

Redistributive Allocation Mechanisms*

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

December 11, 2020

Abstract

Many scarce public resources are allocated at below-market-clearing prices, and sometimes for free. Such “non-market” mechanisms necessarily sacrifice some surplus, yet they can potentially improve equity. In this paper, we develop a model of mechanism design with redistributive concerns. Agents are characterized by a privately observed social welfare weight and willingness to pay for quality, as well as a publicly observed label. A market designer controls allocation and pricing of a set of objects of heterogeneous quality, and maximizes the expectation of a welfare function defined by the social welfare weights. We derive structural insights about the form of the optimal mechanism, leading to a framework for determining how and when to use non-market mechanisms. The key determinant is the strength of the statistical correlation of the unobserved social welfare weights with the label and the willingness to pay that the designer can, respectively, directly observe or elicit through the mechanism.

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, Andrew Kahn, Simon Loertscher, Ellen Muir, Alessandro Pavan, Stefanie Stantcheva, Bruno Strulovici, 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; GRAPE Research Center; 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, school seats, national park permits, and certain types of healthcare—are allocated at below-market-clearing prices, and sometimes for free. Such “non-market” mechanisms naturally raise concerns among economists because they necessarily sacrifice some allocative surplus by failing to allocate resources to those with the highest willingness to pay. At the same time, however, policy-makers often justify non-market mechanisms on fairness grounds: If resources were allocated using market-clearing prices, they argue, 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, health problems, or unemployment—marketplace designers may be naturally concerned about the welfare of such agents. But how should we think about the resulting efficiency–equity trade-off? In this paper, we develop a market-design approach to this question.

We study a model in which a market designer or policymaker must decide about the allocation of a fixed supply of goods with heterogeneous quality. Each agent’s utility is linear in the quality of the received good and in monetary transfers—allowing us to parameterize agents’ preferences by a single parameter called willingness to pay. Besides the privately observed willingness to pay for quality, each agent is characterized by a publicly observed label, and an unobserved social welfare weight. We characterize the optimal incentive-compatible and individually-rational mechanism for a designer who seeks to maximize the expectation of the welfare function, given by the sum of agents’ utilities weighted by their social welfare weights.

While optimal redistribution has been extensively in the context of public finance, here we take a complementary market-design approach. The designer decides about the allocation of a single type of good without considering the interaction of this allocation process with macro-level redistribution. The supply of goods and the social welfare weights are thus modeled as exogenous. While these assumptions are limiting in some contexts, they are natural descriptions of many relevant policy problems. For a recent example, consider the question of how to allocate vaccines during the Covid-19 pandemic. A state government may face an inelastic supply of vaccines in the short run. Quality in this example may be understood as time priority (higher quality means getting the vaccine sooner). While state governments can effect redistribution by setting taxes, they have also been separately considering redistributive consequences of allocating vaccines¹—for instance, setting a high

¹See for example [Schmidt et al. \(2020\)](#), [Pathak et al. \(2020\)](#), and the references therein.

price for the vaccine would exclude large parts of poorer populations from getting it.

Our results show how the form of the optimal mechanism depends on the relationship between agents’ observed and unobserved characteristics. Our analysis thus provides a framework for determining when and how non-market allocation is socially beneficial. In particular, we show that the key parameters that determine the optimal redistributive scheme are as follows.

Social preferences and the information available to the policymaker. The underlying premise of our analysis is that in the presence of perfect information the market designer would know exactly what social welfare weight to attach to every individual by taking into account all of her relevant characteristics.² However, in practice information is not perfect. Instead, the designer has access to two types of information. First, she can observe some publicly available (or verifiable) information summarized by the agent’s label—examples could include employment status, tax data on income, or marital status. Second, the designer can elicit information through the mechanism—we show that she can additionally infer the willingness to pay of each agent (but nothing else). As a result, the designer must form Bayesian beliefs about the remaining unobservable characteristics that could be relevant for the welfare weights—such as the detailed financial situation, or consumption and health shocks. Formally, when maximizing the welfare function, the designer attaches to each agent a *Pareto weight* equal to the *expectation* of the unobserved welfare weight conditional on the publicly observed label and the elicited willingness to pay. (Hereafter, to avoid any confusion, we will use the term “welfare weights” to refer to the underlying social preferences, and “Pareto weights” to refer to the expected weights in the designer’s objective function.)

We show how the shape of the Pareto weights—as a function of the label and willingness to pay—translates into the properties of the optimal mechanism. The dependence of Pareto weights on publicly available information about individuals allows us to capture applications like affirmative action or preferential treatment of low-income households. The dependence of Pareto weights on willingness to pay is due to a more subtle statistical link between unobservable characteristics relevant for social preferences and the agent’s behavior in the mechanism—this effect also shapes the form of optimal redistribution by complementing the use of verifiable information.

²A social welfare weight has a natural interpretation in our quasi-linear model: It is the social value of giving a dollar to a given agent. See [Saez and Stantcheva \(2016\)](#) for a general treatment.

The importance and role of raising revenue. Unless all goods are allocated for free, the mechanism generates some strictly positive revenue; the optimal allocation scheme naturally depends on whether this revenue is used redistributively or not.³ In our welfare function, we introduce a parameter α that reflects the marginal value, or opportunity cost, of a dollar of revenue. Effectively, α is a Pareto weight placed on the beneficiary of the monetary surplus generated by the mechanism.

If the designer “rebates” the revenue directly and frictionlessly to the target population, then α equals the average Pareto weight in the population (possibly conditioned on labels). An α below the average Pareto weight allows us to capture the case in which there are frictions associated with direct transfers, for example reflecting administrative costs or unmodeled inefficiencies associated with giving cash to agents suffering from behavioral biases.⁴ Finally, α above the average Pareto weight implies that the designer subsidizes a socially valuable outside cause.

The type of goods being allocated. Intuitively, in-kind transfers are typically associated with goods that are deemed “essential” in the sense that every agent must consume them (e.g., housing and food). An essential good can be captured in our model by specifying that the willingness to pay for it in the population is bounded away from zero; we explain how the essentiality of the good influences the amount of in-kind redistribution that emerges in the optimal mechanism.

Preview of results

We construct an optimal mechanism for our setting in two steps: First, we solve the optimal allocation problem for each group of agents with the same observable characteristics (i.e., the same label) separately, conditional on a given allocation of goods to that group. Then, we show how that within-group solution gives rise to a simple greedy algorithm for optimally allocating goods across the groups.

For a given set of goods to allocate to a group of agents with the same label, the optimal within-group mechanism partitions agents into blocks according to their willingness to pay, and within each block allocates the good either (i) uniformly at random or (ii) assortatively, with higher-quality goods assigned to agents with higher willingness to pay. Which allocative

³For example, the proceeds from running a public housing program can be used to finance infrastructure investments, or subsidize the city budget. Alternatively, revenue can be returned to the target population, in the form of direct subsidies or via a reduction in taxes.

⁴An extreme case of $\alpha = 0$ corresponds to the case when agents “pay” for the goods not using money but by engaging in a costly activity such as waiting in line; in that case, revenue is just an aggregate measure of the size of the costly screening activity that has no intrinsic value in itself.

mechanism is used is determined by a trade-off: increasing the amount of randomization relative to assortative matching increases the probability that agents with lower willingness to pay receive the good and lowers prices for all agents—but also reduces allocative inefficiency and lowers revenue.

Given our solution for an optimal mechanism within 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 . We show that the optimal mechanism allocates goods across groups greedily with respect to this “group value” statistic, starting with lowest-quality goods and proceeding to the highest.

The main insight of this paper is that the use of below-market-clearing prices and random allocation (in-kind redistribution) in the optimal mechanism is largely determined by three key factors:

1. the level of dispersion in Pareto weights as a function of willingness to pay (conditional on any fixed label);
2. the relationship between the average Pareto weight (for a given group) and the opportunity cost of revenue; and
3. whether the good is essential or not.

Roughly speaking, a random allocation at below-market prices is more likely to emerge as optimal, *ceteris paribus*, when the opportunity cost of revenue is low, the good is essential, and the level of dispersion in Pareto weights is intermediate. Concretely, we prove the following results on within-group allocation.

A fully random allocation in some group can only be optimal when the average Pareto weight on that group is strictly higher than the opportunity cost of revenue α . In particular, a necessary condition is that the designer *cannot* target a lump-sum transfer to this group of agents (or that this is sufficiently costly). If a lump-sum transfer to the entire population is feasible, then a necessary condition is that this group is “preferred” in the sense that the group has a strictly-higher-than-average Pareto weight (as in affirmative action).

If the good is essential, then being a “preferred group” that cannot receive direct lump-sum transfers is also sufficient for random allocation to be used in the optimal mechanism. Formally, we show that if the weight on revenue α is strictly below the average Pareto weight for some group of agents, then the optimal mechanism allocates some goods for free to those with the lowest willingness to pay within that group.

In contrast, if lump-sum transfers are feasible (so that α is weakly above the average Pareto weight), then as long as Pareto weights are non-increasing in the willingness to pay,

it is optimal to use assortative matching—at least for agents with the highest willingness to pay within a group. Moreover, for non-essential goods and Pareto weights that are not too dispersed as a function of willingness to pay, a fully assortative matching within each group is optimal. Monotonicity and dispersion of Pareto weights are determined by the statistical correlation between the unobserved social welfare weights and willingness to pay: Pareto weights are non-increasing if willingness to pay and social welfare weights are negatively correlated; they have low dispersion, when that correlation is not too strong.

Sometimes, in the optimal mechanism, random allocation to low-willingness-to-pay agents co-exists with assortative matching to high-willingness-to-pay agents. In this case, the effect of dispersion in Pareto weights is non-monotone. That is, assortative matching becomes optimal both when the weights are nearly constant (willingness to pay does not influence the Pareto weight too much) and when the weights become extremely skewed (almost all the weight is attached to the individuals with the lowest willingness to pay). Thus, the use of random allocation is maximized at some intermediate level of dispersion.⁵

Finally, we investigate how the structure of the optimal mechanism within each group influences the allocation of objects across groups. We first show that fully assortative matching and fully random matching lead to two opposite properties of the across-group allocation: When assortative matching is used within each group, the allocation of objects is “horizontal” in the sense that each group receives both the lowest- and highest-quality objects. When a fully random allocation is used within each group, meanwhile, the allocation of objects is “vertical” in the sense that groups can be ordered (by their average Pareto weights), and groups higher in the order receive higher quality.

When allocation is random in some groups and assortative in others, we prove that, in general, the groups with random allocation should receive goods of *intermediate* quality. To understand why, note that assortative matching leads to higher revenue when the quality of the good is more dispersed (so that a relatively high price can be charged for high-quality goods). In contrast, when the good is allocated uniformly at random, both the revenue and the utility of agents only depend on the *average* quality. Thus, goods with extreme quality (both high and low) are typically allocated to groups with assortative matching.

We furthermore discuss the market design consequences of our findings in Section 5.

⁵For intuition, note that when the dispersion in Pareto weights is low, the motive to maximize total surplus dominates the redistribution concerns. On the other extreme, when only the agents with lowest willingness to pay receive large Pareto weights, it becomes optimal to maximize revenue (which is achieved by allocating assortatively) and redistribute via lump-sum subsidies—this is because, by definition, agents with lowest willingness to pay have the highest relative value for receiving money (relative to receiving the good).

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 huge number 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 \(2020\)](#). [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, a continuum of goods and agents, and groups of agents with the same observable 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. Finally, by specializing the designer’s objective to be a weighted sum of agents’ utilities and revenue, our paper has an applied focus and attempts to connect market circumstances with the model’s conclusions about the optimal design.

[Dworczak [®] Kominers [®] Akbarpour \(2020\)](#) (henceforth, DKA) studied a closely related question in the context of buyers and sellers trading a good of homogeneous quality. They show that the optimal mechanism—subject to market-clearing and budget-balance constraints—may in general deviate from competitive trading in two simple ways: via the introduction of lump-sum payments to one side of the market and the use of rationing (offering the good at an attractive price but with an interior probability of transacting). This departure from the conclusion of the second welfare theorem is driven by the fact that the designer maximizes a weighted sum of agents’ utilities with weights that depend on the side of the market and willingness to pay. Thus, the two papers share the interest in investigating properties of social preferences that justify the use of non-market mechanisms. However, the

current paper takes a more practical approach by focusing on the problem of allocating public resources and incorporating a range of features that play a key role for real-life policymakers: heterogeneous quality of objects, richer preferences over revenue, additional observable information about the agents, and a restriction on the use of lump-sum transfers. Heterogeneity of object qualities creates new dimensions of the design problem, with potentially multiple “regions” of assortative versus random matching in the optimal mechanism. 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, 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 is a classical idea in the taxation literature. For example, [Akerlof \(1978\)](#) describes how “tagging” could be 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.

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 welfare weights. This is in contrast to many classical public finance models (such as the canonical framework of [Diamond and Mirrlees, 1971](#), or [Atkinson and Stiglitz, 1976](#)), 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 allocation rules—is distinct and complementary. Because of the special structure, agents’ behavior in our model is described by a bang-bang solution, rather than first-order conditions that are used in much of modern public finance. As a result, the redistribution question in our context is especially tractable, and the optimal mechanism can often be found in closed form.

Our paper relates to a large literature in economics on price controls as a redistributive

tool. [Viscusi, Harrington, and Vernon \(2005\)](#) 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. An extreme version of price control—providing the good for free as an in-kind transfer—has also been studied. The literature on in-kind redistribution already emphasized that individuals’ market behavior can help the planner identify their type. [Besley and Coate \(1991\)](#) analyze the problem of providing a free, public option for a good that is already being provided in the private market. The public option is funded by taxing individuals, who are either rich or poor. They show that providing a “lower quality” public option can benefit the poor, as rich individuals opt for the private option but pay the tax anyways. [Gahvari and Mattos \(2007\)](#) compare in-kind transfers with direct cash as two forms of helping those who are in need. They show that in-kind transfers have the advantage that when “needs” are private information, in-kind transfers can help identify individuals who truly need a good.

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](#)).⁶ 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, Moldovanu, and Sela \(2009\)](#), 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.

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 \(1981\)](#).⁷ Following the intuitive approach to ironing developed by [Bulow and Roberts \(1989\)](#), [Hartline and Roughgarden \(2008\)](#) applied it to a problem with multiple goods, and [Condorelli \(2012\)](#) to multiple goods with heterogeneous quality. In concurrent research, [Muir and Loertscher \(2020\)](#) rely

⁶Similar 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.

⁷See also [Toikka \(2011\)](#).

on similar techniques to solve a problem of a revenue-maximizing seller in the presence of resale; Ashlagi, Monachou, and Nikzad (2020) show that these methods can be also used in designing the optimal dynamic allocation in a multi-good environment by optimizing over how much information is disclosed about different types of objects; finally, Kleiner, Moldovanu, and Strack (2020) demonstrate that all 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 who differ in both their observable and unobservable characteristics. There is a unit mass of agents, with each agent characterized by a type vector (i, r, λ) . The three dimensions of a type vector have a joint distribution in the population that is known to the designer.⁸ The first ingredient of the type vector, called the *label*, takes one of finitely many values from the set I , and is assumed to be publicly observed. Agents with the same label form a *group*; there are $\mu_i > 0$ agents in group i . The parameter $r \in \mathbb{R}_+$ is the *willingness to pay* (for quality) which is privately observed by the agent. Conditional on label i , r has a distribution with cumulative distribution function G_i and continuous density g_i , fully supported on $[r_i, \bar{r}_i]$. Finally, $\lambda \in \mathbb{R}_+$ is the social welfare weight on a given individual, interpreted as the social value of giving that individual one unit of money. Neither r nor λ of any given individual are observed by the designer.⁹

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 willingness to pay r is assigned a good with quality q in exchange for a monetary transfer t , that agent's utility is $rq - t$; if that agent has a social welfare weight λ , the contribution of that individual to the social welfare function will be $\lambda(rq - t)$. 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).

Our framework incorporates a few strong assumptions about the environment. First, we assume that agent's utility is quasi-linear in money. Second, agents differ only in their

⁸We do not introduce notation for the joint distribution because it will not play a direct role in our formal analysis.

⁹As we show, it does not matter whether an individual can observe her own social welfare weight λ .

“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.¹⁰ Fourth, social preferences are captured by weights that are constant—capturing the implicit assumption that the designer does not take into account how her chosen allocation impacts social preferences.

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 willingness to pay r in group i is assigned an object with quality q or less.¹¹ The assignment Γ is *feasible* if

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

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

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 willingness to pay r , assigns objects according to $\Gamma_i(q|r)$ in group i , and charges agents according to the transfer function $t_i(r)$. As it will turn out, we do not have to include the social welfare weight λ in the agent’s report because no incentive-compatible mechanism can improve the social welfare function by trying to elicit this information from agents (see Claim 1 below).

Lump-sum payments to agents may or may not be allowed in different applications of our framework. We use the following modeling approach to accommodate all possible cases. There is no hard budget constraint for the designer but the mechanism must use non-negative transfers: $t_i(r) \geq 0$ for all i and r .¹² However, lump-sum payments to agents may happen

¹⁰We can define the agent’s utility to be $rv(q) - t$ for some concave function v , partly capturing the consequences of risk aversion. In that case, we would define a new random variable $\tilde{q} = v(q)$ with CDF \tilde{F} , called “risk-adjusted quality,” and apply our results with \tilde{F} as the primitive distribution of quality.

¹¹We assume that all agents with the same willingness to pay r are assigned the same lottery over objects.

¹²Because agents are buyers in our framework, this constraint on transfers has no impact on the set of

“outside of the mechanism;” this is captured through the designer’s value for generating monetary surplus in the mechanism (in the objective function that we formally introduce in the next subsection). For example, if the value for generating monetary surplus is 0 in the designer’s objective, the constraint of non-negative transfers is binding and means that lump-sum payments are not allowed. However, if the value for generating monetary surplus is equal to the value of giving a lump-sum payment to all agents, then it is *as if* lump-sum payments to all agents were allowed. We comment on other cases later. For incentive-compatible mechanisms, the condition that transfers are non-negative is equivalent to requiring that for each group i , the utility \underline{U}_i of type \underline{r}_i satisfies $\underline{U}_i \leq Q^{\Gamma_i}(\underline{r}_i)\underline{r}_i$.

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 willingness to pay truthfully:

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

- each agent receives non-negative utility from the mechanism but does not receive a positive money transfer:

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

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

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

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 $\underline{U}_i \in [0, Q^{\Gamma_i}(\underline{r}_i)\underline{r}_i]$.

The objective function. We assume that the designer maximizes the expectation of a weighted sum of revenue and agents’ utilities weighted by their social welfare weights. The following observation—which has been made before in different contexts—implies that our definition of a feasible mechanism is without loss of generality for maximizing this objective.

implementable allocation rules—it only constrains lump-sum transfers.

Claim 1 (Jehiel and Moldovanu (2001); Che et al. (2013)¹³). *The designer cannot increase the expectation of her objective function by using an incentive-compatible mechanism that elicits information about λ .*

Claim 1 is intuitive: Since, conditional on r , λ has no bearing on the individual's preferences over choices offered by any allocation mechanism, no mechanism can truthfully elicit information about λ . As a consequence, the designer must form beliefs about λ based on the information she does observe which is r and i . Define $\lambda_i(r) = \mathbb{E}[\lambda | i, r]$ as the expectation of λ conditional on i and r , under their joint distribution. To distinguish $\lambda_i(r)$ from the underlying social welfare weight λ , we call $\lambda_i(r)$ the *Pareto weight* on an agent with label i and willingness to pay r .

With this, we can write the designer's objective function as

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

where $\alpha \geq 0$ is the weight on revenue. For technical reasons, we assume 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 willingness to pay above r . With fully

¹³See also Dworzak [®] Kominers [®] Akbarpour for the formulation of this claim in an analogous economic context.

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 welfare 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. Claim 1 makes it clear that $\lambda_i(r)$ —the Pareto weight—is effectively a primitive of our model. Indeed, we have not even specified the underlying joint distribution of the three dimensions of the type vector because it is only relevant for determining the shape of $\lambda_i(r)$ and the marginal distributions of (i, r) . Nevertheless, we introduced the unobserved social welfare weights to emphasize the economic forces that give rise to any particular $\lambda_i(r)$.

The average Pareto weights $\bar{\lambda}_i$ and $\bar{\lambda}_j$ differ to the extent that the labels i and j capture observable information that is correlated with the social welfare weights. For example, if tax data allows the designer to determine the income bracket for each agent, then agents associated with lower income brackets might receive a higher average Pareto weight.

Similarly, dispersion in $\lambda_i(r)$ for any given i should be interpreted as residual correlation between willingness to pay and social welfare weights, conditional on i . For a concrete example, suppose first that no observable information is available, but we elicit the willingness to pay of two individuals, A and B , for a high-quality house in an attractive neighborhood. Agent A is willing to pay \$500,000, while agent B is only willing to pay \$50,000. While the differences between A and B ’s willingness to pay may be driven by preferences, they likely also reflect characteristics such as income and opportunity cost of money that could in turn affect the welfare weights. Thus, without observing the characteristics that inform welfare directly, the designer may place a higher Pareto weight on the agent with lower willingness to pay, reflecting her Bayesian belief that this agent is more likely to be poor. Now suppose that the designer additionally has access to tax data, and she knows that these agents A and B have the same income. Conditional on that information, the correlation between willingness to pay and welfare weights will get weaker; willingness to pay originally appeared to be more strongly correlated with welfare weights due to the omission of a rel-

evant variable—income. However, that correlation is likely still negative, as long as other unobserved characteristics—such as health shocks or future job prospects—influence *both* the welfare weights and willingness to pay for a house. More generally, the more informative the label, the less residual correlation one would expect between r and λ . There are also cases when the correlation can be very weak even in the absence of very informative labels. For example, when the good to be allocated is a movie ticket, and agent A is willing to pay \$10, while agent B is only willing to pay \$1, the most likely inference is that agent A enjoys watching movies more than agent B ; not that B is very poor or otherwise socially disadvantaged. Summarizing, $\lambda_i(r)$ naturally depends on the strength of the underlying social preferences and the degree to which they are uncovered by the label i , but also on the characteristics of the good, such as the importance of personal taste for determining the willingness to pay for it.

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 as 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 the revenue 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 possibility is that the designer can spend the monetary surplus generated by the mechanism on a socially valuable outside cause.

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

We identify an optimal mechanism for our setting in two steps:

1. First, the objects are allocated “across” groups: F is split into $|I|$ CDFs F_i^* .
2. Then, 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). When i is fixed, we will sometimes abuse terminology slightly by referring to r as the agent's type.

We say that there is *assortative matching* among types $r \in [a, b]$, if $Q_i^*(r) = F_i^{-1}(G_i(r))$ for all $r \in [a, b]$. To account for the possibility that some objects may remain unallocated, we say that the matching is *effectively assortative* when it is assortative for $r \in [\inf\{r : Q_i^*(r) > 0\}, \bar{r}_i]$. 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 $r \in [a, b]$.¹⁴

¹⁴Throughout, $H^{-1}(x)$ 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 $x \leq \max_y H(y)$.

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 coincide (if F_i is constant in the relevant range). In particular, the two concepts do not 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 .¹⁵

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

As mentioned in Section 1.1, 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.

¹⁵An 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.

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

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 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

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

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

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

Theorem 2 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-decreasing.

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.

4 Economic implications

Proposition 1 (In-kind redistribution). *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 ($r_i > 0$), then there exists $r_i^* > r_i$ 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. The second assumption says that the good is essential, that is, the willingness to pay is bounded away from 0.

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, the good is essential. 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.

Our next result describes conditions under which the 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 .

Proposition 2 (Full randomization). *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 $\bar{\lambda}_i$ be 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. 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 $(\bar{\lambda}_i - \alpha) r_i$ 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: the good is allocated for free to those satisfying certain verifiable eligibility criteria, as

captured by the label i . The result states that for such a form of redistribution to be optimal, several conditions must be met. 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 $(1/2)(\bar{r}_i + \underline{r}_i)\bar{\lambda}_i$, 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 the assumptions of Proposition 2 fail, the optimal mechanism features assortative matching at the “top of the distribution.” Formally, we say that there is *assortative matching at the top* if the mechanism allocates the highest-quality objects assortatively to agents with willingness to pay r above some threshold. For this result, we assume that Pareto weights are non-increasing, an assumption that is justified when willingness to pay is negatively correlated with the unobserved welfare weights. This is a natural case that arises, for example, when the designer has a preference for redistribution towards agents with lower wealth, and willingness to pay increases—everything else being equal—with the wealth level.

Proposition 3 (Assortative matching at the top). *If Pareto weights are non-increasing, and $\alpha \geq \bar{\lambda}_i$, any optimal mechanism features assortative matching at the top within group i .*

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 (Fully assortative matching). *Effectively assortative matching is optimal within group i 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.¹⁶ This requirement is inconsistent with an upward jump of Ψ_i at 0 which is what the first assumption

¹⁶Strictly 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 the case when Ψ_i is only weakly concave.

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.

To show why fully assortative matching often fails to be optimal, note that the second condition in Proposition 4 is violated as long as for some r , (assuming enough differentiability)

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

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, even conditional on i , willingness to pay is strongly correlated with the unobserved social welfare weights. This is more likely to be true when the label i is not very informative of the agents' underlying weights (e.g., when the label does not include any information about the agent's income), forcing the designer to rely on willingness to pay to identify individuals who are most in need.

Dworczak [®] Kominers [®] Akbarpour (2020) 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.

Proposition 4 also relates to 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 non-decreasing.¹⁷ Our paper also points out that this result is true *only* under the assumption that the good is non-essential (that is, $\underline{r}_i = 0$, an assumption that is made in all of the above papers).

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 allocation in the level of inequality.

Proposition 5 (Non-monotonicity in the use of non-market mechanisms). *Suppose that $J'_i(r) \geq \underline{J}_i > 0$, for all r and some \underline{J}_i . Consider any sequence of within-group- i problems indexed by $n \in \mathbb{N}$, differing only in the specification of Pareto weights λ_i^n . Assume that, for all n , $\lambda_i^n(r)$ is non-increasing in r , and $\bar{\lambda}_i^n \leq \alpha$. 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 convergent sequence of optimal allocations converges point-wise to effectively assortative matching.

The first case of Proposition 5 is intuitive because 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 low willingness to pay uniquely identifies the individuals in society that the designer cares about (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

¹⁷Because 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 $\alpha = 0$.

weight stays below the weight on revenue, and that the allocation to the lowest type affects 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 each group i . We now focus on what the above insights, combined with Theorem 2, tell us about the allocation of objects *across* the groups. We begin by characterizing the structure of $\text{supp}(F_i^*)$ —the set of object qualities allocated to group i —in simple cases in which allocation takes the same form in all groups.

Proposition 6 (Across-group allocation with random matching). *Suppose that under the optimal mechanism, 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 $\text{supp}(F_i^*) = [q_i, q_{i+1}] \cap \text{supp}(F)$, 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.¹⁸ 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 (Across-group allocation with assortative matching). *Suppose that under the optimal mechanism, all groups have effectively 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 $\text{supp}(F_i^*) = [\underline{q}_i, \bar{q}_i] \cap \text{supp}(F)$, 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

¹⁸By “uniformly higher quality” we mean that the lowest quality allocated to one group is higher than the highest quality allocated to another group.

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 always gets the highest-quality object even if the designer puts no weight on the welfare of agents in group i , and a high weight 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 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 (Highest quality to a random-matching group). *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. 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 exclude agents with high willingness to pay.

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

Proposition 9 (Intermediate quality to a random-matching group). *Suppose that $|I| = 2$ and that in group 1 there is effectively assortative matching, and in group 2 there is fully random matching. Then, there exist $\underline{q} \leq \bar{q}$ such that $\text{supp}(F_2^*) = [\underline{q}, \bar{q}] \cap \text{supp}(F)$, and $\text{supp}(F_1^*) = ([0, \underline{q}] \cup [\bar{q}, 1]) \cap \text{supp}(F)$.*

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 Market Design Implications

Our analysis leads to a number of general insights on the design of allocation mechanisms under redistributive concerns. The optimal mechanism is always a combination of (i) random matching, which can be seen as a form of in-kind redistribution, and (ii) assortative matching, which is effectively the allocation that would arise in a competitive market equilibrium. Moreover, we can characterize how the underlying social preferences, expressed by the unobserved social welfare weights, filter through the observable information to produce the optimal allocation.

In-kind redistribution. Random allocation can be optimal only if the designer is able to identify the inequalities in the unobserved social welfare weights. The designer can observe, directly or through the mechanism, the label and the willingness to pay. The label and the willingness to pay thus give rise to two distinct reasons to use in-kind redistribution:

1. *Label-revealed inequality:* If some label i identifies a group of agents that have a high welfare weight *on average* (higher than the weight on revenue α), then in-kind redistribution becomes optimal when the good being allocated is *essential* (Proposition 1). In this case, the designer does not attempt to redistribute to any particular subset of agents within group i . Rather, she wants to increase the utility of *all* agents in the group uniformly. In-kind redistribution achieves this goal for essential goods because an essential good has a positive value to all agents, even those with the lowest willingness to pay. Low-WTP agents benefit by getting a low-quality good for free, while high-WTP benefit by paying a lower price for higher quality goods. Sometimes, the trade-off is resolved in favor of a fully-random allocation, as predicted by Proposition 2.

Food stamp programs serve as a good illustration. Group i can be defined by a set of verifiable eligibility criteria, such as low income, that are strongly correlated with what society associates with those most in need. For various reasons, it might be impractical or impossible to give monetary transfers to group i . Then, since food is arguably as essential good, in-kind redistribution can be justified by our Proposition 1. It is even plausible that the assumption of Proposition 2 is satisfied, because condition (4.1) is more likely to hold when the dispersion in willingness to pay is low.¹⁹ In that case, our framework predicts that everyone who is eligible should receive the same food stamp free of charge.

In contrast, consider the example of public housing programs. In case of housing, we can plausibly assume that, even conditional on satisfying some eligibility criteria, peo-

¹⁹Indeed, as \bar{r}_i tends to \underline{r}_i , condition (4.1) “converges” to the requirement that $\alpha < \bar{\lambda}_i$.

ple will differ significantly in their willingness to pay. Then, a fully random allocation is probably suboptimal. A superior solution, based on Proposition 1, is to provide the lowest-quality houses at a minimal price in a lottery, and use a price gradient for granting access to higher-quality housing. Using a low price for the lottery ensures that even those who opt for higher quality can be charged a below-market price. At the same time, the price gradient ensures a more efficient allocation, and raises more revenue.

2. *WTP-revealed inequality*: The second, distinct reason for using random allocation is when the label fails to accurately identify those most in need. This is the case when there is high residual correlation between willingness to pay and the unobserved welfare weight. In these cases, in-kind redistribution is used to specifically *target* those individuals within group i who are likely to have a high welfare weight. The most plausible case is when low willingness to pay reveals a high expected welfare weight, which corresponds to violating the assumption of Proposition 4. In that case, the designer introduces a reduced-price lottery for low-quality objects in order to separate low- and high-WTP agents—and subsidize the former via a reduced price.

Unlike the prior case, this form of in-kind redistribution might be optimal in markets for non-essential goods, if there are sufficient reasons to believe that low willingness to pay is a consequence of unobserved adverse social and economic circumstances, rather than taste or preferences. A good example is health care for non-life-threatening conditions. Many European countries provide a wide range of health services to all citizens at low or zero prices but often at a quality that is significantly lower than the quality of services in the private sector (e.g., because of long waiting times). While there may be many justifications for such systems, one is that using the public health system is likely to be correlated with characteristics of an individual that are not easily observable and yet influence the social welfare weight. In contrast, the US Medicare and Medicaid are probably examples of the first reason for using in-kind redistribution, since they are conditional on satisfying certain observable eligibility criteria.

Market allocation. An assortative matching, which can be seen as a form of market allocation, also arises for two distinct reasons.

1. *The revenue motive*: As predicted by Proposition 3 and 4, assortative matching becomes optimal when α —the weight on revenue—exceeds the average Pareto weight $\bar{\lambda}_i$ in a given group i . This is the case, in particular, when the revenue can be used to subsidize any group of agents via direct lump-sum transfers. But more generally, this

case obtains when the designer can use the revenue to subsidize an outside cause that is valuable from a social perspective. In such situations, it is socially optimal to use assortative matching to maximize revenue in order to subsidize that cause with the resulting monetary surplus.

When the government allocates public goods, such as spectrum licenses or oil and gas leases, to corporations, it is arguably the case that the marginal value of revenue—which funds the government budget—exceeds the weight that the government places on the welfare of the corporations and their owners. In such cases, it is optimal to use auctions that allocate the highest-quality goods to those with highest willingness to pay.

However, the same force behind optimality of assortative matching also applies in any situation in which direct label-specific lump-sum payments are feasible (so that $\alpha \geq \bar{\lambda}_i$ for any group i). For example, if it is feasible to give cash transfers to those eligible for public housing (perhaps in the form of tax credit), then there is an argument against using lotteries to allocate public housing—we can do better by allocating assortatively at least at the top of the distribution of willingness to pay, raising revenue, and using that revenue to fund monetary transfers to all eligible agents (see Proposition 3).

2. *The efficiency motive:* Assortative matching is also optimal for maximizing the efficiency of the allocation. This is the second force that works in favor of a market allocation, even if the weight on revenue α is strictly below the average Pareto weight $\bar{\lambda}_i$. Efficiency becomes the dominant force when Pareto weights do not vary too much with willingness to pay, conditional on some label i . Indeed, Proposition 4 implies that a fully assortative matching becomes optimal when $\alpha \geq \max_r \{\lambda_i(r) - \Lambda_i(r)\}$ which can be true even for very low α when there is little dispersion in $\lambda_i(r)$.²⁰ Low dispersion in $\lambda_i(r)$ can arise in two cases: (i) when the designer does not have strong redistributive preferences to begin with (there is little dispersion in the unobserved welfare weights) or, more interestingly, (ii) when willingness to pay is not correlated with the underlying welfare weights, conditional on the label.

Observation (ii) helps explain why a market allocation is desirable for most goods and services even when the designer has strong preferences for redistribution. Agents' needs are unlikely to be strongly correlated with willingness to pay for goods that are relatively cheap (affordable, at least in small quantities, to most people) and whose value depends heavily on tastes (soft drinks, video games, books etc.) Additionally,

²⁰Recall that $\lambda_i(r) - \Lambda_i(r)$ is the difference between the expected welfare weight on agent (i, r) and the average expected welfare weight on all agents (i, \tilde{r}) for $\tilde{r} \geq r$.

the residual correlation between willingness to pay and the unobserved welfare weights decreases when more information becomes available in the form of labels (see the interpretation of the model in Section 2). For example, if a country provides free health care to eligible citizens, it becomes less likely that the low willingness to pay in the non-eligible group reflects adverse social or economic circumstances, since these circumstances would likely be partly captured by the label. Hence, using in-kind redistribution to address label-reveled inequality should be expected to coexist with a market allocation to agents associated with non-preferential labels.

While it is useful to separately identify various forces behind random and assortative allocation, these forces will typically co-exist, leading to more subtle conclusions. An example is our Proposition 5, which predicts a certain non-monotonicity in the use of random allocation. In the setting of Proposition 5, the weight on revenue is above the average Pareto weight, and we vary the strength of the residual correlation between willingness to pay and the unobserved welfare weights. Thus, there is a trade-off between the revenue motive (a force behind assortative matching) and the WTP-revealed inequality (a force behind random matching). When willingness to pay is relatively uninformative about agents' needs, the revenue motive dominates. As willingness to pay becomes increasingly informative, the designer may opt for a partially random allocation to identify those most in need through the mechanism. Eventually, when only a small fraction of agents receive an increasingly high welfare weight, the use of random allocation becomes negligible because the revenue motive dominates for all remaining agents.

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, then 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 Pareto 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.

6 Conclusion

Using a mechanism design framework, we examined the optimal way for a designer to allocate a set of goods of heterogeneous quality to agents differing in three dimensions: the unobserved social welfare weight, publicly observable label, and willingness to pay for quality. We identified an optimal mechanism, and showed that its form depends on how informative the observable characteristics are about the underlying and unobservable “need” captured by the welfare weight.

Focusing on an objective function that assigns arbitrary welfare weights to market participants sets this work apart from the standard mechanism design paradigm. Indeed, while the mechanism design literature has developed an impressive framework for designing revenue-maximizing auctions and allocatively efficient mechanisms, there has been far less focus on how to use those same tools to understand how the structure of mechanisms responds to redistributive goals. Our paper is thus one of relatively few attempts to use mechanism design to give guidance to real-world market designers about how to optimally structure market-level redistributive systems. We hope to see more work devoted to this problem.

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.
- ASHLAGI, I., F. MONACHOU, AND A. NIKZAD (2020): “Optimal Dynamic Allocation: Simplicity through Information Design,” Available at SSRN.
- ATKINSON, A. B. AND J. E. STIGLITZ (1976): “The design of tax structure: direct versus indirect taxation,” Journal of Public Economics, 6, 55–75.

- BESLEY, T. AND S. COATE (1991): “Public provision of private goods and the redistribution of income,” The American Economic Review, 81, 979–984.
- 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.
- CHE, Y.-K., W. DESSEIN, AND N. KARTIK (2013): “Pandering to Persuade,” American Economic Review, 103, 47–79.
- CONDORELLI, D. (2012): “What money can’t buy: Efficient mechanism design with costly signals,” Games and Economic Behavior, 75, 613 – 624.
- (2013): “Market and non-market mechanisms for the optimal allocation of scarce resources,” Games and Economic Behavior, 82, 582–591.
- 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 (2020): “Redistribution through markets,” Econometrica (forthcoming).
- 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.
- JEHIEL, P. AND B. MOLDOVANU (2001): “Efficient Design with Interdependent Valuations,” Econometrica, 69, 1237–1259.
- 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.
- PATHAK, P. A., T. SÖNMEZ, M. U. ÜNVER, AND M. B. YENMEZ (2020): “Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value Behind in Health Care Rationing,” 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.

SCHMIDT, H., P. PATHAK, T. SÖNMEZ, AND M. U. ÜNVER (2020): “Covid-19: how to prioritize worse-off populations in allocating safe and effective vaccines,” British Medical Journal, 371.

TOIKKA, J. (2011): “Ironing without control,” Journal of Economic Theory, 146, 2510 – 2526.

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 Insight 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 $U(r) + \int_{\underline{r}}^r Q(\tau)d\tau$. 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 $U(\underline{r}) = Q(\underline{r})\underline{r}$.

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.




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) .

B Proofs omitted from the main text

B.1 Proof of Theorem 1

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

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

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

Φ 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 willingness to pay 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 concurrent 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 [Dworczak 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 verifying 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:

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

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

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,²¹ 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

²¹In case there are multiple Lagrange multipliers, we pick the largest one.

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 $S(q) = -L(q)$.

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

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

Thus, if $\alpha \geq \bar{\lambda}_i$, we would get $(\alpha - \bar{\lambda}_i) \bar{r}_i < (\alpha - \bar{\lambda}_i) \underline{r}_i$ 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 $\Psi_i(0) \geq \Psi_i(1) - \Psi'(1)$. Together this implies that $\Psi_i(x) \leq (1 - x)\Psi_i(0)$ for all x . \square

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

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

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 effectively assortative matching, either $|\text{dco}(\Psi_1)'(0)| = 0$ or $\text{dco}(\Psi_1)'(0) = \Psi_1'(0) \geq -\alpha \underline{r}_1$. This yields the condition $\alpha \underline{r}_1 \leq \int_{\underline{r}_2}^{\bar{r}_2} \tau \lambda_2(\tau) dG_2(\tau)$. \square