Optimal Membership Design*

Piotr Dworczak ⓡ Marco Reuter ⓡ Scott Duke Kominers ⓡ Changhwa Lee†

March 17, 2026

Abstract

The value of a good to an individual often depends on *average* characteristics of other agents consuming it. For example, the value of a concert ticket depends on the composition of the audience; the value of attending college depends on the structure of its student body; and the value of a blockchain token depends on the types of owners. We call the sale of such goods *membership design*. We show that the optimal sales mechanism for membership goods offers distinct tiers—differing in prices and level of access—with the number of membership tiers reflecting the number of distinct moments summarizing the distribution of relevant member characteristics. This insight provides market design guidance for selling membership goods and helps explain a number of mechanisms used in such markets in practice.

*An abstract of this work appeared in the *Proceedings of the 26th ACM Conference on Economics and Computation*. We appreciate helpful comments from Shai Bernstein, Joseph Burleson, Sonal Chokshi, Michael Crystal, Kate Dellolio, Tom Eisenmann, Andrea Galeotti, Carl-Christian Groh, Faruk Gul, Kerry Herman, Ebehi Iyoha, Valet Jones, Steve Kaczynski, Ray Kluender, PL, Eddy Lazzarin, Kevin Leyton-Brown, Shengwu Li, Taylor Lundy, Julien Manili, Das Narayandas, Chris Peterson, Narun Raman, Tim Roughgarden, SAFA, Tim Sullivan, Marek Weretka, Liang Wu, Refine.ink, and presentation audiences at a16z crypto, HBS, EC'25 and Tokenomics'23. Dworczak and Lee gratefully acknowledge support received under the ERC Starting grant IMD-101040122. Kominers gratefully acknowledges support from the Digital Data Design (D^3) Institute at Harvard and the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications. Part of this work was conducted during the Simons Laufer Mathematical Sciences Institute Fall 2023 program on the Mathematics and Computer Science of Market and Mechanism Design, which was supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778. Kominers is a Research Partner at a16z crypto, which reviewed a draft of this article for compliance prior to publication and is an investor in crypto projects and online platforms (for general a16z disclosures, see https://www.a16z.com/disclosures/). Notwithstanding, the ideas and opinions expressed herein are those of the authors, rather than of a16z or its affiliates. Kominers also holds digital assets, including both fungible and non-fungible tokens, and advises a number of companies on marketplace and incentive design. The views expressed herein are those of the authors and do not necessarily represent the views of the EU or ERC, or the IMF, the IMF Executive Board, or IMF management. Any errors or omissions remain the sole responsibility of the authors.

†*Dworczak:* Department of Economics, Northwestern University; and Group for Research in Applied Economics. *Reuter:* International Monetary Fund. *Kominers:* Entrepreneurial Management Unit, Harvard Business School; Department of Economics and Center of Mathematical Sciences and Applications, Harvard University; and a16z crypto. *Lee:* University of Bristol.

1 Introduction

Many goods have the property that their value to any individual consumer depends on the average characteristics of other consumers. This is true of academic programs and professional organizations, where peer effects are important; of concerts and other performances, where the atmosphere depends on the audience; and of online platforms and social networks, where user base shapes the experience. Emphasizing that access to such goods entails joining a set of agents whose composition affects the value of participation, we refer to them as *membership goods*.

In this paper, we introduce a simple framework for studying *optimal membership design*—the sale of membership goods—in a setting where agents differ along both observable and unobservable characteristics. Our key modeling 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 types and chooses an individually rational and incentive compatible mechanism to maximize an objective such as revenue or welfare.

We show that an optimal membership sale can always be implemented using a tiered pricing mechanism, with distinct price schedules for different groups (i.e., sets of agents with the same observable characteristics). Tiered pricing allows the designer to achieve the optimal composition of members within groups by using self-selection to screen for unobserved characteristics associated with membership externalities. Moreover, the number of pricing tiers is directly linked to the complexity and structure of externalities.

In particular, a single price is optimal for any group that does not generate allocative externalities. When a group exerts a uniform externality—i.e., when its magnitude is constant across agents' private types—a two-price scheme is optimal: a high price granting full access and a discounted price associated with a lower level of access. Finally, in the general case, in which externalities vary arbitrarily across types and groups, the number of pricing tiers required for optimality may be as large as the number of groups plus one.

As we illustrate through a series of applications, our framework rationalizes a number of mechanisms used in practice to sell membership goods, including lotteries and group-specific discounts. More broadly, it provides market-design guidance on how the optimal membership sales mechanism should adjust to differences in market objectives, buyer characteristics, and the structure of membership externalities.

As a leading example, consider selling tickets to an event. Absent concerns about audience composition, in our setting, it is optimal to charge a single monopoly price. However, when the quality of the event depends on the composition of attendees, tickets become a membership good—and our main result predicts that introducing a second pricing tier might be optimal. To see why, suppose that the participants contributing the most to the quality of the event cannot be readily identified based on observables and have relatively low willingness to pay (WTP). Then, it becomes optimal to allocate a share of tickets at a discounted price with rationing, while charging a higher price for guaranteed access. The lower tier secures—through self-selection—the presence of low-WTP high-externality attendees who enhance the value of the event, while the higher tier extracts that value from participants with the highest WTP.[1] This logic helps explain why pop stars like Taylor Swift and sports organizations like UEFA often allocate a fraction of their tickets via lotteries—appearing to leave money on the table. Our framework shows that such mechanisms may in fact be revenue-maximizing—even in the absence of dynamic or reputational considerations—when the highest-externality participants ("megafans") have low ability to pay (e.g., teenage fans of Taylor Swift or diehard football fans from the local community).

Membership design offers a unifying perspective across a range of additional applications, from college admissions and token-based networks to ad auctions and brand marketing. We briefly discuss the implications of our analysis for these settings in Section 4.

1.1 Related literature

There is a sizable mechanism design literature studying allocation problems with externalities. One class of models, introduced by Jehiel et al. (1996, 1999), studies allocation to a finite number of agents exerting fully flexible individual-specific externalities typically assumed to be unobserved.[2] A more recent class of models—studied, for example, by Kang (2020), Pai and Strack (2022), Ostrizek and Sartori (2023), and Akbarpour ⓡ al. (2024)—focuses instead on a large population of agents contributing (potentially heterogeneously) to a single aggregate externality, which often enters the designer's objective function additively.

The average externalities associated with membership design put our framework in between these two classes. 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), which makes membership design tractable. 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' WTP. The resulting set of applications also differs: the first

[1]By contrast, if observable characteristics identify the relevant externalities, group-specific pricing is optimal; and if agents with highest WTP also exert the highest positive externalities (as is arguably the case, e.g., for upscale events such as gala dinners), screening via a price mechanism is aligned with maximizing the event value, so a single access tier remains optimal.

[2]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.

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 in 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 event tickets, determining access to digital communities, building multi-sided platforms, or designing policies with compositional effects such as school admission.[3]

Allocative externalities similar to those we study appear in the classical literature on two-sided buyer-seller platforms (see Jullien et al. (2021) for a survey) and on sales with network effects (see, e.g., Candogan et al. (2012), Fainmesser and Galeotti (2015), Jadbabaie and Kakhbod (2019), and Carpizo and Crystal (2024)); most papers in these literatures do not adopt a mechanism-design approach, focusing on questions such as equilibrium market dynamics.[4] 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 studying optimal groupings of agents with social status externalities (Moldovanu et al. (2007); Board (2009); Rayo (2013)).

2 Model

2.1 Membership goods

Membership in our setting is represented by owning a particular good, which can be thought of as directly conveying access (like a ticket or membership card) or representing a right to participate (like admission to an academic program). Members (i.e., owners) experience externalities as a function of the average characteristics of the full set of owners.

Each agent is characterized by a *private type* θ, an externality vector e, and a publicly observable *label* i. An agent's *level of access* to the good is described by a variable $x \in [0, 1]$, and the agent can be charged a transfer $t \in \mathbb{R}$ for that access. Our baseline interpretation of x is that it represents the probability of becoming a member—but it could also reflect the quality of access (e.g., which specific membership features are available).

[3] 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.

[4] 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. Jadbabaie and Kakhbod (2019) examined how the presence of network externalities distorts the optimal contract offered by a profit-maximizing seller.

There are finitely many labels $i \in I$, and each group i has mass $\mu_i > 0$, with $\sum_{i \in I} \mu_i = 1$. Conditional on label i, agents' types θ are distributed according to a cumulative distribution function F_i with a continuous density f_i fully supported on $[\underline{\theta}_i, \bar{\theta}_i]$. Conditional on each label i and type θ, the externality vector $e = (e_1, ..., e_{|I|}) \in \mathbb{R}^{|I|}$ is distributed according to $G(e|i, \theta)$, with marginals denoted by $G_j(e_j|i, \theta)$, for each $j \in I$.

The utility an agent receives from membership depends on that agent's type θ and level of access, as well as the average characteristics of the set of members. The simplifying assumption we make is that heterogeneity in WTP for access can be decomposed into a vertical component captured by the type θ and a horizontal component that is the same for all agents with the same label i. Specifically, an agent with label i and type θ has WTP for access equal to θv_i, where v_i is common to all agents in group i.

The value v_i is determined by the members' externalities. An agent with externality vector e exerts an externality e_j on group j (when allocated the object). Concretely, if $x_j(\theta)$ denotes the level of access of agents with label j and type θ, then

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

where $\nu_i : \mathbb{R} \to \mathbb{R}$ is continuous, non-negative, and non-decreasing, for all $i \in I$.[5]

We refer to $e_{j \to i}(\theta)$ as the *expected externality* exerted by agent (j, θ) on group i. Because the membership externality that a given agent exerts on others does not affect that agent's payoff, the true externality vector e cannot be elicited in an incentive-compatible way (even if the agent knows their own externality on others, which we have not assumed).[6] Thus, the statistical object $e_{j \to i}(\theta)$ describes the relevant input into the membership design problem.

Finally, we assume utilities are quasi-linear in the monetary transfer and linear in level of access: Allocation (x, t) gives an agent with (i, θ) utility

$$\theta v_i x - t.$$

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

[6]See Jehiel et al. (1996) for a formal statement and proof of this observation.

2.2 The membership design problem

We consider a *designer* who allocates membership access to agents (*potential members*). To focus on the novel aspects of our framework, we assume that the designer does not incur marginal costs for allowing access and does not face an explicit capacity constraint.[7]

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.[8] She knows the population distribution of agents' characteristics and observes the public label i for each agent (but not the individual type θ or externality vector e). 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:

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

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

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

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

Following the standard approach in mechanism design (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:[9]

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

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

where $W_i(\theta) \equiv \tilde{W}_i(\theta) + \alpha J_i(\theta)$, and $J_i(\theta) \coloneqq \theta - (1 - F_i(\theta))/f_i(\theta)$ is the virtual surplus

[7] These features can be added to the model without changing our qualitative results in any substantial way.

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

[9] 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.

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

The presence of externalities significantly complicates the designer's problem. Keeping the value of membership constant, assigning access $x_i(\theta)$ to an agent allows the designer to realize the payoff $v_i W_i(\theta)$. At the same time, however, allocating access affects the composition of buyers and thus the value of membership. Mathematically, this implies that the problem is non-linear in the allocation rule. The key observation we will use to derive structural properties of the optimal mechanism is that—because externalities affect membership value only via average characteristics—the problem becomes linear after conditioning on (a finite number of) moments of the distribution of member characteristics.

3 Tiered-pricing as an optimal membership mechanism

We now show that an optimal mechanism in our setting always takes the form of a tiered-pricing mechanism.

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

A k-tiered pricing mechanism can be implemented by offering agents in a given group k different prices for k distinct levels of access; we show that the number of tiers in the optimal mechanism depends on the complexity of the externalities, in a precise sense we define next.

3.1 Measuring the complexity of externalities

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

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

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-WTP 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 image of the operator E^i; formally,

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

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

For intuition, consider the case in which 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. If externalities are present but 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 js. 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 observation follows directly from the definition of complexity, and formalizes the preceding intuition.

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

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

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

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 the moments representing the statistical dependence of these externalities on agents' types. The mechanism designer, having access to high-dimensional data on agents' types and externalities (collected, e.g., through A/B testing or other empirical methods), reduces the dimensionality of the data using a regression model to estimate the relationship between types and externalities. The more distinct moments are needed to account for the relevant externalities exerted by group i, the higher the complexity of the operator E^i. For example, performing a regression on two moments summarizing the externalities within a group would result in an estimated operator E^i having complexity 2. Fitting a regression model with k moments 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 Constructing an optimal tiered-pricing mechanism

Theorem 1. *It is optimal to use a k^i-tiered pricing mechanism for each group i, where $k^i \leq \text{cplx}(E^i) + 1$, and the top tier provides full access ($x = 1$).*

The proof of Theorem 1 is presented in Section 5. The argument is mathematically intuitive: While problem (1) faced by the designer is non-linear in the allocation rule x_i for group i, it becomes linear if we fix (and treat 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 extension of Carathéodory's theorem yields that the optimal mechanism is a k^i-tiered pricing mechanism with $k^i \leq \text{cplx}(E^i)+1$ prices. Moreover, the top price can always be assumed to be associated with full access.[11]

The 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, the total number of prices charged by the optimal mechanism can be upper bounded by a number that is (in some cases strictly) smaller than the sum of upper bounds from Theorem 1.

Remark 1. *It is optimal to use a k^i-tiered pricing mechanism for each group i, with*

$$\sum_{i \in I} k^i \leq |I| + \text{cplx}\Big(\sum_{i \in I} \mu_i E^i\Big),$$

where $\sum_{i \in I} \mu_i E^i$ is treated as a function from $[0,1]^{|I|}$ to $\mathbb{R}^{|I|}$.

3.3 Economic implications

Theorem 1 establishes how the structure of the externalities across groups impacts the optimal tiered-pricing mechanism design. Here, we refine the result to clarify the economic

[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. Corollary 2 shows that this possibility may indeed arise at optimum.

intuition behind the theorem.

We call (x, t) a *posted-price mechanism* if it offers a single price for full access, i.e., $x(\theta) = \mathbf{1}_{\theta \geq \theta_0}$, for some θ_0; we call (x, t) a *uniform-lottery mechanism* if it offers the same level of access to all types, i.e., $x(\theta) = x_0$, for some x_0; finally, we say that (x, t) features *rationing* if $x(\theta) = x_0 \in (0, 1)$ for some interval of types θ.

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

Corollary 1 confirms that it is the presence of externalities that motivates the introduction of additional pricing tiers.[12] However, Theorem 1 only provides an upper bound on the number of tiers, without explaining when they are actually needed. To cast more light on this issue, we consider the case of a single group (and thus drop the subscript i). Theorem 1 immediately implies that the optimal mechanism will offer at most two pricing tiers: high-price guaranteed membership and a discounted membership subject to rationing.[13]

Corollary 2. *Suppose that there is only one group, and that ν is differentiable. Then:*

1. *If $W(\theta)$ and $e(\theta)$ are both non-decreasing, a posted-price mechanism is optimal;*
2. *If $W(\theta)$ and $e(\theta)$ are both non-increasing, a uniform-lottery mechanism is optimal;*
3. *If $\theta^\star$ is the cutoff type at the optimal posted-price mechanism (i.e., the interior type that is indifferent between buying the good and not), then introducing a second tier with rationing is optimal if there exists $\theta_0 \neq \theta^\star$ such that*

$$\frac{W(\theta^\star)}{e(\theta^\star)} < \frac{\int_{\theta_0}^{\theta^\star} W(\theta) dF(\theta)}{\int_{\theta_0}^{\theta^\star} e(\theta) dF(\theta)}, \tag{2}$$

with $\int_{\theta_0}^{\theta^\star} e(\theta) dF(\theta) > 0$, or if the reverse inequality holds with $\int_{\theta_0}^{\theta^\star} e(\theta) dF(\theta) < 0$.

The first two cases of Corollary 2 are intuitive, and correspond to scenarios in which the two channels affecting the designer's objective are aligned. If the designer achieves a higher payoff from allocating to higher types (fixing the base value of membership) and the base value of membership is increasing in the types of members, then it is optimal to admit agents with high-enough types—corresponding to a posted-price mechanism. If the opposite

[12] Of course, tiered pricing would also arise if agents' values or the 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.

[13] The corollary is a straightforward consequence of the proof of Theorem 1—for completeness, we provide a proof in Appendix A.3.

is true, then the designer would ideally like to admit agents with low-enough types, and incentive-compatibility then implies that the optimal allocation is constant—corresponding to a uniform lottery. In both instances, a single tier of membership is sufficient.

In general, however, there can be a trade-off between optimal allocation of membership for a fixed base value and maximizing the base value of membership. This trade-off is measured by the ratio $W(\theta)/e(\theta)$ of *the designer's direct payoff from allocating to type θ* to *type θ's externality contribution*. Analysis of first-order conditions reveals that when the ratio is monotone (whenever it is negative), it is still optimal to offer a single membership tier. But when the ratio is non-monotone, introducing a second tier with rationing could relax the trade-off. To see why, observe that regulating access through a price increase has a different effect than regulating access through rationing: An increase in the price excludes the *marginal* attendee, while rationing excludes the *average* attendee. Inequality (2) captures exactly this effect: It implies that the ratio of the designer's payoff W to the externality contribution e can be improved (relative to the optimal posted-price mechanism) by admitting a range of agent types with interior probability—making a second membership tier with rationing optimal.

Because condition (2) involves an endogenous variable $\theta^\star$, we present a simple parametric example to demonstrate that the condition is not vacuous.

Example 1. Suppose that there is only one group; F is the uniform distribution on $[0, 1]$; ν is the identity function; and for some $\epsilon > 0$,

$$e(\theta) = \epsilon + \frac{\mathbf{1}_{\{\theta \leq \epsilon^2\}}}{\epsilon^2}.$$

The objective of the designer is to maximize revenue.

When ϵ is small enough, the optimal posted-price mechanism has a price of $(2/9)\epsilon$, which makes type $\theta^\star = 1/3$ indifferent between buying and not and achieves the revenue of $(4/27)\epsilon$. Intuitively, although agents with low WTP would increase the value of membership for others, including them would require dropping the price to very low levels.

However, a tiered pricing mechanism can achieve a revenue of more than $1/16$ by offering an access level of $1/2$ for free and access level of 1 priced at $1/8$, which already implies that a two-tier mechanism is optimal. Indeed, inequality (2) holds with $\theta^\star = 1/3$ and $\theta_0 = 0$ for ϵ small enough, reflecting that fact that rationed access by low types can increase the value of membership without eroding revenue to zero.[14] $\square$

[14]Since $W(\theta) = 2\theta - 1$, the left side of inequality (2) is equal to $-1/(3\epsilon)$, while the right side evaluated at $\theta_0 = 0$ is equal to $\int_0^{1/3} (2\theta - 1)d\theta / (\epsilon/3 + 1)$, which is larger than $-1/(3\epsilon)$ for sufficiently small ϵ.

4 Applications

Many real-world allocation problems can be seen as membership design. Here, we discuss a number of examples, illustrating how our framework explains and provides market design guidance for the selling mechanisms used in these contexts.

Ticket sales. Our analysis casts light on 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 (Corollary 1). However, for many events, the presence of specific audience members influences the atmosphere—an integral component of the overall experience. For example, the excitement of a live football match is notably heightened by diehard fans cheering for their teams; similarly, "Swifties" are important to the atmosphere at Taylor Swift concerts. (And note that in such contexts, in general the impact depends on the fan composition of the audience rather than which specific individuals attend.) Theorem 1 indicates that it may then be optimal to offer two pricing tiers: A higher price with guaranteed access and a below-market-clearing price subject to rationing. Paradoxically, in this case, selling a significant share of tickets at below-market-clearing prices via lotteries—appearing to leave money on the table—may actually be revenue-maximizing.[15]

Adding the lottery option is optimal when the fans that contribute most positively to the atmosphere have relatively low ability to pay (and cannot easily be identified based on observables). Intuitively, if fans with the most positive externality also have the highest WTP, then a posted-price mechanism remains optimal (part 1 of Corollary 2). If the opposite is true, however, the reduced-price option allows the seller to ensure participation by fans generating positive externalities, which increases the WTP of all audience members—and thus the seller's revenue (part 3 of Corollary 2; see also Appendix B.1 for an extended discussion).

Our findings closely align with the intuitions of some industry practitioners: for example, a US Government Accountability Office (2018, p. 8) report on ticket pricing and resale listed the following reason 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."

College admissions. Another key membership design context, which we develop more fully in Appendix B.2, is in college admissions where students' values for attending depend on

[15] This is true even absent any dynamic or reputational considerations that have been suggested as a possible explanation by previous work; see Appendix B.1 for a discussion.

diversity characteristics of the incoming class. Students differ in unobserved WTP, observed measures of merit (such as test scores), and background characteristics; the college's objective is to maximize a weighted sum of revenue and student welfare.

For intuition, suppose that the college wants to admit a certain number of students coming from economically underprivileged backgrounds. When admissions policies are able to explicitly consider students' backgrounds, the college achieves the diversity objective via targeted admission decisions and scholarships—corresponding to a group-specific posted-price mechanism. But when labels pertaining to background characteristics cannot be used in admissions (as required by recent U.S. court decisions),[16] the college must rely on student self-selection to ensure diversity. Since students coming from economically underprivileged backgrounds tend to have lower WTP due to low ability to pay, the college would have to reduce prices substantially for everyone to ensure enrollment of such students. Our membership design framework implies that it thus may become optimal to bundle admission and funding decisions by offering students the possibility to apply for partially-funded admission that is subject to rationing: When applying, students can be offered a choice between a high-tuition option and a rationed lower-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. In this way, the college relies on correlations between WTP and background characteristics to implement the desired student body composition through self-selection.[17]

NFTs and token-based networks. One of the primary business models around non-fungible tokens (NFTs) leverages those tokens to codify membership in a network associated to a brand or product ecosystem (see Kaczynski and Kominers (November 10, 2021, 2024)). In these contexts, NFT owners gain access to a range of cross-member messaging systems and services, whose value naturally depends on the composition of membership.

Our characterization of the optimal mechanism suggests optimal mechanisms that correspond closely to the primary sales mechanisms seen for these types of NFTs in practice. In particular, when NFTs can be eventually traded in a secondary market, in order to catalyze a membership network with high positive externalities, it may be necessary to initially allocate the tokens at prices significantly below the resale price that would be expected to arise in the case of successful network formation. In Appendix B.3, we show furthermore

[16] See for example *Students for Fair Admissions, Inc. v. President and Fellows of Harvard College* and *v. University of North Carolina* [Supreme Court of the United States, 2023a,b]

[17] In Appendix B.2, we also discuss how these findings help explain recent shifts in college admissions practices, such as the move away from and then back toward standardized testing.

how the pattern of across-group externalities determines whether particular groups behave as substitutes or complements when determining optimal participation.

Blockchain applications often issue fungible crypto tokens (such as Ethereum) to regulate utilization of network services, typically via payment of token-denominated transaction fees. In Appendix B.4, we consider a model where all users generate network congestion, but some users ("potential developers") also exert positive externalities by improving the network's functionalities. If all agents value the token as a function of its future market price, then their beliefs about future value determine WTP for network participation. The (plausible) scenario in which developers with the highest positive externality have intermediate beliefs about future token value is structurally similar to the ticket sales problem—analogously, the optimal mechanism combines lotteries and label-dependent prices.

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.[18]

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. An advertiser's externality typically varies 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.

Our framework 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 their use of LTVs. At the same time, due to Theorem 1, randomization could be useful to ensure a more balanced composition of ads, in lieu of explicit preferential treatment of some bidders.[19]

[18]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)).

[19]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: For example, a Google representative claimed that if the auction did not occasionally flip the winner with the runner-up, then "Amazon would always

Brand image concerns. Membership-like effects arise in sales of branded goods that are consumed in public, such as clothes or cars. Companies that sell such goods must take into account how potential pricing strategies affect the composition of their customer base, as that feeds into the brand perception, which in turn affects consumers' WTP. 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 others, especially those with high status, own) or the "snob effect" (widespread ownership of a good, especially by one's peers, can cause it to lose value).[20]

By Corollary 2, 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). However, by Theorem 1, lotteries at below-market-clearing prices may be optimal if consumers with intermediate WTP contribute most positively to 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")—our framework provides an explanation rationalizing such a sales mechanism.

Additional applications. Composition effects are relevant for the design of affordable housing programs, as the value of a housing unit to potential tenants may depend on the demographic composition of the neighborhood. For example, sales of new apartments in Korea are subject to price caps with an intention of fostering social interactions between residents with various income levels.[21] Furthermore, when direct use of certain demographic characteristics (such as race) is not allowed in social programs, our results suggest a potential role for carefully constructed lotteries.[22]

At an abstract level, allocation of health insurance can also be seen as a problem of membership design. Insured agents exert an externality through their impact on the insurer's budget constraint or availability of services. In such a context, high WTP (reflecting a privately observed high probability of utilizing health care services) may be associated with negative externalities. Theorem 1 predicts that it may then be optimal to pool a subset of agents at an interior level of access (interpreted, for example, as a low quality of service or long waiting time)—resembling some features of public health care programs.

Finally, digital platforms leverage consumer data gathered from their services to enhance

show up on top" (Agius (October 18, 2023)).

[20] See, e.g., Leibenstein (1950), Vigneron and Johnson (1999), and Han et al. (2010), as well as Lundy et al. (2025) for an empirical analysis in the context of "digital fashion" NFTs.

[21] See Yun and Kim (2014).

[22] See Cook et al. (2023) for a discussion of how suboptimally designed rationing schemes may conflict with the designer's redistributive goals.

user experience and increase revenue.[23] In pricing their services, platforms need to consider not only the WTP but also the value of the data that consumers provide. Our results predict that providing randomized access (e.g., via A/B testing) at below-market prices might be optimal if consumers with low or moderate WTP contribute significant data externalities.

5 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:

$$\begin{aligned} \max_{x_i} \quad & \int W_i(\theta) x_i(\theta) dF_i(\theta) \\ \text{s.t.} \quad & c_j = \int e_{i \to j}(\theta) x_i(\theta) \mu_i dF_i(\theta), \quad \forall j \in I, \\ & x_i(\theta) \text{ is non-decreasing.} \end{aligned} \tag{3}$$

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

[23] See, e.g., Bergemann et al. (2022) where externalities arise because consumer data are predictive of others' behavior, or Galperti et al. (2023) where externalities result from optimal information policies used by data platforms.

rules for group i, defined by

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

Then, we can rewrite the first constraint in problem (3) in vector notation as

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

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). We have

$$\dim(\mathrm{Im}(\mathbf{E}^i)) \overset{(1)}{=} \dim(\mathrm{Exm}(\mathrm{Im}(\mathbf{E}^i))) \overset{(2)}{\leq} \dim(\mathbf{E}^i(\mathrm{Exm}(\mathcal{X}))) \overset{(3)}{=} \dim(\mathrm{Im}(E^i)) \overset{(4)}{\leq} \dim(\mathrm{Im}(\mathbf{E}^i)),$$

where (1) follows from the fact that, for any compact convex set A in a Euclidean space, the dimension of A is equal to the dimension of $\mathrm{Exm}(A)$; (2) follows from the fact that—for linear operators—extreme points of the image are a subset of the image of extreme points of the domain (see Lemma 1 in Appendix A.1 for a formal statement and proof); (3) follows from the characterization of extreme points of $\mathcal{X}$ and the definition of the operator E^i; and (4) 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 $\bar{k}^i = \mathrm{cplx}(E^i) + 1 = \mathrm{cplx}(\mathbf{E}^i) + 1$. After a change of basis (multiplying both sides of the equality $c = \mathbf{E}^i(x_i)$ by an appropriate change-of-basis matrix), we can assume that $\mathbf{E}^i$ has an image contained in $\mathbb{R}^{\bar{k}_i-1}$, and that $c \in \mathbb{R}^{\bar{k}_i-1}$ (the remaining coefficients are zero, and 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 $L \subset \mathbb{R}^{k-1}$ 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*

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

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 $c \in \mathbb{R}^{\bar{k}_i - 1}$, 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 $\bar{k}_i$ extreme points of the set of non-decreasing allocation rules; that is, there exists a solution to (3) that takes the form

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

for some weights $\alpha_j \geq 0$ adding up to 1 and cutoff types θ_j, $j = 1, .., \bar{k}_i$. Applying the same argument for all $i \in I$ establishes optimality of a k_i-tiered pricing mechanism for every group, with $k_i \leq \bar{k}_i$. Additionally, the top tier can always be granted full access: If it is not true that $\theta_j < \bar{\theta}_i$ for all j with $\alpha_j > 0$, then we can set $\theta_j = \bar{\theta}_i$ for one such j without affecting the designer's expected payoff.

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.

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.

Bergemann, D., A. Bonatti, and T. Gan (2022): "The economics of social data," *The RAND Journal of Economics*, 53, 263–296.

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, 142–170.

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," *Operations Research*, 60, 883–905.

Carpizo, L. and M. Crystal (2024): "Resource Allocation on Networks of Externalities," Working Paper.

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.

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," *Review of Eco-*

nomic 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.

Jadbabaie, A. and A. Kakhbod (2019): "Optimal Contracting in Networks," *Journal of Economic Theory*, 183, 1094–1153.

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.

Lundy, T., N. Raman, S. D. Kominers, and K. Leyton-Brown (2025): "NFTs as a Data-Rich Test Bed: Conspicuous Consumption and its Determinants," in *Proceedings of* *the ACM Web Conference 2025*, 3451–3465.

Markovich, S. and L. Rayo (2024): "What's on the Menu? Consumption Externalities & Product Quality," Tech. rep., Paris and London, CEPR Discussion Paper No. 19792.

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.

Supreme Court of the United States (2023a): "Students for Fair Admissions, Inc. v. President and Fellows of Harvard College," 600 U.S. 181, no. 20-1199.

——— (2023b): "Students for Fair Admissions, Inc. v. University of North Carolina," 600 U.S. 181, no. 21-707.

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," in *International Conference on Economics, Education and Humanities*, 40–43.

Online Appendix

A Proofs omitted from the main text

A.1 Extreme points and linear mappings

We formalize a step from the proof of Theorem 1.

Lemma 1. *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* $\operatorname{Exm}(A)$ *denote the set of extreme points of set* A*. Then,*

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

Proof. Suppose the contrary; then, there exists an extreme point e of $\operatorname{Im}(\mathbf{E})$ that is not an image of any extreme point of $\mathcal{X}$. Since $e \in \operatorname{Im}(\mathbf{E})$, 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 $\operatorname{Exm}(\mathcal{X})$. 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 attach probability 1 to extreme points y such that $\mathbf{E}(y) = e$; in particular, $e = \mathbf{E}(y)$ for some extreme point y of $\mathcal{X}$. A contradiction. □

A.2 Proof of Remark 1

Recall that $N = |I|$ and let $K = \operatorname{cplx}(\sum_{i \in I} \mu_i E^i)$. By the proof of Theorem 1, we know that there exists an optimal solution in which each allocation rule x_i can be written as $x_i(\theta) = \sum_{j=1}^{N} \alpha_i^j \mathbf{1}_{\{\theta \geq \theta_j^i\}}$, for some weights $\alpha_i^j \geq 0$ (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 θ_j^i, we can consider an auxiliary optimization problem of optimizing over the weights α_i^j subject to preserving the values for all the groups; under the maintained assumptions, this problem can be written as

$$\max_{(\alpha_i^j)_{j=1,\ldots,N+1, i \in I}} \sum_{i \in I} \sum_{j=1}^{N+1} \alpha_i^j \, v_i \, \tilde{W}_j^i,$$

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

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

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

for some constants c_k, $\tilde{W}_j^i$, and $\tilde{e}_{i,j,k}$ that do not depend on the weights α_i^j. The key simplification—the fact that we can keep the value of membership constant for all groups by imposing K equalities—follows from the definition of the constant K (and a change of basis). A solution to the system exists (because there exists an optimal solution to the designer's problem that takes this form). The linear program has $N+K$ equality constraints. It is then a consequence of Carathéodory's theorem that there exists a solution 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.3 Proof of Corollary 2

We first prove parts 1 and 2 of the corollary. Consider the auxiliary problem of reoptimizing over the allocation rule x subject to keeping the expected externality 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 the function ν is non-decreasing and the designer's payoff must be non-decreasing in the value v at any candidate optimal mechanism (since $\int W(\theta)x(\theta)dF(\theta) \geq 0$ at any optimal mechanism). The auxiliary problem thus becomes:

$$\max_{x} \quad \int W(\theta)x(\theta)dF(\theta) \tag{4}$$

$$\text{s.t.}\ \ c \leq \int e(\theta)x(\theta)dF(\theta), \tag{5}$$

$$x(\theta) \text{ is non-decreasing;}$$

for some constant c. We can assign a Lagrange multiplier $\lambda \geq 0$ to the inequality constraint (5),[24] and then the problem is to maximize a Lagrangian

$$\int \left(W(\theta) + \lambda e(\theta)\right) x(\theta)dF(\theta).$$

In the first case, when $W(\theta)$ and $e(\theta)$ are both non-decreasing in θ, $W(\theta)+\lambda e(\theta)$ is non-decreasing in θ for any Lagrange multiplier, so there exists an optimal solution in which $x(\theta)$ is a threshold (one-price) allocation rule. In the second case, when $W(\theta)$ and $e(\theta)$ are both non-increasing in θ, $W(\theta)+\lambda e(\theta)$ is non-increasing in θ for any Lagrange multiplier, so there exists an optimal solution in which $x(\theta)$ is constant. In both cases, by continuity of the

[24] The existence of Lagrange multipliers follows from Theorem 2.165 of Bonnans and Shapiro (2000) as long as c is in the interior of all possible values that the right side can take as $x(\theta)$ varies; whenever this is not the case, $x(\theta)$ must take the form of a threshold allocation rule when $e(\theta)$ is non-decreasing, and a constant function when $e(\theta)$ is non-increasing, which is what we are trying to prove for parts 1 and 2, respectively.

constraint (5), it is possible to choose the threshold in x (in the first case) or the constant level of x (in the second case) to satisfy (5) and the complementarity-slackness conditions.

Finally, we prove the third part of Corollary 2. Recall that $\theta^\star$ is defined as the cutoff type that is indifferent between buying the good and not buying the good at the optimal mechanism. Thus, the first-order conditions for optimality of $\theta^\star$ must hold:

$$\frac{W(\theta^\star)}{-e(\theta^\star)} = \frac{W^\star \nu'(E^\star)}{v(E^\star)},$$

where

$$E^\star = \int_{\theta^\star}^{\bar{\theta}} e(\theta)dF(\theta) \text{ and } W^\star = \int_{\theta^\star}^{\bar{\theta}} W(\theta)dF(\theta).$$

We now consider how the expected payoff of the designer changes under the following perturbation of the mechanism: For some type $\theta_0 < \theta^\star$, allocate access to types $\theta \in [\theta_0, \theta^\star]$ with some small probability $\epsilon > 0$; this perturbation is profitable for small enough ϵ if

$$\int_{\theta_0}^{\theta^\star} e(\theta)dF(\theta) \cdot W^\star \nu'(E^\star) + v(E^\star) \cdot \int_{\theta_0}^{\theta^\star} W(\theta)dF(\theta) > 0.$$

Relying on the first-order condition for $\theta^\star$ to substitute $W^\star \nu'(E^\star)/v(E^\star)$ and rearranging, we obtain the condition from the corollary. For the case when $\theta_0 > \theta^\star$, we obtain the same condition by considering a perturbation under which types $\theta \in [\theta^\star, \theta_0]$ receive allocation $1 - \epsilon$ for some small $\epsilon > 0$.

B Applications

We first state a generalization of the first part of Corollary 2 that we will use in some applications.

Proposition 1. *Suppose that the designer's objective function $W_i(\theta)$ is non-decreasing in θ, and that $\int_{\theta}^{\bar{\theta}_j} W_j(\tau)dF_j(\tau) \geq 0$, for all θ and $j \in I$.*[25] *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.*

Proof. The proof is virtually the same as the proof of the first part of Corollary 2, except that

[25] 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.

the fact that the designer's payoff is non-decreasing in the values v_j for all j now follows from the assumption that $\int W_j(\theta)x_j(\theta)dF_j(\theta) \geq 0$, for all non-decreasing allocation rules x_j. $\square$

B.1 Ticket sales

In this appendix, 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.[26] 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.[27] 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 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

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

[26] 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).

[27] 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.

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 WTP.

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:[28]

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 WTP 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 k-tiered pricing mechanisms with $k \leq 2$. The seller is thus maximizing

$$\underbrace{\left(v_0 + x_0 \int_{\theta_0}^{\theta_1} e(\theta)dF(\theta) + \int_{\theta_1}^{\bar{\theta}} e(\theta)dF(\theta)\right)}_{V} \underbrace{\left(x_0 \int_{\theta_0}^{\theta_1} J(\theta)dF(\theta) + \int_{\theta_1}^{\bar{\theta}} J(\theta)dF(\theta)\right)}_{R} \tag{6}$$

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 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*

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

[28] The calculations underlying the results presented in this section can be found in Appendix B.1.1.

where R and V are functions of $\theta_0^\star$ and $x_0^\star$ as defined in (6) *(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 monotonicity 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.[29] 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 WTP 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 WTP 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 $J(\theta)/(-e(\theta))$: 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* (6) *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*

$$\frac{J(\theta_1^\star)}{-e(\theta_1^\star)} < \frac{\int_{\theta_0}^{\theta_1^\star} J(\theta)dF(\theta)}{\int_{\theta_0}^{\theta_1^\star}(-e(\theta))dF(\theta)}. \tag{7}$$

[29] 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 WTP of the highest type $\bar{\theta}$)—this is why Result 2 is consistent with Theorem 1.

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 (6).

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 WTP.

To understand condition (7), 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 WTP. Then, condition (7) 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 WTP 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 WTP.

The puzzle of why performers often sell tickets to their events at below-market-clearing prices has attracted considerable attention in economics.[30] Becker (1991) put forward an

[30] Budish and Bhave (2023) provide a discussion and an empirical application (see also, e.g., Baliga and

explanation by arguing that WTP 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;[31] 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.

B.1.1 Proofs for Section B.1

First, we note that the first-order conditions for optimality of $\theta_0^\star$ and $\theta_1^\star$ for problem (6) 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 (6) when θ_1 is set to $\bar{\theta}$. The derivative of objective (6) 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

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

Ely (May 8, 2013); Arslan et al. (2023)).

[31] 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).

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

B.2 College admissions

In this appendix, we develop an application of our framework to admission policy in the context of a college or other academic program in which students' value for participating (and hence, their WTP) 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 interest.[32]

In this setting, on top of the unobserved heterogeneity in WTP 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 WTP 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.[33]

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 $h_i(\theta) = (1 - F_i(\theta))/f_i(\theta)$

[32]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.

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

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 B.2.1, 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.

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, WTP 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 WTP (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 possible, the college achieves that via targeted admission decisions and scholarships. But when such a policy is not available, the college must rely on self-selection to ensure diversity. Since students coming from economically underprivileged backgrounds tend to have low WTP (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)).[34]

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 nu-

[34] 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 examined these distributional impacts empirically (see, e.g., Heller (2006); Biró et al. (2020); Heo (2023)).

merous extracurriculars.[35] 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.[36] If we think of the score t as comprising two dimensions—a standardized 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.

B.2.1 Proofs for Section B.2

To prove Results 4 and 5, we 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: $v = \nu(s_1, \dots, s_K)$, 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 (3) 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.

B.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

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

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

(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:

$$p_i = \gamma_i v_i - (1-\gamma_i) \sum_{j \in I} e_{i \to j} \mu_j R_j, \tag{8}$$

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 (8) 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 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 WTP 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 (8) 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 externalities.

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 (8) 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.[37]

Result 7. *Suppose that $\{\theta_j\}_{j\in I}$ are the threshold types determined by the solution to the system of equations (8). The participation of group i is a substitute for the participation of group j if*

$$e_{i\to j} \cdot J_j(\theta_j) - \frac{1}{v_i} \cdot e_{j\to i} \cdot e_{i\to R} < 0, \tag{9}$$

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 (9) 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 (9)—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 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 (9)—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

[37]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 (8). See Appendix B.3.1 for the derivations.

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 (9) 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)).[38]

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. 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.[39] 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,

[38]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).

[39]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).

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).

B.3.1 Proofs for Section B.3

To derive the pricing formula (8), 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 (8): $p_i = \gamma_i v_i - (1 - \gamma_i) \sum_{j \in I} e_{i \to j} \mu_j R_j$.

To determine when participation by group i is a substitute or a complement for participation 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 (8) 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,[40] 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 $e_{j\to R}$.

B.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 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

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

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.[41] 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 WTP 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 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

[41] 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)).

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

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

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*

$$\exists \theta^\star \geq 1 \text{ such that } \int_{\theta^\star}^{\bar{\theta}_D} e_D(\theta) dF_D(\theta) > 0. \tag{10}$$

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*

$$\exists \theta^\star \geq 0 \text{ such that } \int_{\theta^\star}^{\bar{\theta}_D} e_D(\theta) dF_D(\theta) > 0. \tag{11}$$

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 (10) fails but condition (11) 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 (10) and (11), 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 B.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.[42] 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.

B.4.1 Proofs for Section B.4

To prove Result 8, note that since the general public only generates congestion, $e_P(\theta) = -1$ for all θ, any allocation to group P weakly decreases the resale value. Hence the token has strictly positive resale price if and only if there exists a feasible allocation to developers alone such that

$$\int e_D(\theta)x_D(\theta)\, dF_D(\theta) > 0.$$

Under the constraint $x_D(\theta) = 0$ for $\theta < 1$ and monotonicity of the allocation rule, this is possible if and only if there exists a cutoff $\theta^\star \geq 1$ such that allocating to all developer types above $\theta^\star$ yields positive total externality (it suffices to consider extreme points of the set of feasible allocation rules):

$$\int_{\theta^\star}^{\bar{\theta}_D} e_D(\theta)\, dF_D(\theta) > 0.$$

Group P has type-independent externality, so by Proposition 1 it is optimal to use a posted price for the general public. For developers, complexity of their externalities is at most one; hence by Theorem 1, an optimal mechanism for developers uses at most 2 tiers.

To prove Result 9, note that with vesting we may choose *any* monotone allocation rule for developers. Therefore the token has strictly positive resale price if and only if there exists an allocation to developers such that

$$\int e_D(\theta)x_D(\theta)\, dF_D(\theta) > 0.$$

Since there is now no restriction excluding types below 1, this holds if and only if there exists

[42] 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).

a cutoff $\theta^\star \geq 0$ such that

$$\int_{\theta^\star}^{\bar{\theta}_D} e_D(\theta)\, dF_D(\theta) > 0.$$

The second part of Result 9 follows from direct inspection of the two conditions (10) and (11).