, it may suffice that $\bar{\lambda}_i$ is only slightly above α .

Next, we show that when a fully random allocation fails to be optimal, assortative matching often becomes optimal at the top of the distribution of willingness to pay, that is, for high enough types r . For this result, we assume that Pareto weights are non-increasing, an assumption that can be justified if the designer views willingness to pay as being positively correlated with wealth, and has preference for redistribution towards poorer agents (see DKA for a detailed discussion).

Proposition 3. *If Pareto weights are non-increasing, and $\alpha \geq \bar{\lambda}_i$, then there is assortative matching at the top.*

The result is intuitive: Non-increasing Pareto weights along with the assumption that the weight on revenue is weakly larger than the average Pareto weight, imply that the weight on revenue is larger than the weight on the utility of agents with high willingness to pay. Since assortative matching is optimal for revenue maximization at the top of the distribution (the so-called “no distortion at the top” result), it dominates random allocation for high enough r .

The condition $\alpha \geq \bar{\lambda}_i$ rules out the force behind random matching in Proposition 1. However, we show that even though assortative allocation may be optimal at the top of the distribution, it will still be often optimal to use a random matching for some set of types.

Proposition 4. *Assortative matching (of objects that are allocated) is optimal if and only if (i) either $\alpha \geq \bar{\lambda}_i$ or $\underline{r}_i = 0$ and (ii) $\alpha J_i(r) + \Lambda_i(r)h_i(r)$ is non-decreasing.*

By Theorem 1, assortative matching is optimal when Ψ_i is a concave function.⁷ This requirement is inconsistent with an upward jump of Ψ_i at 0 which is what the first assumption rules out. The second condition says that the derivative of Ψ_i is non-increasing. Together with Proposition 2, we obtain the conclusion that $\bar{\lambda}_i > \alpha$ is necessary for fully random allocation while $\bar{\lambda}_i \leq \alpha$ is necessary for a fully assortative allocation when goods are essential.

Proposition 4 relates to multiple results from the literature (see Hartline and Roughgarden, 2008, Condorelli, 2012, Chakravarty and Kaplan, 2013) that predict that in the costly-signaling setting ($\alpha = 0$), assortative allocation maximizes unweighted agent surplus when the inverse hazard rate is non-decreasing. Proposition 4 extends this condition to the case when surplus is weighted by the Pareto weights—it is now required that the product $\Lambda_i(r)h_i(r)$ of the inverse hazard rate at r and the average Pareto weight on all types above r is

⁷Strictly speaking, Theorem 1 only implies this conclusion when Ψ_i and hence $\text{co}(\Psi_i)$ are strictly concave; however, when $\text{co}(\Psi_i)$ is linear and equal to Ψ_i on an interval, then both assortative and random matching are optimal on that interval. Thus, the conclusion extends to weak concavity.

non-decreasing.⁸ Our paper also points out that this result is only true under the assumption $\underline{r}_i = 0$ (which is made in all of the above papers).

To show why fully assortative matching will often fail to be optimal, note that the second condition in Proposition 4 is violated (and hence some random allocation is optimal) as long as for some r ,

$$\alpha + \Lambda'_i(r)h_i(r) + (\Lambda_i(r) - \alpha)h'_i(r) < 0.$$

Assume that the inverse hazard rate is non-increasing; this assumption is satisfied by many popular distributions with bounded support, and implies that assortative matching maximizes revenue. Fixing G_i , assortative matching will fail to be optimal for agents with willingness to pay close to r if either (i) the average Pareto weight on types above r is sufficiently greater than the weight on revenue, or (ii) the Pareto weights are declining sufficiently fast with r . That last condition can be interpreted as saying that the designer has strong redistributive preferences within group i in that she puts a significantly higher weight on agents with low willingness to pay.

DKA show that when Pareto weights are decreasing in r , competitive pricing (which is a special case of our assortative matching) is optimal when the dispersion in Pareto weights is low (there is *low inequality*), and rationing may become optimal when it is high. They allow for lump-sum transfers, but not for exogenous value for revenue, and do not consider labels—this amounts to assuming that $\alpha = \bar{\lambda}_i$. In our more general setting, under the assumption of non-increasing inverse hazard rate, non-increasing Pareto weights, and $\bar{\lambda}_i \leq \alpha$, observe that

$$\frac{d}{dr} [\alpha J_i(r) + \Lambda_i(r)h_i(r)] = \alpha + \Lambda_i(r) - \lambda_i(r) + (\Lambda_i(r) - \alpha)h'_i(r) \geq \alpha + \Lambda_i(r) - \lambda_i(r),$$

and thus a sufficient condition for assortative matching is that $\alpha \geq \max_r \{\lambda_i(r) - \Lambda_i(r)\}$. Note that $\max_r \{\lambda_i(r) - \Lambda_i(r)\}$ measures both the level and dispersion of Pareto weights; for example, when the weights are constant, it is equal to 0, and hence assortative matching is optimal. In the special case when $\max_r \{\lambda_i(r) - \Lambda_i(r)\} = \lambda_i(\underline{r}_i) - \Lambda_i(\underline{r}_i)$ and $\alpha = \bar{\lambda}_i$, this condition reduces to the low-inequality condition from DKA.

So far, we have argued that a dispersion in Pareto weights is likely to lead to suboptimality of fully assortative matching but we have not analyzed how far the optimal mechanism deviates from being fully assortative. Our next result establishes a (perhaps surprising) conclusion that, when $\alpha \geq \bar{\lambda}_i$, the set of types for which the allocation is random must shrink both when the weights are approximately constant and when they get extremely skewed towards low types. Thus, there is a certain non-monotonicity of the use of random

⁸Because we assume bounded support of willingness to pay, we have $h_i(\bar{r}_i) = 0$, and thus assortative matching is never optimal in our model when $\alpha = 0$.

allocation in the level of inequality.

Proposition 5. *Consider a sequence of non-increasing Pareto weights $\lambda_i^n(r)$ such that $\alpha \geq \bar{\lambda}_i^n$, and assume that $J_i'(r) \geq \underline{J}_i > 0$. If either*

- *for all r and $\epsilon > 0$, $|\lambda_i^n(r) - \bar{\lambda}_i^n| < \epsilon$ for large enough n , or*
- *for all $r > \underline{r}$ and $\epsilon > 0$, $\lambda_i^n(r) < \epsilon$ for large enough n ,*

then any optimal solution converges point-wise to assortative matching (of objects that are allocated).

The first case is not surprising since we know from the preceding discussion that when $\alpha \geq \bar{\lambda}_i$, a constant Pareto weight leads to optimality of assortative matching (however, Proposition 5 relaxes the assumption that the inverse hazard rate is non-increasing to a weaker assumption that the virtual surplus function is strictly increasing).

The second case looks at the opposite extreme in which the designer is concerned exclusively with the welfare of agents with the lowest willingness to pay (while the average weight does not increase above α). While the optimal mechanism may (and typically will) use a random allocation for the lowest types, the randomization region vanishes as the weights become increasingly skewed. The utility of the lowest type converges to the utility she receives in the assortative matching (which is 0 if $\alpha > \bar{\lambda}_i$) even though her Pareto weight diverges to infinity. The key to understanding this result is to recall that the *average* Pareto weight stays below the weight on revenue, and that the allocation to the lowest type constrains the allocation that can be given to all types, because of IC constraints. While the designer could maximize the welfare of the lowest type by giving her a random good at a price of 0, this would necessarily decrease revenue to 0. As the weight on the lowest types increases, the weight on all higher types converges to 0 which makes assortative matching (which maximizes revenue) increasingly attractive, and optimal in the limit.

Proposition 5 suggests that the use of random allocation is maximized (under $\alpha \geq \bar{\lambda}_i$) when the dispersion in Pareto weights is intermediate.

So far we have focused on allocation within the groups. We now focus on what the above insights, combined with Theorem 2, tell us about the allocation of objects *across* the groups. We begin with simple cases in which allocation takes the same form in all groups.

Proposition 6. *Suppose all groups have fully random matching. Relabel the groups so that lower $i = 1, \dots, |I|$ corresponds to lower $\int_{\underline{r}_i}^{\bar{r}_i} \tau \lambda_i(\tau) dG_i(\tau)$. Then, there exists an optimal mechanism in which group i receives qualities belonging to $[q_i, q_{i+1}] \cap \text{supp}(F)$, where $q_1 = \min \text{supp}(F)$ and $q_{I+1} = \max \text{supp}(F)$.*

The result is straightforward. Mathematically, if there is fully random matching within groups, the slope of $\text{dco}(\Psi_i)$ is constant for each i and equal to $\Psi_i(0)$. Thus, groups with higher $\Psi_i(0)$ receive uniformly higher quality objects. Economically, the conclusion follows from the observation that the designer's marginal value from allocating a unit of quality to group i is constant and equal to $\int_{\underline{r}_i}^{\bar{r}_i} \tau \lambda_i(\tau) dG_i(\tau)$.

Proposition 7. *Suppose all groups have assortative matching. Relabel the groups so that lower $i = 1, \dots, |I|$ corresponds to lower $\alpha \bar{r}_i$. Then there exists an optimal mechanism in which group i receives qualities $[\underline{q}_i, \bar{q}_i] \cap \text{supp}(F)$, where $\bar{q}_i \leq \bar{q}_{i+1}$. Moreover, if $\underline{r}_i = 0$ and $\alpha > \bar{\lambda}_i$ for all i , then all $\underline{q}_i = \min \text{supp}(F)$.*

Note that in the special case of the result when all groups have the same support of willingness to pay, with lower bound equal to 0, then they also have the same support of quality in the optimal allocation. This is in some sense the opposite of the conclusion for random matching, where the quality levels allocated across groups are almost disjoint.

Mathematically, the result follows from Theorem 2 and a few observations. First, $\text{dco}(\Psi_i)$ is a concave function with a continuous derivative, so it follows immediately that in the greedy algorithm based on Theorem 2, the quality levels allocated to a given group are an interval within the support of the distribution of quality F . Moreover, the lower and upper bounds of the interval are pinned down by the slope of $\text{dco}(\Psi_i)$ at 0 and at 1, respectively. Second, when there is assortative matching, the function Ψ_i coincides with its decreasing concave closure, except potentially in some initial interval. Because the derivative of Ψ_i at 1 is equal to $-\alpha \bar{r}_i$, \bar{r}_i determines the allocation of the “last” (highest-quality) good allocated in the greedy algorithm. Finally, if $\underline{r}_i = 0$ and $\alpha > \bar{\lambda}_i$ (note that at least one of these conditions must hold, by Proposition 4), then in all groups some objects are not allocated, and thus all $\text{dco}(\Psi_i)$ have a zero slope in some initial interval. Thus, all groups are allocated the lowest quality objects.

The conclusion about the upper bound on the quality allocated to each group continues to hold as long as there is assortative matching at the top. That is, if there is assortative matching at the top in all groups, the highest-quality object is allocated to the group with the highest maximal willingness to pay.

It may seem surprising that the allocation of the highest-quality object does not depend on the Pareto weights within group i , even though these Pareto weights (including on the highest type) could be as large as the weight on revenue. For example, if \bar{r}_i is only slightly higher than \bar{r}_j , then group i still gets the highest-quality object even if the designer puts no weight on the welfare of agents in group i , and puts weight α on the welfare of agents in group j . The resolution of this puzzle lies in the realization that the utility of the highest

type \bar{r}_i is fully pinned down by the allocation to lower types $r < \bar{r}_i$ within her group (see the envelope formula (2.5)). Since there is assortative matching at the top, when the highest-quality object is allocated, it always goes to the agent with the highest r within her group; however, whatever the quality of this object is, the highest- r type's utility is the same – higher quality simply translates to a higher price. This implies that the allocation at the top of the distribution only affects the designer's revenue, and hence the highest-quality object is allocated to the group with the highest maximal type \bar{r}_i .

A more interesting case arises when some groups have a fully random allocation, while others have assortative matching at the top.

Proposition 8. *Suppose that all groups have the same upper bound \bar{r}_i of the support of willingness to pay. Then, any group with fully random allocation gets uniformly higher quality than any group for which there is assortative matching at the top.*

The results follows from Theorem 2 and two observations. First, if there is assortative matching at the top in group i , then the slope of $|\text{dco}(\Psi_i)'(1)|$ is equal to $\alpha\bar{r}_i$, and since $\text{dco}(\Psi_i)$ is concave, this is the highest slope for group i . If there is fully random allocation in group j , then the slope of $|\text{dco}(\Psi_j)'(1)|$ must be weakly higher than $\alpha\bar{r}_j = \alpha\bar{r}_i$, and since the slope of $\text{dco}(\Psi_j)$ is constant, it is uniformly higher than the slope of $\text{dco}(\Psi_i)$. Hence, group i gets uniformly higher quality in the optimal allocation.

In practice, when a certain group of agents receives the good for free (as in some social housing or food stamps programs), typically the quality of these goods is lower than the quality in the “market.” Thus, we view the value of Proposition 8 as indicating that its assumption of equal upper bounds of willingness to pay should be violated in order to avoid its extreme conclusion (of course, there may be many outside-of-the-model reasons for why the conclusion of Proposition 8 could fail). Loosely interpreted, this indicates that if the designer defines the label i of agents eligible for a free allocation of the good, then this category should not include agents with high willingness to pay.

Our final result indicates, however, that groups of agents that receive the object for free should rather be allocated intermediate-quality goods in general.

Proposition 9. *Suppose that there are two groups and that in group 1 there is assortative matching (of allocated objects), and in group 2 there is fully random matching. Then, there exist $\underline{q} \leq \bar{q}$ such that group 2 gets objects of quality $[\underline{q}, \bar{q}]$, and group 1 gets objects of quality $[0, \underline{q}] \cup (\bar{q}, 1]$.*

Moreover, $\bar{q} < 1$ if $\alpha\bar{r}_1 \geq \int_{r_2}^{\bar{r}_2} \tau \lambda_2(\tau) dG_2(\tau)$, and $\underline{q} > 0$ when $\alpha r_{-1} \leq \int_{r_2}^{\bar{r}_2} \tau \lambda_2(\tau) dG_2(\tau)$.

Proposition 9 follows immediately from Theorem 2 by observing that the slope of $\text{dco}(\Psi_2)$ is constant, while the (absolute value of the) slope of $\text{dco}(\Psi_1)(q)$ is increasing in q . The last

part of the result gives sufficient conditions for the (constant) slope of $|\text{dco}(\Psi_2)|$ to lie strictly between $|\text{dco}(\Psi_1)'(0)|$ and $|\text{dco}(\Psi_1)'(1)|$.

The economic intuition behind the result is as follows. When a random-allocation mechanism is used for group 2, the designer’s payoff depends only on expected quality allocated to that group—this is a consequence of the fact that the price does not depend on the quality in this case (it is 0). In contrast, when assortative matching is used, the designer’s payoff depends crucially on the dispersion in quality— that is why she allocates both the lowest- and highest-quality objects to group 1. This is particularly intuitive in the context of revenue maximization in which the seller lowers the allocation of low types in order to decrease the information rents of the high types (see [Myerson, 1981](#)). In fact, it will often be optimal not to allocate some objects in group 1, in which case the marginal value of quality is 0 up to some point (in the greedy algorithm described in [Theorem 2](#)). However, the marginal value of quality allocated to agents with high willingness to pay in group 1 may be large, especially if \bar{r}_1 is high (in line with [Proposition 8](#)).

5 Discussion

Our analysis leads to a number of general insights on the design of allocation mechanisms under redistributive concerns. Specifically, we can characterize how different types of social preferences influence the optimal design. In the discussion below, we assume that willingness to pay is positively correlated with wealth. We will abuse terminology slightly by referring to agents with low willingness to pay as “poor,” and agents with high willingness to pay as “wealthy.” While this assumption may not hold in various settings (for example, if the good to be allocated is an inferior good), it is fairly natural and will make our arguments more intuitive.

1. *In-kind redistribution and the impact of the weight on revenue, α* : When a dollar of revenue has welfare value equal to (or greater than) that of transferring a dollar to agents directly (which is the case, in particular, when lump-sum payments are feasible), raising revenue is often a more preferred method of redistribution relative to random allocation of goods; this drives towards assortative matching ([Proposition 4](#)). Thus, for example, if the market is mediated by a government agency (or heavily taxed), we may prefer assortative matching in order to maximize the available revenue for redistribution. We always, for example, want assortative matching at the top for this reason, so long as we value revenue sufficiently and weight the welfare of the poor more than that of the rich ([Proposition 3](#)).

By contrast, when the weight on revenue is low relative to the average welfare weight—in particular, when lump-sum payments to the target group are not feasible—full randomization may be optimal as a way of redistributing in-kind to the poor (Proposition 2). That said, the conditions for optimality of full randomization are quite strong. Often, even when actively redistributing to a specific group, the designer might still want to introduce a price gradient, both to ensure a more efficient matching between agents and object qualities, and to raise some revenue.

2. *Essential vs. nonessential goods markets:* In essential goods markets—that is, settings in which the lowest willingness to pay in a group \underline{r} is positive—it is possible to raise the utility of the “poorest” agents by allocating objects randomly for free. Even when full randomization is not optimal, as long as the designer prefers to redistribute to that group (that is, the average Pareto weight is above the weight attached to revenue), in essential goods markets there should always be some amount of random allocation to the poorest agents in the market (Proposition 1). This allows the designer to lower prices for all agents in that group, effectively circumventing the constraint that she cannot give money to agents directly. More broadly, the optimal mechanism uses lotteries more when \underline{r} is higher. Thus, we may think of random allocation as a more appropriate policy response in essential goods settings such as markets for housing, food, and healthcare.
3. *Non-monotonicity of the share of objects allocated randomly:* In order to satisfy our incentive constraints, the rents enjoyed by higher types must always be greater than the rents enjoyed by lower types. As a consequence, it is never possible to redistribute *only* to the poorest buyers through allocation of the good. Thus, somewhat paradoxically, as our welfare weights skew more and more in favor of the lowest types, the optimal mechanism tends towards fully assortative matching (so long as we value revenue sufficiently): even though we would like to redistribute to the lower types, giving away the good at a price of 0 would necessarily bring revenue to 0 because of the incentive constraints (Proposition 5).

Put more simply, there is non-monotonicity: When we do not weight the welfare of the poor highly relative to that of the rich, assortative matching tends to be optimal as a way of increasing the welfare of the wealthy. As we increase the weight on the poor, we shift towards random allocation of the good as a way of increasing the poor’s welfare; but then, eventually, tend back to assortative matching to raise revenue once that becomes a more effective way of redistributing.

4. *Allocation across different groups*: The form of allocation within groups influences the optimal structure of allocation across groups. In particular, if objects are matched randomly within groups, then the optimal mechanism gives the different groups essentially disjoint levels of quality (Proposition 6). In contrast, if objects are matched assortatively within groups, the optimal mechanism involves giving the different groups objects of overlapping quality, with the maximum quality level determined solely by the maximum willingness to pay in each group (Proposition 7). More generally, among groups with assortative matching at the top, the highest-quality goods are allocated to the group with the highest maximum willingness to pay. Relatedly, if all groups have the same maximum willingness to pay, then groups that receive fully random allocation must also receive higher-quality objects because the use of random allocation implies a higher welfare weight on that group (Proposition 8). Conversely, if we introduce a policy under which some group of agents is eligible to receive the good for free—entertainment ticket giveaways for fans, for example—then we probably do not want that group to include agents with the highest willingness to pay.

Last, in allocation programs in which one group is matched assortatively and another group receives allocation by lottery, the optimal mechanism allocates *middle-quality goods* in the lottery (Proposition 9). While this may seem counterintuitive at first, it has a natural explanation: by keeping some low-quality goods in the pool used for assortative matching, the designer increases competition for the high quality goods, which increases the revenue raised through the assortative matching process. A loose implication, for example, is that it may be optimal for states to provide scholarships throughout their university systems rather than just for lower-tier schools, as the resulting constraints on the supply of top-tier positions increases the revenue that can be raised from students with high willingness to pay.

References

- AKERLOF, G. A. (1978): “The Economics of ”Tagging” as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning,” American Economic Review, 68, 8–19.
- BESLEY, T. AND S. COATE (1991): “Public provision of private goods and the redistribution of income,” American Economic Review, 81, 979–984.
- BLACKORBY, C. AND D. DONALDSON (1988): “Cash versus kind, self-selection, and efficient transfers,” American Economic Review, 78, 691–700.

- BODOH-CREED, A. L. AND B. R. HICKMAN (2018): “College assignment as a large contest,” Journal of Economic Theory, 175, 88 – 126.
- BROOKS, B. AND S. DU (2019): “Optimal auction design with common values: An informationally-robust approach,” Working Paper.
- BULOW, J. AND P. KLEMPERER (2012): “Regulated prices, rent seeking, and consumer surplus,” Journal of Political Economy, 120, 160–186.
- BULOW, J. AND J. ROBERTS (1989): “The Simple Economics of Optimal Auctions,” Journal of Political Economy, 97, 1060–1090.
- CHAKRAVARTY, S. AND T. R. KAPLAN (2013): “Optimal allocation without transfer payments,” Games and Economic Behavior, 77, 1 – 20.
- CONDORELLI, D. (2012): “What money can’t buy: Efficient mechanism design with costly signals,” Games and Economic Behavior, 75, 613 – 624.
- (2013): “Market and non-market mechanisms for the optimal allocation of scarce resources,” Games and Economic Behavior, 82, 582–591.
- DAMIANO, E. AND H. LI (2007): “Price discrimination and efficient matching,” Economic Theory, 30, 243–263.
- DIAMOND, P. A. AND J. A. MIRRLEES (1971): “Optimal taxation and public production I: Production efficiency,” American Economic Review, 61, 8–27.
- DWORCZAK, P. AND G. MARTINI (2019): “The Simple Economics of Optimal Persuasion,” Journal of Political Economy, 127, 1993–2048.
- DWORCZAK, P. (R) S. D. KOMINERS (R) M. AKBARPOUR (2019): “Redistribution through markets,” Becker Friedman Institute for Research in Economics Working Paper.
- GAHVARI, F. AND E. MATTOS (2007): “Conditional cash transfers, public provision of private goods, and income redistribution,” American Economic Review, 97, 491–502.
- GOMES, R. AND A. PAVAN (2016): “Many-to-many matching and price discrimination,” Theoretical Economics, 11, 1005–1052.
- (2018): “Price Customization and Targeting in Platform Markets,” Working Paper.

- HARTLINE, J. D. AND T. ROUGHGARDEN (2008): “Optimal mechanism design and money burning,” in Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, 75–84.
- HOPPE, H. C., B. MOLDOVANU, AND A. SELA (2009): “The Theory of Assortative Matching Based on Costly Signals,” Review of Economic Studies, 76, 253–281.
- KLEINER, A., B. MOLDOVANU, AND P. STRACK (2020): “Extreme Points and Majorization: Economic Applications,” Working Paper.
- KOLOTILIN, A. (2018): “Optimal Information Disclosure: A Linear Programming Approach,” Theoretical Economics, 13, 607–636.
- MCAFEE, R. P. AND J. MCMILLAN (1992): “Bidding Rings,” American Economic Review, 82, 579–599.
- MUIR, E. V. AND S. LOERTSCHER (2020): “Monopoly pricing, optimal randomization, and resale,” Working Paper.
- MÜLLER, A. AND D. STOYAN (2002): Comparison Methods for Stochastic Models and Risks, Wiley.
- MYERSON, R. B. (1981): “Optimal auction design,” Mathematics of Operations Research, 6, 58–73.
- OLSZEWSKI, W. AND R. SIEGEL (2019): “Pareto Improvements in the Contest for College Admissions,” Working Paper.
- RAY, D. (R) A. ROBSON (2018): “Certified Random: A New Order for Coauthorship,” American Economic Review, 108, 489–520.
- SAEZ, E. AND S. STANTCHEVA (2016): “Generalized social marginal welfare weights for optimal tax theory,” American Economic Review, 106, 24–45.
- VISCUSI, W. K., J. E. HARRINGTON, AND J. M. VERNON (2005): Economics of Regulation and Antitrust, MIT press.
- WEITZMAN, M. L. (1977): “Is the price system or rationing more effective in getting a commodity to those who need it most?” Bell Journal of Economics, 8, 517–524.

A A precise intuition for Proposition 1

For a more precise intuition for Proposition 1, consider Figure A.1 (we drop the subscript i to simplify notation). Suppose that the expected quality schedule $Q(r)$ is strictly increasing. Recall that utility of type r can be expressed as $U(\underline{r}) + \int_{\underline{r}}^r Q(\tau) d\tau$. Under the assumption $\bar{\lambda} > \alpha$, the designer wants to minimize prices subject to the constraint $t(r) \geq 0$, and hence we can assume $t(\underline{r}) = 0$ and $U(\underline{r}) = Q(\underline{r})\underline{r}$.

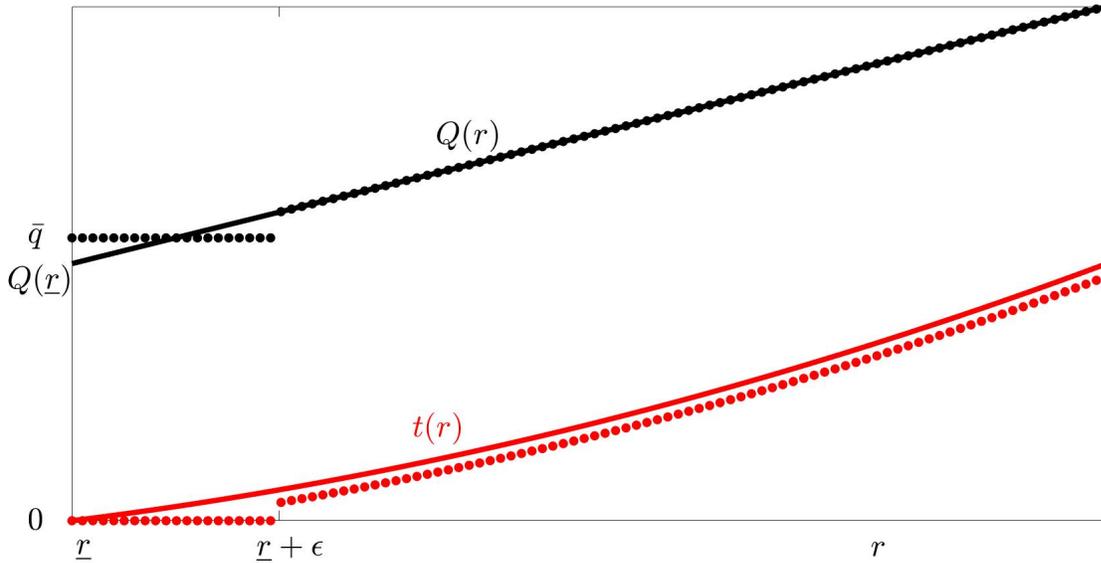


Figure A.1: An expected quality schedule $Q(r)$ and the corresponding payment rule $t(r)$ (solid lines). Dotted lines indicate the perturbation of the mechanism (Q, t) .

We will show that the designer's objective is increased by a perturbation of the mechanism that allocates objects at random and for free to some small set of types $[\underline{r}, \underline{r} + \epsilon]$. This allows her to decrease prices for everyone else (as long as $\underline{r} > 0$) while only causing a second-order distortion to allocative efficiency. Let q^ϵ denote the expected quality of objects allocated to types $[\underline{r}, \underline{r} + \epsilon]$ under Q , and let $Q^\epsilon(r) = q^\epsilon$ for $r \leq \underline{r} + \epsilon$ and $Q^\epsilon(r) = Q(r)$ otherwise. Setting $t^\epsilon(\underline{r}) = 0$ yields $U^\epsilon(\underline{r}) = q^\epsilon \underline{r}$. The associated change in utility for type r equals

$$[U^\epsilon(\underline{r}) + \int_{\underline{r}}^r Q^\epsilon(\tau) d\tau] - [U(\underline{r}) + \int_{\underline{r}}^r Q(\tau) d\tau] = (q^\epsilon - Q(\underline{r}))\underline{r} + \int_{\underline{r}}^{\min\{r, \underline{r} + \epsilon\}} (q^\epsilon - Q(\tau)) d\tau$$

The first term is first-order in ϵ and captures the increase in utility due to the increase in utility for the lowest type \underline{r} (which happens as long as that type values the increase in quality, that is, $\underline{r} > 0$). For types above $\underline{r} + \epsilon$, this increase in utility is achieved by a price discount which is possible when type's $\underline{r} + \epsilon$ allocation is decreased, relaxing the IC constraints for all

higher types. The second term is second-order in ϵ and captures the welfare effects of the distortion in allocation.

B Missing proofs

B.1 Proof of Theorem 1

We prove the theorem under the assumption that the designer maximizes a general objective function of the form

$$\int_{\underline{r}}^{\bar{r}} V(r)Q^F(r)dG(r) + v \underline{U} \tag{B.1}$$

for some upper semi-continuous functions $V : [\underline{r}, \bar{r}] \rightarrow \mathbb{R}$, and some constant $v \in \mathbb{R}$ (we dropped the subscripts i to simplify notation).

Given two non-decreasing functions $F, G : [a, b] \rightarrow [c, d]$ that are 0 at a and coincide at b , we will say that F is a *mean-preserving spread* of G if

$$\int_a^t F(x)dx \geq \int_a^t G(x)dx, \forall t \in [a, b],$$

with equality for $t = b$. We will say that F *first-order stochastically dominates* G if $F(x) \leq G(x)$ for all $x \in [a, b]$.

The following lemma describes all feasible expected quality assignments for a given distribution of quality F , assuming no free disposal.

Lemma 1. *If F is the CDF of available qualities, then $Q(r)$ is a feasible assignment of expected qualities (with no free disposal) if and only if $Q(r) = \Phi(G(r))$, where $\Phi : [0, 1] \rightarrow [0, 1]$ is non-decreasing, left-continuous, and*

$$\Phi \text{ is a mean-preserving spread of } F^{-1}.$$

Proof. Since $F(q)$ is a CDF, we can apply Strassen's Theorem (see Theorem 3.4.2(a) in Müller and Stoyan, 2002): A CDF $\bar{F}(q)$ is a distribution of posterior means of a random variable distributed according to F if and only if F is a mean-preserving spread of \bar{F} . Moreover, by the usual argument, the IC constraint (2.3) implies that the assignment of expected qualities must be non-decreasing. This monotonicity condition uniquely pins down $Q(r)$ given \bar{F} and G : We know that $\bar{F}(q)$ is the (normalized) mass of objects of quality q or less available to agents. This mass must be allocated to agents with some rate r or lower. Therefore, for any

q , there exists r such that $\bar{F}(q) = G(r)$, and it follows that

$$Q(r) = \bar{F}^{-1}(G(r)).$$

We claim that a function Φ is equal to \bar{F}^{-1} for some feasible \bar{F} if and only if Φ satisfies the conditions of the Lemma. That is,

$$\begin{aligned} \bar{F} \text{ is a CDF on } [0, 1] \text{ and } F \text{ is a MPS of } \bar{F} &\iff \\ \bar{F}^{-1} : [0, 1] \rightarrow [0, 1] \text{ is non-decreasing, left-continuous, and } \bar{F}^{-1} \text{ is a MPS of } F. &\quad (\text{B.2}) \end{aligned}$$

This follows from Lemma 1 found in [Brooks and Du \(2019\)](#). □

Intuitively, the proof of Lemma 1 can be understood through its connection to information design: We can treat F as the prior distribution of a random variable X (quality); Strassen's Theorem implies that a distribution \bar{F} of posterior means of X can be induced from the prior F (under some signal when X is treated as a state variable) if and only if F is a mean-preserving spread of \bar{F} . Hence, in our assignment problem, mean-preserving contractions of the distribution F describe all feasible distributions of expected qualities available to agents. Moreover, incentive-compatibility constraints imply that there is a unique assignment of expected qualities to types because the assignment must be monotone in the rate of substitution r .

Because the function $\Phi(q)$ from Lemma 1 is left-continuous, its value at 0 is not pinned down. This is a reflection of the fact that the payoff from the mechanism does not depend on the allocation for a measure-zero set of types, in particular, on the allocation for type \underline{r} . However, the allocation for type \underline{r} , $Q(\underline{r})$, appears in the constraint defining the non-negative transfers condition. It is clear that this constraint is least binding when $Q(\underline{r})$ is set to its maximal feasible level which is $Q(\underline{r}^+)$ (since Q must be non-decreasing). (Here, and thereafter, we denote $f(x^+) = \lim_{y \searrow x} f(y)$.) Because it is convenient to keep Φ left-continuous also at 0, we will extend the function Φ by assuming that $\Phi(x) = 0$ for all $x \leq 0$, and then the non-negative transfers condition becomes $\underline{U} \leq \underline{r}\Phi(0^+)$.

Given Lemma 1, we can write the problem of maximizing (B.1) under no-free-disposal as

$$\max_{\Phi} \int_{\underline{r}}^{\bar{r}} V(r)\Phi(G(r))dG(r) + \max\{0, v\}\underline{r}\Phi(0^+)$$

subject to

$$\Phi \text{ is a MPS of } F^{-1}.$$

Indeed, notice that when $v \leq 0$, it is optimal to choose \underline{U} as low as possible, and hence $\underline{U} = 0$ in the optimal mechanism ($\underline{U} \geq 0$ by individual rationality). In contrast, when $v > 0$, the non-negative transfers condition implies that it is optimal to set \underline{U} to its maximal feasible level $\underline{r} \Phi(0^+)$.

Integration by parts and by substitution yields

$$\int_{\underline{r}}^{\bar{r}} V(r) \Phi(G(r)) dG(r) = \int_0^1 \left(\int_t^1 V(G^{-1}(x)) dx \right) d\Phi(t).$$

Whenever we write $\int f(x) d\Phi(x)$ for some measurable function f , we mean the Lebesgue integral with respect to the σ -additive measure μ_Φ defined by $\mu_\Phi([a, b]) = \Phi(b^+) - \Phi(a)$, in particular, $\mu_\Phi(\{a\}) = \Phi(a^+) - \Phi(a)$. Under this convention, and recalling that $\Phi(x) = 0$ for $x \leq 0$, we can also write

$$\Phi(0^+) = \int_0^1 \mathbf{1}_{\{t=0\}} d\Phi(t).$$

Then, we can write (B.1) as

$$\int_0^1 \left(\int_t^1 V(G^{-1}(x)) dx + \max\{0, v\} \underline{r} \mathbf{1}_{\{t=0\}} \right) d\Phi(t).$$

Therefore, using the definition of Ψ from Theorem 1, we obtain an objective function $\int_0^1 \Psi(x) d\Phi(x)$. Next, we show that problems of this form admit an easy-to-describe solution.

Lemma 2. *Consider the problem*

$$\max_{\Phi: \Phi \text{ is a MPS of } \Phi_0} \int_0^1 \Psi(x) d\Phi(x),$$

where $\Psi(x)$ is an upper semi-continuous function and Φ_0 is given. Then, the value of the problem is $\int_0^1 \text{co}(\Psi)(x) d\Phi_0(x)$, and the solution is given by

$$\Phi^*(x) = \begin{cases} \frac{\int_a^b \Phi_0(x) dx}{b-a} & \text{if } x \in (a, b) \text{ and } (a, b) \text{ is a maximal interval on which } \text{co}(\Psi) \text{ is linear,} \\ \Phi_0(x) & \text{otherwise,} \end{cases}$$

for almost all x .

Proof. For any Φ , we have

$$\int_0^1 \Psi(x) d\Phi(x) \leq \int_0^1 \text{co}(\Psi)(x) d\Phi(x).$$

Moreover, the function on the right hand side of the inequality is maximized at $\Phi = \Phi_0$

because $\text{co}(\Psi)(x)$ is a concave function. It follows that the value of the problem in the lemma is bounded by $\int_0^1 \text{co}(\Psi)(x)d\Phi_0(x)$. We show that this upper bound can be achieved. Consider the candidate solution $\Phi^*(x)$ from the statement of the lemma. First, this function is feasible (by Kamenica Gentzkow 2016). Moreover, $\text{supp}(\Phi^*) \subseteq \{x : \Psi(x) = \text{co}(\Psi)(x)\}$ and on that set, $\Phi^* = \Phi_0$. Thus, $\int_0^1 \Psi(x)d\Phi^*(x) = \int_0^1 \text{co}(\Psi)(x)d\Phi_0(x)$. \square

The form of the solution is consistent with the recent findings of Kleiner et al. (2020) who derive general properties of extreme points that emerge as solutions to problems of the form considered in the lemma. The maximization problem in Lemma 2 can also be seen analogous to a Bayesian persuasion problem in which the designer’s preferences over posterior beliefs depend only on the posterior mean (see Kolotilin, 2018, and Dworzak and Martini, 2019) with a key difference: The MPS condition is flipped, requiring the solution Φ to be a mean-preserving spread (rather than a mean-preserving contraction) of the prior Φ_0 . This makes the problem very easy to solve by finding a concave closure of the objective function.

Lemmas 1 and 2 immediately imply that the value of the maximization problem under no-free-disposal is given by

$$\int_0^1 \text{co}(\Psi)(x)dF^{-1}(x) = \int_0^1 \text{co}(\Psi)(F(q))dq,$$

where the equality follows from integration by substitution. Moreover, a solution is given by $Q^*(r) = \Phi^*(G(r))$, where Φ^* is described in Lemma 2.

We can now derive Theorem 1. Allowing for free disposal is equivalent to allowing for “downward” first-order stochastic dominance shifts in the distribution of expected quality allocated to agents. That is, $Q(r)$ is a feasible expected-quality schedule with free disposal if $Q(r) = \bar{\Phi}(G(r))$ for some $\bar{\Phi} \leq \Phi$, where Φ is a mean-preserving spread of F^{-1} (see Lemma 1). Note that $\bar{\Phi}$ dominates Φ in the FOSD order because the FOSD relation is reversed by taking the inverse of the CDFs (and both $\bar{\Phi}$ and Φ are inverses of the CDFs of the expected quality).

Therefore, to derive the optimal expected-quality schedule under free disposal from the corresponding solution without free disposal, it is enough to solve an optimization problem of the following form:

Lemma 3. *Consider the problem*

$$\max_{\Phi} \int_0^1 \text{co}(\Psi)(x)d\Phi(x)$$

subject to

$$\Phi(x) \leq \Phi^*(x),$$

where $\Phi^*(x)$ is the solution given in Lemma 1. The value of the problem is $\int_0^1 \text{dco}(\Psi)(x)d\Phi^*(x)$, and the solution is given by

$$\Phi^{**}(x) = \Phi^*(x)\mathbf{1}_{\{x \geq x^*\}}$$

for almost all x , where $[0, x^*]$ is the maximal interval on which the concave decreasing function $\text{dco}(\Psi)$ is constant.

Proof. By definition of x^* , the function $\text{dco}(\Psi)(x)$ is constant and equal to $\text{co}(\Psi)(x^*)$ on $[0, x^*]$ and coincides with $\text{co}(\Psi)(x)$ otherwise. On one hand, we have for any feasible Φ ,

$$\int_0^1 \text{co}(\Psi)(x)d\Phi(x) \leq \int_0^1 \text{dco}(\Psi)(x)d\Phi(x) \leq \int_0^1 \text{dco}(\Psi)(x)d\Phi^*(x),$$

where the first inequality follows from the fact that $\text{co}(\Psi) \leq \text{dco}(\Psi)$, and the second from the fact that $\text{dco}(\Psi)$ is non-increasing and Φ dominates Φ^* in the FOSD order. On the other hand, if we define Φ^{**} as in the statement of the lemma, then we have

$$\begin{aligned} \int_0^1 \text{co}(\Psi)(x)d\Phi^{**}(x) &= \int_0^{x^*} \text{co}(\Psi)(x)d\Phi^{**}(x) + \int_{x^*}^1 \text{co}(\Psi)(x)d\Phi^{**}(x) \\ &= \text{co}(\Psi)(x^*)\Phi^*(x^*) + \int_{x^*}^1 \text{dco}(\Psi)(x)d\Phi^*(x) = \int_0^1 \text{dco}(\Psi)(x)d\Phi^*(x), \end{aligned}$$

by the properties of $\text{co}(\Psi)$, $\text{dco}(\Psi)$ and $\Phi^{**}(x)$. Thus, Φ^{**} achieves the upper bound and hence is a solution to the problem described in Lemma 3. \square

With Lemma 3, Theorem 1 follows directly from Lemma 1: The value of the problem is

$$\int_0^1 \text{dco}(\Psi)(x)d\Phi^*(x) = \int_0^1 \text{dco}(\Psi)(x)dF^{-1}(x) = \int_0^1 \text{dco}(\Psi)(F(q))dq$$

where the last equality follows from integration by substitution. The optimal solution is given by an expected-quality schedule

$$Q^*(r) = \Phi^{**}(G(r)) = \Phi^*(G(r))\mathbf{1}_{\{G(r) \geq x^*\}} = \Phi^*(G(r))\mathbf{1}_{\{r \geq G^{-1}(x^*)\}},$$

where Φ^* is described in Lemma 1. Finally, the choice of the optimal \underline{U} was described in the reasoning leading up to Lemma 2.

B.2 Proof of Theorem 2

We solve the program (3.2)–(3.3) by solving a relaxed problem in which the constraint that $F_i(q)$ is a CDF is dropped, and then verify that the solution of the relaxed program is feasible.

The relaxed program is to solve for the optimal $F_i(q)$ for every $q \in Q$ separately:

$$\max_{0 \leq x_i \leq 1} \sum_{i \in I} \mu_i \text{dco}(\Psi_i)(x_i), \quad (\text{B.3})$$

$$\text{s.t. } \sum_{i \in I} \mu_i x_i = F(q). \quad (\text{B.4})$$

This program can be solved using standard Lagrangian techniques (concavity of $\text{dco}(\Psi_i)$ guarantees their validity). Fix q . There exists a Lagrange multiplier,⁹ that we will denote by $L(q)$, such that the optimal x_i^* maximizes $\sum_{i \in I} \mu_i [\text{dco}(\Psi_i)(x_i) - L(q)x_i]$ while satisfying the constraint. Because the Lagrangian is concave, the first-order condition is both necessary and sufficient. Let $X_i^*(q)$ be the set of points satisfying the first-order condition: $X_i^*(q) = \{x : \text{dco}(\Psi_i)'(x) = L(q)\}$ whenever this set is non-empty, and otherwise $X_i^*(q) = \{0\}$ if $\text{dco}(\Psi_i)'(0) < L(q)$ and $X_i^*(q) = \{1\}$ if $\text{dco}(\Psi_i)'(1) > L(q)$. By the above, we know that there exists a selection $x_i^* \in X_i^*(q)$ such that (B.4) holds. Moreover, because each $\text{dco}(\Psi_i)$ is concave and continuous, we know that each $X_i^*(q)$ is a closed interval (potentially a singleton).

To prove the theorem, it remains to show that there exists a selection $F_i^*(q)$ from each $X_i^*(q)$ that is non-decreasing (then, it can be modified on a measure-zero set of points to make it into a CDF; notice that it is guaranteed by the constraint (B.4) that each F_i^* is 0 at 0 and 1 at 1).

Because the constraint in (B.4) is increasing in q , it follows that the Lagrange multiplier $L(q)$ is a non-increasing function of q . Moreover, the sets $X_i^*(q)$ are non-decreasing in the strong set order by concavity of $\text{dco}(\Psi_i)$. Define a vector function

$$C(q, \alpha) = [(1 - \alpha) \min X_1^*(q) + \alpha \max X_1^*(q), \dots, (1 - \alpha) \min X_I^*(q) + \alpha \max X_I^*(q)].$$

By definition, for each q , $\sum_i C_i(q, 0) \leq F(q)$ while $\sum_i C_i(q, 1) \geq F(q)$. By continuity, there exists $\alpha^*(q)$ such that $\sum_i C_i(q, \alpha^*(q)) = F(q)$ (moreover, the values of $C_i(q, \alpha^*(q))$ are uniquely pinned down, even if $\alpha^*(q)$ is not). We can now define $F_i^*(q)$ as $C_i(q, \alpha^*(q))$. By direct inspection and the strong-set order property of $X_i^*(q)$, each $F_i^*(q)$ is non-decreasing, which finishes the proof once we set $S(q) = -L(q)$.

⁹In case there are multiple Lagrange multipliers, we pick the largest one.

B.3 Proofs of results in Section 4

Proof of Proposition 1. The proof is immediate from Theorem 1. The assumptions of Proposition 1 ensure that there is an upward jump at 0 in Ψ_i , and therefore $\text{dco}(\Psi_i)(x)$ must be linear for small enough x . (Of course, when $\text{dco}(\Psi_i)(x)$ is constant for small x , it is possible that types $r \leq r_i^*$ do not receive any objects; however, we still call such allocation random, in line with Remark 1.) □

Proof of Proposition 2. By Theorem 1, full randomization is optimal if and only if $\text{dco}(\Psi_i)$ is linear which is true if and only if

$$\Psi_i(x) \leq (1-x)\Psi_i(0) + x\Psi_i(1),$$

for all $x > 0$. We have

$$\Psi_i(0) = \max\{0, \alpha - \bar{\lambda}_i\}r_i + \int_{r_i}^{\bar{r}_i} \tau \lambda_i(\tau) dG_i(\tau).$$

Using the fact that $\Psi_i(1) = 0$, we can write the condition as, for all $r > r_i$,

$$\Psi_i(G_i(r)) \leq (1 - G_i(r)) \left[\max\{0, \alpha - \bar{\lambda}_i\}r_i + \int_{r_i}^{\bar{r}_i} \tau \lambda_i(\tau) dG_i(\tau) \right]. \quad (\text{B.5})$$

To see that this implies $\alpha < \bar{\lambda}_i$, note that by dividing both sides by $1 - G_i(r)$, using the expression (3.1), and taking the limit as $r \rightarrow \bar{r}_i$, we get

$$\alpha \bar{r}_i \leq \max\{0, \alpha - \bar{\lambda}_i\}r_i + \int_{r_i}^{\bar{r}_i} \tau \lambda_i(\tau) dG_i(\tau) < \max\{0, \alpha - \bar{\lambda}_i\}r_i + \bar{r}_i \bar{\lambda}_i. \quad (\text{B.6})$$

Thus, if $\alpha \geq \bar{\lambda}_i$, we would get $(\alpha - \bar{\lambda}_i)\bar{r}_i < (\alpha - \bar{\lambda}_i)r_i$ which is a contradiction. Using this observation in (B.6), we obtain the necessary conditions from Proposition 2, and equation (B.5) gives us the necessary and sufficient condition states before Proposition 2.

Finally, suppose that $\alpha J_i(r) + \Lambda_i(r)h_i(r)$ is quasi-convex. This implies that Ψ_i is first convex and then concave on $(0, 1]$. The necessary condition implies that $\Psi_i(0) \geq \Psi_i(1) - \Psi_i'(1)$. Together this implies that $\Psi_i(x) \leq (1-x)\Psi_i(0)$ for all x . □

Proof of Proposition 3. Suppose that there is random allocation at the top, that is, $\Psi_i(x)$ is linear for $x \in [\underline{x}, 1]$ for some \underline{x} . Take \underline{x} so that this is the maximal random-allocation region. There are two cases to consider. If $\underline{x} = 0$, then we have to rule out that $\Psi_i(0) \geq \Psi_i(1) - \Psi_i'(1)$;

if $\underline{x} > 0$, and then it suffices to rule out that $\Psi'_i(\underline{x}) \geq \Psi'_i(1)$ (if $\underline{x} > 0$ is the beginning of the maximal interval of a random allocation, then the slope of $\text{dco}(\Psi_i)$ at \underline{x} must be equal to the slope of Ψ_i , and that slope must be larger than the slope of Ψ_i at 1 since $\text{dco}(\Psi_i) \geq \Psi_i$ with an equality at 1). Because G_i has a bounded support, its inverse hazard rate is 0 at the upper bound; thus, $\Psi'_i(1) = -\alpha\bar{r}_i$. Thus, the first possibility can be ruled out if

$$\alpha\bar{r}_i > (\alpha - \bar{\lambda}_i)\underline{r}_i + \int_{\underline{r}_i}^{\bar{r}_i} \tau\lambda_i(\tau)dG_i(\tau) \iff \alpha(\bar{r}_i - \underline{r}_i) > -\bar{\lambda}_i\underline{r}_i + \int_{\underline{r}_i}^{\bar{r}_i} \tau\lambda_i(\tau)dG_i(\tau).$$

But we have

$$\int_{\underline{r}_i}^{\bar{r}_i} \tau\lambda_i(\tau)dG_i(\tau) - \bar{\lambda}_i\underline{r}_i < \bar{r}_i \int_{\underline{r}_i}^{\bar{r}_i} \lambda_i(\tau)dG_i(\tau) - \bar{\lambda}_i\underline{r}_i = \bar{\lambda}_i(\bar{r}_i - \underline{r}_i),$$

thus this can be ruled out by $\alpha \geq \bar{\lambda}_i$. The second possibility can be ruled out if for all r ,

$$\alpha r - (\alpha - \Lambda_i(r))h_i(r) < \alpha\bar{r},$$

which clearly holds as long as $\alpha \geq \Lambda_i(r)$ which is true by the fact that $\lambda_i(r)$ is non-increasing and $\bar{\lambda}_i \leq \alpha$. \square

Proof of Proposition 4. The proof is immediate from Theorem 1. \square

Proof of Proposition 5. Throughout the proof, we drop the dependence on n from the notation.

Consider the first case first. We know that $J'_i(r) = 1 - h'_i(r) \geq \underline{J}_i$, so $h'_i(r) \leq 1 - \underline{J}_i$. By Proposition 4, since we know that $\alpha \geq \bar{\lambda}_i$ for all n , to prove that assortative matching is optimal, it is enough to prove that the second derivative of Ψ_i is non-positive. The sign of the second derivative of Ψ_i is opposite to the sign of the following expression:

$$\alpha + \Lambda_i(r) - \lambda_i(r) + \underbrace{(\Lambda_i(r) - \alpha)}_{\leq 0} h'_i(r) \geq \alpha - 2\epsilon + (\Lambda_i(r) - \alpha)(1 - \underline{J}_i) = -2\epsilon + \Lambda_i(r) + (\alpha - \Lambda_i(r))\underline{J}_i \geq 0,$$

for all $2\epsilon < \alpha \min\{1, \underline{J}_i\} \leq \Lambda_i(r) + (\alpha - \Lambda_i(r))\underline{J}_i$. Thus, by taking ϵ satisfying that last condition, and n large enough, we conclude that the solution to the problem is assortative matching (this conclusion is stronger than that of Proposition 4 in that assortative matching is exactly optimal for n large enough.)

Now consider the second case. By the same calculation as before, for any $x > 0$, there exists a large enough n so that Ψ_i is strictly concave on $[x, 1]$. This means that if there is a random-allocation region that does not vanish in the limit as $n \rightarrow \infty$, then it must take

the form of $[x_0, x_1]$ with $x_0 \rightarrow 0$ and $x_1 > \underline{x} > 0$ as $n \rightarrow \infty$, where \underline{x} does not depend on n . (Intuitively, while Ψ_i is concave on $[x, 1]$ for any x if n is large enough, it could be the case that the concave closure of Ψ_i is supported at a point x_0 that converges to 0, and some other point—bounded away from 0—that lies in the region where Ψ_i is concave.) We will show that this leads to a contradiction.

First, it is convenient to decompose

$$\Psi_i(x) = \underbrace{\int_x^1 J_i(G_i^{-1}(x)) dx}_{\Psi_i^R} + \underbrace{\int_x^1 \Lambda_i(G_i^{-1}(x)) h_i(G_i^{-1}(x)) dx}_{\Psi_i^W}.$$

Our strategy is to show that, for large enough n , Ψ_i^R is strictly concave (with a second derivative bounded away from 0), while Ψ_i^W and its derivative are arbitrarily small, and thus they cannot change the shape of Ψ_i in the limit.

Note that there exists $m > 0$ such that

$$(\Psi_i^R)''(x) = -\frac{J_i'(G_i^{-1}(x))}{g_i(G_i^{-1}(x))} < -m < 0,$$

by assumption that the derivative of J_i is lower bounded, and that the density g_i is continuous on its support (so it has an upper bound). Also note that for any $\epsilon > 0$, and x such that $G_i^{-1}(x) < \epsilon$, for large enough n , we have

$$\Psi_i^W(x) \leq \Psi_i^W(0) = \int_0^1 \left(\int_{G_i^{-1}(x)}^{\bar{r}_i} \lambda_i(\tau) dG_i(\tau) \right) dr \leq \bar{\lambda}_i G_i^{-1}(x) + \epsilon(\bar{r}_i - \underline{r}_i) \leq \epsilon \cdot M,$$

where the second to last inequality uses the assumption that Pareto weights are below ϵ for large enough n , and M is some constant. By the same assumption, for any $\epsilon > 0$, $x > 0$, and large enough n ,

$$|(\Psi_i^W)'(y)| \leq |-\Lambda_i(G_i^{-1}(y)) h_i(G_i^{-1}(y))| \leq \epsilon,$$

for any $y \geq x$.

We are ready to obtain a contradiction. A necessary condition for $\text{dco}(\Psi_i)$ to be linear on $[x_0, x_1]$ is that

$$\Psi_i(x_1) - \Psi_i'(x_1)(x_1 - x_0) - \Psi_i(x_0) \leq 0. \tag{B.7}$$

Note, however, that

$$\Psi_i^R(x_1) - (\Psi_i^R)'(x_1)(x_1 - x_0) - \Psi_i^R(x_0) = - \int_{x_0}^{x_1} y(\Psi_i^R)''(y) dy \geq \frac{1}{2}m(x_1 - x_0)^2. \quad (\text{B.8})$$

Since $x_1 \geq \underline{x} > 0$ for all n , and \underline{x} does not depend on n , this expression is bounded away from 0. Yet, by the inequalities established above on Ψ_i^W and $(\Psi_i^W)'$, we have

$$|(\Psi_i(x_1) - \Psi_i'(x_1)(x_1 - x_0) - \Psi_i(x_0)) - (\Psi_i^R(x_1) - (\Psi_i^R)'(x_1)(x_1 - x_0) - \Psi_i^R(x_0))| \leq \epsilon \cdot \tilde{M},$$

for some constant \tilde{M} . For large enough n , we can take ϵ small enough so that $\epsilon \cdot \tilde{M} < \frac{1}{2}m(x_1 - x_0)^2$ which is inconsistent with (B.7) and (B.8), a contradiction. \square

Proof of Proposition 6. The conclusion is immediate from Theorem 2. When there is fully random matching in group i , the function $\text{dco}(\Psi_i)$ is linear, and thus its slope is constant, equal to $\Psi_i(0)$ (since $\Psi_i(1) = 0$). By Proposition 2, fully random matching requires that $\bar{\lambda}_i > \alpha$, and under this inequality, we have that $\Psi_i(0) = \int_{\underline{r}_i}^{\bar{r}_i} r \lambda_i(r) dG_i(r)$. \square

Proof of Proposition 7. By the assumption that assortative matching is optimal, we must have $\text{dco}(\Psi_i)(x) = \Psi_i(x)$, except possible for $x \leq x_i^*$ if $\text{dco}(\Psi_i)(x)$ is constant on $[0, x_i^*]$. By direct calculation (and using the fact that for bounded support distributions, the inverse hazard rate is 0 at the upper bound), we obtain $\Psi_i'(1) = -\alpha \bar{r}_i$. The conclusion follows directly from Theorem 2, and the observation that Ψ_i has a continuous derivative (by assumption that $g_i(r)$ and $\lambda_i(r)$ are continuous). When $\alpha > \bar{\lambda}_i$ and $\underline{r}_i = 0$, we have that

$$\Psi_i'(0) = -\alpha \left(\underline{r}_i - \frac{1}{g_i(\underline{r}_i)} \right) - \bar{\lambda}_i \frac{1}{g_i(\underline{r}_i)} > 0,$$

and hence Ψ_i is increasing in the neighborhood of 0. Thus, $\text{dco}(\Psi_i)$ is constant in some initial interval, and hence has a zero slope. When $\alpha > \bar{\lambda}_i$ and $\underline{r}_i = 0$ for all i , by Theorem 2, all groups are allocated the lowest-quality objects. \square

Proof of Proposition 8. The proof of this result was presented after the statement in the main text. \square

Proof of Proposition 9. The proof of the first part of this result was presented after the statement in the main text.

We prove the second part. For $\bar{q} < 1$, we need that $|\text{dco}(\Psi_1)'(1)| > |\text{dco}(\Psi_2)'(1)| = \Psi_2(0)$, by Theorem 2. This yields the condition $\alpha \bar{r}_1 \geq \int_{\underline{r}_2}^{\bar{r}_2} \tau \lambda_2(\tau) dG_2(\tau)$. For $\underline{q} > 0$, we need that $|\text{dco}(\Psi_1)'(0)| < |\text{dco}(\Psi_2)'(0)| = \Psi_2(0)$. Since group 1 features assortative matching (of

allocated objects), either $|\text{dco}(\Psi_1)'(0)| = 0$ or $\text{dco}(\Psi_1)'(0) = \Psi_1'(0) \geq -\alpha r_1$. This yields the condition $\alpha r_1 \leq \int_{r_2}^{\bar{r}_2} \tau \lambda_2(\tau) dG_2(\tau)$. \square