Inequality and Market Design

PIOTR DWORCZAK
Northwestern University & Group for Research in Applied Economics

Policymakers are often concerned about inequalities in the markets they control. In this letter, I argue that mechanism design has not responded sufficiently to the need for a comprehensive theory of inequality-aware market design. I review some of my recent work trying to fill this gap and identify research directions where input from computer scientists would be particularly useful.

Categories and Subject Descriptors: Computer Applications [**Social and Behavioral Sciences**]: Economics

General Terms: Design, Economics, Theory

Additional Key Words and Phrases: Market Design, Mechanism Design, Inequality

1. INTRODUCTION

Within both economics and computer science, mechanism design has been successfully employed as a methodology to understand the optimal design of markets, auctions, trading platforms, and other economic systems. Following the seminal work of Myerson [1981], thousands of papers have been written on revenue-maximizing mechanisms in a variety of contexts and under various constraints. The question of efficient design received a similar level of attention: The classical results about the implementability of efficient outcomes (Vickrey [1961], Clarke [1971], Groves [1973]) have been examined from every possible angle.

Policymakers, though, often attempt to maximize an objective function that evidently differs from either revenue or allocative efficiency. Around the world, local governments of large cities impose rent control policies and run affordable housing programs. Developed and developing countries alike provide food at subsidized prices to some populations (e.g., food stamps in the US, or in-kind provision of rice and wheat in India). Citizens of many countries can use publicly provided (but typically low-quality) basic health care. Access to roads is free (or relatively cheap), leading to inefficient levels of congestion in most large agglomerations. Covid-19 vaccines were distributed (effectively) free of charge, despite their initial scarcity. Following the spike in energy prices caused by Russia's invasion of Ukraine, virtually all European countries began subsidizing or directly controlling electricity prices faced by households. National parks tend to allocate hiking permits by lottery rather than by charging a market-clearing (or revenue-maximizing) price. The list goes on.

Author's address: piotr.dworczak@northwestern.edu. The letter summarizes the research progress in my long-term collaboration with Mohammad Akbarpour and Scott Duke Kominers on Inequality-aware Market Design. I thank Andrzej Skrzypacz for helpful comments. I gratefully acknowledge support received under the ERC Starting grant IMD-101040122. Views and opinions expressed are those of the author only and do not necessarily reflect those of the European Union or the European Research Council.

In short, policymakers often care about inequalities in the markets they control—their objective function is *redistributive*, in the sense that they have preferences over splits of surplus across market participants, and are especially concerned about the welfare of the poorest (or otherwise disadvantaged) individuals.

My impression is that economic theorists have not paid enough attention to these kinds of problems because classical economic intuition—based on the celebrated welfare theorems—suggests that policies such as in-kind transfers, subsidies, or price controls are simply a mistake. According to welfare theorems, we should "redistribute endowments," and then let markets do their job (we would only intervene if markets are not competitive or feature externalities). However, the classical intuition is incomplete, to say the least. The second welfare theorem, in particular, breaks down under asymmetric information. In all the above examples, information available to policymakers is imperfect at best. And mechanism design provides tools to address the imperfect information problem head-on.

This is the subject of the research agenda on inequality-aware market design that I have been pursuing with Mohammad Akbarpour and Scott Duke Kominers (as well as several new co-authors more recently). We employ a mechanism-design framework to understand the optimal design of markets under imperfect information when the designer has a redistributive objective function. As it turns out, in-kind provision, subsidies, and even price controls may be part of the optimal design. But—as always—understanding whether and how we want to use these policy tools is a complex question that no single paper can address.[1]

My primary goal in this letter is to introduce the readers to a simple framework of inequality-aware market design, based on our two papers Dworczak ⓡ Kominers ⓡ Akbarpour [2021] and Akbarpour ⓡ Dworczak ⓡ Kominers [2024]. The interactions of economists and computer scientists, especially recently, led to a proliferation (and popularization) of new results and approaches in mechanism design. Issues related to robustness, complexity, and computational constraints are becoming part of mainstream mechanism design *across* the two fields. My hope is that some of the same ideas applied to inequality-aware market design can generate a comprehensive understanding of when and how markets should be regulated by policymakers concerned about inequality.

2. FRAMEWORK

Let us consider a simple setting of allocating a set of homogeneous indivisible goods (e.g., identical houses) to agents differing in their privately observed values for the good and social welfare weights. Welfare weights are a classical way of capturing the designer's preferences over how surplus is split: Higher welfare weights can be placed on agents who are poor, disadvantaged, or deserve preferential treatment for other social and moral reasons. In a reduced-form way, dispersion in welfare weights can thus capture various inequalities among market participants (under the implicit assumption that these inequalities motivate the designer's redistributive preferences).

[1]This letter is not intended to serve as a guide to the literature on the topic. Nevertheless, a shoutout to a few key contributions is in place. In particular, our work was preceded by the important papers of Weitzman [1977], Spence [1977], and Condorelli [2013].

There is a unit mass of agents, with each agent characterized by a type (r, λ), where r is the willingness to pay (WTP) for the good and λ is the social welfare weight. The designer knows the joint distribution $F(r, \lambda)$ of WTP and welfare weights but does not observe individual realizations. The unobservability of λ is a key assumption: The designer's redistributive preferences depend on characteristics (such as wealth, income, life circumstances, and health status) that may be difficult to observe or condition on.

An agent with type (r, λ) who gets the good with probability x and pays t gets utility $rx-t$ and contributes $\lambda(rx-t)$ to the social welfare function. Quasi-linearity of preferences gives λ a nice interpretation: It is the social value of giving an agent a dollar. The social welfare function is the sum of $\lambda(rx-t)$ across all agents; additionally, the designer attaches a weight $\alpha \geq 0$ to the revenue generated by the mechanism. Since α can be flexibly specified, it is without loss of generality to assume that transfers in the mechanism are non-negative: If the designer wants to give agents a lump-sum transfer (using the revenue generated by the mechanism), it suffices to set α to be equal to the average social welfare weight.

The mechanism chosen by the designer—specifying the allocation x and transfer t for each type (r, λ)—must be individually-rational and incentive-compatible. The designer can choose a total quantity Q of goods to allocate at a total cost given by $C(Q) \geq 0$. The problem of the designer is to maximize social welfare over all feasible mechanisms:

$$\max_{\substack{x:\,\mathrm{supp}(F)\to[0,\,1]\\ t:\,\mathrm{supp}(F)\to\mathbb{R}_+}} \int \left(\lambda\left[x(r,\lambda)r - t(r,\lambda)\right] + \alpha t(r,\lambda)\right) dF(r,\lambda) - C\left(\int x(r,\lambda)dF(r,\lambda)\right)$$

subject to

$$x(r,\lambda)r - t(r,\lambda) \geq x(r',\lambda')r - t(r',\lambda'), \qquad \forall (r,\lambda),\,(r',\lambda') \in \mathrm{supp}(F),$$
$$x(r,\lambda)r - t(r,\lambda) \geq 0, \qquad \forall (r,\lambda) \in \mathrm{supp}(F).$$

The problem can be solved by adapting a few standard steps.[2]

(1) Even though the designer would like to condition the allocation on realized welfare weights λ, she cannot; this is because λ does not affect agents' individual payoffs, and hence agents would not report λ truthfully if it led to better outcomes conditional on some reports. As long as r is continuously distributed,[3] it is without loss of optimality for the designer to condition the allocation and transfers on the reported WTP only. I will abuse notation slightly and denote by $F(r)$ the marginal distribution of r. I also assume F has a continuous density f with support $[0, \bar{r}]$.[4]

[2]See Dworczak ⓡ Kominers ⓡ Akbarpour [2021] and Akbarpour ⓡ Dworczak ⓡ Kominers [2024] for details and missing steps.
[3]With a discrete type space, the designer could elicit some information about λ relying on agents' indifferences; see Ostrizek and Sartori [2023].
[4]That the lower bound of WTP is 0 is not an innocuous assumption; see Akbarpour ⓡ Dworczak ⓡ Kominers [2024] for the analysis of the general case.

(2) As a result, the objective function can be written as

$$\int \left(\lambda(r)\left[x(r)r - t(r)\right]\right) dF(r) + \alpha \int t(r)dF(r) - C\left(\int x(r)dF(r)\right),$$

where $\lambda(r) = \mathbb{E}[\lambda | \, r]$. That is, even though the designer cannot elicit λ truthfully, she can still use its *statistical correlation* with willingness to pay to inform the choice of the optimal mechanism.

(3) Once we have a standard one-dimensional screening problem, we can use the ideas of Myerson [1981]: Incentive-compatibility requires $x(r)$ to be weakly increasing in r and pins down transfers uniquely, up to the utility $\underline{U}$ of the agent with the lowest WTP. Because I assumed that $\min(\text{supp}(F)) = 0$ and transfers are non-negative, we must in fact have $\underline{U} = 0$. Then, we know that

$$t(r) = x(r)r - \int_0^r x(\tau)d\tau, \forall r.$$

(4) A simple sequence of standard transformations applied to the social welfare function yields that the objective function (gross of costs) is given by

$$\int \left(\Lambda(r)h(r) + \alpha J(r)\right) x(r)dF(r),$$

where $\Lambda(r) = \mathbb{E}[\lambda(\tilde{r}) | \, \tilde{r} \geq r]$ is the average welfare weight on agents with WTP above r, $h(r) = (1 - F(r))/f(r)$ is the inverse hazard rate, and $J(r) = r - h(r)$ is the virtual surplus function. Let $V(r) \equiv \Lambda(r)h(r) + \alpha J(r)$ and note that this value function collapses to well-known objects in special cases: willingness to pay itself if the designer has no redistributive preferences ($\lambda \equiv 1 = \alpha$); virtual surplus if the designer does not care about agent welfare ($\lambda \equiv 0$); and the inverse hazard rate in the case when the designer does not have redistributive preferences and screens with a costly-screening instrument rather than with money ($\lambda \equiv 1$ and $\alpha = 0$).

(5) Whenever the objective function $V(r)$ fails to be non-decreasing, the constraint requiring the allocation rule $x(r)$ to be non-decreasing might bind at the optimal solution. I will therefore apply a variant of Myerson's ironing procedure. Let us first rewrite the problem in the quantile space:

$$\int_0^{\bar{r}} V(r)x(r)dF(r) = \int_0^1 V(F^{-1}(q))x(F^{-1}(q))dq.$$

We can optimize over non-decreasing (quantile) allocation rules $\tilde{x} : [0, 1] \to [0, 1]$ given by $\tilde{x}(q) = x(F^{-1}(q))$. Since each candidate $\tilde{x}(r)$ is a weakly-increasing function from [0, 1] to [0, 1], we can formally treat it as a cumulative distribution function (it is without loss of generality to assume that $x(0) = 0$ and $x(1) = 1$). Moreover, if Q is the total quantity of objects, then $1 - Q$ is the mean of the distribution $\tilde{x}$:

$$1 - Q = 1 - \int_0^{\bar{r}} x(r)dF(r) = 1 - \int_0^1 \tilde{x}(t)dt = \int_0^1 t \, d\tilde{x}(t).$$

An analogous transformation of the objective function (using integration by parts) yields an equivalent representation of the problem, for a fixed Q:

$$\max_{\text{cdf } \tilde{x}} \quad \int_0^1 \left(\int_t^1 V(F^{-1}(q))dq \right) d\tilde{x}(t) \tag{1}$$

$$\text{subject to} \quad 1 - Q = \int_0^1 t \, d\tilde{x}(t). \tag{2}$$

Thus, for a fixed quantity Q, the problem is to choose a distribution $\tilde{x}$ to maximize the objective function subject to preserving a given mean of the distribution.[5]

(6) Let $\mathcal{V}(t) = \int_t^1 V(F^{-1}(q))dq$, and let co$\mathcal{V}$ denote the concave closure of $\mathcal{V}$ (the smallest concave function that lies above $\mathcal{V}$). The value of problem (1)-(2) is co$\mathcal{V}(1-Q)$. Moreover, there always exists an optimal distribution $\tilde{x}$ that has at most two points in its support, which implies that the optimal allocation rule $x(r)$ takes the form

$$x(r) = \begin{cases} 0 & r < r_0, \\ x_0 & r \in [r_0, r_1), \\ 1 & r \geq r_1, \end{cases}$$

for some $0 \leq r_0 \leq r_1 \leq 1$, and $x_0 \in (0, 1)$. In particular, a posted-price mechanism is optimal if and only if the optimal $\tilde{x}$ is a degenerate distribution, which is the case if and only if co$\mathcal{V}(1-Q) = \mathcal{V}(1-Q)$. Rationing appears in the optimal mechanism if and only if co$\mathcal{V}(1-Q) > \mathcal{V}(1-Q)$, that is, whenever ironing of the objective function is required given the chosen quantity Q.

(7) To solve the final problem, all we have to do is maximize the total objective co$\mathcal{V}(1-Q) - C(Q)$ over possible quantity levels $Q \in [0, 1]$.

3. ECONOMIC IMPLICATIONS

For the remainder of the note, I will explore the economic consequences of the derivation. To focus on the role of welfare weights, I assume that $F(r)$ is the uniform distribution on $[0, 1]$. Let us also assume that $\alpha = \mathbb{E}[\lambda] = 1$, meaning that all revenue is redistributed back to agents as a lump-sum transfer. Under these assumptions, we obtain

$$V(r) = \int_r^1 \lambda(\tau)d\tau + 2r - 1. \tag{3}$$

To gain some intuition for the shape of the function $\lambda(r)$, recall that $\lambda(r)$ is the conditional expectation of the unobserved welfare weight λ conditional on the observed WTP r. WTP is shaped by both the taste for the good and the "opportunity cost of money"—it is natural to expect that wealthier agents will be able to

[5] This problem is mathematically equivalent to a Bayesian persuasion problem with a binary state: Choose a distribution of posterior beliefs $\tilde{x}$ to maximize an objective function (1) that depends on the posterior belief, subject to a Bayes-plausibility constraint (2). By Aumann and Maschler [1995] and Kamenica and Gentzkow [2011], the value of such a problem is equal to the concave closure of the objective function at the prior (here, $1 - Q$), and the optimal distribution has at most two points in its support.

pay more for things they need. Thus, as long as the designer has a preference for redistribution towards agents with a higher opportunity cost of money, we might expect a negative correlation of λ and r, resulting in a decreasing function $\lambda(r)$. The strength of the correlation depends on whether heterogeneity in WTP is driven more by heterogeneity in tastes or by heterogeneity in ability to pay (the negative correlation is perhaps strongest for expensive goods and services satisfying basic needs, such as health care or housing).

When is it optimal to use a price mechanism?

When inequality—as measured by the dispersion in welfare weights—is not too high, we should use a price mechanism:

PROPOSITION 3.1. *If $\lambda(r) \leq 2$ for all r, then it is optimal to use a price mechanism.*

Proposition 3.1 follows from noticing that when the expected welfare weights never exceed the average welfare weight (normalized to 1) by more than a factor of 2,[6] then the objective function $V(r)$ is non-decreasing—ironing is not required, so a posted-price mechanism is optimal. Nevertheless, the price in the mechanism will typically be "distorted" relative to the efficient mechanism. To see that, suppose that $C(Q) = cQ$, so that the good can be produced at constant marginal cost $c > 0$. If all $\lambda \equiv 1$, it is optimal to set the price equal to marginal cost. This is not the case when the designer has redistributive preferences:

PROPOSITION 3.2. *Under the assumptions of Proposition 3.1, if $C(Q) = cQ$ for $c > 0$, the optimal price $p^\star$ satisfies*

$$p^\star = c + (1 - \Lambda(p^\star))(1 - p^\star).$$

It follows that the designer optimally imposes a tax (the price $p^\star$ is higher than the marginal cost c) if and only if the average welfare weight $\Lambda(p^\star)$ on agents who buy the good is lower than the average welfare weight (normalized to 1). This allows the designer to increase the revenue and hence the lump-sum transfer. This case arises when $\lambda(r)$ is decreasing in r. Otherwise, when $\Lambda(p^\star) > 1$, the designer uses a subsidy ($p^\star < c$): This allows the designer to transfer utility from the average agent to an agent who buys the good. Only in the knife-edge case $\Lambda(p^\star) = 1$ is the efficient mechanism (pricing at marginal cost) optimal.

When is it optimal to ration?

If inequality in the market is large, it may be optimal to complement a posted price with a lower price but subject to rationing ($x_0 < 1$):

PROPOSITION 3.3. *Suppose that $\lambda(r) > 2$ for some r. Then, there exists a cost function $C(Q)$ such that the optimal mechanism features rationing.*

The assumption that $\lambda(r) > 2$ for some r implies that the objective function $V(r)$ is decreasing in some region, and hence ironing may be needed (ironing is used if the

[6]The reader may reasonably ask: "Why 2?" There is a good reason which I explain in Dworczak [2023] (see also the Online Appendix of Dworczak ⓡ Kominers ⓡ Akbarpour [2021]).

cost function makes the optimal quantity Q fall within the ironing region). In our context, ironing means that the optimal mechanism offers a rationed option with a lower per-unit price. The role of rationing is to shift more gains towards agents with intermediate values of r: Agents with the highest r self-select into the non-rationed high-price option, which allows the designer to subsidize agents choosing the rationed option by reducing its price. However, because rationing is inefficient, the welfare weights on agents who benefit from it must be sufficiently large.

At least among economists, Myerson's ironing is mostly perceived as an "annoying detail" that can be formally ignored if we make appropriate regularity conditions. With redistributive preferences, however, ironing is no longer a technicality. It underlies the key economic result: Inefficient rationing may be part of the optimal mechanism.

The assumption that the expected welfare weight $\lambda(r)$ exceeds 2 for some willingness to pay r has an important economic interpretation. It states that the designer has strong redistributive preferences to begin with (a necessary condition, obviously, is that there are at least some agents with $\lambda > 2$) *and* that market behavior reveals information about that underlying inequality (i.e., r is sufficiently informative about λ). Intuitively, even if the underlying inequality is substantial, rationing is suboptimal if willingness to pay is driven more by tastes than ability to pay.

4. DIRECTIONS FOR FUTURE RESEARCH

Relative to the simple framework presented here, recent work incorporated several realistic features such as *(i)* two-sided markets with both buyers and sellers (Dworczak ⓡ Kominers ⓡ Akbarpour [2021]), *(ii)* heterogeneous object qualities and partially informative signals available to the designer (Akbarpour ⓡ Dworczak ⓡ Kominers [2024]), *(iii)* allocative externalities (Kang [2024], Pai and Strack [2022], Akbarpour ⓡ Budish ⓡ Dworczak ⓡ Kominers [2024]), *(iv)* co-existence of regulated and private markets (Kang [2023]), *(v)* endogeneity of buyer and seller roles (Kang and Zheng [2023]), *(vi)* models with finitely many agents (Reuter and Groh [2020]), *(vii)* the role of costly screening devices (Dworczak [2023], Yang et al. [2024]), and *(viii)* the use of information design and market segmentation (Barreto et al. [2022], Arya and Malhotra [2022]). However, all of these papers adopt a standard Bayesian optimization framework. I conclude this letter by identifying a few research directions along which the expertise of the Econ CS community would be particularly valuable.

(1) **Price of anarchy**. A classical question asked in the computer science literature concerns the performance of a system relative to its optimal design, evaluated against the worst-case scenario (see, e.g., Koutsoupias and Papadimitriou [2009]; Roughgarden and Tardos [2000]; Andelman et al. [2009]). One can ask a version of that question when performance (e.g., of a market) is measured in an inequality-sensitive way, for example, using a redistributive objective function. Just how bad can an unregulated market be in achieving an equitable outcome? Are there circumstances under which we should be particularly concerned about inequality in markets?

(2) **Performance of simple mechanisms**. A closely related question is that of the performance of simple (non-optimal) mechanisms (in the context of allocating goods, see, e.g., Hartline and Roughgarden [2009]; Chawla et al. [2010]; Feldman et al. [2015]; Babaioff et al. [2020]; Feldman et al. [2020]). In the above framework, the optimal mechanism is relatively simple, in that it involves at most two prices (one at which agents can buy for sure, and one with rationing). Even there, however, one can ask: what fraction of the optimal welfare can the designer achieve by relying only on posted-price mechanisms? In other words, is the use of rationing essential to achieving a good welfare guarantee? In more complex settings, the optimal mechanism can become very complicated, and finding simpler mechanisms that achieve good welfare guarantees (in terms of the redistributive objective function) is of first-order importance for practical purposes.

(3) **Robustness**. The computer scientists' insistence on robustness has undoubtedly (at least partially) inspired the distributionally-robust mechanism design literature that is flourishing in economics (see, e.g., Frankel [2014], Carroll [2017], Brooks and Du [2021], Suzdaltsev [2022], He and Li [2022], among many others). Whether rationing is used in the above framework depends on the joint distribution of willingness to pay and welfare weights: r and λ must be sufficiently strongly correlated. But what if the designer does not know their exact joint distribution? What if she considers a robust approach of maximizing her (redistributive) objective against the worst-case joint distribution consistent with some known moments?

(4) **Multidimensional screening**. One of the most spectacular advancements brought together by the interaction of economics and computer science is in the area of multi-dimensional screening (see, e.g., Daskalakis et al. [2013, 2017]; Haghpanah and Hartline [2020]). The framework I considered here features two-dimensional types but the underlying screening problem ends up being one-dimensional. Many natural extensions of the framework, however, feature multi-dimensional information that cannot be reduced to a one-dimensional statistic. For example, the designer could be controlling multiple markets, or have access to multiple redistributive instruments. The question of optimal design with an inequality-sensitive objective in such environments is wide open.

Most likely, the best future research directions are not on the above list. The beauty of good ideas is that they cannot be anticipated. And we need more good ideas to provide smart policy guidance for dealing with inequality (policymakers will try to deal with it, with or without our input!).

REFERENCES

Akbarpour, M. ⓡ E. Budish ⓡ P. Dworczak ⓡ S. D. Kominers. 2024. An economic framework for vaccine prioritization. *Quarterly Journal of Economics 139,* 1, 359–417.

Akbarpour, M. ⓡ P. Dworczak ⓡ S. D. Kominers. 2024. Redistributive allocation mechanisms. *Journal of Political Economy (forthcoming).*

Andelman, N., Feldman, M., and Mansour, Y. 2009. Strong price of anarchy. *Games and Economic Behavior 65,* 2, 289–317.

Arya, Y. and Malhotra, R. 2022. Rethinking distribution: Introducing market segmentation as a policy instrument. Working Paper.

AUMANN, R. J. AND MASCHLER, M. 1995. *Repeated Games with Incomplete Information.* MIT press.

BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2020. A simple and approximately optimal mechanism for an additive buyer. *J. ACM 67,* 4 (jun).

BARRETO, D. M. A., GHERSENGORIN, A., AND AUGIAS, V. 2022. Price discrimination with redistributive concerns. Working Paper.

BROOKS, B. AND DU, S. 2021. Optimal auction design with common values: An informationally robust approach. *Econometrica 89,* 3, 1313–1360.

CARROLL, G. 2017. Robustness and separation in multidimensional screening. *Econometrica 85,* 2, 453–488.

CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing.* STOC '10. Association for Computing Machinery, New York, NY, USA, 311–320.

CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice 11,* 1, 17–33.

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

DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2013. Mechanism design via optimal transport. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. Association for Computing Machinery, New York, NY, USA, 269–286.

DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2017. Strong duality for a multiple-good monopolist. *Econometrica 85,* 3, 735–767.

DWORCZAK, P. 2023. Equity-efficiency trade-off in quasi-linear environments. Working Paper.

DWORCZAK, P. ⓡ S. D. KOMINERS ⓡ M. AKBARPOUR. 2021. Redistribution through markets. *Econometrica 89,* 4, 1665–1698.

FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2020. Simultaneous auctions without complements are (almost) efficient. *Games and Economic Behavior 123*, 327–341.

FELDMAN, M., GRAVIN, N., AND LUCIER, B. 2015. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms.* SODA '15. Society for Industrial and Applied Mathematics, USA, 123–135.

FRANKEL, A. 2014. Aligned delegation. *American Economic Review 104,* 1 (January), 66–83.

GROVES, T. 1973. Incentives in teams. *Econometrica 41,* 4, 617–631.

HAGHPANAH, N. AND HARTLINE, J. 2020. When is pure bundling optimal? *The Review of Economic Studies 88,* 3 (08), 1127–1156.

HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce.* EC '09. Association for Computing Machinery, New York, NY, USA, 225–234.

HE, W. AND LI, J. 2022. Correlation-robust auction design. *Journal of Economic Theory 200*, 105403.

KAMENICA, E. AND GENTZKOW, M. 2011. Bayesian persuasion. *American Economic Review 101,* 6, 2590–2615.

KANG, M. AND ZHENG, C. Z. 2023. Optimal design for redistributions among endogenous buyers and sellers. *Economic Theory*, 1141–1180.

KANG, Z. Y. 2023. The public option and optimal redistribution. Working Paper.

KANG, Z. Y. 2024. Optimal indirect regulation of externalities. Working Paper.

KOUTSOUPIAS, E. AND PAPADIMITRIOU, C. 2009. Worst-case equilibria. *Computer Science Review 3,* 2, 65–69.

MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1, 58–78.

OSTRIZEK, F. AND SARTORI, E. 2023. Screening while controlling an externality. *Games and Economic Behavior 139*, 26–55.

PAI, M. AND STRACK, P. 2022. Taxing externalities without hurting the poor. Working Paper.

Reuter, M. and Groh, C.-C. 2020. Mechanism design for unequal societies. *Working Paper*.

Roughgarden, T. and Tardos, É. 2000. How bad is selfish routing? j acm. *Foundations of Computer Science, 1975., 16th Annual Symposium on 49*, 93–102.

Spence, M. 1977. Nonlinear prices and welfare. *Journal of Public Economics 8,* 1, 1–18.

Suzdaltsev, A. 2022. Distributionally robust pricing in independent private value auctions. *Journal of Economic Theory 206*, 105555.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance 16,* 1, 8–37.

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,* 2, 517–524.

Yang, F., Dworczak, P., and Akbarpour, M. 2024. Comparison of screening devices. Working Paper.