Optimal Membership Design[*]

Piotr Dworczak $^{\text{Ⓡ}}$ Marco Reuter $^{\text{Ⓡ}}$ Scott Duke Kominers $^{\text{Ⓡ}}$ Changhwa Lee[†]

January 21, 2025

Abstract

Membership design involves allocating an economic good whose value to any individual depends on who else receives it. We introduce a framework for optimal membership design by combining an otherwise standard mechanism-design model with allocative externalities that depend flexibly on agents' observable and unobservable characteristics. Our main technical result demonstrates that the optimal mechanism offers distinct membership tiers—differing in prices and level of access—and that the number of membership tiers is increasing in the complexity of the externalities. This insight may help explain a number of mechanisms used in practice to sell membership goods, including musical artists charging below-market-clearing prices for concert tickets, certain admission procedures used by colleges concerned about the diversity of the student body, heterogeneous pricing tiers for access to digital communities, and the use of vesting and free allocation in the distribution of network tokens.

1 Introduction

Many allocation problems can be recast as designing *membership*. The defining feature of membership as an economic good is that its value for a given agent depends on who else is a member. For example, a ticket gives its owner the right to attend a concert, but the value of attending may depend on the concert atmosphere—which is influenced by the composition of the audience. In the context of blockchain-based networks, digital tokens are often required for access and utilization; but it is often the interactions among token owners that determine the value of being part of the network in the first place. Finally, designing rules for admission to a school or an academic program such as executive education requires understanding how prospective students' values for enrolling depend on peer effects and opportunities for networking.

In short, membership goods differ from conventional goods in that they feature potentially complex externalities between their owners. In this paper, we provide a simple framework for studying optimal design of membership that allows for a rich pattern of externalities that can depend on both observable and unobservable agent characteristics. We show how viewing a diverse set of allocation problems as membership design leads to a unifying explanation for a number of real-life selling practices such as lotteries and group-specific price discounts. Finally, our framework can be used to guide the design of a membership system that must account for the interactions between various members to correctly price access.

Our model combines a standard mechanism-design framework—featuring a large population of agents, one-dimensional private information, quasi-linear utilities, and labels capturing observable information—with allocative externalities. We introduce allocative externalities by assuming that the value of membership for an agent is obtained as a product of her type with a term determined by the distribution of other members' characteristics. The main simplifying assumption is that—conditional on observables—agents care about the *average* characteristics of other members. The designer is assumed to know the distribution of agent characteristics and chooses an individually-rational and incentive-compatible mechanism to maximize an objective such as revenue or welfare.

We show that optimal membership in our setting can be implemented by a mechanism with label-dependent tiered pricing. Dependence on labels means that different groups—i.e., sets of agents sharing the same observable characteristics—could face different price schedules, allowing incentives to join to vary with group identity and the associated externalities. Tiered pricing means that agents within a group are offered a menu of different levels of access at different prices. Level of access can simply be the probability of becoming a member but—depending on the application—it can also capture factors such as a higher

quality of service or additional membership privileges. The key idea behind tiered pricing in our framework is that it allows the designer to achieve a more valuable composition of members by relying on self-selection to screen for unobserved characteristics associated with membership externalities.

Our main technical result connects the number of distinct membership tiers offered by the optimal mechanism to the complexity of the structure of the externalities. We measure the complexity of externalities exerted by a group as the dimension of the image of an operator relating an allocation rule for that group to the value generated for each of the other groups. In particular, a single price is optimal for a group that does not exert any allocative externalities. When a group exerts a uniform externality (i.e., its magnitude is the same for all agents), two prices may be needed for that group: a high price with the highest level of access, and a discounted price with a lower level of access. Finally, when agents' externalities within a given group vary arbitrarily with their types and have heterogeneous effects on other groups, the number of membership tiers required for optimality may be as high as the number of groups plus one.

The main contribution of our paper is to argue that membership design provides a unifying lens through which a diverse set of allocation problems can be analyzed and understood. First, we explore the problem faced by an artist or sports club selling tickets to their events. Without concerns about the composition of the audience, it is optimal to charge a single monopoly price. However, if the atmosphere at the event depends on the presence of the most dedicated enthusiasts (e.g., screaming teenagers at a Taylor Swift concert, or diehard fans cheering for their team at a football game), then our framework indicates that it may become optimal to sell at a below-market-clearing price with rationing. We show that using lotteries is optimal when the fans that contribute the most to the atmosphere cannot be easily identified based on observables and have relatively low ability to pay. This may help explain why pop stars like Taylor Swift and sports associations like UEFA sometimes sell tickets to their events at below-market-clearing prices via lotteries—appearing to leave money on the table. Our explanation suggests that such sales mechanisms may actually be revenue-maximizing, even absent any dynamic or reputational considerations.

Second, we consider the problem faced by a college designing its admissions procedure in the face of diversity considerations. We show that affirmative action (conditioning on the applicants' background characteristics) allows the college to implement the optimal student body composition via a system of merit and diversity scholarships. However, if affirmative action is not allowed, it may become optimal to bundle admission and funding decisions by forcing students to choose between non-funded traditional admission versus admission that is partially- or fully-funded but rationed. In this way, the college relies on correlations

3

between willingness to pay and background characteristics to implement the desired student body composition through self-selection.

Third, we analyze the primary sales process of membership in a community, such as a social organization, professional association, nonprofit society, or fan club. Sales mechanisms of the type we discuss here have been used historically, but they have become especially prevalent recently in the context of digital communities whose membership is managed through ownership of non-fungible tokens (NFTs). Branded NFTs are often intentionally sold at low prices or given away for free, while later reaching prices in the hundreds or even thousands of dollars. We provide an explanation for this phenomenon by recognizing the role of externalities among NFT owners in determining the value of community membership. One challenge in designing NFT communities is that prospective members have a particularly broad set of labels drawing upon range of on-chain data, with rich implications for the shape of their externalities. We thus show how the pattern of across-group externalities determines whether particular groups behave as substitutes or complements when determining optimal participation. Additionally, our analysis applies to traditional clubs and professional associations that often have simple tiered pricing schemes, and lets us recover some of the classical insights about pricing access to traditional buyer–seller platforms.

Finally, we consider a common case in the context of blockchain infrastructure and protocol applications: the design of a network where a fungible crypto token (such as Ethereum) is required for utilization, typically via payment of token-denominated transaction fees. We examine the tension between using the token to secure participation by developers who improve the network's functionalities and limiting the congestion externality generated by the network's users. We show that when the tokens can be eventually traded in a secondary market, getting the network off the ground may require allocating tokens at prices significantly below the expected resale price—sometimes, taking the form of what are called "free mints"—but with vesting.

Beyond the ability to explain certain real-life phenomena, our model provides a simple but flexible conceptual framework for thinking about optimal design and policy in contexts in which externalities due to membership structure play an important role.

1.1 Related literature

Our paper is most closely related to a sizable literature within mechanism design studying allocative externalities. Broadly, two classes of models have been considered. The first class of models, introduced by Jehiel et al. (1996, 1999), studies allocation to a finite number of agents exerting fully flexible individual-specific externalities that are typically assumed to

be unobserved.[1] Such a model is challenging to analyze and often leads to negative results (see, e.g., Jehiel and Moldovanu (2001)). A more recent class of models—studied by Kang (2020), Pai and Strack (2022), Ostrizek and Sartori (2023), and Akbarpour ⓣ al. (2024)— focuses instead on a large population of agents contributing (in a potentially heterogeneous way) to a single aggregate externality (which often enters the designer's objective function additively). Our framework sits in between these two approaches: Relative to the first class of models, we impose the simplifying assumption that externalities depend only on compositional effects (in a model with a continuum of agents). Relative to the second class of models, we allow for a much richer pattern of externalities that depend on group identity and interact with individuals' willingness to pay. The resulting set of applications also differs: the first class of models is ideal for studying a small number of strategically interacting agents (e.g., the problem of selling nuclear weapons discussed by Jehiel et al. (1996)); the second class of models is ideal for studying problems such as regulating the sale of goods that generate externalities affecting *all* agents (e.g., goods whose consumption contributes to climate change, as the settings of Kang (2020), and Pai and Strack (2022)); and our framework is geared towards settings in which agents care about the characteristics of other agents receiving the allocation, such as in the context of selling tickets to events, determining access to digital communities, building multi-sided platforms, or designing policies with compositional effects such as school admission.[2]

Allocative externalities similar to the ones we study appear in the classical literature on two-sided buyer-seller platforms (see Jullien et al. (2021) for a comprehensive survey) and on sales with network effects (see, e.g., Candogan et al. (2012), Fainmesser and Galeotti (2015), and Jadbabaiea and Kakhbod (2019)); most papers in these literatures do not adopt a mechanism-design approach, focusing instead on questions such as equilibrium market dynamics.[3] To our knowledge, our work is the first to explicitly characterize the role of rationing (or lotteries) in the optimal mechanism design for network membership. More broadly, our paper is related to the literature trying to understand optimal groupings of agents in the presence of social status externalities (Moldovanu et al. (2007); Rayo (2013); Board (2009)).

---

[1] See also the work of Segal (1999) who introduced a general model of contracting with externalities without asymmetric information, and Bernstein and Winter (2012) for an application of this framework to a problem similar to membership design.

[2] Meisner and Pillath (2024) study the sale of digital content with network effects, but restrict the externalities to be a function of the number of consumers.

[3] Some notable exceptions include the work of Damiano and Li (2007) and Gomes and Pavan (2021); however, these papers focus on the optimal matching between the two sides of the market, rather than pricing policy. Jadbabaiea and Kakhbod (2019) examined how the presence of network externalities distorts the optimal contract offered by a profit-maximizing seller.

We comment further on connections to various applied literatures (including the work on ticket pricing, fungible and non-fungible tokens, and matching with diversity considerations) in the context of our discussions of applications in Section 4.

2   Model

A *designer* allocates access to a platform or community to a unit mass of agents (potential members). Each agent is characterized by a *private type* θ and a publicly observable *label i*. An agent's *level of access* to the platform is described by a variable $x \in [0,1]$, and the agent can be charged a transfer $t \in \mathbb{R}$. Our baseline interpretation of x is that it represents the probability of becoming a member but it could also reflect the quality of access to the platform (e.g., which features are available to the agent). To focus on the novel aspects of our framework, we assume that the designer does not incur marginal costs for allowing agents' access and does not face an explicit capacity constraint—an assumption that is most appropriate for the design of digital communities.[4]

Conditional on label i, agents' types have a distribution with a cumulative distribution function F_i and a continuous density f_i fully supported on an interval $[\underline{\theta}_i, \bar{\theta}_i]$. There are finitely many labels $i \in I$, and each group i has mass $\mu_i > 0$, with $\sum_{i \in I} \mu_i = 1$.

The utility an agent receives from access depends on that agent's type and the *composition* of other agents using the platform. Specifically, an agent with label i and type θ has willingness to pay for membership equal to θv_i, where v_i is common to all agents in group i. The simplifying assumption here is that heterogeneity in willingness to pay can be decomposed into a vertical component captured by the type θ and a horizontal component that only depends on the label i. The value v_i is determined by the expected externalities exerted by all members of the community. Conceptually, each member has some unobserved "true" externality on agents in group i. Because an externality an agent exerts on others does not affect that agent's payoff, it cannot be elicited by any incentive-compatible mechanism.[5] However, the designer can form a "forecast" of the true externality based on characteristics that are observable (the label i) or elicitable by a mechanism with transfers (the type θ). Thus, we let $e_{j \to i}(\theta)$ denote the *expected externality* exerted by an agent with label j and type θ onto any agent in group i. The function $e_{j \to i}(\theta)$ is assumed to be continuous in θ for all $i, j \in I$. Finally, if $x_j(\theta)$ denotes the probability of access for an agent (j, θ), we assume

---

[4]These features could be added to the model without changing our results in any substantial way.

[5]Jehiel et al. (1996) provide a formal statement and proof of this observation.

that

$$v_i \equiv \nu_i \left(\sum_{j \in I} \int e_{j \to i}(\theta) x_j(\theta) \mu_j dF_j(\theta) \right),$$

where, for all $i \in I$, $\nu_i : \mathbb{R} \to \mathbb{R}$ is upper semi-continuous, non-negative, and non-decreasing.[6] We put no sign restrictions on $e_{j \to i}(\theta)$, i.e., agents can exert both positive and negative externalities on one another, depending on group identity and type. We assume utilities are quasi-linear in the monetary transfer and linear in probability of access: Allocation (x, t) gives an (i, θ) agent utility

$$\theta v_i x - t.$$

The designer maximizes an objective function—such as revenue or agents' welfare—over all direct-revelation mechanisms, subject to individual-rationality and incentive-compatibility constraints. Formally, for a non-negative weight $\alpha \geq 0$ and continuous functions $\tilde{W}_i : [\underline{\theta}_i, \bar{\theta}_i] \to \mathbb{R}$, the designer solves the following problem:

$$\max_{x_1,...,x_{|I|}, t_1,...,t_{|I|}} \left\{ \sum_{i \in I} \int \left(v_i \tilde{W}_i(\theta) x_i(\theta) + \alpha t_i(\theta) \right) \mu_i dF_i(\theta) \right\}$$

$$\text{s.t. } v_i = \nu_i \left(\sum_{j \in I} \int e_{j \to i}(\theta) x_j(\theta) \mu_j dF_j(\theta) \right),$$

$$\theta v_i x_i(\theta) - t_i(\theta) \geq 0, \qquad\qquad \forall i \in I, \forall \theta \in [\underline{\theta}_i, \bar{\theta}_i],$$

$$\theta v_i x_i(\theta) - t_i(\theta) \geq \theta v_i x_i(\theta') - t_i(\theta'), \qquad \forall i \in I, \forall \theta, \theta' \in [\underline{\theta}_i, \bar{\theta}_i].$$

Following the standard approach in the mechanism design literature (Myerson (1981)), we can characterize feasible mechanisms through the envelope formula for transfers and monotonicity conditions on the allocation rules, giving rise to the following equivalent optimization problem for the designer:[7]

$$\max_{x_1,...,x_{|I|}} \left\{ \sum_{i \in I} \int \nu_i \left(\sum_{j \in I} \int e_{j \to i}(\theta) x_j(\theta) \mu_j dF_j(\theta) \right) W_i(\theta) x_i(\theta) \mu_i dF_i(\theta) \right\} \qquad (1)$$

$$\text{s.t. } x_i(\theta) \text{ is non-decreasing in } \theta, \quad \forall i \in I,$$

---

[6] Our results extend with minimal modifications to the case when values can become negative, as long as the designer can pay the agents for joining the platform; however, covering this more general case is notationally cumbersome, so we restrict attention to non-negative values.

[7] Strictly speaking, incentive compatibility does not impose any restrictions on the allocation rule x_i in the degenerate case $v_i = 0$; we nevertheless assume that allocation rules are always non-decreasing so that the problem is economically and mathematically well-behaved.

where $W_i(\theta) \equiv \tilde{W}_i(\theta) + \alpha J_i(\theta)$, and $J_i(\theta) := \theta - (1 - F_i(\theta))/f_i(\theta)$ is the virtual surplus function. For example, setting $\alpha = 1$ and $\tilde{W}_i \equiv 0$ corresponds to maximizing revenue, while setting $\alpha = 0$ and $\tilde{W}_i(\theta) \equiv h_i(\theta) := (1 - F_i(\theta))/f_i(\theta)$ (the inverse hazard rate) corresponds to maximizing agents' total surplus.

The presence of externalities significantly complicates the designer's problem. Assigning access $x_i(\theta)$ to an agent allows the designer to realize the payoff $v_i W_i(\theta)$ from allocating to that agent. At the same time, however, it affects the composition of agents on the platform and thus potentially changes the values that all other agents have for joining. Mathematically, the problem is non-linear in the allocation rule. We purposefully constructed our framework to have a very simple optimal mechanisms when there are no externalities. Indeed, if $\nu_i(\cdot)$ is constant for all $i \in I$, it is well-known that, regardless of the properties of the function $W_i(\theta)$, the designer can achieve optimality by simply posting a price p_i for allocation $x = 1$ to each group i. Thus, tiered membership can only arise in our setting due to allocative externalities.[8]

Our restriction to direct-revelation mechanisms implicitly assumes that the designer can select her preferred equilibrium in case of multiplicity (otherwise the revelation principle does not apply). The issue of robustness to equilibrium selection received considerable attention in the literature on network externalities and is beyond the scope of this paper.[9]

3   Results

3.1   Measuring the complexity of externalities

We define an operator E^i from the unit interval $[0, 1]$ to $\mathbb{R}^{|I|}$ by

$$E^i_j(q) := \int_{\theta_q}^{\bar{\theta}_i} e_{i \to j}(\theta) dF_i(\theta), \quad j = 1, \dots, |I|,$$

where θ_q is the q-quantile of the distribution F_i. That is, for a given $q \in [0, 1]$, the operator E^i computes the expected externality exerted on each group j by the fraction $(1 - q)$ of highest-type agents from group i. As q varies, the operator traces out how the externality contributions from group i change as more or fewer agents from group i participate. We quantify the complexity of these externality contributions by using the dimension of the

---

[8]Of course, tiered pricing also arises if agent's values or seller's costs are non-linear in the allocation x, as in the celebrated product-quality model of Mussa and Rosen (1978). We turned off these forces in order to focus on the role of externalities.

[9]See Rohlfs (1974) for a classical reference, Segal (2003) and Winter (2004) for the first mechanism-design formulations, and Halac et al. (2020) and Halac et al. (2024) for recent developments.

image of the operator E^i; formally,

$$\mathrm{cplx}(E^i) := \dim(\mathrm{Im}(E^i)).$$

That is, the complexity of the operator E^i is the dimension of its image as a set in the Euclidean space $\mathbb{R}^{|I|}$.[10]

For intuition, suppose that agents in group i do not exert any externalities on agents in any group j: i.e., $e_{i \to j}(\theta) = 0$ for all j and θ. Then, E^i has complexity 0. Alternatively if externalities do not depend on an agent's type—$e_{i \to j}(\theta) = e_j$, for all j and θ—then E^i has complexity 1. More generally, the complexity of externalities exerted by members drawn from group i is *low* when the externality contributions $e_{i \to j}(\theta)$ exhibit a structurally similar pattern of dependence on θ across different j's. For example, the complexity of E^i is at most 2 if the dependence on θ is linear in each group j. On the other hand, the complexity of externalities exerted by members from group i is *high* when the externality contributions $e_{i \to j}(\theta)$ depend on θ in a different way for each group j. Finally, note that the complexity of E^i is bounded above by the number of groups $|I|$. The following simple lemma follows directly from the definition of complexity, and formalizes the preceding intuition.

Lemma 1. *Suppose that, fixing i, we can write*

$$e_{i \to j}(\theta) = \sum_{k=1}^{K} \beta_{i,j,k}\, g_k(\theta), \quad \forall j = 1, \dots, |I|,$$

for some $\beta_{i,j,k} \in \mathbb{R}$ and $g_k : [\underline{\theta}_i, \bar{\theta}_i] \to \mathbb{R}$, $k = 1, \dots, K$. Then, $\mathrm{cplx}(E^i) \le K$.

Returning to our interpretation of $e_{i \to j}(\theta)$ as the *expected* externality that an agent with type θ in group i exerts on agents in group j, we can think of the functions g_k as principal components representing the statistical dependence of these externalities on agents' types. The mechanism designer, having access to high-dimensional data on agents' types and externalities (possibly collected through A/B testing or other empirical methods), reduces the dimensionality of the data using principal component analysis to estimate the relationship between types and externalities. The more components are needed to account for all the externalities exerted by group i, the higher the complexity of the operator E^i. For example, performing a linear regression with two principal components on the externalities within a group would result in an estimated operator E^i having complexity 2. Fitting a linear regression with k principal components would yield an operator E^i with complexity k, while a fully non-parametric regression would most likely result in a full-dimensional E^i.

---

[10]The *dimension* of a set in a Euclidean space is the number of basis vectors generating the smallest affine subspace containing that set.

3.2 Tiered pricing mechanisms

We will show that the complexity of E^i is closely connected to *how many prices* the optimal mechanism may be required to use. To that end, we define a class of mechanisms that we refer to as k-tiered pricing mechanisms:

Definition 1. An allocation rule $x_i(\theta)$ for group i is a *k-tiered pricing mechanism* if $x_i(\theta)$ takes on at most k non-zero values on $\text{supp}(F_i)$.

A k-tiered pricing mechanism can be implemented by offering agents in a given group at most k different prices for k different levels of access. For example, a 1-tiered pricing mechanism could be a posted-price mechanism, which we define as offering a single price for full access ($x = 1$). But a 1-tiered pricing mechanism could also allow agents to access the platform for free (or at a positive price) with some interior probability. Similarly, a 2-tiered pricing mechanism could allow agents to choose between a low price that offers a low level of access and a higher price that offers full access.

Our main technical result establishes that the optimal mechanism is a k-tiered pricing mechanism for every group, where k can be related to the complexity of the externalities.

Theorem 1. *It is optimal to use a k^i-tiered pricing mechanism for each group i, where $k^i = \text{cplx}(E^i) + 1$, and the top tier provides full access ($x = 1$). Moreover, the optimal mechanism uses at most $|I| + K$ prices in total, where $K = \text{cplx}(\sum_{i \in I} \mu_i E^i) + 1$.*

The proof of Theorem 1 can be found in Appendix A. The intuition for the result is simple: While problem (1) faced by the designer is non-linear in the allocation rule x_i for group i, it would become linear if we fixed (and treated as exogenous) the values v_j for all groups. Moreover, fixing the value v_j while letting $x_i(\theta)$ vary corresponds to imposing a single linear constraint on $x_i(\theta)$ since, by assumption, the value v_j only depends on the average externality exerted by members from group i. As a result, the auxiliary problem in which we only optimize over the allocation rule $x_i(\theta)$ for a single group subject to inducing the same target values v_j for all groups j is a linear problem with $|I|$ linear constraints. Furthermore, the complexity of the operator E^i determines whether some of these constraints are co-linear and hence can be dropped. Once we represent the auxiliary problem as a linear program with $\text{cplx}(E^i)$ active linear constraints, an application of an extension of the Carathéodory's theorem yields that the optimal mechanism is a k^i-tiered pricing mechanism with $k^i = \text{cplx}(E^i)+1$ prices. Moreover, the top price can always be assumed to be associated with full access.[11]

---

[11]Except for the case $k^i = 1$, the formulation of Theorem 1 does not preclude the possibility that all agents choose an interior level of access; this is because the price of full access can be set to be so high that no one chooses it. In Section 4.1, we show that this possibility may indeed arise at optimum.

The second part of Theorem 1 imposes an upper bound on the total number of prices used by an optimal mechanism. The reason why the total bound is typically much lower than the sum of bounds for all groups i is that our proof strategy of looking at a single group at a time leads to "multiple counting" of the constraints in the auxiliary problem. When all allocation rules can be flexibly chosen, achieving a target value v_j for group j does not require fixing separately the average contributions from each group i; rather, it requires fixing the average contribution *across* all groups, which is a single linear constraint. Thus, we can find an optimal mechanism that charges at most $K = \text{cplx}(\sum_{i \in I} \mu_i E^i) + 1$ additional prices (on top of charging a price to each group).

Theorem 1 establishes our main general result which is that the number of prices offered by an optimal mechanism depends on the complexity of externalities. In the remainder of this section, we explore a few special cases of that insight.

First, we confirm an earlier statement that a posted price would be optimal in the absence of externalities.

Corollary 1. *It is optimal to sell to group i at a posted price if group i does not exert externalities, i.e., $e_{i \to j}(\theta) = 0$ for all $j \in I$ and θ.*

The absence of externalities exerted by agents from a certain group can be seen as an extreme case; next, we offer a more permissive condition under which the optimal mechanism still sells access at a posted price.

Proposition 1. *Suppose that the designer's objective function $W_i(\theta)$ is strictly increasing in θ, and that $\int_{\theta}^{\bar{\theta}_j} W_j(\tau) dF_j(\tau) \geq 0$, for all θ and $j \in I$.[12] Then, it is optimal to sell to group i at a posted price if the externality that group i exerts on any group j, i.e., $e_{i \to j}(\theta)$, is non-decreasing in θ.*

Under the conditions of Proposition 1, there is an alignment between the designer's objective (for a fixed profile of values v_j) and how the values v_j vary with the allocation—both considerations point towards optimality of a threshold allocation rule that grants access to all agents with a sufficiently high type. Proposition 1 thus implies that the additional prices predicted by Theorem 1 will only be needed when there is a conflict between the designer's objective for fixed v_j's and the desire to maximize v_j's over the composition of members.

---

[12]The last assumption always holds for revenue and surplus maximization when types are non-negative, and could be weakened to requiring that the inequality holds at the optimal mechanism. Even that weaker assumption could fail if the designer wanted to allocate membership to some group j to generate positive externalities for other groups at the cost of achieving a negative payoff from group j; in such a case, the designer's objective would become *decreasing* in the value v_j.

Finally, we state a simple corollary of Theorem 1 identifying two cases in which two prices—a high price with full access and a low price with lower level of access—are sufficient for achieving optimality.

Corollary 2. *It is optimal to use a 2-tiered pricing mechanism in group i (with the higher price providing full access) if either (1) there is only one group, or (2) group i exerts a uniform externality: $e_{i \to j}(\theta) = e_i(\theta)$ for all $j \in I$.*

4 Applications

4.1 Ticket sales

We apply our framework to the problem of an artist selling concert tickets or a sports team selling tickets to the game, aiming to maximize revenue. At such events, the atmosphere is often a crucial part of the experience. For example, screaming teenagers are somewhat essential for the atmosphere at a Taylor Swift concert, while the atmosphere at a football match depends on the most dedicated fans cheering for their teams.

We assume that there is a single group of agents (and hence drop the subscript i throughout), reflecting an assumption that the seller cannot easily identify the most enthusiastic fans based on observables.[13] It is then natural to assume that the lower bound of the support of types θ is zero. To model how the atmosphere at an event is generated, we use the externalities that agents exert on others. We assume that the value of externalities is captured by

$$v = v_0 + \int_{\underline{\theta}}^{\overline{\theta}} e(\theta) x(\theta) dF(\theta).$$

We do not explicitly model capacity constraints, such as limited seating availability, but assume that $e(\theta)$ contains a congestion externality that serves as a proxy for capacity considerations.[14] Even though our analysis is valid for a general function $e(\theta)$, we provide the following simple way of deriving $e(\theta)$ from more primitive elements. Suppose that agents are either high-value or low-value audience members (e.g., high-value members cheer for teams at a football game while low-value members sit quietly), and this is not observed. All members generate a congestion externality of -1. However, high-value members additionally generate

---

[13]Our results go through qualitatively if the seller can observe a noisy signal. Ticketmaster uses the "Verified Fan" program to sell some concert tickets but the "verification" part is mostly making sure that the potential buyer is a human being (not a bot).

[14]By specializing the congestion externality to take the form of a Lagrange multiplier on the capacity constraint, we can effectively replicate the case of fixed supply of seats.

a positive externality $\Delta > 0$. The high-value members' type is drawn from $F_H(\theta)$, while the low-value members' type is drawn from $F_L(\theta)$, with respective densities $f_H(\theta)$ and $f_L(\theta)$. A fraction μ of all agents are high-value. Then, by a simple application of Bayes' rule, we have

$$v = v_0 + \int_{\underline{\theta}}^{\bar{\theta}} \left(\underbrace{\frac{\mu f_H(\theta)}{\mu f_H(\theta) + (1-\mu) f_L(\theta)} \Delta - 1}_{e(\theta)} \right) x(\theta) dF(\theta).$$

Intuitively, the seller is trying to infer the unobserved value of a fan by conditioning on her type. As a result, the shape of the externality $e(\theta)$ depends on the joint distribution of the unobserved value and willingness to pay.

To simplify exposition, we assume that the virtual surplus function $J(\theta)$ is strictly increasing, and that $e(\theta) < 0$, that is, the negative congestion externality dominates the positive externality even for the most dedicated fans (corresponding to $\Delta < 1$ in the above example). The case when some members have a positive net contribution would make our key insights *easier* to obtain (we comment on this further at the end of the section).

First, relying on Proposition 1, we show the following result:[15]

Result 1. *A posted price is optimal when $e(\theta)$ is non-decreasing.*

When the externality $e(\theta)$ is increasing, selecting the most valuable audience members is aligned with the direct revenue motive of the seller. In our example, the externality is increasing if high-value members have a higher willingness to pay than low-value members (in the monotone likelihood ratio order). This scenario is empirically plausible for cases such as selling tickets to a networking event, where it is natural to expect that higher willingness-to-pay individuals are also more likely to hold professional positions that make them attractive to connect with.

In general, applying Corollary 2 allows us to restrict attention to 2-tiered pricing mechanisms. The seller is thus maximizing

$$\underbrace{\left(v_0 + x_0 \int_{\theta_0}^{\theta_1} e(\theta) dF(\theta) + \int_{\theta_1}^{\bar{\theta}} e(\theta) dF(\theta) \right)}_{V} \underbrace{\left(x_0 \int_{\theta_0}^{\theta_1} J(\theta) dF(\theta) + \int_{\theta_1}^{\bar{\theta}} J(\theta) dF(\theta) \right)}_{R} \quad (2)$$

over $x_0 \in [0, 1]$ and $\theta_0, \theta_1 \in [\underline{\theta}, \bar{\theta}]$, where V represents the base value of the event given the audience composition, and R represents the normalized revenue from selling tickets. Because we assumed that externalities $e(\theta)$ are negative, the term V is *decreasing* in the allocation, while the term R is *increasing* in the allocation (for types for whom virtual surplus is

---

[15]The calculations underlying the results presented in this section can be found in Appendix A.3.

positive)—leading to a trade-off. Our next observation is that a single tier is sufficient if that trade-off admits a certain monotone structure.

Result 2. *If $J(\theta)/(-e(\theta))$ is strictly increasing, then it is optimal to offer a 1-tiered pricing mechanism that allocates level of access $x_0^\star$ to all types above $\theta_0^\star$, where $(x_0^\star, \theta_0^\star)$ satisfy*

$$\frac{J(\theta_0^\star)}{-e(\theta_0^\star)} = \frac{R}{V} \quad and \quad x_0^\star = \min\left\{1, \frac{v_0}{2\int_{\theta_0^\star}^{\bar{\theta}}(-e(\theta))dF(\theta)}\right\},$$

where R and V are functions of $\theta_0^\star$ and $x_0^\star$ as defined in (2) *(with $\theta_1^\star = \bar{\theta}$).*

The ratio of virtual surplus to (the negative of) the externality, $J(\theta)/(-e(\theta))$, can be seen as measuring the trade-off between extracting revenue and preserving the value of the event: Ideally, the seller wants to target agents with high virtual surplus $J(\theta)$ and low (in magnitude) negative externality $e(\theta)$. Under the monotoniticy assumption made in Result 2, the trade-off is always resolved in favor of allocating to all agents with a type above a certain threshold $\theta_0^\star$, pinned down by the first-order condition. Unlike in Result 1, however, it may be optimal to ration access to the event, especially when the negative externalities are substantial.[16] Intuitively, a reduction in congestion can be achieved either through random rationing or through a price increase. These two tools, however, are not perfect substitutes: While an increase in price excludes the *marginal* attendee, rationing excludes the *average* attendee. Therefore, the optimal combination of price and rationing balances out the ratio of revenue to externalities of the marginal attendee (excluded by the price) against the ratio of revenue to externalities of the average attendee (excluded by the lottery). Regulating access via lotteries is thus useful precisely when average and marginal characteristics of potential attendees are not aligned: for example, when some lower-willingness-to-pay individuals exert a smaller negative externality than some higher-willingness-to-pay individuals.

The most interesting case arises when the ratio $J(\theta)/(-e(\theta))$ is non-monotone. This can happen when—in the language of our example—the probability of being a high-value member is strongly negatively correlated with willingness to pay in some range of the distribution. For example, blue collar workers who often make for the most dedicated fans of local soccer clubs may have lower ability to pay than wealthy tourists but they are essential to creating the atmosphere at the game. Mathematically, if a majority of high-value members have willingness to pay bounded above by a relatively low number, there might be a sharp drop in $e(\theta)$ around that threshold. This could in turn introduce a non-monotonicity in

---

[16]Note that a 1-tiered pricing mechanism with rationing is payoff-equivalent to a 2-tiered pricing mechanism in which the top tier has full access but is degenerate (the price is equal to the willingness to pay of the highest type $\bar{\theta}$)—this is why Result 2 is consistent with Theorem 1.

$J(\theta)/(-e(\theta))$: The ratio could be large for small θ due to low $(-e(\theta))$ and large for high θ due to high $J(\theta)$, but relatively low for intermediate θ. The following result describes the optimal mechanism in this case:

Result 3. *Suppose that $\theta_1^\star$ maximizes (2) when x_0 is constrained to be zero. It is then profitable to introduce a second tier with rationing if there exists $\theta_0 < \theta_1^\star$ such that*

$$\frac{J(\theta_1^\star)}{-e(\theta_1^\star)} < \frac{\int_{\theta_0}^{\theta_1^\star} J(\theta)dF(\theta)}{\int_{\theta_0}^{\theta_1^\star}(-e(\theta))dF(\theta)}. \tag{3}$$

If such θ_0 exists, then, in the optimal mechanism, assuming an interior solution,

$$\frac{J(\theta_0^\star)}{-e(\theta_0^\star)} = \frac{J(\theta_1^\star)}{-e(\theta_1^\star)} = \frac{R}{V},$$

where R and V are functions of $\theta_0^\star$, $\theta_1^\star$, and $x_0^\star$ as defined in (2).

Result 3 predicts that it may be optimal to offer two prices: a low price with rationing and a high price with full access. Intuitively, the low-price option is designed to attract agents who contribute to the atmosphere of the event but do not have a very high willingness to pay.

To understand condition (3), recall that the seller is trying to target agents with a high ratio of virtual surplus to (the negative of) the externality. If the seller is constrained to choosing a single price, she will choose a cutoff type $\theta_1^\star$ at which the ratio $J(\theta_1^\star)/(-e(\theta_1^\star))$ equates R/V, as in Result 2. Moreover, the second-order condition implies that the ratio must be locally increasing around $\theta_1^\star$. However, if $J(\theta)/(-e(\theta))$ is non-monotone, there may exist θ_0 such that agents with types between θ_0 and $\theta_1^\star$ achieve an even higher ratio on average, on the margin (i.e., when they enter with small probability). Thus, the seller finds it optimal to let these types in—with appropriate rationing to limit the congestion externality. At the optimal $\theta_0^\star$, the ratio of virtual surplus to (the negative of) the externality must be the same at $\theta_0^\star$ and $\theta_1^\star$, implying that the seller does not want to further adjust the composition of the rationed and full-access groups.

Because we assumed that externalities are always negative, both cutoffs $\theta_0^\star$ and $\theta_1^\star$ must belong to the region in which virtual surplus is positive. Suppose, instead, that some low-willingness-to-pay agents have a positive net externality—in the context of our example, this corresponds to assuming that $\Delta > 1$ and there is negative correlation between being a high-value member and willingness to pay. Then, condition (3) could be satisfied at θ_0 with $J(\theta_0) < 0$ and $e(\theta_0) > 0$. That is, the seller could decide to use a lottery to allocate to agents who contribute negatively to profits as long as they contribute positively to the willingness

15

to pay of agents selecting the full-access option.

Overall, Results 2 and 3 predict that performers selling tickets to their events may find it optimal to charge below-market-clearing prices and use lotteries to ration. The revenue-maximizing mechanism may take the form of a single lottery, or a combination of a lottery with a high-price full-access option. In the latter case, the lottery allows the seller to strike a better balance between improving the atmosphere at the event (by securing participation of high-value members) and extracting more revenue from agents with high willingness to pay.

The puzzle of why performers often sell tickets to their events at below-market-clearing prices has attracted considerable attention in economics.[17] Becker (1991) put forward an explanation by arguing that willingness to pay for goods consumed in public (such as going to a concert) may directly depend on demand of others (e.g., via image concerns); Mortimer et al. (2012) observed that artists may be motivated by generating demand for complementary products (e.g., recorded music); Che et al. (2013) showed that rationing arises when agents are budget-constrained and the seller cares about their utility;[18] and Loertscher and Muir (2022) demonstrated that lotteries may arise in a standard monopolist's problem with capacity constraints whenever ironing would be used in the canonical Myerson (1981) model. To the best of our knowledge, our explanation in terms of allocative externalities and associated concerns related to the composition of the audience is new to this literature.[19]

4.2 College admissions

Our next application is to design admission policy in the context of a college or other academic program in which students' value for participating (and hence, their willingness to pay) depends on the composition of the student body. Specifically, we assume that students' values for attending are higher when the incoming class is balanced across various characteristics such as socioeconomic background, minority status, geography, or field of

---

[17]Budish and Bhave (2023) provide a discussion and an empirical application (see also, e.g., Baliga and Ely (May 8, 2013); Arslan et al. (2023)).

[18]Rationing may also emerge if agents are not budget-constrained but the seller puts different welfare weights on different agents, as in the frameworks of Dworczak ⓣ al. (2021) and Akbarpour ⓣ al. (2024).

[19]That said, the idea that ticket pricing impacts the composition of attendees is well recognized in the industry; for example, a US Government Accountability Office (2018, p. 8) report on ticket pricing and resale listed the following among the reasons for ticket underpricing (emphasis in original):

> "**Audience mix.** Some artists prefer to have the most enthusiastic fans at their shows, rather than just those able to pay the most, especially in the front rows, where tickets are generally the most expensive."

interest.[20]

In this setting, on top of the unobserved heterogeneity in willingness to pay parameterized by the type θ, each (prospective) student has a label $i = (b, t)$, where $b \in \mathcal{B}$ captures diversity characteristics (that we refer to as "background") and $t \in \mathcal{T}$ captures observable measures of talent or ability (that we refer to as "test score"). We assume that a student with characteristics (b, t) exerts an externality $e_{(b,t)}$ on all other students—for example, students could value interacting with talented peers. Additionally, students are assumed to value diversity: For any b, let

$$s_b = \sum_{t \in \mathcal{T}} \mu_{(b,t)} e_{(b,t)} \int x_{(b,t)}(\theta) dF_{(b,t)}(\theta),$$

be the mass of admitted students with background b, weighted by their externality. Then, a student's willingness to pay for attending the college is θv, where $v = \nu(s_1, \ldots, s_{|\mathcal{B}|})$ satisfies the usual Inada conditions in each argument. These assumptions reflect the idea that—everything else being equal—there is a higher marginal value for adding a student from a background group that is underrepresented. For example, setting $\nu(s_1, \ldots, s_{|\mathcal{B}|}) = \prod_b s_b^{\mu_b}$ corresponds to assuming that the ideal composition of the student body mimics the population shares μ_b of the respective groups.[21]

The designer (the college) maximizes a weighted sum of revenue (with weight $\alpha \geq 1/2$) and student welfare: $W_i(\theta) = \alpha J_i(\theta) + (1-\alpha)h_i(\theta)$, where recall that $h_i(\theta) = (1 - F_i(\theta))/f_i(\theta)$ is the inverse hazard rate, which we assume is non-increasing.

The example is a special case of our baseline model except that the value v is now a function of multiple arguments; in Appendix A.4, we explain how the proofs of Theorem 1 and Proposition 1 can be modified to cover this case.

When affirmative action is allowed. When the college is allowed to condition admission decisions on all observable characteristics, we obtain the following result as a corollary of Proposition 1.

---

[20]Our model here complements the large market design literature on matching with diversity considerations. The frameworks of Abdulkadiroğlu and Sönmez (2003), Abdulkadiroğlu (2005), Hafalir et al. (2013), Echenique and Yenmez (2015), Wang et al. (2019), and Sönmez and Yenmez (2022), for example, provide different ways of designing matching mechanisms to implement diversity objectives when the relevant dimensions of identity are fully observable and the objectives can be stated in terms of target allocation or priority rules. By contrast, we take the potential externalities from diverse membership as a primitive and ask how to design a mechanism that produces the optimal selection on diversity status within a single school, in a setting with private information.

[21]See Strack and Yang (2023) for a related analysis of admission policies where compositional considerations are modeled as a hard constraint.

Result 4. *When admission decisions can explicitly condition on background characteristics, it is optimal to use a posted-price mechanism for every group i, with prices varying across the labels i.*

According to Result 4, the optimal admission policy is relatively simple: Each student faces a price for admission that depends on both the measure of their talent t and their background b. This mechanism could be implemented through a constant tuition and a combination of merit and diversity scholarships, along with prohibitively high prices for candidates with low ability t.

When affirmative action is not allowed. Next, we investigate the optimal admission mechanism when the college is not allowed to explicitly condition on the background characteristics $b \in \mathcal{B}$. Mathematically, this means that the set of labels is taken to be just the values $t \in \mathcal{T}$; the allocation rule $x_t(\theta)$ cannot depend on b directly; and thus

$$s_b = \sum_{t \in \mathcal{T}} \mu_t \int \underbrace{e_{(b,t)} \frac{\mu_{(b,t)} f_{(b,t)}(\theta)}{\mu_t f_t(\theta)}}_{e_{(b,t)}(\theta)} x_t(\theta) dF_t(\theta),$$

where we now treat $e_{(b,t)}(\theta)$ as the effective externality of group (b, t). The key observation is that the "composition externalities" are now a function of the type θ. While θ does not matter for the actual externality, when the college cannot condition admission decisions on background characteristics, willingness to pay becomes useful in inferring a candidate's background. Indeed, $e_{(b,t)}(\theta)$ is proportional to the share of candidates with background b among candidates with test score t. As a result, the expected externality may naturally be decreasing in θ for groups who tend to have low willingness to pay (e.g., students coming from disadvantaged backgrounds). This changes the optimal admission policy:

Result 5. *When admission decisions cannot explicitly condition on background characteristics, it is optimal to use a k^t-tiered pricing mechanism for group t, where $k^t \leq |\mathcal{B}| + 1$.*

The main insight of Result 5 is that randomized admission at lower prices may become optimal when affirmative action is not allowed. Moreover, the number of distinct admission options (differing in the price and probability of acceptance) is upper-bounded by the number of distinct background characteristics that matter for assessing diversity.

The intuition for the result is straightforward. Suppose that the college must admit a certain mass of students coming from economically underprivileged backgrounds to ensure a high value v. When affirmative action is available, the college achieves that via targeted admission decisions and scholarships. But when such a policy is not available, the college

18

must rely on self-selection to ensure diversity. Since students coming from economically underprivileged backgrounds tend to have low willingness to pay (due to low ability to pay), the college would have to reduce prices substantially for everyone to ensure enrollment of such students. Randomization relaxes that trade-off: When applying, students can be offered a choice between a high-tuition option and a lottery-assigned low-tuition option. Higher willingness-to-pay students would choose to avoid the additional admission risk and select into the high-tuition option, while low willingness-to-pay students could be admitted via the rationed option with sufficient probability to ensure the desired level of diversity. Note that such a policy necessarily bundles together admission and funding decisions.

This mechanism is similar in spirit to the work of Aygün and Bó (2021) illustrating how to design an affirmative action policy that incentivizes applicants to reveal their minority status—except that in our setting, agents are revealing their diversity status endogenously through their choice of admission option, rather than explicitly revealing it to the mechanism upfront. Meanwhile, our observation that the optimal mechanism for implementing a diverse membership outcome depends heavily on whether explicitly conditioning on group labels is permitted echoes an insight of Fryer Jr. et al. (2008) that implementing "color-blind" (versus "sighted") affirmative action to some degree requires flattening the allocation distribution with regards to signals such as test scores that are not directly linked to group identity, but may be correlated with it (see also Fryer Jr. and Loury (2013), and Heo and Park (2023)).[22]

Our results may help explain the recent admission trends among US colleges. In the last few years, many colleges moved to stop conditioning admission decisions on standardized test scores because of concerns that this was advantaging more affluent and otherwise better-resourced students. More recently, a number of colleges have reversed course on this, reinstating standardized test score consideration because they have found that test scores were less susceptible to socioeconomic bias than other factors such as participation in numerous extracurriculars.[23] Our framework demonstrates that if the distribution of scores t across groups b is heavily skewed towards more advantaged students (e.g., because they have access to private tutoring), then admitting more students with higher t undermines the self-selection mechanism. Thus, the college may find it optimal not to condition admission decisions on t.[24] If we think of the score t as comprising two dimensions—a standardized

---

[22]Our observations on how need- and merit-based financial aid shift the distribution of students selecting into admission also relate to work in public finance, market design, and higher education more broadly that has has examined these distributional impacts empirically (see, e.g., Heller (2006); Biró et al. (2020); Heo (2023)).

[23]See, e.g., Wren (March 28, 2022), Korn (February 5, 2024), and the discussion therein.

[24]In our model, it may in fact be optimal to give higher priority to certain intermediate or even low scores t if they are particularly prevalent within some underrepresented groups b. However, (unmodeled) moral-hazard considerations may prevent colleges from implementing admission rules that are non-monotone in

test score along with a second component accounting for extracurriculars, essays, and so forth—then the preceding logic suggests that colleges would want to adjust the weights on these two components to reflect how they correlate with background characteristics.

4.3 Membership communities

Next, we apply our model to the problem of designing membership communities, including clubs and affinity organizations, as well as digital communities where membership is managed via non-fungible tokens (NFTs).

We assume that there is a set I of groups of potential members, which one might think of as content creators, marquee members (e.g., celebrities), the general public, advertisers (i.e., people who wish to sell products to community members), and so forth. We assume that, for each group, types θ are distributed according to a cumulative distribution function $F_i(\theta) = 1 - (1-\theta)^{(1-\gamma_i)/\gamma_i}$ on the interval $[0,1]$, for $\gamma_i \in (0,1)$; this family of distributions is conveniently parameterized so that γ_i is the optimal monopoly price in the model without externalities (i.e., when $v_i \equiv 1$). To focus on across-group interactions, we assume that externalities are constant in agent's type: $e_{i \to j}(\theta) = e_{i \to j}$, for all $i, j \in I$. Furthermore, we assume that the designer maximizes revenue, and that

$$v_i = \max \left\{ 0,\, v_i^0 + \sum_{j \in I} \int_{\underline{\theta}_j}^{\bar{\theta}_j} e_{j \to i}\, x_j(\theta) \mu_j dF_j(\theta) \right\},$$

where $v_i^0 \in \mathbb{R}$ is the base value of access (which could be 0).

By Proposition 1, the optimal mechanism within each group is a posted price. Assuming an interior solution, so that first-order conditions are valid, we can derive the following formula for the optimal posted prices p_i:

$$p_i = \gamma_i v_i - (1 - \gamma_i) \sum_{j \in I} e_{i \to j} \mu_j R_j, \tag{4}$$

where $R_j := \frac{p_j}{v_j} \left(1 - F_j(\frac{p_j}{v_j}) \right)$ is the normalized revenue generated from selling to group j. That is, holding fixed the cutoff type $\theta_j := p_j/v_j$ buying at price p_j, R_j would be the revenue raised by selling to types above θ_j if $v_j \equiv 1$. The pricing formula (4) quantifies the trade-off between two motives of the designer setting the optimal price p_i for group i: *revenue-extraction* and *value-creation*. On the one hand, fixing v_i, the designer can extract the most revenue from group i by posting the monopoly price $\gamma_i v_i$; on the other hand, the

---

test scores (see Sönmez (2013) for discussion of an instance when precisely this moral hazard issue arose in practice).

price p_i influences the values v_j created for all groups j in proportion to the externality $e_{i \to j}$. For further intuition, note that—in line with Corollary 1—if group i does not exert any externalities, then it is charged the monopoly price $\gamma_i v_i$. On the other extreme, suppose group i exerts a positive externality on all other groups, i.e., $e_{i \to j} > 0$ for all $j \in I$, and either there is a group j that generates large revenue $\mu_j R_j$, or willingness to pay in group i is low (i.e., γ_i is close to 0). Then, the optimal price p_i charged to group i can become 0 (corresponding to a corner solution)—it is optimal to allow group i to join the community for free.

For the remainder of this subsection, our goal is to understand the economic consequences of formula (4) for the level of participation of various groups under the optimal mechanism. The level of participation by group i can be identified with the threshold type θ_i that joins the community given the optimal price p_i, as defined in the preceding paragraph. We refer to participation of all types above γ_i as the *level of participation absent externalities*. It is a classical insight of the two-sided platform literature (see, e.g., Caillaud and Jullien (2001); Rochet and Tirole (2002); Eisenmann et al. (2006)) and network literature (see, e.g., Candogan et al. (2012); Fainmesser and Galeotti (2015)) that groups exerting a positive externality should be incentivized to join via reduced prices or fees. However, it is not clear what the correct measure of the externality is in a setting with multiple groups that potentially interact in complicated ways. The following result identifies the key statistic in our framework.

Result 6. *Let*

$$e_{i \to R} := \sum_{j \in I} e_{i \to j} \mu_j R_j$$

denote the aggregate revenue externality exerted by group i. If group i exerts a positive (negative) aggregate revenue externality, then its optimal level of participation is higher (lower) than the optimal level of participation absent exernalities.

Therefore, whether group i should be subsidized or "taxed" relative to the optimal monopoly price $\gamma_i v_i$ depends exclusively on the weighted sum of normalized profits from all groups, weighted by the externalities $e_{i \to j}$ that group i exerts on each group $j \in I$. Note that profits R_j are normalized, that is, they are not scaled by the value v_j that group j has for joining the community. Intuitively, what matters is how participation by group i affects the value v_j on the margin (which is captured by the term $e_{i \to j}$), not the absolute value of v_j.

In practice, there are groups that are particularly important for the success of a membership community. For example, creators contribute to the community by providing content

for members to consume; celebrities may generate value to other members who are interested in interacting with them; and particularly enthusiastic consumers (as identified, e.g., through their prior engagement in similar communities) may serve as brand evangelists who drive up the visibility and value of membership. As a result, participation of certain groups may be incentivized through low or even zero prices. On the other hand, groups that exert negative externalities, such as advertisers, may be charged higher prices in order to reduce their level of participation.

Result 6 shows that the optimal participation of group i is particularly sensitive to externalities exerted on groups j with high normalized profits $\mu_j R_j$. However, because equation (4) determines optimal participation via a complicated fixed-point problem, it is unclear precisely how changes in participation by group j affect the optimal participation by group i. To cast light on this question, we can perturb the threshold type θ_j that joins the community from group j and ask how it affects the optimal threshold type θ_i in group i, holding fixed the level of participation of all other groups. We say that the participation of group i is a *substitute* for the participation of group j if a small upward perturbation in the level of participation for group j leads to decreased participation of group i; participation of group i is a *complement* to the participation of group j if a small upward perturbation of the level of participation for group j leads to increased participation of group i.[25]

Result 7. *Suppose that $\{\theta_j\}_{j \in I}$ are the threshold types determined by the solution to the system of equations (4). The participation of group i is a substitute for the participation of group j if*

$$e_{i \to j} \cdot J_j(\theta_j) - \frac{1}{v_i} \cdot e_{j \to i} \cdot e_{i \to R} < 0, \tag{5}$$

where $J_j(\theta_j) := \theta_j - (1 - F(\theta_j))/f(\theta_j)$ is the virtual surplus at the optimal allocation (which has the opposite sign to $e_{j \to R}$). The participation of group i is a complement to the participation of group j if the inequality sign in inequality (5) is reversed.

Intuitively, an exogenous increase in the participation of group j impacts the relative importance of the revenue-extraction and value-creation channels for group i. There are two effects of adding a new member from group j. The first effect—corresponding to the first term in inequality (5)—is that an increase in participation of group j changes the normalized revenue R_j that is raised from group j. Revenue raised from group j increases if and only if the virtual surplus $J_j(\theta_j)$ at the initial threshold type θ_j is positive—that is, when participation by group j is lower than it would be absent externalities. (In light of

---

[25]Formally, we ask about the sign of the partial derivative of θ_i with respect to θ_j when θ_i is pinned down by equation (4). See Appendix A.5 for the derivations.

Result 6, this is the case when group j exerts a negative aggregate revenue externality.) If the normalized revenue R_j goes up, the designer cares more about the value-creation motive through the externality that group i exerts on group j—thus, she wants to add more members from group i if $e_{i \to j}$ is positive (and vice versa if $e_{i \to j}$ is negative). If the normalized revenue R_j goes down (i.e., group j exerts a positive aggregate revenue externality), these incentives are reversed. The second effect—corresponding to the second term in inequality (5)—is that an increase in participation of group j changes the value v_i of group i from joining the community. The value is higher if and only if group j exerts a positive externality, $e_{j \to i} > 0$. Suppose that v_i indeed goes up. Then, the designer cares more about the revenue-extraction motive when pricing access for group i. Thus, she wants to add more members from group i if and only if membership was restricted compared to the optimal level absent externalities (i.e., if group i exerts a negative aggregate revenue externality); she wants to remove members if membership was higher than the optimal level absent externalities. These incentives are reversed if the value v_i goes down when an additional agent in group j enters.

Suppose that both groups i and j exert a positive aggregate revenue externality and positive externalities on one another. A casual intuition would suggest that participation by these two groups should be complementary. As Result 7 shows, however, this is not necessarily the case. When a new member from group j enters, the two forces work in opposite directions. The intuitive force is that joining is now more valuable for agents in group i, and the designer can thus benefit by increasing participation by group i (the first term in expression (5) is positive). On the other hand, however, group j already had excessive participation (compared to the benchmark without externalities) and hence adding one more member j drives down the normalized profits from group j. This in turn means that the designer now puts relatively more weight on revenue extracted from group i than on how it affects the value for group j. Since participation by group i is higher than would be optimal for revenue extraction, the designer has an incentive to *exclude* members from group i. The net effect depends on which of these two forces is stronger.

In addition to extending the classic intuitions about dynamics of multi-sided platforms, our analysis is particularly relevant to the emergence of new online community formats, such as those with membership gated by NFT ownership (see, e.g., Kaczynski and Kominers (November 10, 2021, 2024); Cassatt (2023); Kraski and Shenkarow (2023)).[26]

NFTs are digital records (akin to "digital deeds") that can be used to define and exchange ownership of media such as imagery or song files, as well as other digital or physical goods.

---

[26]Our perspective on NFTs as gateways to a community of owners follows a common interpretation—and associated business models—in the NFT industry, and is complementary to the Veblen-good view of NFTs explored in the work of Oh et al. (2023).

NFTs by nature link their owners into a network, and through software integrations can be used to gate access to online spaces. They are thus often used to establish communities of enthusiasts for a media creator and/or an associated brand.[27] For example, a number of NFTs have been used to anchor community membership, engagement, and co-creation around character brands such as Bored Ape Yacht Club, Goblintown, Pudgy Penguins, SupDucks, Tycoon Tigers, and 1337 Skulls (see Kaczynski and Kominers (2024)).

Our framework both provides practical advice for the design of NFT primary sales and helps explain some NFT allocation mechanisms already commonly used in practice. Branded NFTs are often intentionally sold at low prices or given away for free, while later reaching prices in the hundreds or even thousands of dollars. Recognizing the role of NFTs as gateways to membership communities provides an explanation for this pricing phenomenon: Often, a substantial share of an NFT's equilibrium value is determined by the set of community members and their mutual externalities; this justifies a low price ex ante (when the community is in effect empty and membership is uncertain), and a higher equilibrium price in the event that the community takes off. Variation in externalities among prospective NFT community members implies that different groups of agents should be charged different prices depending on their desirability as potential members, as measured by their revenue externality (Result 6). And indeed, NFT primary sales are often managed through some mixture of an open public sale and a private sale for people on an "allow list." The allow-list sale is typically conducted at a lower price than the public sale—and the allow list comprises prospective buyers who the founders believe are likely to be particularly value-generating for the community, often assessed as a function of participation in similar communities, or through a survey or other costly and informative signal of interest (see, e.g., Kominers and Roughgarden (September 10, 2022); Kominers and The 1337 Skulls Sers (April 6, 2023)). Our analysis implies that the optimal design of allow lists must take into account the substitutability and complementarity between different categories of potential members, and highlights the importance of expanding or contracting any given allow list depending on the takeup from other groups (Result 7).

4.4 Token-based networks

In blockchain-based protocol networks, a common design pattern is to manage access and utilization through the ownership of a digital token. For example, in the Ethereum blockchain, the associated Ether token is used to pay for computing capacity (Ethereum Foundation (accessed November 3, 2024)); similarly, the oracle service Chainlink uses LINK tokens to pay

---

[27]For discussion of the function of NFTs and their role in establishing and managing digital brand communities, see, e.g., Kraski and Shenkarow (2023) and Kaczynski and Kominers (2024).

for data provision (Chainlink DevHub (accessed November 3, 2024)), and storage services such as Filecoin use their own tokens to pay for storage (Filecoin Docs (accessed November 3, 2024)). These tokens are often at least partially distributed in advance of the network's launch, which means that agents may receive tokens before those tokens' full functionality is realized.

In this section, we investigate the optimal design for a blockchain project that allocates tokens in the primary market to raise revenue and regulate access to the network. We focus on conditions under which the token can have any value at all in equilibrium (as measured by the price it achieves in the secondary market), the role of vesting in getting the network off the ground, and optimal pricing.

We consider two groups of agents. First, there are potential developers, denoted D, who might actively contribute to building or improving the network. Second, there are agents who are part of the "general public," denoted P, who may decide to hold the token for speculative or investment purposes or use the network's functionality but do not contribute to its development. We assume that agents belonging to the general public exert a network congestion externality

$$e_P(\theta) = -1,$$

which can be interpreted as occupying the network's computational resources.[28] The developers potentially contribute to the network development, in addition to exerting a congestion externality:

$$e_D(\theta) = -1 + e(\theta),$$

where the dependence of $e(\theta)$ on θ reflects potential correlation between willingness to pay and the likelihood of contributing. Finally, we interpret the value v as the resale price of the token that will arise in the secondary market once it opens (hence, v is not indexed by group label i). We assume that the value is given by

$$v = \max\left\{ \sum_{i\in\{D,P\}} \int e_i(\theta)x_i(\theta)\mu_i dF_i(\theta),\, 0 \right\}.$$

Intuitively, the value of the token is a reflection of the network functionality—given by the total expected non-congestion externalities of all developers who hold tokens—net of the

---

[28]For example, congestion leads to increased transaction fees on the Bitcoin blockchain (Huberman et al. (2021)), as well as on the Ethereum network (Ethereum Foundation (accessed November 3, 2024)).

network congestion. Since only developers can provide positive externalities, it is crucial to incentivize them to join. At the same time, giving the token to too many agents undermines the value of the network through congestion—a network that is too congested (relative to its functionalities) becomes useless and has no value.

We interpret θ as capturing an agent's belief about the (unmodeled) long-term resale value of the token. It is then natural to assume that the median of the unconditional distribution of θ is 1. If we assume that tokens can be costlessly liquidated at v once the secondary market opens, then agents with $\theta \geq 1$ are the ones who believe the token will appreciate in value and hence decide to hold it. In contrast, agents with $\theta < 1$ want to sell as soon as the secondary market opens.

We assume that the designer is the network founder who controls the allocation of tokens in the primary market (with total quantity normalized to 1). The founder thus maximizes

$$\sum_{i \in \{D, P\}} \int \left(v J_i(\theta) x_i(\theta) + v(1 - x_i(\theta)) \right) \mu_i dF_i(\theta),$$

under the assumption that the tokens kept by the founder are valued at v (and do not contribute to the network congestion). The allocation rule $x_i(\theta)$ represents holding tokens, which is required for access to the network. In particular, only agents with $x_i(\theta) > 0$ can contribute to network development.

Solution without vesting. Under our interpretation of θ as the agent's optimism about the future value of the network (relative to the resale price of the token), the designer is constrained to allocation rules satisfying $x_i(\theta) = 0$ for $\theta < 1$ (since agents with $\theta < 1$ will not hold tokens once the secondary market opens). We can then make the following simple observation:

Result 8. *The network token has a strictly positive resale price if and only if*

$$\exists \theta^\star \geq 1 \text{ such that } \int_{\theta^\star}^{\bar{\theta}_D} e_D(\theta) dF_D(\theta) > 0. \tag{6}$$

It is optimal to allocate tokens by posting a single price for the general public, and at most two prices for developers (where there could be rationing at the lower price).

A strictly positive value of the network can be generated if there exists a price weakly above the resale value with the property that—among the developers who decide to buy and then hold the tokens—the average positive externality $e(\theta)$ exceeds the congestion externality 1. In other words, the founder must attract the high-value developers through

self-selection based on optimism about the network value. We view Result 8 as a negative result. Compared to the general public, the beliefs of the well-informed high-skill developers are more likely to be concentrated around 1. Moreover, developers with unrealistically high expectations may not be the best contributors—$e(\theta)$ may be decreasing in θ above $\theta = 1$. At the same time, the founder cannot engage the potentially valuable developers with $\theta < 1$ because they would prefer to sell their tokens as soon as the secondary market opens.

Solution with vesting. Next, we observe that if we allow for vesting (an allocation of tokens with a temporary restriction on reselling them in the secondary market), then the condition from Result 8 can be relaxed. We model vesting simply by allowing the designer to choose any allocation rule, where $x_i(\theta) > 0$ for $\theta < 1$ is now possible due to vesting.

Result 9. *With vesting, the network token has a strictly positive resale price if and only if*

$$\exists \theta^\star \geq 0 \ such \ that \ \int_{\theta^\star}^{\bar{\theta}_D} e_D(\theta) dF_D(\theta) > 0. \tag{7}$$

It is optimal to sell tokens to the general public at a single price that strictly exceeds v and to developers at at most two prices; if condition (6) fails but condition (7) holds, then at least one of these prices is below v, involves vesting, and may involve rationing.

With vesting, the designer can access the potentially high-externality developers with lower beliefs θ about the network value. This expands the set of parameters for which the token can have a strictly positive resale price. In practice, such a mechanism may take the form of promising token grants to developers who commit to building around a new blockchain network in advance of launch.

Rationing schemes do not play a role in determining whether the token can have positive value—in conditions (6) and (7), it is enough to consider the effect of charging different posted prices to developers. However, rationing may play an important role in the revenue-maximizing mechanism. Rationing at a low price with vesting may allow the designer to ensure sufficient participation by high-externality developers with relatively low θ, without putting too much downward pressure on the price she charges for the non-rationed option to more optimistic developers. Unlike in the ticket sales example in Section 4.1, it may be optimal to have a rationing scheme with a price of zero for developers, as the designer can still make money by selling to the general public (or by keeping the tokens and reselling them in the secondary market once it opens). In practice, such mechanisms sometimes take the form of "free mints." Free mints may be optimal when $e(\theta)$ exceeds the congestion externality even for θ close to zero. Such a case is plausible if the network founders believe

that initially skeptical developers are likely to engage with the network once they gain access to it.

In token-based networks where the goal is to create a decentralized software ecosystem—such as with open-source projects—our results suggest it makes sense to reserve a share of the initial tokens for developers, and implement vesting schedules to incentivize active participation. While our model is static, intuitively, vesting should be calibrated so that token holders' ability to sell kicks in after the ecosystem value has already been well established; moreover, it is particularly important to do this for large token holders who do not directly contribute to the growth of the software ecosystem.[29] And indeed, in practice we often see token-based ecosystems implementing vesting schedules for early participants for precisely this reason (see, e.g., Jennings et al. (August 8, 2024) for discussion), and often explicitly tying those schedules to project milestones. For example, Howell et al. (2020) and Davydiuk et al. (2023) document that vesting has been widely used in initial coin offerings; vesting has also been common for "Layer 2" blockchain ecosystems such as those of the Arbitrum DAO (accessed March 17, 2024) and Optimism Collective (December 18, 2023). Singh (December 16, 2023) provides an overview of mechanisms for token vesting used in practice.

4.5   Additional applications

To close, we briefly comment on other potential applications.

Ad auctions.   A recent US Department of Justice antitrust case revealed that Google had been relying on a Randomized Generalized Second Price auction (RGSP) to sell ad spots to advertisers. The RGSP used by Google differed from a classical generalized second-price auction in two ways. First, the competitiveness of the bid depended not only on the bid amount but also on the campaign's long-term value (LTV), determined by factors such as ad quality and the context of the search query. Second, even conditional on LTV, Google would sometimes randomly permute the top bidders if their bids were close enough.[30]

Setting aside the question of whether RGSP was anti-competitive, our framework provides a simple explanation for why Google might have preferred the RGSP design to a simpler auction format. Effectively, Google is selling access to its multi-sided platform: Allocation of ad slots generates (mostly negative) externalities because ads shown as search results affect consumers' search experience (and hence their willingness to engage with the search engine).

---

[29]For analysis of the dynamics of token-based networks see for example Bakos and Halaburda (2022), Cong et al. (2022), Sockin and Xiong (2023), and Reuter (2024).

[30]For details, see Weitzman (October 5, 2023). Randomization was also used by Yahoo! under the name of "squashing;" see, e.g., Agius (October 18, 2023).

A bidder's externality would typically vary with characteristics that are observable to Google (e.g., the fit between the ad and the consumer's query); however, it could also depend on characteristics that are unobserved or cannot be conditioned on. For example, a consumer's search experience might deteriorate if they are shown an ad by the same advertiser over and over again; yet, Google might be unable to explicitly down-weight bidders who bid aggressively on particular search queries.

Adopting a model similar to the one we developed in Section 4.3 demonstrates that Google would have an incentive to use heterogeneous prices for bidders differing in their observed externalities (e.g., a higher price for campaigns with low ad quality)—as captured by the use of LTVs. At the same time, reasoning analogous to that used in Section 4.2 suggests that randomization could be used to ensure a more balanced composition of ads, in lieu of explicit preferential treatment of some bidders. While other reasons for randomization could be possible—e.g., the classical ironing solution of Myerson (1981) provides a justification if the distribution of bidders' values is sufficiently irregular—some of Google's own testimony at trial was suggestive of an externality story: A Google representative claimed that if the auction did not occasionally flip the winner with the runner-up, then "Amazon would always show up on top" (Agius (October 18, 2023)).

Brand image concerns. Membership-like effects arise in sales of branded goods that are consumed in public, such as clothes, cars, or even certain types of NFTs. Companies that sell such goods must take into account how potential pricing strategies affect the composition of their customer base; this feeds into the brand perception, which in turn affects consumers' willingness to pay. Our framework can be used to quantitatively study membership design in the presence of well-known effects such as the "bandwagon effect" (people seek to own goods that other individuals, especially those with high status, own) or the "snob effect" (widespread ownership of a good, especially by a consumer's peers, can cause it to lose value).[31] In the language of our model, these effects can be captured by externalities (positive or negative) that potentially depend on consumer characteristics that are not directly observed by the seller (e.g., socioeconomic status).

The intuition from Section 4.1 suggests that a posted-price mechanism maximizes sales revenue for goods whose value is increasing in the wealth and status of the consumers that own them (e.g., luxury fashion and flashy cars). As in Section 4.3, however, it may be optimal to offer such goods at low or zero prices to selected individuals (identified by observables) whose public consumption has a particularly strong effect on the perceived value of the good (e.g., to celebrities or social media influencers). Meanwhile, our model predicts that

---

[31]See, e.g., Leibenstein (1950), Vigneron and Johnson (1999), and Han et al. (2010).

lotteries at below-market-clearing prices may be optimal if consumers with intermediate willingness to pay contribute most positively to the brand value; certain brands may aim to maintain scarcity without creating the perception of being out of reach for the middle class. For example, Nike sells some collections of its shoes at prices that create excess demand that is then rationed via an explicit lottery (referred to as the "SNKRS drawing"; see Nike (accessed March 17, 2024))—our framework provides an explanation rationalizing such a sales mechanism.

Affordable housing. Composition effects—similar to the ones we analyzed in Section 4.2—are relevant for the design of affordable housing programs (such as LIHTC in the US). On the one hand, the value of a housing unit to potential tenants may depend on the demographic composition of the neighborhood; on the other hand, allocation of an affordable housing unit often creates externalities for the current residents (see, e.g., Diamond and McQuade (2019)). This makes housing in a given area a type of a membership good whose optimal allocation can be analyzed using our framework. For example, sales of new apartments in Korea are subject to price ceilings below the market-clearing price with an intention of fostering social interactions between residents with various income levels (see Yun and Kim (2014)). Furthermore, when direct use of certain demographic characteristics (such as race) is not allowed in social programs, our results (in particular from Section 4.2) suggest a potential role for carefully constructed lotteries (see Cook et al. (2023) for a discussion of how suboptimally designed rationing schemes may conflict with the designer's redistributive goals).

Consumer data. Digital platforms leverage consumer data gathered from their services to enhance user experience and increase revenue. This data, while specific to individual users, also provides insights into other consumers with similar traits and behaviors, influencing the services tailored to them and their willingness to pay (see, e.g., Galperti et al. (2023)). In pricing their services, platforms need to consider not only the willingness to pay of each consumer group but also the value of the data these groups provide about others. Our model proposes that services should be offered at low or no cost to consumers who generate exceptionally valuable data, such as early adopters or those with rare but important characteristics. Additionally, our analysis would recommend implementing A/B testing or limited-time discount offers that provide randomized access at below-market prices, for consumers with moderate willingness to pay who nonetheless contribute significant externalities (see Ascarza (2018) for a discussion of how platforms can benefit from A/B testing targeted to specific subgroups).

References

ABDULKADIROĞLU, A. (2005): "College admissions with affirmative action," *International Journal of Game Theory*, 33, 535–549.

ABDULKADIROĞLU, A. AND T. SÖNMEZ (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

AGIUS, N. (October 18, 2023): "What is RGSP? Google's Randomized Generalized Second-Price ad auctions explained," *Search Engine Land*, https://searchengineland.com/google-rgsp-randomized-generalized-second-price-ad-auctions-explained-433053.

AKBARPOUR, M. ⓡ E. BUDISH ⓡ P. DWORCZAK ⓡ S. D. KOMINERS (2024): "An Economic Framework for Vaccine Prioritization," *Quarterly Journal of Economics*, 139, 359–417.

AKBARPOUR, M. ⓡ P. DWORCZAK ⓡ S. D. KOMINERS (2024): "Redistributive Allocation Mechanisms," *Journal of Political Economy*, 132, 1831–1875.

ARBITRUM DAO (accessed March 17, 2024): "Governance docs: $ARB airdrop eligibility and distribution specifications," https://docs.arbitrum.foundation/airdrop-eligibility-distribution.

ARSLAN, H. A., N. TEREYAĞOĞLU, AND Ö. YILMAZ (2023): "Scoring a touchdown with variable pricing: Evidence from a quasi-experiment in the NFL ticket markets," *Management Science*, 69, 4435–4456.

ASCARZA, E. (2018): "Retention Futility: Targeting High-Risk Customers Might be Ineffective," *Journal of Marketing Research*, 55.

AYGÜN, O. AND I. BÓ (2021): "College admission with multidimensional privileges: The Brazilian affirmative action case," *American Economic Journal: Microeconomics*, 13, 1–28.

BAKOS, Y. AND H. HALABURDA (2022): "Overcoming the coordination problem in new marketplaces via cryptographic tokens," *Information Systems Research*, 33, 1368–1385.

BALIGA, S. AND J. ELY (May 8, 2013): "The Power of Purple Pricing," *Harvard Business Review*.

BAUER, H. (1958): "Minimalstellen von Funktionen und Extremalpunkte," *Archiv der Mathematik*, 9, 389–393.

BECKER, G. S. (1991): "A Note on Restaurant Pricing and Other Examples of Social Influences on Price," *Journal of Political Economy*, 99, 1109–1116.

BERNSTEIN, S. AND E. WINTER (2012): "Contracting with Heterogeneous Externalities," *American Economic Journal: Microeconomics*, 4, 50–76.

BIRÓ, P., A. HASSIDIM, A. ROMM, R. I. SHORRER, AND S. SÓVÁGÓ (2020): "The Large Core of College Admission Markets: Theory and Evidence," Working Paper.

BOARD, S. (2009): "Monopolistic group design with peer effects," *Theoretical Economics*, 4, 89–125.

BONNANS, J. F. AND A. SHAPIRO (2000): *Perturbation Analysis of Optimization Problems*, Springer.

BUDISH, E. AND A. BHAVE (2023): "Primary-Market Auctions for Event Tickets: Eliminating the Rents of 'Bob the Broker'?" *American Economic Journal: Microeconomics*, 15, 142–1170.

CAILLAUD, B. AND B. JULLIEN (2001): "Competing cybermediaries," *European Economic Review*, 45, 797–808, 15th Annual Congress of the European Economic Association.

CANDOGAN, O., K. BIMPIKIS, AND A. OZDAGLAR (2012): "Optimal Pricing in Networks with Externalities," *Optimal Pricing in Networks with Externalities*, 60, 883–905.

CASSATT, A. (2023): *Web3 Marketing: A Handbook for the Next Internet Revolution*, Wiley.

CHAINLINK DEVHUB (accessed November 3, 2024): "LINK Token Contracts," https://docs.chain.link/resources/link-token-contracts.

CHE, Y.-K., I. GALE, AND J. KIM (2013): "Assigning Resources to Budget-Constrained Agents," *Review of Economic Studies*, 80, 73–107.

CONG, L. W., Y. LI, AND N. WANG (2022): "Token-based platform finance," *Journal of Financial Economics*, 144, 972–991.

COOK, C., P. Z. LI, AND A. J. BINDER (2023): "Where to Build Affordable Housing? Evaluating the Tradeoffs of Location," Working Paper.

DAMIANO, E. AND H. LI (2007): "Price Discrimination and Efficient Matching," *Economic Theory*, 30, 243–263.

DAVYDIUK, T., D. GUPTA, AND S. ROSEN (2023): "De-crypto-ing signals in initial coin offerings: Evidence of rational token retention," *Management Science*, 69, 6584–6624.

DIAMOND, R. AND T. MCQUADE (2019): "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development," *Journal of Political Economy*, 127, 1063–1117.

DUBINS, L. E. (1962): "On extreme points of convex sets," *Journal of Mathematical Analysis and Applications*, 5, 237–244.

DWORCZAK, P. ⓡ S. D. KOMINERS ⓡ M. AKBARPOUR (2021): "Redistribution through Markets," *Econometrica*, 89, 1665–1698.

ECHENIQUE, F. AND M. B. YENMEZ (2015): "How to control controlled school choice," *American Economic Review*, 105, 2679–2694.

EISENMANN, T., G. PARKER, AND M. W. VAN ALSTYNE (2006): "Strategies for Two-Sided Markets," *Harvard Business Review*, 84, 92–101.

ETHEREUM FOUNDATION (accessed November 3, 2024): "Gas and Fees," https://ethereum.org/en/developers/docs/gas/.

FAINMESSER, I. P. AND A. GALEOTTI (2015): "Pricing Network Effects," *The Review of Economic Studies*, 83, 165–198.

FILECOIN DOCS (accessed November 3, 2024): "The FIL token," https://docs.filecoin.io/basics/assets/the-fil-token.

FRYER JR., R. G. AND G. C. LOURY (2013): "Valuing diversity," *Journal of Political Economy*, 121, 747–774.

FRYER JR., R. G., G. C. LOURY, AND T. YURET (2008): "An economic analysis of color-blind affirmative action," *Journal of Law, Economics, & Organization*, 24, 319–355.

GALPERTI, S., A. LEVKUN, AND J. PEREGO (2023): "The Value of Data Records," *The Review of Economic Studies*, 91, 1007–1038.

GOMES, R. AND A. PAVAN (2021): "Two-sided Markets, Platforms, and Matching Design," in *Online and Matching-Based Market Design*, Cambridge University Press.

HAFALIR, I. E., M. B. YENMEZ, AND M. A. YILDIRIM (2013): "Effective Affirmative Action in School Choice," *Theoretical Economics*, 8, 325–363.

HALAC, M., I. KREMER, AND E. WINTER (2020): "Raising Capital from Heterogeneous Investors," *American Economic Review*, 110, 889–921.

HALAC, M., E. LIPNOWSKI, AND D. RAPPOPORT (2024): "Pricing for Coordination,"

Working Paper.

HAN, Y. J., J. C. NUNES, AND X. DRÈZE (2010): "Signaling status with luxury goods: The role of brand prominence," *Journal of Marketing*, 74, 15–30.

HELLER, D. E. (2006): "Merit Aid and College Access," in *Symposium on the Consequences of Merit-Based Student Aid.*

HEO, E. J. (2023): "Financial aid in college admissions: Need-based versus merit-based," *Social Choice and Welfare*, 60, 265–297.

HEO, E. J. AND I.-U. PARK (2023): "Need-blind versus Need-aware College Admissions, Financial Aid and Student Loan Programs," Working Paper.

HOWELL, S. T., M. NIESSNER, AND D. YERMACK (2020): "Initial coin offerings: Financing growth with cryptocurrency token sales," *Review of Financial Studies*, 33, 3925–3974.

HUBERMAN, G., J. D. LESHNO, AND C. MOALLEMI (2021): "Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System," *Review of Economic Studies*, 88, 3011–3040.

JADBABAIEA, A. AND A. KAKHBOD (2019): "Optimal Contracting in Networks," *Journal of Economic Theory*, 183, 1094–1153.

JEHIEL, P. AND B. MOLDOVANU (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237–1259.

JEHIEL, P., B. MOLDOVANU, AND E. STACCHETTI (1996): "How (Not) to Sell Nuclear Weapons," *American Economic Review*, 814–829.

———— (1999): "Multidimensional Mechanism Design for Auctions with Externalities," *Journal of Economic Theory*, 85, 258–293.

JENNINGS, M., J. BURLESON, AND Z. GRAY (August 8, 2024): "Token rights in term sheets: How to avoid predatory deals," *a16z crypto.*

JULLIEN, B., A. PAVAN, AND M. RYSMAN (2021): "Two-sided Markets, Pricing, and Network Effects," in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, Elsevier, vol. 4, 485–592.

KACZYNSKI, S. AND S. D. KOMINERS (2024): *The Everything Token: How NFTs and Web3 Will Transform the Way We Buy, Sell, and Create*, Portfolio.

———— (November 10, 2021): "How NFTs Create Value," *Harvard Business Review.*

KANG, Z. Y. (2020): "Optimal Indirect Regulation of Externalities," Working Paper.

——— (2023): "The Public Option and Optimal Redistribution," Working Paper.

KOMINERS, S. D. AND T. ROUGHGARDEN (September 10, 2022): "NFT Sales: Clearing the Market, Avoiding Gas Wars," *a16z crypto*.

KOMINERS, S. D. AND THE 1337 SKULLS SERS (April 6, 2023): "A new NFT launch strategy: The wave mint," *a16z crypto*.

KORN, M. (February 5, 2024): "Dartmouth Reinstates SAT Requirement in First for Ivy League," *Wall Street Journal*.

KRASKI, J. AND J. SHENKAROW (2023): *The Future of Community: How to Leverage Web3 Technologies to Grow Your Business*, Wiley.

LEIBENSTEIN, H. (1950): "Bandwagon, snob, and Veblen effects in the theory of consumers' demand," *Quarterly Journal of Economics*, 64, 183–207.

LOERTSCHER, S. AND E. V. MUIR (2022): "Monopoly Pricing, Optimal Randomization, and Resale," *Journal of Political Economy*, 130, 566–635.

MEISNER, V. AND P. PILLATH (2024): "Monetizing digital content with network effects: A mechanism-design approach," *arXiv preprint arXiv:2408.15196*.

MOLDOVANU, B., A. SELA, AND X. SHI (2007): "Contests for status," *Journal of Political Economy*, 115, 338–363.

MORTIMER, J. H., C. NOSKO, AND A. SORENSEN (2012): "Supply responses to digital distribution: Recorded music and live performances," *Information Economics and Policy*, 24, 3–14.

MUSSA, M. AND S. ROSEN (1978): "Monopoly and product quality," *Journal of Economic Theory*, 18, 301–317.

MYERSON, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–78.

NIKE (accessed March 17, 2024): "How can I join Nike SNKRS drawings?" https://www.nike.com/help/a/nike-snkrs-drawings.

OH, S., S. ROSEN, AND A. L. ZHANG (2023): "Digital Veblen Goods," Working Paper.

OPTIMISM COLLECTIVE (December 18, 2023): "What is the Optimism Collective?: OP Allocations," https://community.optimism.io/docs/governance/allocations/.

OSTRIZEK, F. AND E. SARTORI (2023): "Screening While Controlling an Externality," *Games and Economic Behavior*, 139, 26–55.

PAI, M. AND P. STRACK (2022): "Taxing Externalities Without Hurting the Poor," Working Paper.

RAYO, L. (2013): "Monopolistic Signal Provision†," *The B.E. Journal of Theoretical Economics*, 13, 27–58.

REUTER, M. (2024): "Platform Precommitment via Decentralization," IMF Working Paper 2024/028.

ROCHET, J.-C. AND J. TIROLE (2002): "Cooperation among Competitors: Some Economics of Payment Card Associations," *RAND Journal of Economics*, 33, 549–570.

ROHLFS, J. (1974): "A Theory of Interdependent Demand for a Communications Service," *The Bell Journal of Economics and Management Science*, 5, 16–37.

SEGAL, I. (1999): "Contracting with Externalities," *Quarterly Journal of Economics*, 114, 337–388.

——— (2003): "Coordination and discrimination in contracting with externalities: divide and conquer?" *Journal of Economic Theory*, 113, 147–181.

SINGH, J. (December 16, 2023): "Vesting in crypto, explained," *Cointelegraph*.

SOCKIN, M. AND W. XIONG (2023): "Decentralization through tokenization," *Journal of Finance*, 78, 247–299.

SÖNMEZ, T. (2013): "Bidding for army career specialties: Improving the ROTC branching mechanism," *Journal of Political Economy*, 121, 186–219.

SÖNMEZ, T. AND M. B. YENMEZ (2022): "Affirmative action in India via vertical, horizontal, and overlapping reservations," *Econometrica*, 90, 1143–1176.

STRACK, P. AND K. H. YANG (2023): "Privacy Preserving Signals," Working Paper.

SZAPIEL, W. (1975): "Points Extrémaux dans les Ensembles Convexes(I). Théorie Générale," *Bulletin de l'Académie Polonaise des Sciences. Sériedes Sciences Mathématiques, Astronomiques et Physiques*, 23, 389–393.

US GOVERNMENT ACCOUNTABILITY OFFICE (2018): "Event Ticket Sales: Market Characteristics and Consumer Protection Issues," https://www.gao.gov/assets/gao-18-347.pdf.

VIGNERON, F. AND L. W. JOHNSON (1999): "A review and a conceptual framework of prestige-seeking consumer behavior," *Academy of Marketing Science Review*, 1, 1–15.

WANG, F., R. JAGADEESAN, AND S. D. KOMINERS (2019): "Optimizing reserves in school choice: A dynamic programming approach," *Operations Research Letters*, 47, 438–446.

WEITZMAN, Y. (October 5, 2023): "A look inside Google's advertising auctions," https://www.bigtechontrial.com/p/a-look-inside-googles-advertising.

WINTER, E. (2004): "Incentives and Discrimination," *American Economic Review*, 94, 764–773.

WREN, K. (March 28, 2022): "Q&A: Stuart Schmill on MIT's decision to reinstate the SAT/ACT requirement," *MIT News Office*.

YUN, S. AND K. KIM (2014): "The Effects of the Price Ceiling Regulation on the New Apartment Price on Characteristics of Housing in Korea," *International Conference on Economics, Education and Humanities*.

A Proofs Omitted from the Main Text

A.1 Proof of Theorem 1

We first note that an optimal solution to the designer's problem exists, by a standard argument (the objective function is upper semi-continuous on a compact set, when we endow the space of non-decreasing allocation rules with the L^1 topology). Our proof strategy is to take an arbitrary mechanism, and then show that this mechanism can be weakly improved upon by a modified mechanism that satisfies the properties listed in Theorem 1. This will establish that there exists an optimal mechanism with these properties.

Let $N = |I|$. We consider an auxiliary problem for the designer of reoptimizing her objective over the allocation rule $x_i(\theta)$ for group i only, keeping the other allocation rules fixed, subject to preserving the values v_j for all groups $j \in I$. Value preservation requires

$$v_j = \nu_j \left(\sum_{k \in I} \int e_{k \to j}(\theta) x_k(\theta) \mu_k dF_k(\theta) \right), \quad \forall j \in I.$$

Because we are keeping fixed all allocation rules other than the one for group i, we can

impose a stronger condition that implies value preservation:

$$c_j = \int e_{i \to j}(\theta) x_i(\theta) \mu_i dF_i(\theta), \quad \forall j \in I,$$

for some constants c_j (pinned down by the initial mechanism). Formally, the auxiliary problem becomes:

$$\max_{x_i} \quad \int W_i(\theta) x_i(\theta) dF_i(\theta) \tag{8}$$

$$\text{s.t.} \quad c_j = \int e_{i \to j}(\theta) x_i(\theta) \mu_i dF_i(\theta), \quad \forall j \in I,$$

$$x_i(\theta) \text{ is non-decreasing.}$$

Let us introduce an auxiliary *linear* operator $\mathbf{E}^i$ on the space of non-decreasing allocation rules for group i, defined by

$$\mathbf{E}^i_j(x_i) := \int e_{i \to j}(\theta) x_i(\theta) \mu_i dF_i(\theta), \quad \forall j = 1, \dots, |I|.$$

Then, we can rewrite the first constraint in problem (8) in vector notation as

$$c = \mathbf{E}^i(x_i).$$

The following simple lemma will allow us to connect the complexity of the non-linear operator E^i with the complexity of the linear operator $\mathbf{E}^i$.

Lemma 2. *Suppose that* $\mathbf{E} : \mathcal{X} \to \mathcal{Y}$ *is a linear operator, where* $\mathcal{X}$ *and* Y *are convex compact subsets of vector spaces, and let* $\mathrm{Exm}(A)$ *denote the set of extreme points of set A. Then,*

$$\mathrm{Exm}(\mathrm{Im}(\mathbf{E})) \subset \mathbf{E}(\mathrm{Exm}(\mathcal{X})).$$

Proof. Suppose not; then, there exists an extreme point e of $\mathrm{Im}(\mathbf{E})$ that is not an image of any extreme point of $\mathcal{X}$. Since $e \in \mathrm{Im}(\mathbf{E})$, let us write $e = \mathbf{E}(x)$ for some x that is not extreme in $\mathcal{X}$. Since $\mathcal{X}$ is convex and compact, by Choquet's theorem, we can write $x = \int y d\omega(y)$, where ω is a probability measure supported on $\mathrm{Exm}(\mathcal{X})$. But then, by linearity of $\mathbf{E}$, we have $e = \int \mathbf{E}(y) d\omega(y)$. Since e is an extreme point, the measure ω must be degenerate and $e = \mathbf{E}(y)$ for some extreme point y of $\mathcal{X}$. A contradiction. $\square$

Take $\mathcal{X}$ to be the set of non-decreasing allocation rules on $[\underline{\theta}_i, \bar{\theta}_i]$. This is a convex and compact subset of a vector space (when we endow it with the L^1 topology) with extreme

points that are cutoff allocation rules: $x(\theta) = \mathbf{1}_{\{\theta \geq \theta_q\}}$ for $q \in [0, 1]$, where θ_q denotes the q-quantile of the distribution F_i (using the fact that F_i is a continuous distribution). For a compact convex set A in a Euclidean space, the dimension of A is equal to the dimension of $\mathrm{Exm}(A)$. Hence, we have

$$\dim(\mathrm{Im}(\mathbf{E}^i)) = \dim(\mathrm{Exm}(\mathrm{Im}(\mathbf{E}^i))) \leq \dim(\mathbf{E}^i(\mathrm{Exm}(\mathcal{X}))) = \dim(\mathrm{Im}(E^i)) \leq \dim(\mathrm{Im}(\mathbf{E}^i)),$$

where the first inequality follows from Lemma 2, the second equality follows from the characterization of extreme points of $\mathcal{X}$ and the definition of the operator E^i, and the second inequality follows from the fact that the operator E^i is effectively a restriction of the operator $\mathbf{E}^i$ to cutoff allocation rules. Thus, we have $\mathrm{cplx}(E^i) = \mathrm{cplx}(\mathbf{E}^i)$.

Let $k^i = \mathrm{cplx}(E^i) + 1 = \mathrm{cplx}(\mathbf{E}^i) + 1$. After a change of basis (multiplying the left- and right-hand side of the above equality by an appropriate matrix), we can assume that $\mathbf{E}^i$ has an image contained in $\mathbb{R}^{k^i-1}$, and that $c \in \mathbb{R}^{k^i-1}$ (the remaining coefficients are always zero, so they can be dropped).

We now restate (in a slightly modified form) a result from Kang (2023), which can be seen as an extension of Carathéodory's Theorem to infinite-dimensional spaces based on the contributions of Bauer (1958) and Szapiel (1975) (see also Dubins (1962)):

Theorem 0. *Let $\mathcal{X}$ be a convex, compact set in a locally convex Hausdorff space, and let $\mathbf{E} : \mathcal{X} \to \mathbb{R}^{k-1}$ and $W : \mathcal{X} \to \mathbb{R}$ be continuous affine functions. Suppose that $L \subset \mathbb{R}^k$ is closed and convex, and that $\mathbf{E}(\mathcal{X}) \cap L \neq \emptyset$. Then, there exists $x^\star \in \mathcal{X}$ such that $W(x^\star) = \max_{x \in \mathcal{X}: \mathbf{E}(x) \in L} W(x)$ and*

$$x^\star = \sum_{j=1}^k \alpha_j x_j, \quad \sum_{j=1}^k \alpha_j = 1, \quad \alpha_j \geq 0, \quad and \ x_j \in \mathrm{Exm}(\mathcal{X}), \ \forall j = 1, \ldots, k,$$

where $\mathrm{Exm}(\mathcal{X})$ denotes the set of extreme points of $\mathcal{X}$.

Intuitively, Theorem 0 states that the problem of maximizing a linear objective function over a compact convex set $\mathcal{X}$ subject to $k - 1$ linear constraints admits a solution that is a convex combination of at most k extreme points of $\mathcal{X}$.

We can apply Theorem 0 to our problem by letting $\mathcal{X}$ be the set of non-decreasing allocation rules, $\mathbf{E}$ be $\mathbf{E}^i$, L to consist of the single point $c \in \mathbb{R}^{k^i-1}$, and $W(x) = \int W_i(\theta)x(\theta)dF_i(\theta)$. By Theorem 0, there exists a solution to the auxiliary problem that is a convex combination of at most k^i extreme points of the set of non-decreasing allocation rules; that is, there exists

a solution to (8) that takes the form

$$x_i(\theta) = \sum_{j=1}^{k^i} \alpha_j \mathbf{1}_{\{\theta \geq \theta_j\}},$$

for some weights $\alpha_j \geq 0$ adding up to 1 and cutoff types θ_j, $j = 1, .., k^i$. Applying the same argument for all $i \in I$ establishes existence of a solution in which there is a k^i-tiered pricing mechanism for every group. Additionally, the top tier is granted full access: $x_i(\theta) = 1$ for $\theta \geq \theta_j$ (which does not preclude the possibility that $\theta_j = \bar{\theta}$ in which case removing the top tier from the mechanism would not affect the designer's expected payoff).

To prove the second part of the theorem, we consider an auxiliary problem in which we fix all the externality values v_i but reoptimize over all allocation rules jointly:

$$\max_{x_1, \ldots, x_N} \quad \sum_{i \in I} v_i \int W_i(\theta) x_i(\theta) \mu_i dF_i(\theta)$$

$$\text{s.t. } v_i = \nu_i \left(\sum_{j \in I} \int e_{j \to i}(\theta) x_j(\theta) \mu_j dF_j(\theta) \right), \qquad \forall i \in I,$$

$$x_i(\theta) \text{ is non-decreasing}, \qquad \forall i \in I.$$

The problem can be rewritten, for some vector of constants $c \in \mathbb{R}^N$, as

$$\max_{x_1, \ldots, x_N} \quad \sum_{i \in I} v_i \int W_i(\theta) x_i(\theta) \mu_i dF_i(\theta)$$

$$\text{s.t. } c = \sum_{j \in I} \mu_j \mathbf{E}^j(x_j),$$

$$x_i(\theta) \text{ is non-decreasing}, \qquad \forall i \in I.$$

By assumption and Lemma 2 (applied in the same way as in the previous case), we have that $K = \dim(\text{Im}(\tilde{\mathbf{E}})) + 1$ where $\tilde{\mathbf{E}}(x_1, \ldots, x_N) := \sum_{j \in I} \mu_j \mathbf{E}^j(x_j)$. After a change of basis, we can assume that $\tilde{\mathbf{E}}$ has an image contained in $\mathbb{R}^{K-1}$ and $c \in \mathbb{R}^{K-1}$.

We know from the first part of the proof that there exists an optimal solution in which each allocation rule x_i can be written as $x_i(\theta) = \sum_{j=1}^N \alpha_j^i \mathbf{1}_{\{\theta \geq \theta_j^i\}}$, for some weights $\alpha_j^i \geq 0$ (adding up to 1) and cutoff types θ_j^i, $j = 1, .., N + 1$ (where we have used the fact that $k^i \leq N + 1$). Fixing the cutoff types, and the value $\bar{V}$ of the optimal payoff for the designer,

the optimal weights α_j^i must satisfy the following system of inequalities:

$$\bar{V} = \sum_{i \in I} \sum_{j=1}^{N+1} \alpha_j^i \, v_i \, \tilde{W}_j^i,$$

$$c_k = \sum_{i \in I} \sum_{j=1}^{N+1} \mu_i \, \alpha_j^i \, \tilde{e}_{i,j,k}, \quad \forall k = 1, \ldots, K-1,$$

$$1 = \sum_{j=1}^{N+1} \alpha_j^i, \quad \forall i \in I,$$

$$0 \leq \alpha_j^i, \quad \forall j = 1, \ldots, N+1, \, i \in I,$$

for some constants $\tilde{W}_j^i$ and $\tilde{e}_{i,j,k}$ that do not depend on the weights α_j^i. Moreover, any set of weights that satisfies this system of equations defines an optimal solution. A solution to the system exists (because there exists an optimal solution to the designer's problem that takes this form). The system of linear equations has $N(N+1)$ variables and $N+K$ equality constraints. By the fundamental theorem of linear algebra, there exists a solution to the system of equations in which (at most) $N+K$ weights are non-zero. It follows that setting $N+K$ cutoff types (equivalently, prices) optimally allows the designer to achieve optimality in her optimization problem.

A.2 Proof of Proposition 1

Consider the auxiliary problem of reoptimizing over x_i subject to keeping the expected externalities exerted by group i on other groups at least as large as in some initial mechanism. Imposing a weak inequality (as opposed to an equality, as in the proof of Theorem 1) is sufficient because we assumed that *(i)* all functions ν_j are non-decreasing, and *(ii)* the designer's payoff is non-decreasing in the values v_j (since $\int W_j(\theta) x_j(\theta) dF_j(\theta) \geq 0$ for all non-decreasing allocation rules x_j). The auxiliary problem thus becomes:

$$\max_{x_i} \quad \int W_i(\theta) x_i(\theta) \mu_i dF_i(\theta) \tag{9}$$

$$\text{s.t. } c_j \leq \int e_{i \to j}(\theta) x_i(\theta) \mu_i dF_i(\theta), \quad \forall j \in I,$$

$$x_i(\theta) \text{ is non-decreasing,}$$

for some constants c_j. We can assign Lagrange multipliers $\lambda_j \geq 0$ to the inequality constraints,[32] and then the problem is to maximize a Lagrangian

$$\int \left(W_i(\theta) + \sum_{j \in I} \lambda_j e_{i \to j}(\theta) \right) x_i(\theta) \mu_i dF_i(\theta).$$

By assumption that $W_i(\theta)$ is strictly increasing and $e_{i \to j}(\theta)$ is non-decreasing in θ, for all j, $W_i(\theta) + \sum_{j \in I} \lambda_j e_{i \to j}(\theta)$ is strictly increasing in θ for any set of Lagrange multipliers. It follows that there exists an optimal solution in which $x_i(\theta)$ is a threshold (one-price) allocation rule.

A.3 Supplementary Materials for Section 4.1

In this appendix, we provide supplementary calculations for Results 2 and 3.

First, we note that the first-order conditions for optimality of $\theta_0^\star$ and $\theta_1^\star$ for problem (2) read

$$\frac{J(\theta_0^\star)}{-e(\theta_0^\star)} = \frac{J(\theta_1^\star)}{-e(\theta_1^\star)} = \frac{R}{V}.$$

Thus, we cannot have two non-degenerate membership tiers when $J(\theta)/(-e(\theta))$ is strictly increasing. In particular, this shows that a 1-tiered pricing mechanism is optimal under the assumption of Result 2. The rest of Result 2 follows from standard analysis of first-order conditions for maximizing objective (2) when θ_1 is set to $\bar\theta$. The derivative of objective (2) with respect to x_0 is

$$x_0 \int_{\theta_0}^{\bar\theta} e(\theta) dF(\theta) \int_{\theta_0}^{\bar\theta} J(\theta) dF(\theta) + \left(v_0 + x_0 \int_{\theta_0}^{\bar\theta} e(\theta) dF(\theta) \right) \int_{\theta_0}^{\bar\theta} J(\theta) dF(\theta).$$

Solving for x_0 yields the second condition characterizing the solution in Result 2.

To obtain Result 3, suppose that $\theta_1^\star$ is the optimal cutoff type if the designer restricts attention to posted-price mechanisms. We can then consider how the expected payoff of the designer changes under the following perturbation of the mechanism: For some type $\theta_0 < \theta_1^\star$, allocate access to types $\theta \in [\theta_0, \theta_1^\star]$ with some small probability $\epsilon > 0$. This perturbation is profitable for small enough ϵ if

$$\int_{\theta_0}^{\theta_1^\star} e(\theta) dF(\theta) \cdot R + V \cdot \int_{\theta_0}^{\theta_1^\star} J(\theta) dF(\theta) > 0 \iff \frac{J(\theta_1^\star)}{-e(\theta_1^\star)} < \frac{\int_{\theta_0}^{\theta_1^\star} J(\theta) dF(\theta)}{\int_{\theta_0}^{\theta_1^\star} (-e(\theta)) dF(\theta)},$$

---

[32]The existence of Lagrange multipliers follows from Theorem 2.165 in Bonnans and Shapiro (2000) as long as c_j, for each $j \in I$, is in the interior of all possible values that the right-hand side can take as $x_i(\theta)$ varies. Whenever this is not the case, $x_i(\theta)$ must take the form of a threshold allocation rule, which is what we want to show.

where we have relied on the first-order condition for $\theta_1^\star$ to substitute $R/V = J(\theta_1^\star)/(-e(\theta_1^\star))$.

A.4 Supplementary Materials for Section 4.2

In this appendix, we briefly explain how our arguments extend to the case when the value functions ν_i are common to all groups, $\nu_i = \nu$ for all $i \in I$, but ν is a function of multiple arguments: $v = \nu(s_1, \ldots, s_K)$, where each s_k depends linearly on the allocation rule to each group.

To extend the first part of Theorem 1 to this case (in order to obtain Result 5), note that we can modify the proof of Theorem 1 by fixing the values of each of the inputs s_k. Doing so implies that the auxiliary problem (8) of optimizing over an allocation rule for a single group will have K linear constraints. Thus, by the same argument as in the proof of Theorem 1, we need at most $K + 1$ prices in the optimal mechanism. This gives us Result 5.

Result 4 is implied by an analogous modification applied to the proof of Proposition 1 (instead of imposing a constraint on the values v_j for all groups j, we impose a constraint on the inputs s_k to the single value function ν). Note that the designer's payoff is non-decreasing in the value v because—under the objective of weighted revenue and surplus—the designer's payoff from every group is always non-negative.

A.5 Supplementary Materials for Section 4.3

To derive the pricing formula (4), we first rewrite the designer's maximization problem relying on the fact that—by Proposition 1—optimal allocation rules take the form $x_i(\theta) = \mathbf{1}_{\{\theta \geq \theta_i\}}$ for some cutoff types θ_i, for all $i \in I$:

$$\max_{\{\theta_i\}_{i \in I}} \sum_{i \in I} \left(\mu_i \left(v_i^0 + \sum_{j \in I} \mu_j \int_{\theta_j}^1 e_{j \to i} dF_j(\theta) \right) \left(\int_{\theta_i}^1 J_i(\theta) dF_i(\theta) \right) \right).$$

Assuming the first-order approach is valid (i.e., that we have an interior solution), the first-order condition with respect to θ_i yields

$$-\mu_i v_i J_i(\theta_i) f_i(\theta_i) - \sum_{j \in I} \mu_j \mu_i e_{i \to j} f_i(\theta_i) R_j = 0,$$

where v_i and R_j are defined as in the main body of the paper. Rearranging and using the equality $p_i = \theta_i v_i$ yields the pricing equation (4): $p_i = \gamma_i v_i - (1 - \gamma_i) \sum_{j \in I} e_{i \to j} \mu_j R_j$.

To determine when participation by group i is a substitute or a complement for participa-

tion by group j, we employ the implicit function theorem to calculate the partial derivative of θ_i with respect to θ_j, when θ_i satisfies the pricing formula (4) and θ_j is treated as a free parameter. Holding fixed θ_k for $k \neq \{i, j\}$, denote the objective function in the designer's maximization problem by $\mathcal{W}(\theta_i, \theta_j)$. The accompanying first-order condition for θ_i can then be rewritten as $\mathcal{W}'_{\theta_i}(\theta_i, \theta_j) = 0$. Then, using the implicit function theorem,[33] we get that

$$
\begin{aligned}
\frac{\partial \theta_i}{\partial \theta_j} &= -\frac{-(1-\gamma_i)\left(-\frac{\partial v_i}{\partial \theta_j}\frac{1}{v_i^2}\sum_k e_{i \to k}\mu_k R_k + \frac{1}{v_i}e_{i \to j}\mu_j \frac{\partial R_j}{\partial \theta_j}\right)}{\mathcal{W}''_{\theta_i \theta_i}(\theta_i, \theta_j)} \\
&= -\frac{(1-\gamma_i)\mu_j f_j(\theta_j)\left(-\frac{1}{v_i^2}e_{j \to i}e_{i \to R} + \frac{1}{v_i}e_{i \to j}J_j(\theta_j)\right)}{\mathcal{W}''_{\theta_i \theta_i}(\theta_i, \theta_j)}.
\end{aligned}
$$

Note that the second-order condition (which must hold at an interior solution that we are assuming) implies that the denominator of the preceding expression, $\mathcal{W}''_{\theta_i \theta_i}(\theta_i, \theta_j)$, is negative. Therefore, we get that participation by group i is a substitute for participation by group j if and only if

$$
e_{i \to j} \cdot J_j(\theta_j) - \frac{1}{v_i} \cdot e_{j \to i} \cdot e_{i \to R} < 0.
$$

Additionally, since $J_j(\theta)$ is non-decreasing, we have that $J_j(\theta) \leq 0$ if and only if θ_j is below the revenue-maximizing level γ_j (absent externalities). Thus, by Result 6, the sign of $J_j(\theta)$ is opposite to the sign of the aggregate revenue externality $e_{j \to R}$.

---

[33]Formally, in order to do so, we need to assume that $\mathcal{W}''_{\theta_i \theta_i}(\theta_i, \theta_j) \neq 0$.