Redistribution through Markets

Piotr Dworczak ® Scott Duke Kominers ® Mohammad Akbarpour *

First Version: February 7, 2018
This Version: May 26, 2018

Abstract

When macroeconomic tools fail to respond to wealth inequality optimally, regulators can still seek to mitigate inequality within individual markets. A social planner with distributional preferences might distort allocative efficiency to achieve a more desirable split of surplus, for example, by setting higher prices when sellers are poor—effectively, using the market as a redistributive tool.

In this paper, we seek to understand how to design goods markets optimally in the presence of inequality. Using a mechanism design approach, we uncover the constrained Pareto frontier by identifying the optimal trade-off between allocative efficiency and redistribution in a setting where the second welfare theorem fails because of private information and participation constraints. We find that competitive equilibrium allocation is not always optimal. Instead, when there is substantial inequality across sides of the market, the optimal design uses a tax-like mechanism, introducing a wedge between the buyer and seller prices, and redistributing the resulting surplus to the poorer side of the market via lump-sum payments. When there is significant within-side inequality, meanwhile, it may be optimal to impose price controls even though doing so induces rationing.

Keywords: optimal mechanism design, redistribution, inequality, welfare theorems
JEL codes: D47, D61, D63, D82, H21

*The "®" symbol indicates that the authors' names are in certified random order, as described by Ray ® Robson (2018). The authors thank Benjamin Brooks, Eric Budish, Raj Chetty, Pawel Doligalski, Steven Durlauf, Mehmet Ekmekci, Abigail Fradkin, Alex Frankel, Jonathan Gould, Michael Grubb, Ravi Jagadeesan, Emir Kamenica, Paul Kominers, Moritz Lenel, Shengwu Li, Elliot Lipnowski, Giorgio Martini, Paul Milgrom, Jeff Miron, Roger Myerson, Eric Nelson, Alessandro Pavan, Canice Prendergast, Doron Ravid, Phil Reny, Marzena Rostek, Alvin Roth, Emma Rothschild, Larry Samuelson, Amartya Sen, Andy Skrzypacz, Tayfun Sönmez, Stefanie Stantcheva, Philipp Strack, Alex Teytelboym, Utku Ünver, Marek Weretka, Bob Wilson, Eric Zwick, and numerous seminar audiences for helpful comments. We also thank the anonymous referees of the Economics and Computations conference for valuable feedback. All three authors 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, and the Human Capital and Economic Opportunity Working Group sponsored by the Institute for New Economic Thinking.

Dworczak: Becker Friedman Institute, University of Chicago; dworczak@uchicago.edu. *Kominers:* Entrepreneurial Management Unit, Harvard Business School; Department of Economics, Center of Mathematical Sciences and Applications, and Center for Research on Computation and Society, Harvard University; and National Bureau of Economic Research; kominers@fas.harvard.edu. *Akbarpour:* Graduate School of Business, Stanford University; mohamwad@stanford.edu.

1 Introduction

In many markets, there is systematic wealth inequality among participants. When global income redistribution is not feasible, market designers can still seek to mitigate inequality within individual markets. For example, property market regulators frequently use tools like rent control in response to the wealth disparities between renters and owners. In the only legal marketplace for kidneys—the one in Iran—the government sets a price floor partly because it is concerned about welfare of organ donors, who tend to come from low-income households.

In this paper, we seek to understand how to optimally design marketplaces, taking wealth inequality as given. When the designer has distributional preferences, competitive equilibrium pricing is not necessarily optimal. For example, if sellers in the market are poorer than buyers, they tend to value money more relative to the good being traded, and hence will be willing to sell at lower prices, all else equal. However, precisely because sellers have relatively high values for money, a designer who cares about inequality might prefer to set the price higher—effectively, using the market as a redistributive tool.

Our framework is as follows. There is a market for a single indivisible good, with a large number of prospective buyers and sellers. Each agent is characterized by *how much she values the good relative to a monetary transfer*—a ratio which we refer to as the *rate of substitution*. A market designer chooses a mechanism to maximize a social objective, in a manner we explain soon. The designer knows the distribution of agents' characteristics but does not observe any individual agents' rates of substitution. The mechanism must clear the market, maintain budget balance, and respect both incentive-compatibility and individual participation constraints.

In settings like ours, finding the optimal marketplace design is straightforward in a quasi-linear framework with perfectly transferable utility—that is, if monetary transfers between agents have no impact on the social objective function. Because of its tractability, the perfectly transferable utility framework has become the dominant model in auction and mechanism design theory; however, that framework does not allow us to study wealth inequality, as it implicitly embeds the assumption that each agent values money equally. Our work exploits the observation that while quasi-linearity is key for tractability, the assumption of perfectly-transferable utility is logically independent and can be relaxed. Specifically, in our setting we maintain the assumption that each agent has a constant rate of substitution between the good and money (independent of the wealth) but we allow the possibility that money is more valuable to some individuals (from a social perspective).[1] The designer

[1]The dispersion in marginal values for money has been used extensively to model inequality in the public

maximizes the total value of allocating the good and money to agents. Effectively, we are accounting for macro-level wealth differences (which generate differences in agents' values for money) while using a local first-order approximation (quasi-linearity) to preserve tractability.[2] We emphasize that the notion of a "value for money" only makes sense in the context of social preferences—individuals' preferences are fully described by their (constant) rate of substitution between the good and the monetary transfer.

We show formally that the above approach is equivalent to maximizing a weighted sum of agents' utilities, for an appropriately chosen set of Pareto weights, in an otherwise standard quasi-linear model. Using the Pareto weight interpretation, the quasi-linear model with perfectly transferable utility can be seen as making the (strong!) assumption that the designer's Pareto weights are independent of agents' willingness to pay (and whether she is a buyer or a seller). As a consequence, the standard framework "forces" the designer to pick a particular point on the Pareto frontier, the same that would be selected by the so-called "Kaldor-Hicks" criterion. In our setting, by contrast, the designer can express preferences over the entire Pareto frontier.

We thus seek to characterize the Pareto frontier in the presence of incentive-compatibility (IC) and individual-rationality (IR) constraints. Under IC and IR constraints—that is, even when agents' values are not observed and participation in the marketplace is voluntary—the competitive equilibrium mechanism is feasible in our large-market setting. By the first welfare theorem, competitive equilibrium allocation is Pareto efficient; due to quasi-linearity of agents' preferences, this is equivalent to maximizing total willingness to pay (or, again equivalently, maximizing total utility with equal Pareto weights). If private information and voluntary participation were to be ignored, we would have from the second welfare theorem that any split of the maximized surplus between agents can be achieved by redistribution of agents' endowments before trading—so that *(i)* there is no trade-off between maximizing total willingness to pay (allocative efficiency) and achieving the desired distributional outcome (under any Pareto weights), and *(ii)* the only optimal mechanism is the competitive equilibrium mechanism. However, the second welfare theorem fails in our framework because redistribution of endowments prior to trading would typically violate both the IC and IR constraints. As a consequence, *(i)* there is a trade-off between allocative efficiency and the distribution of surplus, and *(ii)* the competitive equilibrium mechanism is sometimes suboptimal.

To illustrate the preceding discussion point, in Figure 1.1 we depict the hypothetical

finance literature, which we review in Section 1.1.

[2]This approach means that we implicitly assume that the market under consideration is a small enough part of the economy that transactions do not substantially change agents' wealth levels.

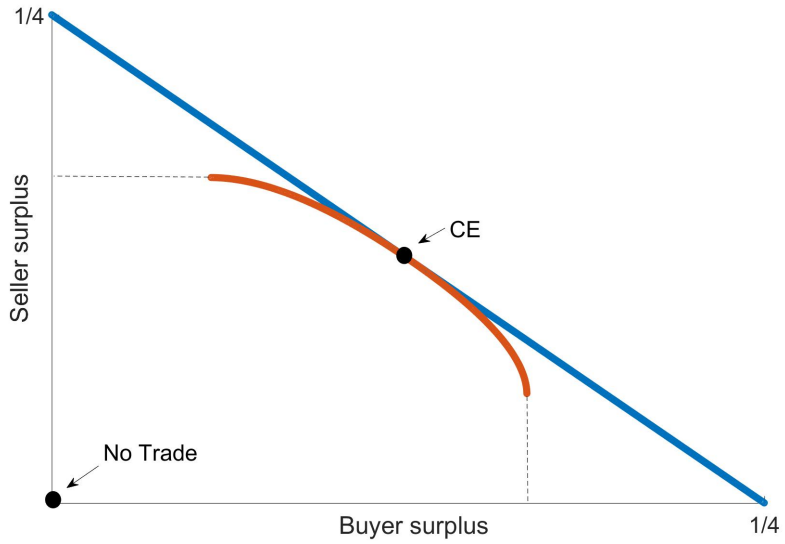


Figure 1.1: The unconstrained Pareto curve (blue line) and the constrained Pareto curve (red curve)

Pareto frontier that would arise in a marketplace with an equal mass of buyers and sellers and a uniform distribution of rates of substitution if we relaxed the IC and IR constraints (blue curve). As we expect from the second welfare theorem, unconstrained Pareto frontier is linear because agents' preferences are quasi-linear. However, the constrained Pareto curve that the designer can achieve (red curve) is strictly concave. The two frontiers coincide only at the competitive equilibrium mechanism. If the social objective function puts a sufficiently high Pareto weight on one side of the market (for example sellers), the optimal mechanism will differ from the competitive equilibrium mechanism. Intuitively, the frontier becomes strictly concave because giving more surplus to one side of the market requires distorting allocative efficiency to preserve the IC and IR constraints: For example, giving more money to sellers requires raising more revenue from the buyer side which, in the presence of IC and IR constraints, can only be achieved by limiting supply. Studying this optimal trade-off is the focus of this paper.

Our first result gives a general characterization of the class of mechanisms that generate the (constrained) Pareto frontier. For an arbitrary set of Pareto weights, in the optimal mechanism, each seller can choose to sell at a low price with probability 1, or sell at a high price with probability determined by rationing. Symmetrically, buyers can choose between a high price with guaranteed purchase and a low price with rationing. Additionally, the mechanism redistributes the monetary surplus (if there is any) as a lump-sum payment to the side of the market that has a higher average Pareto weight in the social objective function.

Consequently, computing the optimal mechanism requires optimizing over six parameters subject to the market-clearing constraint. Thus, the constrained Pareto frontier (e.g. the one depicted in Figure 1.1) is generated by a relatively small set of simple mechanisms. The simple form of optimal mechanisms in our framework stems from our large-market assumption. We notice that any mechanism can be represented as a lottery over quantities, and hence the market-clearing constraint boils down to an equal-means constraint—the average quantity sold to sellers must equal the average quantity sold to buyers. Thus, the optimal value is obtained by concavifying the objective function at the equilibrium trade volume. Concavification implies that the optimal scheme is a lottery over at most two points for each side, yielding the two-price characterization presented above.

Given our general characterization of optimal mechanisms, we attempt to understand which mechanisms are optimal depending on the characteristics of market participants. To this end, we use an interpretation of our model with agents being characterized by two-dimensional types—a value for the good and a marginal value for money—and the designer maximizing the total value of allocating the good and money. (As noted earlier, this is equivalent to agents being characterized by a one-dimensional rate of substitution, and the designer maximizing a Pareto-weighted sum of utilities.) We show that the form of the optimal mechanism depends on the type of inequality present in the market. Cross-side inequality, which measures average disparity between buyers and sellers, determines the direction of the lump-sum payments—the surplus is redistributed to the side of the market with a higher average value for money. Same-side inequality, which measures dispersion within each side of the market, decides about the use of rationing.

Under certain regularity conditions, when same-side inequality is not too large—regardless of cross-side inequality—the optimal mechanism is a price mechanism and there is no rationing. However, the designer may impose a wedge between the buyer and seller prices. The degree of cross-side inequality determines the magnitude of this tax, and hence the size of the lump-sum transfer to the "poorer" side of the market. On the other hand, when same-side inequality is substantial on the seller side, the optimal mechanism may ration the sellers. Rationing, relative to lump-sum redistribution, allows the designer to reach the "poorest" sellers by raising the price that they receive above the market-clearing level. In such cases, the designer uses the redistributive power of the market: willingness to sell at a given price can be used to select sellers with relatively higher values for money. However, as long as our regularity conditions hold, rationing is never optimal on the buyer side, even in the presence of very strong same-side buyer inequality. This is because the decision to trade identifies buyers with relatively low values for money. Poorer buyers are precisely those that do not participate, and hence the only available tool to increase their wealth is a lump-sum trans-

fer. In this sense, our analysis uncovers a fundamental (albeit ex-post intuitive) asymmetry between buyers and sellers with respect to the redistributive role of markets.

While the optimal mechanism mitigates cross-side inequality with lump-sum transfers, we can easily imagine contexts in which lump-sum redistribution is not possible. For example, if we pay a constant amount to sellers regardless of whether they trade, then all sellers get strictly positive utility from participating in the marketplace; this creates an incentive for excess entry of agents, which in turn undermines overall budget balance. We thus consider a restriction of our setting in which lump-sum redistribution is ruled out. We find there that rationing can sometimes emerge as an optimal response to cross-side inequality (even if there is no same-side inequality), and in particular can be used on the buyer side (as it is then essentially the only tool available to the designer).

Our analysis indicates that there might be substantial scope for market design to improve outcomes in contexts with underlying inequality. Moreover, our results give some guidance as to the types of mechanisms that market designers seeking to mitigate inequality might use. We may think of markets as serving two purposes simultaneously: they both allocate objects and transfer money among participants. From a social welfare perspective, sometimes it is worth distorting the allocative role to make better use of the transfer role.

Additionally, our results may help explain the widespread use of price controls and other market-distorting regulations in settings with inequality. Philosophers (Satz (2010); Sandel (2012)) and policymakers (Roth (2007)) often speak of markets as having the power to "exploit" participants through prices. The possibility that prices could somehow take advantage of individuals who act according to revealed preference seems fundamentally unnatural to an economist. Yet our framework illustrates at least one sense in which the idea has a precise economic meaning: as inequality increases, the competitive equilibrium price shifts in response to some market participants' relatively stronger desire for money, leaving more of the surplus with the other agents. At the same time, however, our work shows that the proper social response is not banning or eliminating markets—as Sandel (2012) and others suggest—but rather designing the market-clearing mechanism in a way that directly attends to inequality. A welfare-maximizing social planner might prefer to "redistribute through the market" by choosing a market design that gives up some allocative efficiency in exchange for creating more surplus for poorer market participants.

The remainder of this paper is organized as follows. Section 1.1 reviews the related literature in mechanism design, public finance, and other areas. Then, Section 2 presents a simple example illustrating how inequality may lead a welfare-maximizing social planner to distort away from competitive equilibrium allocation. Section 3 lays out our full model of mechanism design under social preferences. In Section 4, we characterize optimal mechanisms

in the general case; then, in Section 5, we use our characterization result to understand optimal design in the presence of inequality. Section 6 contains the analysis under the additional constraint that lump-sum transfers are not feasible. Section 7 contains a numerical example and Section 8 concludes.

1.1 Related work

For frictionless markets, competitive equilibrium allocation (colloquially, "market pricing") is focal in economics because of its efficiency properties (Arrow and Debreu (1954); McKenzie (1959); see also Fleiner et al. (2017)). Even so, economists have long recognized that alternative mechanisms may be superior—even in markets without transaction frictions. Weitzman (1977), for example, showed that rationing can perform better than the price system when agents' needs are not well expressed by willingness to pay. Likewise, optimal allocation schemes diverge from classical intuitions when agents have budget constraints (see, e.g., Laffont and Robert (1996); Che and Gale (1998); Fernandez and Gali (1999); Che and Gale (2000); Che et al. (2012); Dobzinski et al. (2012); Pai and Vohra (2014); Kotowski (2017)), or when agents are risk-averse and have non-linear preferences (Maskin and Riley (1984); Baisa (2017)).

Our paper is also related to studies of price control as a redistributive tool. Viscusi et al. (2005) discuss "allocative costs" of price regulations and Bulow and Klemperer (2012) show when price control can be harmful to all market participants.

Our principal divergence from classical market models—the introduction of heterogeneity in marginal values for money—has antecedents, as well. Condorelli (2013) asks a question similar to ours, in a setting in which a designer wishes to allocate k objects to $n > k$ agents, and in which agents' willingness to pay is not necessarily the characteristic that appears in the designer's objective. Huesmann (2017) studies the problem of allocating an indivisible item to a mass of agents, in which agents have different wealth levels, and non-quasi-linear preferences. Esteban and Ray (2006) study a model of lobbying under inequality in which, similarly to our setting, it is effectively more expensive for less wealthy agents to spend resources in lobbying. More broadly, the idea that it is more costly for low-income individuals to spend money derives from capital market imperfections that impose borrowing constraints on low-wealth individuals; such constraints are ubiquitous throughout economics (see, e.g., Loury (1981); Aghion and Bolton (1997); McKinnon (2010)).

Meanwhile, our attention to a planner with distributional preferences is closely similar to the view taken in the public finance literature, which seeks the socially optimal tax schedule (Diamond and Mirrlees (1971); Atkinson and Stiglitz (1976); Piketty and Saez

(2013); Saez and Stantcheva (2016, 2017)). The idea of using public provision of goods as a form of redistribution (which is inefficient from an optimal taxation perspective) has also been examined (see, e.g., Besley and Coate (1991); Blackorby and Donaldson (1988); Gahvari and Mattos (2007)).

Unlike our work—which considers a two-sided market in which buyers and sellers trade—both optimal allocation and public finance settings typically consider efficiency, fairness, and other design goals in single-sided market contexts. Additionally, our work specifically complements the broad literature on optimal taxation by considering mechanisms for settings in which global redistribution of wealth is infeasible, and the designer must respect a participation constraint.

We find that suitably designed market mechanisms (if we may stretch the term slightly beyond its standard usage) can themselves be used as redistributive tools. In this light, our work also has kinship with the broad and growing literature within market design that shows how variants of market mechanisms can achieve fairness and other distributional goals in settings that (unlike ours) do not allow transfers (see, e.g., Hylland and Zeckhauser (1979); Bogomolnaia and Moulin (2001); Budish (2011); Prendergast (2017)).

2 Simple Example

In this section, we introduce a simple example that illustrates how a welfare-maximizing utilitarian social planner would set prices in the presence of inequality in the marginal utility for money. Without inequality, a competitive equilibrium price is optimal. However, when inequality is substantial, the welfare-maximizing prices are far from the competitive equilibrium level.

To fix ideas, consider a market with a unit mass of buyers and a unit mass of sellers. Each seller owns one unit of an indivisible good, and each buyer demands one unit. The value of the object for any agent in the market is drawn (almost) independently and uniformly at random from $[0, 1]$. We assume preferences are quasi-linear in money, but we relax the assumption that all agents value money equally. Specifically, we assume that a buyer with value v who purchases an object at price p receives utility $v - p$, while a seller with value v who sells an object at price p receives utility $mp - v$, where $m > 1$. Intuitively, one unit of money is worth more—in utility terms—for sellers than for buyers. Consequently, all else equal, a utilitarian social planner would prefer to transfer money from buyers to sellers.

2.1 Setting the socially-optimal price

Suppose that the social planner chooses a price to maximize the sum of agents' utilities. Since no seller would sell at price 0 and all sellers would sell at price $1/m$, for now we limit the choice of price to $p \in [0, 1/m]$. With price p, the social welfare function is:

$$W(p) = \min\left\{1, \frac{mp}{1-p}\right\} \int_p^1 (v-p)dv + \min\left\{1, \frac{1-p}{mp}\right\} \int_0^{mp} (mp - v)dv. \tag{2.1}$$

We note here (and establish a general result later on) that we can equivalently think of sellers as having value for money normalized to 1, and m being instead the Pareto weight attached to seller surplus in the social objective functions: Indeed, (2.1) is equivalent to

$$W(p) = \min\left\{1, \frac{mp}{1-p}\right\} \int_p^1 (v-p)dv + m \min\left\{1, \frac{1-p}{mp}\right\} \int_0^{mp} \left(p - \frac{v}{m}\right) dv,$$

where v/m is the marginal rate of substitution for a seller with value v.

By a straightforward calculation based on a first-order condition, we can characterize the optimal price $p^\star$ (see Figure 2.1). When $m \leq 3$, it is optimal to use a competitive equilibrium price. When $m > 3$, it is optimal to ration the sellers. More precisely:

$$p^\star = \begin{cases} p^{\mathsf{CE}} & m \leq 3, \\ \frac{m-2}{2m-2} & 3 < m < 2+\sqrt{2}, \\ \frac{1}{m} & m \geq 2+\sqrt{2}. \end{cases} \tag{2.2}$$

When $m \geq 2+\sqrt{2}$, all sellers want to sell, and there is a uniform lottery in which the probability that each seller trades is $1 - 1/m$.

To understand the involved trade-off, note that when the price is above the competitive level, rationing the sellers becomes necessary. As a consequence, it is no longer true that sellers with the lowest rates of substitution are always the ones that trade; this reduces allocative efficiency. On the other hand, social welfare is increased because, conditional on selling, sellers receive a larger transfer. When inequality between buyers and sellers is high enough ($m \geq 3$), the second effect dominates.

One way to interpret the result is to notice that as the value for money increases for an individual seller, that seller's benefit from participating in the marketplace also increases. However, when the value for money gets higher for many sellers, there is a certain "poverty externality:" because sellers really want to sell, there is downward pressure on the price, and this reduces sellers' benefits from participating. The example illustrates that a proper

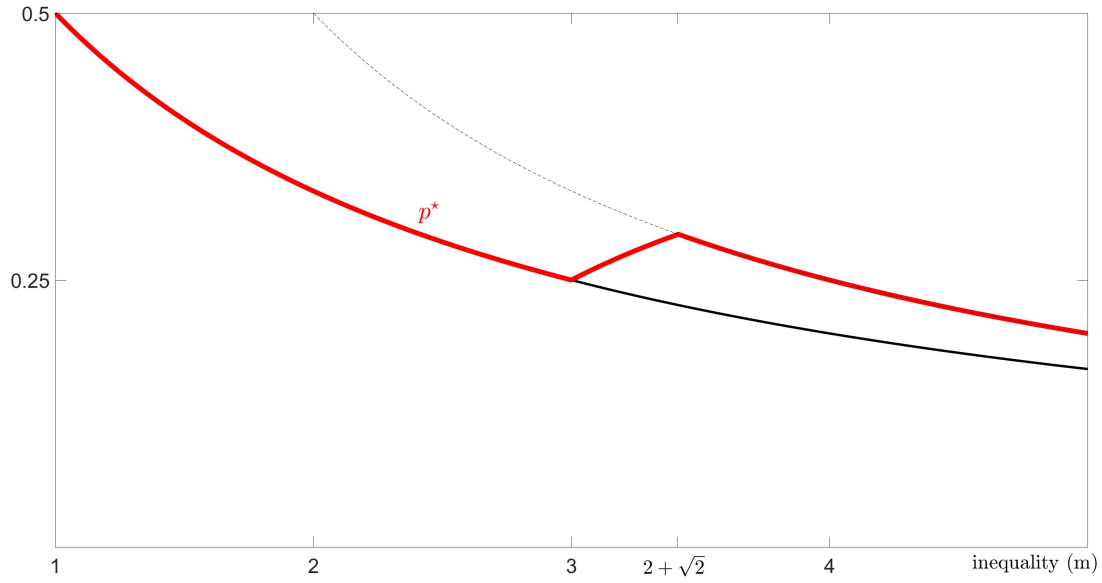


Figure 2.1: The optimal price (thick red line). The solid black line is the competitive equilibrium price, and the dotted black line is the maximum price allowed (the price at which all sellers sell).

regulatory response might be to impose a price floor.

2.2 Improving welfare with the socially optimal price wedge

In the preceding analysis, we restricted attention to single-price mechanisms. Now, we suppose instead that the planner has the option of inserting a "wedge" between buyer and seller prices, choosing a price $p_B \geq 0$ for buyers and a (potentially) different price $p_S \geq 0$ for sellers. The gap between p_B and p_S generates some "revenue" (i.e., monetary surplus), which the planner then redistributes to the sellers. It turns out that the optimal buyer and seller prices are given by $p_B = 1/2$ and $p_S = 1/(2m)$, respectively, as depicted in Figure 2.2. Intuitively, as sellers are uniformly poorer than buyers, the social planner wishes to transfer as much surplus as possible from buyers to sellers. This is achieved at the "monopoly" price, $p_B = 1/2$, at which half of the buyers buy. At price $p_S = 1/(2m)$, half of the sellers sell, and the market clears. (Once we allow a price wedge, as we discuss in detail in Section 5, there is no rationing in this case.) Hence, a total transfer $\underline{U}_S = (1/2)(p_B - p_S) = (m-1)/(4m)$ is redistributed to the seller side in the form of lump-sum payments.[3] Our analysis implies

[3]Figure 2.2 shows the effective price received by sellers who decide to sell (green line), but we emphasize that the transfer $\underline{U}_S$ is also received by remaining sellers who do not sell; lump-sum transfers constitute a lower bound on the utility achieved by the side of the market that receives them, hence the notation "$\underline{U}_S$".

that the competitive equilibrium mechanism ceases to be optimal as soon as there is any inequality, i.e., for any $m > 1$—although the lump-sum transfer is small when inequality is modest.

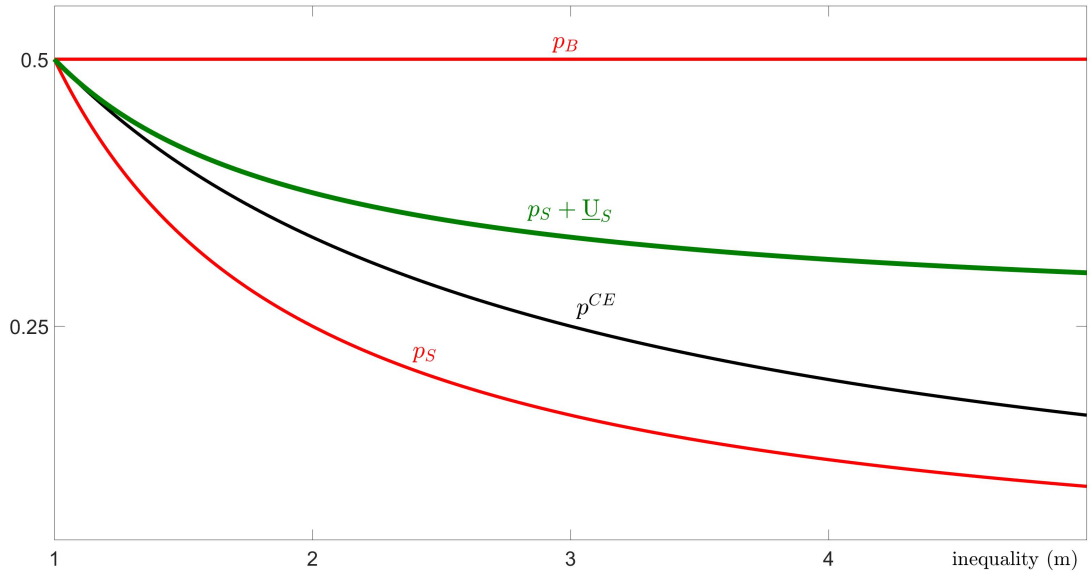


Figure 2.2: The optimal price (thick line) for buyers (p_B) and sellers (p_S). The solid black line is the competitive equilibrium price, and the green line is the total transfer to sellers.

Of course, the space of possible mechanisms is much richer than the one-price or two-price mechanisms just described. Nevertheless, in Section 5 we prove a result implying that the two-price mechanism depicted in Figure 2.2 is optimal (welfare-maximizing) for our example setting. In Section 6 we prove that if lump-sum redistribution is not feasible, then the single-price mechanism depicted in Figure 2.1 is optimal.

3 General Model

In Section 2, we focused on a simple case in which the planner is constrained to choose the same price for buyers and sellers. Here, we relax the requirement of a unique price, and instead allow the planner to choose a general market-clearing mechanism to maximize a weighted sum of agents' utilities.

There is a unit mass of *owners*, and a mass μ of *non-owners* in the market for two *goods*, K and M. All agents can hold at most one unit of K but can hold an arbitrary amount of good M (which can be thought of as money). Owners possess one unit of good K; non-owners have no units of K. Because of the unit-supply/demand assumption, we refer to

owners as *sellers* (S), and to non-owners as *buyers* (B).

In a market where goods K can be exchanged for good M, a parameter that fully describes the behavior of an agent (apart from her ownership type) is the rate of substitution between the goods K and M. In general, the rates of substitution may depend on agents' holdings of the two goods. It is typical in mechanism design to assume quasi-linear and fully transferable utility, under which the numeraire is continuously transferable and all agents value it equally. This has two implications: *(i)* the rate of substitution is constant for any agent, i.e., it does not depend on the endowment of good M, and *(ii)* all agents value good M equally. Components *(i)* and *(ii)* are very different in nature: The former restricts the behavior of each individual agent, while the latter imposes a strong restriction on social preferences. In our analysis, we maintain assumption *(i)* but relax assumption *(ii)* (see Section 3.2).

Under *(i)*, each agent in the economy is characterized by a two-dimensional type (j, r), where $j \in \{B, S\}$, and $r \in \mathbb{R}_+$. If (x_K, x_M) denotes the holdings of goods K and M, then type r's preferences over (x_K, x_M) are induced by a utility function

$$r \cdot x_K + x_M; \tag{3.1}$$

moreover, agents are expected-utility maximizers. We emphasize that the utility function (3.1) only captures the preferences of an individual, and hence it is without loss of generality to normalize the coefficient on x_M to 1. Relaxing assumption *(ii)* means that we cannot compute social welfare by summing (3.1) across all agents.

We denote by $G_j(r)$ the cumulative distribution function of the rate of substitution on side j of the market.[4] We assume that G_j has full support and admits a density g_j on $[\underline{r}_j, \bar{r}_j]$, where $0 \leq \underline{r}_j < \bar{r}_j$, $j \in \{B, S\}$. Moreover, the supports intersect non-trivially so that the equation $\mu(1 - G_B(r)) = G_S(r)$ has a unique solution, implying existence and uniqueness of a competitive equilibrium with strictly positive volume of trade. We denote by $J_S(r) = r + G_S(r)/g_S(r)$ and $J_B(r) = r - (1 - G_B(r))/g_B(r)$ the virtual surplus functions for sellers and buyers, respectively.

3.1 Mechanisms

A designer organizes a marketplace/exchange subject to three constraints:

1. **Anonymity/ Incentive Compatibility** – The designer knows the ownership type of every agent and the distribution of types but does not observe any individuals's rate of substitution.

[4]We adopt the convention that the measures G_j are probability measures, so that to obtain the actual mass of buyers with given characteristics we must multiply by the total mass μ of buyers.

2. **Voluntary Participation/Individual Rationality** – Each agent must weakly prefer the outcome of the mechanism to her status quo.

3. **Market-Clearing** – The designer does not have any goods at the beginning of trade, can only obtain goods from the participants, and must distribute all goods acquired through the mechanism.

By the Revelation Principle, we can restrict attention to direct revelation mechanisms. We can represent mechanisms as a tuple (X_B, X_S, T_B, T_S), where $X_j(r)$ is the probability that an agent with type r trades object K (that is, buys good K if $j = B$, and sells good K if $j = S$), and $T_j(r)$ is the net change in the holdings of good M (which we will refer to as a transfer rule). Formally, we have the following definition of a feasible mechanism.

Definition 1. A feasible mechanism (X_B, X_S, T_B, T_S) consists of $X_j : [\underline{r}_j, \bar{r}_j] \to [0, 1]$, $T_j : [\underline{r}_j, \bar{r}_j] \to \mathbb{R}$, $j \in \{B, S\}$, that satisfy the following conditions:

$$X_B(r)r - T_B(r) \geq X_B(\hat{r})r - T_B(\hat{r}), \qquad \forall r, \hat{r} \in [\underline{r}_B, \bar{r}_B], \tag{IC-B}$$

$$-X_S(r)r + T_S(r) \geq -X_S(\hat{r})r + T_S(\hat{r}), \qquad \forall r, \hat{r} \in [\underline{r}_S, \bar{r}_S], \tag{IC-S}$$

$$X_B(r)r - T_B(r) \geq 0, \qquad \forall r \in [\underline{r}_B, \bar{r}_B], \tag{IR-B}$$

$$-X_S(r)r + T_S(r) \geq 0, \qquad \forall r \in [\underline{r}_S, \bar{r}_S], \tag{IR-S}$$

$$\mu \int_{\underline{r}_B}^{\bar{r}_B} X_B(r) dG_B(r) = \int_{\underline{r}_S}^{\bar{r}_S} X_S(r) dG_S(r), \tag{MC}$$

$$\mu \int_{\underline{r}_B}^{\bar{r}_B} T_B(r) dG_B(r) \geq \int_{\underline{r}_S}^{\bar{r}_S} T_S(r) dG_S(r). \tag{BB}$$

Conditions (IC-B) and (IC-S) are the incentive compatibility constraints, (IR-B)-(IR-S) ensure voluntary participation, (MC) is the market-clearing condition, and (BB) is the budget balance constraint.

3.2 Social preferences

Any individual buyer or seller evaluates outcomes of mechanism (X_B, X_S, T_B, T_S) according to her induced utility level:

$$U_B(r) \coloneqq X_B(r)r - T_B(r),$$
$$U_S(r) \coloneqq -X_S(r)r + T_S(r).$$

We assume that the mechanism designer is utilitarian, and hence chooses an outcome that is Pareto efficient. This, however, does not uniquely pin down social preferences when we relax assumption *(ii)* that states that money can be used for interpersonal utility comparisons. Normally, the social objective is to maximize the total surplus (or willingness-to-pay), as under the "Kaldor-Hicks" criterion. Here, we maximize a weighted sum of agents' utilities but the weights can be arbitrary. This leads to the following definition of an optimal mechanism.

Definition 2. A mechanism (X_B, X_S, T_B, T_S) is optimal with respect to weights $\Lambda := \{\lambda_j(r) \geq 0 : r \in [\underline{r}_j, \bar{r}_j], j \in \{B, S\}\}$ if it is feasible and maximizes

$$\text{TS}(\Lambda) := \mu \int_{\underline{r}_B}^{\bar{r}_B} \lambda_B(r) U_B(r) dG_B(r) + \int_{\underline{r}_S}^{\bar{r}_S} \lambda_S(r) U_S(r) dG_S(r) \tag{OBJ}$$

among all feasible mechanisms.

Our formulation implicitly entails two restrictions on the social objective function. First, the designer cannot assign different weights to two individuals that cannot be separated based on their behavior (e.g., two sellers with the same rate of substitution). Second, any measure-zero set of agents has negligible contribution to the social welfare.[5]

4 Optimal Mechanisms – The General Case

In this section, we present and prove our main technical result, which provides a tight restriction on the class of mechanisms that may emerge as optimal in our framework.

Theorem 1. *There exists an optimal mechanism that takes the following form:*

1. *Each buyer can choose to buy good K at a high price p_B^H or enter a lottery in which she is selected with probability $\delta_B < 1$ to buy at a lower price $p_B^L < p_B^H$.*

2. *Each seller can choose to sell good K at a low price p_S^L or enter a lottery in which she is selected with probability $\delta_S < 1$ to sell at a higher price $p_S^H > p_S^L$.*

3. *The non-negative revenue (i.e., surplus of good M) generated by the pricing system is redistributed as a lump-sum payment $\underline{U}_j$ to agents on the side of the market j with higher average Pareto weight $\Lambda_j = \int \lambda_j(r) dG_j(r)$.*

[5] This last assumption simplifies our exposition, but our conclusions can be extended to the case of discrete weights by an appropriate approximation argument.

Theorem 1 narrows down our set of candidate solutions to a class of mechanisms indexed by eight parameters: four prices, two rationing coefficients, and a pair of lump-sum payments.[6] Our theorem identifies two tools that the designer might use to address inequality: lump-sum transfers and rationing; we explore conditions under which these tools should be used in the next section. Given Theorem 1, we can introduce the following terminology.

Definition 3. Let $\mathcal{M} = \{(p_j^H, p_j^L, \delta_j, \underline{U}_j)\}_{j\in\{B,S\}}$ denote a generic mechanism of the form described in Theorem 1.

- We say that mechanism $\mathcal{M}$ is a *price mechanism* if $\delta_S = \delta_B = 0$. A price mechanism is a *competitive equilibrium mechanism* if additionally $p_B^H = p_B^L = p^{\mathsf{CE}}$, where p^{CE} solves $\mu(1 - G_B(p^{\mathsf{CE}})) = G_S(p^{\mathsf{CE}})$.
- The mechanism $\mathcal{M}$ *rations side* j if $\delta_j > 0$ and a non-zero measure of agents on side j choose the lottery.
- The mechanism $\mathcal{M}$ *subsidizes side* j if $\underline{U}_j > 0$.

4.1 Derivation of the optimal mechanism – Proof of Theorem 1

In this section, we sketch the proof of Theorem 1. (A number of details are relegated to the appendix.) A reader not interested in our techniques may skip to Section 5.

4.2 Outline of the proof

The key ideas behind the proof of Theorem 1 are as follows. Any incentive-compatible individually-rational mechanism may be represented as a pair of lotteries over quantities (one for each side of the market). A lottery specifies the probability that a given type trades the good. Because of our large-market assumption, the lottery generates stochastic outcomes for individuals but deterministic outcomes in the aggregate. With the lottery representation of the mechanism, the market-clearing condition (MC) states the the expected quantity must be the same under both lotteries. The objective function can be represented as expectation of a certain function of the realized quantity with respect to the pair of lotteries. If we further incorporate the budget balance constraint (BB) into the objective function by assigning a Lagrange multiplier, it follows that the value of the optimal lottery must be equal to the concave closure of the Lagrangian at the optimal trade volume.[7] A concave closure of a

[6] The mechanism is effectively characterized by six parameters, since one of the lump-sum payments is always 0. Also, market-clearing conditions reduce the degrees of freedom on prices by 1.

[7] This is reminiscent of concavification in Bayesian persuasion models, see Aumann et al. (1995) and Kamenica and Gentzkow (2011).

one-dimensional function is supported by a mixture over at most two points. Hence, we obtain the two-price characterization of the optimal mechanism from Theorem 1.

4.3 Main argument

First, we simplify the problem by applying the canonical method developed by Myerson (1981) that allows us to express feasibility of the mechanism solely through the properties of the allocation rule and the transfer received by the worst type (the proof is skipped).

Claim 1. *A mechanism* (X_B, X_S, T_B, T_S) *is feasible if and only if*

$$X_B(r) \text{ is non-decreasing in } r, \tag{B-Mon}$$

$$X_S(r) \text{ is non-increasing in } r, \tag{S-Mon}$$

$$\mu \int_{\underline{r}_B}^{\bar{r}_B} X_B(r) dG_B(r) = \int_{\underline{r}_S}^{\bar{r}_S} X_S(r) dG_S(r), \tag{MC}$$

$$\mu \int_{\underline{r}_B}^{\bar{r}_B} J_B(r) X_B(r) dG_B(r) - \mu \underline{U}_B \geq \int_{\underline{r}_S}^{\bar{r}_S} J_S(r) X_S(r) dG_S(r) + \underline{U}_S, \tag{BB}$$

where J_B *and* J_S *denote the virtual surplus functions,* $\underline{U}_B$, $\underline{U}_S \geq 0$.

We can recover the transfers associated with a feasible mechanism via the envelope formula:

$$T_B(r) = \underline{U}_B + \int_{\underline{r}_B}^{r} X_B(\tau) d\tau - X_B(r) r,$$

$$T_S(r) = \underline{U}_S + \int_{r}^{\bar{r}_S} X_S(\tau) d\tau + X_S(r) r.$$

Second, using the preceding formulas and integrating by parts, we can show that the objective function (OBJ) also only depends on the allocation rule:

$$\mathrm{TS}(\Lambda) = \mu \Lambda_B \underline{U}_B + \mu \int \Pi_B^{\Lambda}(r) X_B(r) dG_B(r) + \Lambda_S \underline{U}_S + \int \Pi_S^{\Lambda}(r) X_S(r) dG_S(r), \tag{OBJ'}$$

where $\Lambda_j = \int \lambda_j(r) dG_j(r)$ is the average weight assigned to side j and

$$\Pi_B^{\Lambda}(r) := \frac{\int_r^{\bar{r}_B} \lambda_B(r) dG_B(r)}{g_B(r)}, \tag{4.1}$$

$$\Pi_S^{\Lambda}(r) := \frac{\int_{\underline{r}_S}^{r} \lambda_S(r) dG_S(r)}{g_S(r)}. \tag{4.2}$$

We refer to Π_j^Λ as the *preference-weighted information rents* of side j. Note that in the special case of fully transferable utilities, Π_j^Λ boils down to the usual information rent term, that is, $G_S(r)/g_S(r)$ for sellers, and $(1-G_B(r))/g_B(r)$ for buyers.

Third, finding the optimal mechanism is hindered by the fact that the monotonicity constraints (B-Mon) and (S-Mon) may bind ("ironing"); in such cases, it is difficult to employ optimal control techniques. We get around the problem by representing allocation rules as mixtures over quantities; this allows us to optimize in the space of distributions and make use of the concavification approach.[8] Because G_S has full support (it is strictly increasing), we can represent any non-increasing, left-continuous function $X_S(r)$ as

$$X_S(r) = \int_0^1 \mathbf{1}_{\{r \le G_S^{-1}(q)\}} dH_S(q),$$

where H_S is a distribution on $[0,\,1]$. Similarly, we can represent any non-decreasing, right-continuous function $X_B(r)$ as

$$X_B(r) = \int_0^1 \mathbf{1}_{\{r \ge G_B^{-1}(1-q)\}} dH_B(q).$$

Economically, our representation means that we can express a feasible mechanism in the quantile (quantity) space. To buy quantity q from the sellers, the designer has to offer a price of $G_S^{-1}(q)$, because then exactly sellers with $r \le G_S^{-1}(q)$ sell. An appropriate randomization over quantity (equivalently, prices) will replicate an arbitrary feasible quantity schedule X_S. Similarly, to sell quantity q to buyers, the designer has to offer a price $G_B^{-1}(1-q)$, and exactly buyers with $r \ge G_B^{-1}(1-q)$ buy. We have thus shown that it is without loss of generality to optimize over H_S and H_B rather than X_S and X_B in (OBJ').[9]

Fourth, we arrive at the following equivalent formulation of the designer's problem:

$$\max \left\{ \begin{array}{l} \mu \displaystyle\int_0^1 \left(\int_{G_B^{-1}(1-q)}^{\bar{r}_B} \Pi_B^\Lambda(r) g_B(r) dr \right) dH_B(q) \\ + \displaystyle\int_0^1 \left(\int_{\underline{r}_S}^{G_S^{-1}(q)} \Pi_S^\Lambda(r) g_S(r) dr \right) dH_S(q) \\ + \mu \Lambda_B \underline{U}_B + \Lambda_S \underline{U}_S \end{array} \right\} \tag{4.3}$$

[8] We thank Benjamin Brooks and Doron Ravid for teaching us this strategy.

[9] Formally, considering all distributions H_B and H_S is equivalent to considering all feasible right-continuous X_B and left-continuous X_S. The optimal schedules can be assumed right- and left-continuous, respectively, because a monotone function can be made continuous from one side via a modification of a measure-zero set of points (which thus does not change the value of the objective function (OBJ')).

over H_S, $H_B \in \Delta([0,1])$, $\underline{U}_B$, $\underline{U}_S \geq 0$, subject to

$$\mu \int_0^1 q dH_B(q) = \int_0^1 q dH_S(q), \tag{4.4}$$

$$\mu \int_0^1 \left(\int_{G_B^{-1}(1-q)}^{\bar{r}_B} J_B(r) g_B(r) dr \right) dH_B(q) - \mu \underline{U}_B \geq \int_0^1 \left(\int_{\underline{r}_S}^{G_S^{-1}(q)} J_S(r) g_S(r) dr \right) dH_S(q) + \underline{U}_S. \tag{4.5}$$

Fifth, we can incorporate the constraint (4.5) into the objective function using a Lagrange multiplier. Let $\alpha \geq 0$ denote the Lagrange multiplier. We define

$$\phi_B^\alpha(q) \coloneqq \int_{G_B^{-1}(1-q)}^{\bar{r}_B} (\Pi_B^\Lambda(r) + \alpha J_B(r)) g_B(r) dr + (\Lambda_B - \alpha) \underline{U}_B,$$
$$\phi_S^\alpha(q) \coloneqq \int_{\underline{r}_S}^{G_S^{-1}(q)} (\Pi_S^\Lambda(r) - \alpha J_S(r)) g_S(r) dr + (\Lambda_S - \alpha) \underline{U}_S.$$

Then, our problem becomes one of maximizing the expectation of an additive function over two distributions, subject to an inequality ordering the means of those distributions. We can thus employ the concavification approach to simplify the problem.

Lemma 1. *Suppose that there exists $\alpha^\star \geq 0$ and distributions $H_S^\star$ and $H_B^\star$ such that*

$$\int_0^1 \phi_S^{\alpha^\star}(q) dH_S^\star(q) + \mu \int_0^1 \phi_B^{\alpha^\star}(q) dH_B^\star(q) = \max_{Q \in [0, \mu \wedge 1], \underline{U}_B, \underline{U}_S \geq 0} \left\{ \operatorname{co}\left(\phi_S^{\alpha^\star}\right)(Q) + \mu \operatorname{co}\left(\phi_B^{\alpha^\star}\right)(Q/\mu) \right\}, \tag{4.6}$$

with constraints (4.4) and (4.5) holding with equality (where co *denotes the concave closure). Then, $H_S^\star$ and $H_B^\star$ correspond to the optimal mechanism. Conversely, if $H_B^\star$ and $H_S^\star$ are optimal, we can find $\alpha^\star$ such that (4.4), (4.5), and (4.6) hold.*

Finally, because the optimal $H_j^\star$ concavifies a one-dimensional function ϕ_j^α, it is without loss to assume that the lottery induced by $H_j^\star$ has at most two realizations. Therefore, as a corollary of Lemma 1, we obtain Theorem 1. The proof of Lemma 1, as well as the formal proof of the preceding claim, can be found in Appendix B.1 and B.2, respectively.

5 Optimal Mechanism under Inequality-Based Social Preferences

In this section, we characterize the optimal mechanism for the case in which social preferences come from the following model of inequality. For each agent, there exists a pair of values,

v^K and v^M, for units of goods K and M, respectively; these values represent the social surplus created by assigning the goods to that agent, but they are not observed by the designer. The values v^K and v^M also describe the individual's preferences by defining the rate of substitution v^K/v^M between the good and money (given the rate $r \equiv v^K/v^M$, the agent's preferences are defined as in Section 3). The pair (v^K, v^M) is distributed according to a joint distribution $F_S(v^K, v^M)$ for sellers, and $F_B(v^K, v^M)$ for buyers. The distribution F_j is continuous and fully supported on $[\underline{v}_j^K, \bar{v}_j^K] \times [\underline{v}_j^M, \bar{v}_j^M]$, with $\underline{v}_j^K \geq 0$, $\underline{v}_j^M \geq 0$, and $j \in \{B, S\}$. The designer knows the distribution of (v^K, v^M) on both sides of the market and chooses a mechanism to maximize total value. We will refer to the above framework as the *two-dimensional value model*.

In the two-dimensional value model, in general, direct mechanisms should allow agents to report their two-dimensional types. However, as we show in Appendix A.1, it is without loss of generality to assume that agents only report their rates of substitution. Intuitively, reporting rates of substitution suffices because those rates fully describe individual agents' preferences. (The mechanism could elicit information about both values by making agents indifferent between reports—but we show that this cannot raise the surplus achieved by the optimal mechanism.) This implies that we can write the objective function of the designer as

$$\begin{aligned} \text{TV} = \mu \int_{\underline{v}_B^K}^{\bar{v}_B^K} \int_{\underline{v}_B^M}^{\bar{v}_B^M} \left[X_B(v^K/v^M)v^K - T_B(v^K/v^M)v^M \right] dF_B(v^K, v^M) \\ + \int_{\underline{v}_S^K}^{\bar{v}_S^K} \int_{\underline{v}_S^M}^{\bar{v}_S^M} \left[-X_S(v^K/v^M)v^K + T_S(v^K/v^M)v^M \right] dF_S(v^K, v^M); \quad \text{(VAL)} \end{aligned}$$

the set of feasible mechanisms is the same as in Section 3.

Given objective (VAL), it is straightforward to see that the preceding problem is in fact equivalent to solving the problem of maximizing (OBJ) over feasible mechanisms, with $G_j(r)$ derived from the joint distribution $F_j(v^K, v^M)$ given that $r \equiv v^K/v^M$, and the Pareto weights are chosen to be

$$\lambda_j(r) = \mathbb{E}^j \left[v^M \mid \frac{v^K}{v^M} = r \right]. \tag{5.1}$$

That is, the Pareto weight on an agent with rate r is her expected value for money conditional on r. This observation allows us to apply the results of Section 4 by studying the properties of the optimal mechanism under Pareto weights (5.1). In Appendix A.1, we show additionally that *any* set of Pareto weights can be obtained by choosing appropriate joint distributions F_j. Therefore, the Pareto weight model of Section 3 and the two-dimensional value model are equivalent under the mapping defined by (5.1).

5.1 Measures of inequality

Before we describe the optimal mechanism, we formalize the idea of inequality in our model. Define

$$m_j = \mathbb{E}^j[v_M], \tag{5.2}$$

for $j \in \{B, S\}$, as buyers' and sellers' average values for money (M), and let

$$D_j(r) = \frac{\mathbb{E}^j\left[v_M \mid \frac{v^K}{v^M} = r\right]}{m_j} \tag{5.3}$$

be the (normalized) conditional expectation of the value for money when the marginal rate of substitution is r. The normalization means that $D_j(r)$ is equal to 1 on average, by the law of iterated expectations. For clarity of exposition, we assume that $D_j(r)$ is non-increasing in r for $j \in \{B, S\}$.[10]

Definition 4. We have *cross-side inequality* if $m_S \neq m_B$. We have *same-side inequality* (for side j) if D_j is not identically equal to 1.

We say that same-side inequality is *low* for sellers if $D_S(\underline{r}_S) \leq 2$; we say that same-side inequality is *high* for sellers if $D_S(\underline{r}_S) > 2$.

Under the assumption that $D_j(r)$ is decreasing, a seller with the lowest rate of substitution $\underline{r}_S$ is the "poorest" seller that can be identified based on behavior in the marketplace (she has the highest conditional expected value for money).[11] Same-side inequality is low if this "poorest" seller has a conditional expected value for money that does not exceed the average value for money by more than a factor of 2. The opposite case of high same-side inequality implies that the "poorest" seller has a conditional expected value for money that exceeds the average by more than a factor of 2. It turns out that the threshold of 2 delineates qualitatively different solutions to the optimal design problem—we comment on the interpretation of this threshold in Appendix A.2. A similar concept can be defined for buyers.

5.2 Regularity conditions

We impose additional regularity conditions to simplify the characterization of optimal mechanisms. First, we assume that the densities g_j of the distributions G_j of the rates of substitu-

[10] This assumption is not used in any of the proofs but is used in the discussions. Moreover, we later introduce a regularity condition that implicitly makes a weaker version of the above assumption (and which *is* used in the proofs). We make a stronger assumption here to emphasize the economic intuition—see the discussion in the next subsection.

[11] When the support of (v^K, v^M) is a product set which does not contain the origin $(0, 0)$, the lowest rate of substitution identifies a seller with the highest *realized* value for money.

tion are strictly positive and continuously differentiable (in particular continuous) on $[\underline{r}_j, \bar{r}_j]$, and that virtual surplus functions $J_B(r)$ and $J_S(r)$ are non-decreasing. We make the latter assumption to highlight the role that inequality plays in determining whether the optimal mechanism makes use of rationing: With non-monotone virtual surplus functions, rationing (known in this context as "ironing") could arise as a consequence of revenue-maximization motives implicitly present in our model due to the budget balance constraint. We need an even stronger condition to rule out ironing due to irregular local behavior of the densities g_j. Let

$$\Delta_S(p) \coloneqq \frac{\int_{\underline{r}_S}^{p} [D_S(\tau) - 1] g_S(\tau) d\tau}{g_S(p)}, \tag{5.4}$$

$$\Delta_B(p) \coloneqq \frac{\int_{p}^{\bar{r}_B} [1 - D_B(\tau)] g_B(\tau) d\tau}{g_B(p)}. \tag{5.5}$$

Assumption 1. The functions $\Delta_S(p) - p$ and $p + \Delta_B(p)$ are both quasi-concave in p.

A sufficient condition for Assumption 1 to hold is that the functions $\Delta_j(p)$ are concave. Intuitively, concavity of $\Delta_j(p)$ is closely related to non-increasingness of $D_j(r)$ (these two properties become equivalent when g_j is uniform). A non-increasing $D_j(r)$ reflects the belief of the market designer that that agents with lower willingness to pay (lower r) are "poorer" on average, that is, have a higher conditional expected value for money. This assumption is economically intuitive but is restrictive in that it disciplines social preferences— the designer attaches a higher Pareto weight to agents with lower rates of substitution. Specifically, $\lambda_j(r) = m_j D_j(r)$. When $D_j(r)$ is assumed to be decreasing, concavity of Δ_j rules out irregular local behavior of g_j. Each function $\Delta_j(p)$ is 0 at the endpoints $\underline{r}_j$ and $\bar{r}_j$, and non-negative in the interior. There is no same-side inequality if and only if $\Delta_j(p) = 0$ for all p. We show in Appendix A.2 that, more generally, the functions Δ_j measure same-side inequality by quantifying the change in surplus associated with running a one-sided price mechanism with price p (which redistributes money from richer to poorer agents on the same side of the market). There, we also give examples of primitive distributions F_j that satisfy the regularity conditions.

5.3 Solution with inequality

In this section, we show that lump-sum transfers are an optimal response of the market designer when cross-side inequality is significant, and that rationing can occur only when same-side inequality is sufficiently large.

Theorem 2. *Suppose that Assumption 1 holds, and same-side inequality for sellers is low. Then, the optimal mechanism is a price mechanism (with prices p_B and p_S).*

When $m_S \geq m_B$, a competitive equilibrium mechanism is optimal if and only if

$$m_S \Delta_S(p^{\mathsf{CE}}) - m_B \Delta_B(p^{\mathsf{CE}}) \geq (m_S - m_B) \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}. \tag{5.6}$$

When condition (5.6) fails, we have $p_B > p_S$, and prices are determined by market-clearing $\mu(1 - G_B(p_B)) = G_S(p_S)$, and, in the case of an interior solution,[12]

$$p_B - p_S = -\frac{1}{m_S} \left[m_S \Delta_S(p_S) - m_B \Delta_B(p_B) - (m_S - m_B) \frac{1 - G_B(p_B)}{g_B(p_B)} \right]. \tag{5.7}$$

The mechanism subsidizes the sellers. (The case $m_B > m_S$ is described in Appendix B.3).

The intuition behind Theorem 2 is that a price mechanism performs relatively well both in terms of allocative efficiency and in terms of redistribution. Why? A price mechanism induces the planner's desired selection of traders: the poorest sellers sell (and hence receive a transfer), and the poorest buyers do not buy (hence are not deprived of any money). There is, however, a trade-off between allocative efficiency and redistribution, as the marginal rate of substitution is not a perfect signal for the pair of values (v^K, v^M); that trade-off determines whether there is a gap between the buyer and seller prices (in which case there are also lump-sum transfers) or not (in which case the competitive equilibrium mechanism is optimal).

In the special case of no same-side inequality, Assumption 1 is automatically satisfied, equation (5.6) cannot hold unless $m_B = m_S$, and equation (5.7) boils down to

$$p_B - p_S = \left(\frac{m_S - m_B}{m_S} \right) \frac{1 - G_B(p_B)}{g_B(p_B)}.$$

Thus, when there is no same-side inequality, there is a gap between the prices proportional to the size of the cross-side inequality, and the "poorer" side receives a strictly positive subsidy if and only if $m_S \neq m_B$.

Theorem 2 allows us to derive the optimal mechanism for the setting of Section 2—that is, where there is no same-side inequality, buyers have value 1 for each unit of money, and sellers have value for money $m \geq 1$. By Theorem 2, there is no rationing in the optimal mechanism, and the optimal buyer and seller prices are given by $p_B = 1/2$ and $p_S = 1/(2m)$, respectively. The total lump-sum transfer to the seller side is $(m-1)/(4m)$. This is exactly

[12] When no such interior solution exists, one of the prices is equal to the bound of the support: either $p_B = \underline{r}_B$ or $p_S = \bar{r}_S$.

the mechanism described in Section 2.2. It can be shown (see the generalized statement of Theorem 2 in Appendix B.3) that in the opposite case when buyers are poorer ($m < 1$), the optimal mechanism is given by $p_S = 1/2$, $p_B = 1 - m/2$, and it is the buyers who receive the lump-sum transfer.

A disadvantage of the price mechanism is that it is limited in how much wealth can be redistributed to the poorest agents. Indeed, in the case $m_S \geq m_B$ the price received by the sellers is capped by the market-clearing condition, and the revenue is shared by all sellers equally. When same-side inequality is low, a lump-sum transfer to all sellers is a fairly effective redistribution channel. However, when same-side inequality is high, the conclusion of Theorem 2 may fail, as the planner may prefer to use market-clearing to target transfers to poorer sellers.

Theorem 3. *Suppose that Assumption 1 holds, $m_S \geq m_B$, and same-side inequality for sellers is high. Then, if μ is low enough (there are few buyers relative to sellers), the optimal mechanism rations the sellers (and is a price mechanism for the buyer side).*

The intuition behind Theorem 3 is simple. With high same-side inequality, the designer would like to transfer wealth to the poorest sellers. When there are only few buyers, the revenue of the mechanism is small relative to the number of sellers—and because lump-sum payments must be allocated equally across sellers, the lump-sum transfer to the poorest sellers becomes tiny. By abandoning a price mechanism, however, the designer has an additional channel to redistribute wealth: she can set a price for sellers above the market-clearing level, making sure that more wealth goes to the sellers who agree to trade. To clear the market, the designer must then ration the sellers.

Remark 1. We can derive an explicit upper bound for μ in Theorem 3 if m_S is sufficiently higher than m_B (so that competitive equilibrium cannot be optimal) and the density $g_S(r)$ is decreasing. In this case, let $r^\star = D_S^{-1}(2)$, so that $G_S(r^\star)$ is the mass of sellers whose conditional value for money exceeds the average more than twice. Then, if $\mu \leq G_S(r^\star)$, the conclusion of Theorem 3 applies.

We note that Theorem 3 does not have a counterpart for the case in which buyers are poor.

Theorem 4. *Under Assumption 1, no optimal mechanism ever rations the buyers.*

Theorem 4 explains why we have only defined high same-side inequality for the seller side. This difference between buyers and sellers is not an artifact of our modeling approach; rather, it is a consequence of the inherent asymmetry between buyers and sellers in the context of inequality. A market mechanism selects the poorest sellers and the richest buyers

to trade, all else equal. Thus, the desire to trade can be used as a tool to subsidize the poorest sellers—but no such tool can be used to subsidize the poorest buyers, as the poorest buyers are those who *do not trade*. Thus, the only way to transfer money to poorest buyers is through lump-sum transfers.

We present parametric examples illustrating Theorems 2 and 3 in Section 7.

In the special case when there is no wealth inequality—that is, both buyers and sellers have the same constant value for good M, normalized to 1—Assumption 1 is trivially satisfied, and Theorem 2 yields the following corollary.

Corollary 1. *When there is no wealth inequality, the competitive equilibrium mechanism is optimal:* $X_B(r) = \mathbf{1}_{\{r \geq p^{\mathsf{CE}}\}}$ *and* $X_S(r) = \mathbf{1}_{\{r \leq p^{\mathsf{CE}}\}}$.

Corollary 1 can be seen as a version of the First Welfare Theorem. The price p^{CE}, by definition, equalizes supply and demand and leads to an efficient allocation of good K, which—in the world of fully transferable utility—is equivalent to maximizing the sum of agents' utilities.

6 What if lump-sum redistribution is not feasible?

The optimal mechanism (both in full generality à la Theorem 1 and in the cases just discussed) often redistributes wealth through direct lump-sum transfers. While lump-sum transfers may be feasible in some marketplaces, they might be difficult to implement in cases where participation is flexible. Indeed, imagine a mechanism that pays a constant amount to sellers regardless of whether they trade or not. In such mechanisms, the sellers with the highest marginal rates of substitution might get strictly positive utility from participating in the marketplace even though they never trade; this creates an incentive for additional agents to enter the marketplace to reap the free benefit. Excess entry could then undermine the budget balance condition. Consequently, we consider one additional constraint on the set of mechanisms.[13]

Assumption 2 ("No Free Lunch")**.** The participation constraints of the lowest-utility type of buyers $\underline{r}_B$, and of the lowest-utility type sellers $\bar{r}_S$ both bind, that is, types $\underline{r}_B$ and $\bar{r}_S$ are indifferent between participating or not.

[13] One might think that a cleaner way to rule out the problem just described is to assume that a trader can only receive a transfer conditional on trading. Making transfers conditional on trading, however, would not work in our risk-neutral setting because an arbitrary expected transfer T can be paid to an agent by paying her T/ϵ in the ϵ-probability event of trading; because ϵ is arbitrarily small, this causes no distortion in the actual allocation.

Under Assumption 2, the methods of the previous sections immediately apply with the only modification that we set $\underline{U}_S = \underline{U}_B = 0$ in Lemma 1 (and hence those parameters no longer appear as optimization parameters in the Lagrangian (4.6)). Theorem 1 holds but the optimal mechanism cannot redistribute wealth via lump-sum payments (thus, the optimal mechanism is parametrized by at most six parameters). The sufficient condition for optimality of competitive equilibrium from Theorem 2 is still valid, since a competitve equilibrium mechanism does not use lump-sum transfers. However, the optimal mechanism from the second part of Theorem 2 is now ruled out. Another novel aspect of the analysis is that we can no longer assume that the budget balance condition binds at the optimal solution (Lemma 1 needs to be appropriately modified) because it is not feasible to distribute budget surplus using lump-sum transfers.

6.1 Rationing when lump-sum transfers are ruled out

We have seen in Section 5 that rationing can only be optimal when there is significant same-side seller inequality. Yet the discussion in Section 5 relied to a large extent on the fact that cross-side inequality can be addressed by using lump-sum transfers. In this section, we show that when lump-sum transfers are ruled out, rationing may be used to address cross-side inequality. To make this point in a sharp way, for now we shut down same-side inequality completely; we address same-side inequality again in Section 7.2.

We assume that $D_B(r) = D_S(r) = 1$ for all r. Moreover, sellers are poorer on average: $m_S \geq m_B$. We can without loss normalize $m_B = 1$. We impose a stronger regularity assumption on the behavior of virtual surplus functions. Because we now have $G_B(r) = F_B^K(r)$, and $G_S(r) = F_S^K(m_S r)$, we phrase the conditions in terms of the marginal distributions of values for good K.

Assumption 3. The information rent terms for buyers and sellers are monotone: $F_S^K(r)/f_S^K(r)$ is non-decreasing while $(1 - F_B^K(r))/f_B^K(r)$ is non-increasing.

To simplify notation, we let $I_S(r) \coloneqq F_S^K(r)/f_S^K(r)$.

Proposition 1. *Suppose that Assumptions 2 and 3 both hold. A price mechanism is optimal for the buyer side. If*

$$m_S \leq 1 + \frac{1}{\sup_r\{I_S'(r)\}},$$

then the competitive equilibrium mechanism is optimal. On the other hand, if p^{CE} denotes the competitive equilibrium price, $c \coloneqq f_S^K(m_S p^{\mathsf{CE}})/f_B^K(p^{\mathsf{CE}})$, $d \coloneqq I_S'(m_S p^{\mathsf{CE}})$, and

$$m_S > 1 + \frac{1}{d} + \frac{\mu}{cd},$$

then it is optimal to ration the sellers.

Proposition 1 implies that if f_B^K is bounded away from 0 and the information rents $I_S(r)$ for the sellers have a derivative that is finite and bounded away from 0, then a competitive equilibrium mechanism is optimal under small m_S (when the cross-side inequality is relatively small) and rationing on the seller side is optimal under large m_S (when the cross-side inequality is relatively large).

For example, consider the case in which $F_S^K(x) = x^{\alpha_S}$ and $F_B^K(x) = x^{\alpha_B}$ with $\alpha_S \leq \alpha_B$ (buyers have stochastically higher values). In this case we have $\sup_r\{I_S'(r)\} = d = 1/\alpha_S$ and $c \geq \alpha_S m_S^{\alpha_S - 1}/\alpha_B$. Therefore, competitive equilibrium is optimal for

$$m_S \leq 1 + \alpha_S,$$

and rationing on the seller side is optimal for

$$m_S > 1 + \alpha_S + \mu\alpha_B m_S^{1-\alpha_S}.$$

In the simple setting of Section 2 (uniform distributions), we can obtain a stronger result which implies that the single-price mechanism described in Section 2.1 is actually optimal under Assumption 2.

Claim 2. *In the setting of Section 2 with $m \geq 1$, under Assumption 2, the optimal mechanism is to set a single price $p^\star$, where*

$$p^\star = \begin{cases} p^{\mathsf{CE}} & m \leq 3, \\ \frac{m-2}{2m-2} & 3 < m < 2 + \sqrt{2}, \\ \frac{1}{m} & m \geq 2 + \sqrt{2} \end{cases}$$

is as in (2.2). In particular, there is rationing (on the seller side) if and only if $m > 3$.

Finally, we show, by means of example, that rationing can be optimal on the buyer side when lump-sum redistribution is ruled out. We adopt the setting of Section 2 with $m < 1$. The following claim is established in a fashion analogous to Proposition 1 and Claim 2, so we omit the proof.

Claim 3. *In the setting of Section 2 but with $m < 1$, under Assumption 2, the optimal mechanism is to set a single price $p^\star$, where*

$$p^\star = \begin{cases} p^{\mathsf{CE}} & m \geq 1/3, \\ \frac{1}{2-2m} & m < 1/3. \end{cases}$$

In particular, there is rationing on the buyer side if and only if $m < 1/3$.

The intuition behind Claim 3 is straightforward: When buyers cannot be subsidized via lump-sum transfers, the designer can only raise buyer surplus by pushing the price below the market-clearing level (i.e., through a price ceiling/cap); when buyers are sufficiently poor relative to sellers, such a policy becomes optimal.

7 Numerical example

In this section, we consider a simple parametric example that illustrates our main theoretical results. First, we note that the conclusions from preceding sections continue to hold when we relax some of the regularity conditions on the distributions.

Remark 2. Theorem 2 continues to hold if we relax the assumption that densities g_j are strictly positive on $[\underline{r}_j, \bar{r}_j]$. More generally, as long as virtual surplus functions are non-decreasing and Assumption 1 holds, a sufficient condition for Theorem 2 is that $\Delta'_S(\underline{r}_S) \leq 1$. Theorem 3 continues to hold if we relax the assumption that the density g_S is continuous at $\underline{r}_S$ by allowing $g_S(\underline{r}_S) = \infty$.

The preceding remark follows from direct inspection of the proofs of Theorems 2 and 3. Intuitively, positive density at $\underline{r}_S$ is only needed for Theorem 3: It says that there exists a non-negligible mass of sellers with a rate of substitution close to the lowest one. On the other hand, the assumption that the density does not explode at $\underline{r}_S$ is only needed for Theorem 2: It makes sure that there are relatively few sellers with low rates of substitution.

In our numerical example, we keep fixed the distribution of values for the buyer side. We assume that $v_B^K \sim \text{Uniform}[0, 1]$, and $v_B^M \equiv 1$. For the seller side, we consider the following two-dimensional family of distributions: $v_S^K \sim F_S^K$ with $F_S^K(v) = v^\beta$, for some $\beta > 0$, and $v_S^M \sim \text{Pareto}(1, \alpha)$, where 1 is the location parameter and α is the tail parameter. We assume that $\alpha > 1$ so that $m_S = \alpha/(\alpha - 1)$ is finite. Lower α means that the distribution has a thicker tail—this corresponds to the case in which there are relatively many poor sellers. Note that sellers are always poorer than buyers—that is, they always have higher values for money than buyers (the values are equal in the limit as $\alpha \to \infty$). The parameter β allows us to shift the distribution of values for the good—high β corresponds to values concentrated around high levels. When $\beta \leq 1$, sellers have stochastically smaller values for the good than buyers.

We can provide the following foundation for our assumption that the marginal value for

money has a Pareto distribution: If an agent has CRRA utility for wealth[14]

$$u(w) = \frac{w^{1-\eta} - 1}{1-\eta},$$

and wealth w has a Pareto distribution with tail parameter γ, then $u'(w)$ (the marginal value) also has a Pareto distribution (with tail parameter γ/η).

It can be checked that our parametric family satisfies the following regularity conditions: the function $D_S(r)$ is decreasing ($D_B(r) \equiv 1$); virtual surplus functions $J_B(r)$ and $J_S(r)$ are non-decreasing; and Assumption 1 holds (trivially for the buyer side). It is not true in general that the density g_S is strictly positive and bounded on $[\underline{r}_S, \bar{r}_S] = [0, 1]$—in fact, this is only true if $\beta = 1$. When $\beta > 1$, we have $g_S(0) = 0$, and when $\beta < 1$, $g_S(0) = \infty$ (that is why we extended our main results to cover these cases in Remark 2).

7.1 When lump-sum redistribution is allowed

The parametric family allows us to explicitly calculate some key variables that determine the optimal mechanism in the setting of Section 5. For example, we have (when $\alpha \notin \{\beta, \beta+1\}$)

$$D_S(r) = \frac{(\alpha-\beta)(\alpha-1)}{\alpha(\alpha-\beta-1)} \frac{r^{\beta+1} - r^{\alpha}}{r(r^{\beta} - r^{\alpha})};$$

in particular, high same-side inequality is equivalent to the condition

$$\beta > \frac{\alpha^2 - \alpha}{\alpha + 1},$$

which requires α to be low compared to β. Intuitively, a low α and high β mean that there are relatively many sellers with high values for money, and relatively few sellers with low values for the good; in such cases, low rates of substitution identify sellers with high values for money. When β is low, so that most sellers have low values for the good, low rates of substitution do not necessarily indicate high values for money. Indeed, in the extreme case in which all sellers have value 0 for the good (which is true in the limit as $\beta \to 0$), the rate of substitution is equal to 0, and hence is uninformative about sellers' values for money.

We can also directly calculate the derivative of $\Delta_S(r)$ at $r = 0$, which allows us to apply Remark 2:

$$\Delta_S'(0) = \begin{cases} \frac{1}{\alpha(\alpha-\beta-1)} & \text{when } \alpha > \beta + 1 \\ \infty & \text{when } \alpha \leq \beta + 1. \end{cases}$$

[14] In the case $\eta = 1$, $u(w) = \log(w)$.

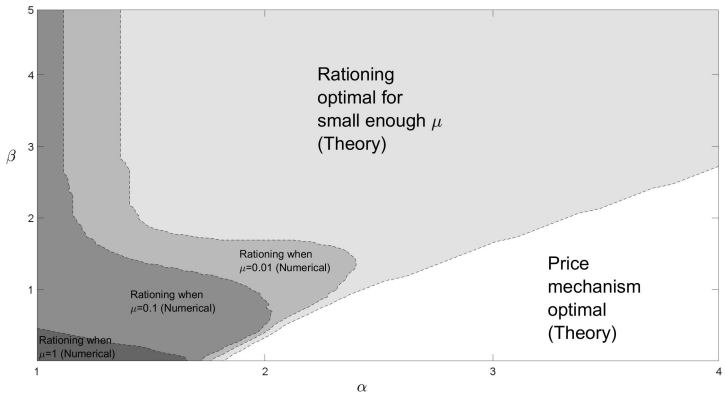


Figure 7.1: Optimal mechanism as a function of parameters (α, β).

It follows that a price mechanism (with a lump-sum transfer to sellers) is optimal whenever $\beta < \alpha - 1 - 1/\alpha$. In the case $\beta > \alpha - 1 - 1/\alpha$, rationing is optimal when μ is small enough. These regions are depicted in Figure 7.1 (labeled "Theory").

Figure 7.1 additionally shows the numerical solution for three fixed levels of μ, starting from a balanced market ($\mu = 1$) and ending at a market in which there is one buyer per 100 sellers ($\mu = 0.01$). As can be seen, the degree of same-side inequality required for rationing to emerge in balanced markets is quite extreme. In particular, the tail of the distribution of values for money must be so thick that the variance of the distribution is infinite. Even when $\mu = 0.01$, rationing still requires a thick tail.

Note that the effect of β on the optimality of rationing is not monotone. On one hand, a relatively large β is needed so that low rates of substitution r reveal a high value for money (see discussion above). On the other hand, when β is too large, the distribution of r shifts to the right, and thus there are relatively few sellers with low r. Therefore, rationing is most likely to occur for intermediate values of β.

The above numerical exercise reveals that same-side inequality must be quite extreme for rationing to be optimal. This is because in most cases inequality concerns can be efficiently addressed through lump-sum transfers. In the next section, we study what happens when such direct transfers are not feasible.

7.2 When lump-sum redistribution is ruled out

We now consider the same numerical example but under our "No Free Lunch" assumption (Assumption 2).

Figure 7.2 depicts the optimal mechanism (computed numerically) for different values

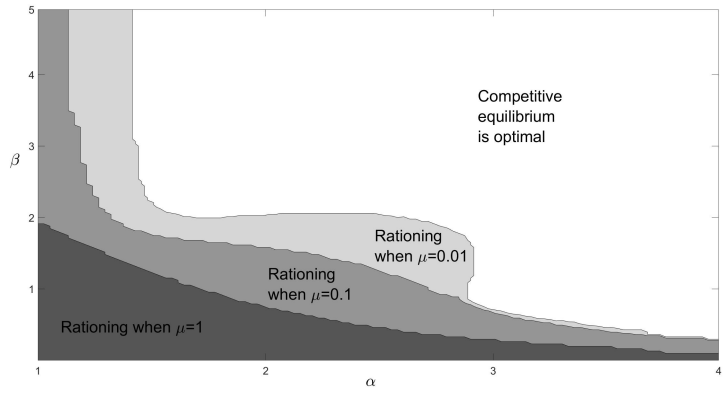


Figure 7.2: Optimal mechanism as a function of parameters (α, β) when lump-sum redistribution is not feasible.

of μ. Outside of the regions where rationing is optimal, the optimal mechanism is a competitive equilibrium. As can be seen, rationing emerges as optimal for a much larger space of parameters, compared to the case in which lump-sum transfers are allowed (pictured in Figure 7.1). Roughly speaking, rationing is optimal whenever there are sufficiently many (identifiable) poor sellers, a property that holds as long as α and β are not too large (a high α implies that sellers have relatively low values for money; a high β implies that there are relatively few sellers with low rates of substitution).

8 Discussion and Conclusion

Wealth inequality is a central and growing problem in modern society. Our work here highlights one way that markets may play a role in the solution: Markets themselves can be used to effect redistribution—particularly when more global redistributive instruments are not available. When there is substantial inequality between buyers and sellers, the optimal mechanism for a planner who cares about inequality imposes a wedge between buyer and seller prices, passing on the resulting surplus to the poorer side of the market. When there is significant within-side inequality, meanwhile, the optimal mechanism imposes price controls even though doing so induces rationing.

Of course, it is a classical result in economics that for any distortionary allocative rule the social planner might choose, there is some system of taxes and redistribution that will implement an outcome in which everyone is (weakly) better off. In practice, however, systemic inequality is hard to completely address through tax system alone, for both methodological and political reasons. Meanwhile, it seems unlikely that the current tax system corresponds to the optimal nonlinear tax schedule that would be set by a fully informed, inequality-aware

social planner. Thus, there might be real scope for market-level redistributive designs of the types we describe here. And indeed, our findings might partially justify the widespread use of price controls and other market-distorting regulations in settings with substantial inequality.

Additionally, our findings suggest that marketplaces that enable lower-income agents to participate in new types of exchanges with higher-income ones—such as ridesharing platforms and online labor marketplaces—may have significant social value beyond their purely allocative impacts. However, our results also imply that to mitigate inequality, such marketplaces may need either active redistribution through a combination of a price wedge and a lump-sum transfer (e.g., in the form of employer-provided health insurance for active employees) or minimum wage-like price controls, or possibly both.

The modeling approach applied here—agents with different marginal values of money—may prove useful for studying inequality in microeconomic contexts. More broadly, there may be value in further reflecting on how underlying macroeconomic issues like inequality should inform market design.

References

Aghion, P. and P. Bolton (1997): "A theory of trickle-down growth and development," Review of Economic Studies, 64, 151–172.

Arrow, K. J. and G. Debreu (1954): "Existence of an equilibrium for a competitive economy," Econometrica, 22, 265–290.

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

Aumann, R. J., M. Maschler, and R. E. Stearns (1995): Repeated Games with Incomplete Information, MIT press.

Baisa, B. (2017): "Auction design without quasilinear preferences," Theoretical Economics, 12, 53–78.

Besley, T. and S. Coate (1991): "Public provision of private goods and the redistribution of income," American Economic Review, 81, 979–984.

Blackorby, C. and D. Donaldson (1988): "Cash versus kind, self-selection, and efficient transfers," American Economic Review, 78, 691–700.

Bogomolnaia, A. and H. Moulin (2001): "A new solution to the random assignment problem," Journal of Economic Theory, 100, 295–328.

Bonnans, J. F. and A. Shapiro (2000): Perturbation Analysis of Optimization Problems, Springer.

Budish, E. (2011): "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes," Journal of Political Economy, 119, 1061–1103.

Bulow, J. and P. Klemperer (2012): "Regulated prices, rent seeking, and consumer surplus," Journal of Political Economy, 120, 160–186.

Che, Y.-K. and I. Gale (1998): "Standard auctions with financially constrained bidders," Review of Economic Studies, 65, 1–21.

——— (2000): "The optimal mechanism for selling to a budget-constrained buyer," Journal of Economic Theory, 92, 198–233.

Che, Y.-K., I. Gale, and J. Kim (2012): "Assigning resources to budget-constrained agents," Review of Economic Studies, 80, 73–107.

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

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

Dobzinski, S., R. Lavi, and N. Nisan (2012): "Multi-unit auctions with budget limits," Games and Economic Behavior, 74, 486–503.

Esteban, J. and D. Ray (2006): "Inequality, lobbying, and resource allocation," American Economic Review, 96, 257–279.

Fernandez, R. and J. Gali (1999): "To each according to...?? Markets, tournaments, and the matching problem with borrowing constraints," Review of Economic Studies, 66, 799–824.

Fleiner, T., R. Jagadeesan, Z. Jankó, and A. Teytelboym (2017): "Trading Networks with Frictions," Oxford University Working Paper.

Gahvari, F. and E. Mattos (2007): "Conditional cash transfers, public provision of private goods, and income redistribution," American Economic Review, 97, 491–502.

Huesmann, K. (2017): "Public Assignment of Scarce Resources when Preferences are Non-Linear," Mimeo.

Hylland, A. and R. Zeckhauser (1979): "The efficient allocation of individuals to positions," Journal of Political Economy, 87, 293–314.

Kamenica, E. and M. Gentzkow (2011): "Bayesian persuasion," American Economic Review, 101, 2590–2615.

Kotowski, M. H. (2017): "First-price auctions with budget constraints," Working Paper.

Laffont, J.-J. and J. Robert (1996): "Optimal auction with financially constrained buyers," Economics Letters, 52, 181–186.

Loury, G. C. (1981): "Intergenerational transfers and the distribution of earnings," Econometrica, 49, 843–867.

Maskin, E. and J. Riley (1984): "Optimal auctions with risk averse buyers," Econometrica, 52, 1473–1518.

McKenzie, L. W. (1959): "On the existence of general equilibrium for a competitive market," Econometrica, 27, 54–71.

McKinnon, R. I. (2010): Money and Capital in Economic Development, Brookings Institution.

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

Pai, M. M. and R. Vohra (2014): "Optimal auctions with financially constrained buyers," Journal of Economic Theory, 150, 383–425.

Piketty, T. and E. Saez (2013): "Optimal labor income taxation," in Handbook of Public Economics, Elsevier, vol. 5, 391–474.

Prendergast, C. (2017): "The Allocation of Food to Food Banks," Booth School of Business Working Paper.

Ray, D. ® A. Robson (2018): "Certified Random: A New Order for Coauthorship," American Economic Review, 108, 489–520.

Roth, A. E. (2007): "Repugnance as a Constraint on Markets," Journal of Economic Perspectives, 21, 37–58.

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

——— (2017): "A simpler theory of optimal capital taxation," Journal of Public Economics.

Sandel, M. J. (2012): What Money Can't Buy: The Moral Limits of Markets, Macmillan.

Satz, D. (2010): Why some things should not be for sale: The moral limits of markets, Oxford University Press.

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 Additional Discussions and Results

A.1 Equivalence between the Pareto weight model and the two-dimensional value model

In this appendix we formally establish the equivalence between the Pareto weight model of Section 3 and the two-dimensional value model of Section 5. Proofs are collected at the end of the section.

Consider the two-dimensional value model. By the Revelation Principle, it is without loss to assume that agents report their two-dimensional types to the designer; we denote a direct mechanism in this setting by $(\bar{X}_B, \bar{X}_S, \bar{T}_B, \bar{T}_S)$. The allocation and transfer rules are defined as in Section 3 but are now functions of the two-dimensional types (v^K, v^M). For example, $\bar{X}_B(v^K, v^M)$ is the probability that a buyer gets the object when she reports (v^K, v^M) to the direct mechanism. The definition of a feasible mechanism is analogous to the one introduced in Section 4 (see Definition 1), except that all constrains are written in terms of the two-dimensional types, and expectations are taken with respect to the multivariate distribution F_j (we omit the formal statement of these definitions). The total-value objective function is defined by

$$\begin{aligned}\text{TV} = \mu \int_{\underline{v}_B^K}^{\bar{v}_B^K} \int_{\underline{v}_B^M}^{\bar{v}_B^M} \left[\bar{X}_B(v^K, v^M)v^K - \bar{T}_B(v^K, v^M)v^M\right] dF_B(v^K, v^M) \\ + \int_{\underline{v}_S^K}^{\bar{v}_S^K} \int_{\underline{v}_S^M}^{\bar{v}_S^M} \left[-\bar{X}_S(v^K, v^M)v^K + \bar{T}_S(v^K, v^M)v^M\right] dF_S(v^K, v^M). \quad \text{(A.1)}\end{aligned}$$

We establish the equivalence to the setting considered in Section 3 in three steps:

1. We show that, without loss of generality, an incentive-compatible mechanism in the two-dimensional value model only elicits information about the rate of substitution, v^K/v^M; thus, the space of feasible mechanisms is the same in both frameworks.

2. We argue that the total value function (VAL) corresponds exactly to the objective function (OBJ) with Pareto weights $\lambda_j(r)$ defined to be the expected value of v^M conditional on observing a rate of substitution $r = v^K/v^M$.

3. As a consequence, if G_j is the distribution of v^K/v^M under F_j, and Pareto weights are defined as in Step 2, the same mechanism is optimal in both settings.

Step 1. We first formalize the idea that although agents have two-dimensional types, it is without loss of generality to consider mechanisms that only elicit information about the rate of substitution.[15]

Lemma 2. *If* $(\bar{X}_B, \bar{X}_S, \bar{T}_B, \bar{T}_S)$ *is an incentive-compatible mechanism, then there exists a mechanism* (X_B, X_S, T_B, T_S) *such that* $\bar{X}_j(v^K, v^M) = X_j(v^K/v^M)$ *and* $\bar{T}_j(v^K, v^M) = T_j(v^K/v^M)$ *for almost all* (v^K, v^M) *and* $j \in \{B, S\}$.

Thanks to Lemma 2, and the assumption that there are no mass points in the distribution of values, we can assume that the space of feasible direct mechanisms in the two-dimensional value setting is the same as the one considered in Section 3. From now on, with slight abuse of notation, we use (X_B, X_S, T_B, T_S) to denote a generic mechanism that elicits one-dimensional reports in both settings (with the convention that in the two-dimensional framework, agents report their rates of substitution v^K/v^M).

Step 2. To define the mapping between the distribution F_j and the Pareto weights $\lambda_j(r)$, define (as in Section 5.1)

$$m_j = \mathbb{E}^j[v_M],$$

for $j \in \{B, S\}$, as the average value for money, and

$$D_j(r) = \frac{\mathbb{E}^j\left[v_M \mid \frac{v^K}{v^M} = r\right]}{m_j}.$$

as the (normalized) conditional expected value for money when the marginal rate of substitution is r. Moreover, let G_j be the distribution of the random variable v^K/v^M when

[15] While it is clear that the rate of substitution fully describes agents' behavior, it could be hypothetically possible that the designer elicits more information by offering different combinations of trade probabilities and transfers among which the agent is indifferent; we show, however, that this is only possible for a measure-zero set of agents' types, and thus cannot strictly improve the designer's objective.

(v^K, v^M) is distributed according to F_j. Then, using Step 1, if we define $\lambda_j(r) = m_j D_j(r)$, the objective functions (OBJ) and (VAL) become identical:

$$\begin{aligned}\text{TV} &= \mu\mathbb{E}^B\left[X_B\left(\frac{v^K}{v^M}\right)v^K - T_B\left(\frac{v^K}{v^M}\right)v^M\right] + \mathbb{E}^S\left[-X_S\left(\frac{v^K}{v^M}\right)v^K + T_S\left(\frac{v^K}{v^M}\right)v^M\right] \\ &= \mu\mathbb{E}^B\left[\underbrace{\mathbb{E}^B\left[v_M \mid r\right]}_{\lambda_B(r)}\underbrace{\left[X_B\left(r\right)r - T_B\left(r\right)\right]}_{U_B(r)}\right] + \mathbb{E}^S\left[\underbrace{\mathbb{E}^S\left[v_M \mid r\right]}_{\lambda_S(r)}\underbrace{\left[-X_S\left(r\right)r + T_S\left(r\right)\right]}_{U_S(r)}\right] = \text{TS}(\Lambda).\end{aligned}$$

Step 3. We can now establish the following result.

Theorem 5. *If a mechanism (X_B, X_S, T_B, T_S) is feasible (optimal) in the setting of Section 5, then it is also feasible (optimal) in the setting of Section 3 with G_j being the distribution of v^K/v^M under F_j, and $\lambda_j(r) = m_j D_j(r)$.*

Conversely, if a mechanism (X_B, X_S, T_B, T_S) is feasible (optimal) in the setting of Section 3, then there exists a joint distribution F_j such that it is also feasible (optimal) in the setting of Section 5, v^K/v^M has distribution G_j, and $m_j D_j(r) = \lambda_j(r)$.

Theorem 5 establishes an equivalence between the two versions of our model. A further implication is that studying the two-dimensional model, for all possible distributions F_j, captures all possible social preferences over the Pareto frontier.

A.1.1 Proof of Lemma 2

We start with the following result that will be a key step in the proof.

Lemma 3. *Let $X(\theta_1, \theta_2)$ be a function defined on $[\underline{\theta}_1, \bar{\theta}_1] \times [\underline{\theta}_2, \bar{\theta}_2]$, with $\underline{\theta}_1, \underline{\theta}_2 \geq 0$, and assume that $X(\theta_1, \theta_2)$ is non-decreasing in θ_1/θ_2, that is*

$$\frac{\theta_1}{\theta_2} > \frac{\theta_1'}{\theta_2'} \implies X(\theta_1, \theta_2) \geq X(\theta_1', \theta_2').$$

Then, there exists a non-decreasing function $x : [\underline{\theta}_1/\bar{\theta}_2, \bar{\theta}_1/\underline{\theta}_2] \to \mathbb{R}$ such that $X(\theta_1, \theta_2) = x(\theta_1/\theta_2)$ almost everywhere.

Proof of Lemma 3. Consider $Y(r, \theta_2) = X(r\theta_2, \theta_2)$. For small enough $\epsilon > 0$ and almost all $r \in [\underline{\theta}_1/\bar{\theta}_2, \bar{\theta}_1/\underline{\theta}_2]$,

$$Y(r + \epsilon, \theta_2) \geq Y(r, \theta_2'), \forall \theta_2, \theta_2',$$

by assumption. Because $Y(r,\, \theta_2)$ is non-decreasing in r for every θ_2, it is continuous in r almost everywhere. Thus, for almost all r we obtain

$$Y(r,\, \theta_2) \geq Y(r,\, \theta_2'),\ \forall \theta_2,\, \theta_2'.$$

This, however, means that $Y(r,\, \theta_2) = x(r)$ for almost all r (does not depend on θ_2), for some function x, that is moreover non-decreasing. Thus, $X(r\theta_1,\, \theta_2) = x(r)$ for almost all r. Therefore,

$$X(\theta_1,\, \theta_2) = X\left(\frac{\theta_1}{\theta_2}\,\theta_2,\, \theta_2\right) = x\left(\frac{\theta_1}{\theta_2}\right)$$

almost everywhere which finishes the proof. □

We will show that incentive compatibility for buyers implies that $\bar{X}_B(v^K,\, v^M) = X_B(v^K/v^M)$ for some non-decreasing X_B. The argument for sellers is identical, and the statement for transfer rules follows immediately from payoff equivalence.

Incentive compatibility means that for all $(v^K,\, v^M)$ and $(\hat{v}^K,\, \hat{v}^M)$ in the support of F_B we have

$$\bar{X}_B(v^K,\, v^M)\frac{v^K}{v^M} - \bar{T}_B(v^K,\, v^M) \geq \bar{X}_B(\hat{v}^K,\, \hat{v}^M)\frac{v^K}{v^M} - \bar{T}_B(\hat{v}^K,\, \hat{v}^M),$$

as well as

$$\bar{X}_B(\hat{v}^K,\, \hat{v}^M)\frac{\hat{v}^K}{\hat{v}^M} - \bar{T}_B(\hat{v}^K,\, \hat{v}^M) \geq \bar{X}_B(v^K,\, v^M)\frac{\hat{v}^K}{\hat{v}^M} - \bar{T}_B(v^K,\, v^M).$$

Putting these two inequalities together,

$$\left(\bar{X}_B(v^K,\, v^M) - \bar{X}_B(\hat{v}^K,\, \hat{v}^M)\right)\left(\frac{v^K}{v^M} - \frac{\hat{v}^K}{\hat{v}^M}\right) \geq 0.$$

It follows that

$$\frac{v^K}{v^M} > \frac{\hat{v}^K}{\hat{v}^M} \implies \bar{X}_B(v^K,\, v^M) \geq \bar{X}_B(\hat{v}^K,\, \hat{v}^M).$$

By Lemma 3, it follows that there exists a non-decreasing $X_B(\cdot)$ such that

$$\bar{X}_B(v^K,\, v^M) = X_B(v^K/v^M)$$

which finishes the proof.

A.1.2 Proof of Theorem 5

Given Lemma 3, the proof of Theorem 5 is a matter of simple accounting. If $(X_B,\, X_S,\, T_B,\, T_S)$ is feasible in the setting of Section 3, then, under the convention that agents report the one-

dimensional rate of substitution, the same mechanism is feasible in the two-dimensional framework. Because we have shown in Section 5 that the objective functions (OBJ) and (VAL) are identical under the assumptions of the theorem, optimality in one framework establishes optimality in the other.

For the converse part, we only have to show that given arbitrary G_j and $\lambda_j(r)$, we can find a joint distribution F_j of (v^K, v^M) that induces them under the mapping described by the theorem. Fixing the random variable r (on some probability space) with distribution $G_j(r)$, define random variables $v^K = r\lambda_j(r)$ and $v^M = \lambda_j(r)$. It is clear that the distribution of v^K/v^M is the same as that of r because these random variables are equal. Moreover, again by definition,

$$m_j D_j(r) = \mathbb{E}^j \left[v^M \mid \frac{v^K}{v^M} = r \right] = \lambda_j(r).$$

This finishes the proof.

A.2 Interpreting our key regularity condition

We have imposed a relatively strong regularity condition (Assumption 1); in this section, we give an interpretation.

Observe that a price mechanism (in which we ignore market-clearing and only look at one side of the market) can be used to measure same-side inequality. Consider the seller side for concreteness, and suppose that we raise the price from p to $p+\epsilon$. The following two terms capture the associated gain in seller surplus:

$$\epsilon \int_{\underline{\tau}_S}^{p} m_S D_S(\tau) g_S(\tau) d\tau + \int_{p}^{p+\epsilon} m_S D_S(\tau)(p+\epsilon-\tau) d\tau.$$

That is, *(1)* sellers who were already selling at price p still sell at price $p+\epsilon$, and hence receive an additional transfer of ϵ, which has expected value $m_S D_S(\tau)$ for a seller with rate τ; and *(2)* sellers with $\tau \in (p,\, p+\epsilon]$ decide to sell receiving the corresponding surplus. The social cost of increasing the price to $p+\epsilon$ is that more revenue must be generated by the mechanism to cover the additional expenditure. If the shadow cost of revenue is α, then that cost is

$$-\alpha \left[\epsilon G(p) + (p+\epsilon)(G(p+\epsilon) - G(p)) \right],$$

where the expression in brackets is equal to the additional monetary transfer to sellers associated with the price increase. Because we are interested in measuring same-side inequality, the relevant shadow cost of revenue comes from charging sellers a lump-sum fee to cover the expenditures—thus, the shadow cost α is equal to the average value for money m_S for sellers.

Dividing by ϵ, and taking $\epsilon \to 0$, yields the local (first-order) net gain in seller surplus at price p:

$$m_S \left[\frac{\int_{\underline{r}_S}^{p} D_S(\tau) g_S(\tau) d\tau}{g_S(p)} - \left(p + \frac{G_S(p)}{g_S(p)} \right) \right] g_S(p). \tag{A.2}$$

The first bracketed term in (A.2) is typically increasing, especially for small p when $D_S(p)$ is relatively high, while the second bracketed term—minus the virtual surplus—is decreasing. Assumption 1 states that the bracketed expression in (A.2) can change from increasing to decreasing at most once.[16] Indeed, note that

$$\frac{\int_{\underline{r}_S}^{p} D_S(\tau) g_S(\tau) d\tau}{g_S(p)} - \left(p + \frac{G_S(p)}{g_S(p)} \right) = \Delta_S(p) - p.$$

The function Δ_S measures same-side inequality by quantifying the reduction in inequality associated with a price mechanism with price p. The function Δ_S is 0 at the extremes $p \in \{\underline{r}_S, \bar{r}_S\}$ (by definition of D_S), and Δ_S is non-negative because D_S is decreasing. This is because a price mechanism with both the maximal (at which all sellers sell) and the minimal price (at which no sellers sell) fails to reduce inequality but any interior price mechanism induces a monetary transfer to relatively poorer sellers. The area below the graph of Δ_S (and above 0) can thus be seen as a compact measure of same-side inequality. In particular, $\Delta_S \equiv 0$ if and only if there is no same-side inequality. Our regularity assumption means that as the price moves between its two extreme values at which inequality is unchanged, the reduction in inequality changes in a monotone fashion.

The distinction between high and low same-side inequality comes from studying the behavior of Δ_S near $\underline{r}_S$, that is, for the poorest sellers. For p close to $\underline{r}_S$ we have

$$\Delta_S(p) - p \approx \Delta_S'(\underline{r}_S)p - p = (D_S(\underline{r}_S) - 2)p.$$

Thus, when same-side inequality is high ($D_S(\underline{r}_S) > 2$), even a one-sided price mechanism leads to an improvement in the social objective function when the price is set so that only the poorest sellers trade. When same-side inequality is low ($D_S(\underline{r}_S) \leq 2$), this is no longer the case, and, intuitively, the transfer received by poorest sellers must come from an exogenous source (the buyer side).

Figure A.1 depicts the function $\Delta_S(p)$ derived under the assumption that v_k is uniform on [0, 1] and v_M follows a Pareto distribution with tail parameter $\alpha > 2$, for different values

[16] The reason why the density $g_S(p)$ was pulled out from the expression in brackets is that in the full analysis we have to account for market-clearing, and thus the bracketed expression is weighted by the mass of sellers at a given price p.

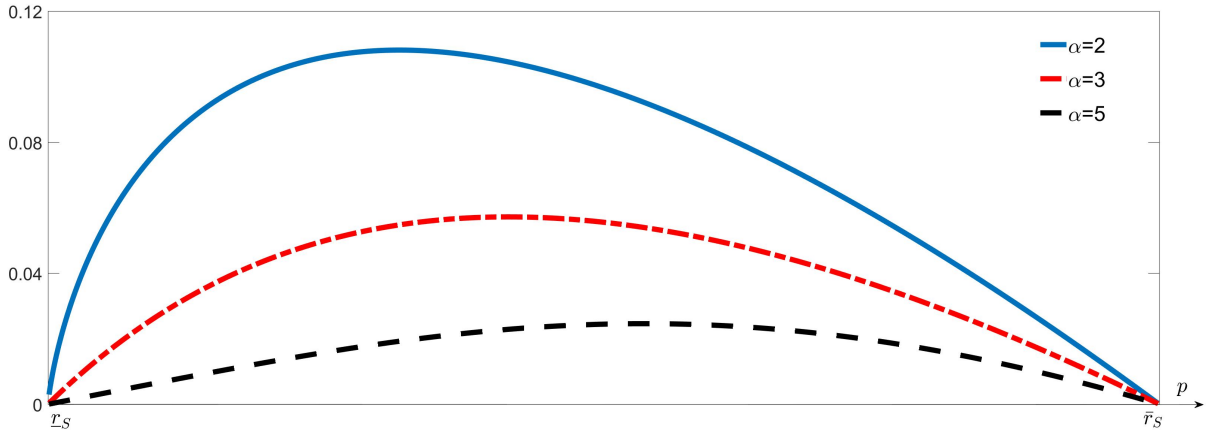


Figure A.1: $\Delta_S(p)$ when v_M is distributed according to Pareto distribution with parameter α.

of α. The functions are concave which is a sufficient condition for Assumption 1. The area below the curve is larger when α is smaller—the fatter the tail of agents with high value for money, the larger the inequality. It can be shown more generally that $\Delta_S(p)$ is concave when v_k follows a distribution with CDF $F_S^K(v) = v^\beta$ for $\beta > 0$ and $\alpha > \beta$. Another sufficient condition for concavity of $\Delta_j(r)$, for $j \in \{B, S\}$, is that $g_B(p)$ is non-increasing, same-side inequality is low for sellers, and $g_S(p)$ is non-decreasing.

B Proofs Omitted from the Main Text

B.1 Proof of Lemma 1

The constraint (4.5) satisfies the generalized Slater condition (see, e.g., Proposition 2.106 and Theorem 3.4 of Bonnans and Shapiro (2000)), so an approach based on putting a Lagrange multiplier $\alpha \geq 0$ on the constraint (4.5) is always valid (strong duality holds). Moreover, we can assume without loss of generality that constraint (4.5) will bind at the optimal solution (because G_j admits a density, it follows that there exists a positive measure of buyers and sellers with strictly positive value for good M). This means that the problem (4.3)–(4.5) is equivalent to the following statement: There exists $\alpha^\star \geq 0$ and the solution to the problem

$$\max \left\{ \int_0^1 \phi_B^{\alpha^\star}(q) d(\mu H_B(q)) + \int_0^1 \phi_S^{\alpha^\star}(q) dH_S(q) \right\} \tag{B.1}$$

over H_S, $H_B \in \Delta([0,1])$, $\underline{U}_B$, $\underline{U}_S \geq 0$, subject to

$$\int_0^1 q d(\mu H_B(q)) = \int_0^1 q dH_S(q) \tag{B.2}$$

satisfies constraint (4.5) with equality.

The value of the problem (B.1)–(B.2) can be computed by parameterizing $Q = \int_0^1 q dH_S(q)$, and noticing that for a fixed Q, the optimal distributions $H_S^\star$ yields the value of the concave closure of $\phi_S^{\alpha^\star}$ at Q. Similarly, the optimal distribution $H_B^\star$ yields the value of the concave closure of $\mu\phi_B^{\alpha^\star}$ at Q/μ. Optimizing the value of the unconstrained problem

$$\operatorname{co}\left(\phi_S^{\alpha^\star}\right)(Q) + \mu \operatorname{co}\left(\phi_B^{\alpha^\star}\right)(Q/\mu)$$

over Q, $\underline{U}_B$, $\underline{U}_S \geq 0$ yields the optimal solution to the original problem if constraint (4.5) holds with equality at that solution. This gives the first part of the lemma.

Conversely, if $H_B^\star$ and $H_S^\star$ are optimal, then constraints (4.4)—(4.5) must bind. Optimality of $H_j^\star$ implies that the value of ϕ_j^α at the optimum must be equal to its concave closure. We can define $Q = \int_0^1 q dH_S(q)$, and it must be that there exists $\alpha^\star \geq 0$ such that Q maximizes (B.1). This yields the second part of the lemma.

B.2 Completion of the proof of Theorem 1

The representation of the optimal mechanism follows directly from Lemma 1. Take for example the optimal $H_S^\star$. Since $H_S^\star$ concavifies a one-dimensional function ϕ_S^α, we can assume without loss of generality that it is supported at at most two points, that is,

$$H_S^\star(q) = (1-\eta)\mathbf{1}_{\{q \geq q_L\}} + \eta \mathbf{1}_{\{q \geq q_H\}}$$

for some $q_L < q_H$ and $\eta \in [0,1]$. Let $r_L = G_S^{-1}(q_L)$ and $r_H = G_S^{-1}(q_H)$. Then, the corresponding optimal $X_S(r)$ is given by

$$X_S(r) = (1-\eta)\mathbf{1}_{\{r \leq r_L\}} + \eta \mathbf{1}_{\{r \leq r_H\}}.$$

From this we can compute that

$$T_S(r) = \begin{cases} \underline{U}_S & r > r_H \\ \underline{U}_S + \eta r_H & r_H \geq r > r_L \\ \underline{U}_S + \eta(r_H - r_L) + r_L & r_L \geq r \end{cases}.$$

Therefore, the mechanism for the seller side can be implemented in the following way. Each seller gets a lump-sum subsidy of $\underline{U}_S$ units of good M. A seller can choose to sell a single good K for a price of $\eta(r_H - r_L) + r_L$ or choose to enter a lottery in which she sells good K with probability η at a price of r_H. Then, sellers with types below r_L choose the first option, types between r_L and r_H choose the second option and remaining types refuse to trade. We can parametrize the problem differently by defining $p_S^H = r_H$, $p_S^L = \delta_S(r_H - r_L) + r_L$, where $\delta_S = \eta$ is the rationing coefficient.

Analogously, for the buyer side, there is a subsidy $\underline{U}_B$, and then the buyer can choose to buy for sure at a price of p_B^H or with probability δ_B at a lower price p_B^L.

The only claim that we have to prove is that the optimal pricing system generates a non-negative surplus of good M, and that the surplus is redistributed to the side of the market j that has a higher average Pareto weight Λ_j (that is, $\underline{U}_j = 0$ for the side of the market with lower Λ_j).

First, notice that the pricing system $(p_B^L,\, p_B^H,\, \delta_B,\, p_S^L,\, p_S^H,\, \delta_S)$ cannot generate a deficit of good M because in such a case the mechanism could not satisfy the budget balance condition (4.5) together with the constraints $\underline{U}_B \geq 0, \underline{U}_S \geq 0$. Second, notice that Lemma 1 requires that the problem

$$\max_{Q\in[0,\,1],\,\underline{U}_B,\,\underline{U}_S\geq 0} \left\{ \operatorname{co}\left(\phi_S^{\alpha^\star}\right)(Q) + \mu \operatorname{co}\left(\phi_B^{\alpha^\star}\right)(Q/\mu) \right\}$$

has a solution, and this restricts the Lagrange multiplier to satisfy $\alpha^\star \geq \max\{\Lambda_S,\, \Lambda_B\}$. Indeed, in the opposite case, it would be optimal to set $\underline{U}_j = \infty$ for some j and this would clearly violate constraint (4.5). When $\Lambda_B = \Lambda_S$, it is either optimal to set $\alpha^\star > \Lambda_S = \Lambda_B$ and satisfy (4.5) with equality and $\underline{U}_S = \underline{U}_B = 0$ (in which case there is no revenue and no redistribution), or to set $\alpha^\star = \Lambda_S = \Lambda_B$ and $\underline{U}_S = 0$ and choose $\underline{U}_B \geq 0$ to satisfy (4.5) with equality (in which case the revenue is redistributed to buyers as an equal lump-sum payment).[17] When $\Lambda_B > \Lambda_S$, by similar reasoning, $\underline{U}_S$ must be 0, and $\underline{U}_B \geq 0$ is chosen to satisfy (4.5). When $\Lambda_S > \Lambda_B$, it is the seller side that obtains the lump-sum payment that balances the budget (4.5).

[17] Of course, in this case, the surplus can also be redistributed to sellers, or to both sides of the market, as long as condition (4.5) holds.

B.3 Proof of Theorem 2

Theorem 2 (generalized statement)**.** *A competitive equilibrium mechanism is optimal if and only if*

$$m_S\Delta_S(p^{\mathsf{CE}}) - m_B\Delta_B(p^{\mathsf{CE}}) \geq \begin{cases} (m_S - m_B)\frac{1-G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})} & \text{if } m_S \geq m_B \\ (m_B - m_S)\frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})} & \text{if } m_B \geq m_S. \end{cases} \tag{B.3}$$

When condition (B.3) fails, we have $p_B > p_S$, and prices are determined by market-clearing $\mu(1 - G_B(p_B)) = G_S(p_S)$, and, in the case of an interior solution,[18]

$$p_B - p_S = \begin{cases} -\frac{1}{m_S}\left[m_S\Delta_S(p_S) - m_B\Delta_B(p_B) - (m_S - m_B)\frac{1-G_B(p_B)}{g_B(p_B)}\right] & \text{if } m_S \geq m_B \\ -\frac{1}{m_B}\left[m_S\Delta_S(p_S) - m_B\Delta_B(p_B) - (m_B - m_S)\frac{G_S(p_S)}{g_S(p_S)}\right] & \text{if } m_B \geq m_S. \end{cases} \tag{B.4}$$

The mechanism subsidizes the sellers when $m_S > m_B$, and the buyers when $m_B > m_S$.

Proof. We will show that under the assumptions of the theorem, the functions $\phi_j^{\alpha^\star}$ are concave with the optimal Lagrange multiplier $\alpha^\star$. This is sufficient to prove optimality of a price mechanism because of Lemma 1—when the objective function is concave, it coincides with its concave closure, and thus the optimal distribution of quantities is degenerate, corresponding to posted-price mechanism.

As argued in the proof of Theorem 1, we must have $\alpha^\star \geq \max\{m_S, m_B\}$. Then, the derivative of the function $\phi_S^{\alpha^\star}(q)$ takes the form

$$(\phi_S^{\alpha^\star})'(q) = \Pi_S^\Lambda(G_S^{-1}(q)) - \alpha^\star J_S(G_S^{-1}(q)),$$

so it is enough to prove that $\Pi_S^\Lambda(r) - \alpha^\star J_S(r)$ is non-increasing in r. Rewriting,

$$\Pi_S^\Lambda(r) - \alpha^\star J_S(r) = m_S \underbrace{\left[\frac{\int_{\underline{r}_S}^r [D_S(\tau) - 1]g_S(\tau)d\tau}{g_S(r)} - r\right]}_{\Delta_S(r) - r} - (\alpha^\star - m_S)J_S(r)$$

Virtual surplus $J_S(r)$ is non-decreasing, and $\alpha^\star \geq m_S$, so it is enough to prove that the first term is non-increasing. The function $\Delta_S(r) - r$ is quasi-concave by Assumption 1, so to prove monotonicity on the entire domain, it is enough to show that the derivative at $r = \underline{r}_S$ is negative. We have

$$\frac{d}{dr}[\Delta_S(r) - r]|_{r=\underline{r}_S} = D_S(\underline{r}_S) - 2 \leq 0,$$

[18] When no such interior solution exists, one of the prices is equal to the bound of the support: either $p_B = \underline{r}_B$ or $p_S = \bar{r}_S$.

where the last inequality follows from the assumption that same-side inequality is low.

We now show that $\phi_B^{\alpha^\star}(q)$ is also concave:

$$(\phi_B^{\alpha^\star})'(q) = \Pi_B^\Lambda(G_B^{-1}(1-q)) + \alpha^\star J_B(G_B^{-1}(1-q)),$$

so it is enough to show that $\Pi_B^\Lambda(r) + \alpha^\star J_B(r)$ is non-decreasing. Rewriting,

$$\Pi_B^\Lambda(r) + \alpha^\star J_B(r) = m_B\left[\Delta_B(r) + r\right] + (\alpha^\star - m_B)J_B(r).$$

Because the virtual surplus function $J_B(r)$ is non-decreasing, and $\alpha^\star \geq m_B$, by assumption, it is enough to prove that $\Delta_B(r)+r$ is non-decreasing. Because this function is quasi-concave by Assumption 1, it is enough to prove that the derivative is non-negative at the end point $r = \bar{r}_B$:

$$\frac{d}{dr}[\Delta_B(r) + r]|_{r=\bar{r}_B} = D_B(\bar{r}_B) \geq 0,$$

which is trivially satisfied. Thus, we have proven that both functions $\phi_j^{\alpha^\star}$ are concave.

It follows that a price mechanism with no rationing is optimal for both sides of the market, and the revenue (if strictly positive) is redistributed to the sellers if $m_S \geq m_B$, and to the buyers otherwise (see Theorem 1). Concavity of $\phi_j^{\alpha^\star}$ implies that the first-order condition in problem (4.6) has to hold and is sufficient for optimality. This means that the optimal volume of trade $Q^\star \in [0,\, \mu \wedge 1]$ (the maximizer of the right hand side of (4.6)) satisfies

$$\Pi_S^\Lambda(G_S^{-1}(Q^\star)) - \alpha^\star J_S(G_S^{-1}(Q^\star)) + \Pi_B^\Lambda(G_B^{-1}(1-\frac{Q^\star}{\mu})) + \alpha^\star J_B(G_B^{-1}(1-\frac{Q^\star}{\mu})) \geq 0$$
$$(= 0 \text{ when } Q^\star = \mu \wedge 1). \quad \text{(B.5)}$$

Rewriting, and noting that $p_S = G_S^{-1}(Q^\star)$ and $p_B = G_B^{-1}(1-\frac{Q^\star}{\mu})$,

$$m_S[\Delta_S(p_S) - p_S] - (\alpha^\star - m_S)J_S(p_S) + m_B[p_B - \Delta_B(p_B)] + (\alpha^\star - m_B)J_B(p_B) \geq 0$$
$$(= 0 \text{ when } Q^\star = \mu \wedge 1). \quad \text{(B.6)}$$

Moreover, prices p_B, p_S have to satisfy $p_B \geq p_S$ and clear the market:

$$\mu(1 - G_B(p_B)) = G_S(p_S). \quad \text{(B.7)}$$

First, assume that (B.3) holds at the competitive equilibrium price p^{CE}. We will show that in this case, competitive equilibrium is optimal. At p^{CE}, market-clearing and budget balance (with $\underline{U}_S = \underline{U}_B = 0$) hold, by construction. Therefore, we only have to prove existence of

$\alpha^\star \geq \max\{m_S,\, m_B\}$ such that the first-order condition holds:

$$m_S[\Delta_S(p^{\mathsf{CE}}) - p^{\mathsf{CE}}] - (\alpha^\star - m_S)J_S(p^{\mathsf{CE}}) + m_B[p^{\mathsf{CE}} - \Delta_B(p^{\mathsf{CE}})] + (\alpha^\star - m_B)J_B(p^{\mathsf{CE}}) \geq 0$$

with equality when the solution is interior: $p^{CE} \in (\underline{r}_S,\, \bar{r}_B)$. Simplifying the above expression:

$$m_S \underbrace{\left[\Delta_S(p^{\mathsf{CE}}) + \frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})}\right]}_{\geq 0} - m_B \underbrace{\left[\Delta_B(p^{\mathsf{CE}}) - \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}\right]}_{\leq 0} \geq \alpha^\star \left[\frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})} + \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}\right]$$

with equality when $p^{CE} \in (\underline{r}_S,\, \bar{r}_B)$. Since the left hand side is non-negative, such a solution $\alpha^\star \geq \max\{m_S,\, m_B\}$ exists if and only if we have an inequality at the minimal possible $\alpha^\star$, that is, $\alpha^\star = \max\{m_S,\, m_B\}$:

$$m_S \left[\Delta_S(p^{\mathsf{CE}}) + \frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})}\right] - m_B \left[\Delta_B(p^{\mathsf{CE}}) - \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}\right] \geq \max\{m_S,\, m_B\} \left[\frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})} + \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}\right].$$

Simplifying the above expression shows that it is equivalent to condition (B.3).

It remains to show how the solution looks like when condition (B.3) fails. A competitive equilibrium cannot be optimal in this case because there does not exist a $\alpha^\star$ under which the corresponding quantity maximizes the Lagrangian (4.6) in Lemma 1. This means that $p_B > p_S$, and, in light of Theorem 1, there will be a strictly positive lump-sum payment for the "poorer" side of the market: $\underline{U}_S > 0$ when $m_S \geq m_B$ and $\underline{U}_B > 0$ when $m_B > m_S$. This implies that we must have $\alpha^\star = \max\{m_S,\, m_B\}$. Subsequently, the optimal prices p_B and p_S are pinned down by market-clearing (B.7) and the first-order condition (B.6) which becomes, assuming that an interior solution exists,

$$m_S(p_B - p_S) = -m_S\Delta_S(p_S) + m_B\Delta_B(p_B) + (m_S - m_B)\frac{1 - G_B(p_B)}{g_B(p_B)}$$

when $m_S \geq m_B$, and

$$m_B(p_B - p_S) = -m_S\Delta_S(p_S) + m_B\Delta_B(p_B) + (m_B - m_S)\frac{G_S(p_S)}{g_S(p_S)}$$

otherwise. When there is no interior solution, one of the prices is equal to the bound of the support of rates of substitution, and the other price is determined by the market-clearing condition. This finishes the proof of the theorem. □

B.4 Proof of Theorem 3

We start by proving a lemma.

Lemma 4. *There exist $\hat{q} > 0$ and $\bar{\alpha} > m_S$ such that if $\alpha < \bar{\alpha}$, then $\phi_S^\alpha(q)$ is strictly convex on $[0, \hat{q}]$.*

Proof. The derivative of $\phi_S^\alpha(q)$ is $\Pi_S^\Lambda(G_S^{-1}(q)) - \alpha J_S(G_S^{-1}(q))$. Because the function G_S^{-1} is strictly increasing, it is enough to prove that $\Pi_S^\Lambda(r) - \alpha J_S(r)$ is strictly increasing for some $r \in [\underline{r}_S, \hat{r}]$ (we then set $\hat{q} = G_S(\hat{r})$). Taking a derivative again, and rearranging, yields the following sufficient condition: for $r \in [\underline{r}_S, \hat{r}]$,

$$D_S(r) > 2 + \frac{g_S'(r)}{g_S(r)}\Delta_S(r) + \frac{\alpha - m_S}{m_S}\left[2 - \frac{g_S'(r)G_S(r)}{g_S^2(r)}\right].$$

Because g_S was assumed continuously differentiable and strictly positive, including at $r = \underline{r}_S$, we can put a uniform (across r) bound $M < \infty$ on $\frac{g_S'(r)}{g_S(r)}$ and $2 - \frac{g_S'(r)G_S(r)}{g_S^2(r)}$. This means that it is enough that

$$D_S(r) > 2 + M\Delta_S(r) + \frac{\alpha - m_S}{m_S}M.$$

Continuity of $D_S(r)$ and the assumption that same-side inequality for sellers is high imply that $D_S(r) > 2 + \epsilon$ for $r \in [\underline{r}_S, \underline{r}_S + \delta]$ for some $\delta > 0$. Continuity of $\Delta_S(r)$ and the fact that $\Delta_S(\underline{r}_S) = 0$ imply that $\Delta_S(r) < \epsilon/(3M)$ for all $r \in [\underline{r}_S, \underline{r}_S + \nu]$ for some $\nu > 0$. Finally, there exists a $\bar{\alpha} > m_S$ such that for all $\alpha < \bar{\alpha}$, we have $(\alpha - m_S)/m_S < \epsilon/(3M)$. Then, for all $r \in [\underline{r}_S, \underline{r}_S + \min\{\delta, \nu\}]$, $\alpha < \bar{\alpha}$,

$$D_S(r) > 2 + \epsilon > 2 + M\underbrace{\Delta_S(r)}_{\leq \epsilon/(3M)} + \underbrace{\frac{\alpha - m_S}{m_S}}_{\leq \epsilon/(3M)}M.$$

The proof is finished by setting $\hat{r} = \underline{r}_S + \min\{\delta, \nu\}$. □

We now prove Theorem 3. The optimal mechanism for the buyer side is a price mechanism—this follows from the proof of Theorem 2. Suppose, towards a contradiction, that the optimal mechanism for sellers is also a price mechanism. Then, there are two possibilities. Either, (1) a competitive equilibrium mechanism is optimal, or (2) $\alpha^\star = m_S$.[19]

Consider case (2) first. We can invoke Lemma 4. Because $\alpha^\star = m_S$, there exists $\hat{q} > 0$ such that $\phi_S^{\alpha^\star}(q)$ is strictly convex on $[0, \hat{q}]$. Suppose that $\mu < \hat{q}$. In this case, at the optimal quantity $0 < Q^\star < \mu$, $\phi_S^{\alpha^\star}(q)$ cannot be equal to its concave closure. Thus, the optimal value

[19] Indeed, when a price mechanism (p_B, p_S) is not a competitive equilibrium, we have $p_B > p_S$, and thus $\underline{U}_j > 0$ for some j to ensure budget balance. But then, in light of Lemma 1, we need $\alpha^\star = m_S$.

is obtained by randomizing over two distinct quantities q_1 and q_2 under $H_S^\star$. This means that the optimal mechanism for sellers necessarily involves a non-zero-measure region $[r_1, r_2]$ in which $X^S(r) \in (0, 1)$ is interior, i.e., the optimal mechanism rations the sellers.

Consider case (1). We know that the first-order condition of the problem (4.6) must hold at the optimal $Q^\star$ and $\alpha^\star$. Because the functions $\phi_j^{\alpha^\star}$ are differentiable, and coincide with their concave closure at the optimum (only then a price mechanism can be optimal), the derivative of ϕ_j^α must be equal to the derivative of its concave closure. Therefore, the first-order condition implies that

$$(\phi_S^{\alpha^\star})'(Q^\star) + (\phi_B^{\alpha^\star})'(Q^\star/\mu) \geq 0$$

with equality for $Q^\star < \mu$. As in the proof of Theorem 2, we can write this condition as

$$m_S \underbrace{\left[\Delta_S(p^{\mathsf{CE}}) + \frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})}\right]}_{\geq 0} - m_B \underbrace{\left[\Delta_B(p^{\mathsf{CE}}) - \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}\right]}_{\leq 0} \geq \alpha^\star \left[\frac{G_S(p^{\mathsf{CE}})}{g_S(p^{\mathsf{CE}})} + \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}\right] \tag{B.8}$$

with equality for $Q^\star < \mu$. Since the left hand side is non-negative, a solution $\alpha^\star \geq m_S$ can only exist if we have an inequality at $\alpha^\star = m_S$, that is, condition (5.6) must hold

$$m_S \Delta_S(p^{\mathsf{CE}}) - m_B \Delta_B(p^{\mathsf{CE}}) \geq (m_S - m_B) \frac{1 - G_B(p^{\mathsf{CE}})}{g_B(p^{\mathsf{CE}})}. \tag{B.9}$$

To get a contradiction, it is enough to prove that this condition fails when μ is low enough. Note that the competitive equilibrium price, which we will now index by μ, satisfies

$$G_S(p_\mu^{\mathsf{CE}}) = \mu(1 - G_B(p_\mu^{\mathsf{CE}})).$$

Thus, $p_\mu^{\mathsf{CE}} \to \underline{r}_S$ as $\mu \to 0$. When $p_\mu^{\mathsf{CE}} \to 0$, the left hand side of (B.9) converges to a non-positive number (because $\Delta_S(\underline{r}_S) = 0$). The right hand side converges to $(m_S - m_B)(1 - G_B(\underline{r}_S))/g_B(\underline{r}_S)$ which is strictly positive (possibly ∞) as long as $m_S > m_B$. Thus, for the case $m_S > m_B$, by continuity, for small enough μ, condition (B.9) is violated, contradicting the optimality of a price mechanism.

In the case when $m_S = m_B$, a different argument is needed. The first-order condition (B.8) can only hold for $\mu \to 0$ if $\alpha_\mu^\star \to m_S$. However, fixing $\hat{q}$ and $\bar{\alpha}$ from Lemma 4, we get that for sufficiently small μ, $\alpha_\mu^\star < \bar{\alpha}$ and thus $\phi_S^{\alpha_\mu^\star}(q)$ is strictly convex on $[0, \hat{q}]$. This, however, leads to a contradiction because once $\mu < \hat{q}$, the optimal mechanism cannot be a price mechanism, since $\phi_S^{\alpha_\mu^\star}(q)$ cannot be equal to its concave closure at any $q \leq \mu$.

Finally, we prove Remark 1 stated after Theorem 3. When m_S is sufficiently higher than m_B so that competitive equilibrium is not optimal, we only have to consider the case $\alpha^\star = m_S$. When $g_S(p)$ is decreasing, it is easy to notice that $\Pi_S^\Lambda(r) - \alpha^\star J_S(r)$ is increasing for all $r \leq r^\star$, and hence $\phi_S^{\alpha^\star}(q)$ is convex for all $q \leq G_s(r^\star)$. It follows that if the equilibrium volume of trade is bounded by $\mu \leq G_S(r^\star)$, the concave closure of $\phi_S^{\alpha^\star}$ at $Q^\star$ is obtained by mixing between the quantity of 0, and some quantity $\bar{q} \geq Q^\star$. The conclusion follows.

B.5 Proof of Theorem 4

The proof follows directly from the arguments used in the proof of Theorem 2. Under Assumption 1, the function $\phi_B^{\alpha^\star}$ is concave, and thus the optimal mechanism for the buyer side is always a price mechanism.

B.6 Proof of Proposition 1

The function $\phi_S^\alpha(q)$ is concave if and only if

$$(m_S - \alpha)\frac{G_S(r)}{g_S(r)} - \alpha r \tag{B.10}$$

is decreasing. Similarly, the function $\phi_B^\alpha(q)$ is concave if and only if

$$-(1-\alpha)\frac{1 - G_B(r)}{g_B(r)} - \alpha r \tag{B.11}$$

is decreasing. We will prove that the optimal Lagrange multiplier α must lie between 1 and m_S.

Lemma 5. *At the optimal solution, $\alpha \in [1, m_S]$.*

Proof. Let us consider all possible cases given that $\alpha \geq 0$ and $m_S \geq 1 = (m_B)$.

1. $\alpha > m_S$: Both functions (B.10) and (B.11) are decreasing, by Assumption 3. Therefore, both ϕ_j^α are concave, and the optimal mechanism is a price mechanism. Given that lump-sum transfers are now ruled out by the no-free-lunch assumption, the only solution in this case can be a competitive equilibrium.

2. $m_S \geq \alpha \geq 1$: The function (B.11) is decreasing, so a price mechanism is optimal for the buyer side.

3. $1 > \alpha$: There could potentially be rationing on both sides of the market.

We will show that cases 1 and 3 cannot take place. Consider case 1. In this case, we know that the first-order condition (B.6) has to hold at the optimal quantity $Q^\star$. Simplifying using the assumptions of the current setting, we obtain

$$\mu \frac{m_S - \alpha}{g_S(p^{\mathsf{CE}})} = \frac{\alpha - 1}{g_B(p^{\mathsf{CE}})}. \tag{B.12}$$

This condition can only hold when $\alpha \in [1, m_S]$ (and can be strengthened to $\alpha \in (1, m_S]$ when $m_S > 1$).

Now consider case 3. We will rely on the following observation. We know that the optimal schedules X_S and X_B can be supported as randomization over at most two quantities for each side: $H_S^\star$ is supported on $q_S^L \le q_S^H$, and $H_B^\star$ is supported on $q_B^L \le q_B^H$. The derivative of the concave closure of ϕ_j^α at the optimal quantity $Q^\star$ must be equal to the derivative of ϕ_j^α at both q_j^H and q_j^L as long as there are interior points. If $q_j^L = 0$, then the derivative of the concave closure at $Q^\star$ is weakly larger than the derivative of ϕ_j^α at q_j^L, and if $q_j^H = 1$, then the derivative of the concave closure at $Q^\star$ is weakly smaller than the derivative of ϕ_j^α at q_j^H. Formally, we obtain

$$(m_S - \alpha)\frac{q_S^L}{g_S(G_S^{-1}(q_S^L))} - \alpha G_S^{-1}(q_S^L) \le (m_S - \alpha)\frac{q_S^H}{g_S(G_S^{-1}(q_S^H))} - \alpha G_S^{-1}(q_S^H)$$

(with equality if both points q_S^L and q_S^H are interior), and

$$(1-\alpha)\frac{q_B^L}{g_B(G_B^{-1}(1 - q_B^L/\mu))} + \alpha G_B^{-1}(1 - q_B^L/\mu) \le (1-\alpha)\frac{q_B^H}{g_B(G_B^{-1}(1 - q_B^H/\mu))} + \alpha G_B^{-1}(1 - q_B^H/\mu)$$

(with equality if both points q_B^L and q_B^H are interior). Because the first-order condition must hold at the optimal expected quantity $Q^\star$ which lies between q_j^L and q_j^H, we obtain

$$\begin{aligned} 0 &\ge (m_S - \alpha)\frac{q_S^L}{g_S(G_S^{-1}(q_S^L))} - \alpha G_S^{-1}(q_S^L) + (1-\alpha)\frac{q_B^L}{g_B(G_B^{-1}(1 - q_B^L/\mu))} + \alpha G_B^{-1}(1 - q_B^L/\mu) \\ &= \underbrace{(m_S - \alpha)\frac{q_S^L}{g_S(G_S^{-1}(q_S^L))} + (1-\alpha)\frac{q_B^L}{g_B(G_B^{-1}(1 - q_B^L/\mu))}}_{>0} - \alpha\left(G_S^{-1}(q_S^L) - G_B^{-1}(1 - q_B^L/\mu)\right). \end{aligned} \tag{B.13}$$

We will argue that $G_S^{-1}(q_S^L) \le G_B^{-1}(1 - q_B^L/\mu)$ which will finish the proof of the lemma. Suppose to the contrary that $G_S^{-1}(q_S^L) > G_B^{-1}(1 - q_B^L/\mu)$. Then we have

$$G_S^{-1}(q_S^H) \ge G_S^{-1}(q_S^L) > G_B^{-1}(1 - q_B^L/\mu) \ge G_B^{-1}(1 - q_B^H/\mu).$$

This however, clearly contradicts the budget balance condition because both prices offered to sellers are strictly higher than the price offered to buyers. Thus, the lemma is proven.

□

We now prove the proposition. First, the lemma implies that a price mechanism is optimal for the buyer side. Second, we rewrite the probability distributions in terms of marginal distributions for good K: $G_B(p) = F_B^K(p)$, $G_S(p) = F_S^K(m_S p)$, $g_S(p) = m_S f_S^K(m_S p)$. Third, a sufficient condition for optimality of competitive equilibrium is that $\phi_S^{\alpha^\star}(q)$ is concave, that is,

$$\frac{m_S - \alpha^\star}{m_S} \frac{d}{dr} \left(\frac{F_S^K(m_S r)}{f_S^K(m_S r)} \right) \leq \alpha^\star.$$

Using the lemma, we know that $\alpha \geq 1$, so it is enough that, for all r,

$$\frac{d}{dr} \left(\frac{F_S^K(r)}{f_S^K(r)} \right) \leq \frac{1}{m_S - 1} \leq \frac{\alpha^\star}{m_S - \alpha^\star}.$$

Finally, a necessary condition for optimality of competitive equilibrium is that the first-order condition (B.12) holds. Rewriting in terms of the distribution of values for good K, we obtain:

$$\frac{m_S - \alpha^\star}{\alpha^\star - 1} = \frac{1}{\mu} \frac{m_S f_S^K(m_S p^{\mathsf{CE}})}{f_B^K(p^{\mathsf{CE}})}$$

Take c defined in the statements of the proposition. Then,

$$\alpha^\star = \frac{m_S(\mu + c)}{\mu + m_S c}.$$

We will show that this implies (under the assumptions of the proposition) that $\phi_S^{\alpha^\star}(q)$ is strictly convex in the neighborhood of the corresponding optimal quantity, leading to a contradiction. By continuity of the derivatives, it is enough to show that

$$(m_S - \alpha^\star)d > \alpha^\star,$$

where d has been defined in the statement of the proposition. Equivalently,

$$m_S > \frac{m_S(\mu + c)}{\mu + m_S c} \left(1 + \frac{1}{d} \right).$$

Rewriting,

$$m_S > 1 + \frac{1}{d} + \frac{\mu}{cd}.$$

B.7 Proof of Claim 2

We will apply a modified version of Lemma 1. The modification, described in Section 6, is that under Assumption 2 we have $\underline{U}_B = \underline{U}_S = 0$, and thus these variables do not appear in the maximization problem (4.6). Under the assumptions of Section 2,

$$\phi_S^\alpha(q) = \frac{m - 2\alpha}{2m} q^2,$$

and

$$\phi_B^\alpha(q) = q\left[\alpha - \frac{2\alpha - 1}{2} q\right].$$

We also know from the proof of Proposition 1 that $\alpha \in [1, m]$. Thus, $\phi_B^\alpha(q)$ is concave.

First, we consider the case $m \leq 3$. We conjecture that $\alpha \geq m/2$, so that $\phi_S^\alpha(q)$ is concave. Then, to show optimality of competitive equilibrium, it is enough to find $\alpha \geq m/2$ such that the first-order condition corresponding to (4.6) holds at the competitive-equilibrium volume of trade $Q^\star = m/(m+1)$. The first-order condition is

$$\left[(2\alpha - 1) - \frac{m - 2\alpha}{m}\right] Q^\star = \alpha$$

which yields $\alpha^\star = (2m)/(m+1)$. For $m \leq 3$, we have $\alpha^\star \geq 1/m$.

Second, we consider the case $3 < m < 2 + \sqrt{2}$. We conjecture that $\alpha^\star = m/2$. This means that $\phi_S^{\alpha^\star}(q) = 0$ for all q. The first-order condition yields

$$Q^\star = \frac{\alpha^\star}{2\alpha^\star - 1} = \frac{m}{2(m-1)}$$

which corresponds to a price $(m-2)/(2m-2)$ for buyers. Because the function $\phi_S^{\alpha^\star}(q)$ is linear, any mechanism that yields the market-clearing volume is optimal for the seller side. In particular, we can specify that the price for sellers is $(m-2)/(2m-2)$ with a rationing coefficient δ_S that leads to market-clearing:

$$\delta_S G_S(p^\star) = Q^\star \implies \delta_S = \frac{1}{m-2}.$$

We only have to verify that $p^\star \leq 1/m$, so that the mechanism is feasible (does not violate Assumption 2). Solving $(m-2)/(2m-2) \leq 1/m$ yields the restriction $m \leq 2 + \sqrt{2}$.

Third, we consider the case $m \geq 2 + \sqrt{2}$. We conjecture that $\alpha \leq m/2$, so that the function $\phi_S^\alpha(q)$ is convex. Convexity of $\phi_S^\alpha(q)$ implies that the optimal mechanism on the seller side is to offer the maximal price $1/m$ that induces all sellers to sell, and ration to

achieve market-clearing. It also follows that

$$\mathrm{co}(\phi_S^\alpha)(q) = \frac{m-2\alpha}{2m}q.$$

Thus, the first-order condition for the problem (4.6) is

$$\frac{m-2\alpha}{2m} + \alpha = (2\alpha - 1)Q.$$

We want to show that $Q^\star = (m-1)/m$ is optimal (this is the volume corresponding to the mechanism described in the proposition). This requires

$$\alpha^\star = \frac{3m-2}{2(m-1)}.$$

We only have to check that $\alpha^\star \leq m/2$. This is indeed true when $m \geq 2+\sqrt{2}$.