Reduced Forms: Feasibility, Extremality, Optimality

Pasha Andreyanov[†] Ilia Krasikov[‡] Alex Suzdaltsev[§]

February 2026

Abstract

We study independent private values auction environments in which the auctioneer's revenue depends nonlinearly on bidders' interim winning probabilities. Our framework accommodates heterogeneity among bidders and places no ad hoc constraints on the auctions available to the auctioneer. Within this general setting, we show that feasibility of interim winning probabilities can be tested along a unidimensional curve—the *principal curve*. We construct this curve in closed form through a novel transformation of the interim winning probabilities. We use this insight to explicitly characterize the extreme points of the feasible set and the allocations that induce them. We then combine our results on feasibility and extremality to solve for the optimal auction under a natural regularity condition. The optimal auction equalizes bidders' marginal revenues along its principal curve and awards the good to the bidder whose type is reached by the curve the earliest. We show that whether and how low types are excluded depends on the curvature of the auctioneer's revenue—at the intensive margin, by scaling winning probabilities down, or at the extensive margin, by cutting low types off. We apply our approach to the classical linear model, settings with endogenous valuations due to ex ante investments, and settings with non-expected-utility preferences. In these settings, previous results were largely limited either to symmetric environments with symmetric allocations or to two-bidder environments.

1 Introduction

Auctions are among the most widely used market institutions. Governments use them to allocate spectrum licenses and procurement contracts, firms employ them to acquire or liquidate inputs, and millions of individuals participate daily on online platforms. A central question in auction theory is how to design auctions that maximize an auctioneer's objective such as expected revenue.

The benchmark case of independent private values (IPV) with linear utilities has been well understood since at least Myerson [1981]. Myerson's celebrated analysis shows that, in this environment, the optimal auction awards the good to the bidder with the highest positive virtual value under mild regularity conditions. A key feature that makes this analysis tractable is that

[†]HSE University, Moscow. Email: pandreyanov@gmail.com.
[‡]Arizona State University. Email: krasikovis.main@gmail.com.
[§]HSE University, Saint Petersburg. Email: asuzdaltsev@gmail.com.

expected revenue is linear in bidders' interim winning probabilities. This linearity allows the auctioneer's problem to be solved pointwise over feasible allocations in regular environments when ironing is not needed.

Much less is known once the auctioneer's objective departs from linearity. Nonlinear dependence on interim winning probabilities arises naturally in economically important settings. For instance, it arises when bidders undertake ex ante investments that affect valuations, as in models of endogenous valuations [Gershkov et al., 2021], or when bidders have non-expected-utility preferences such as constant risk aversion (CRA) [Gershkov et al., 2022]. In these environments, because the auctioneer's revenue depends nonlinearly on interim winning probabilities, the convenient pointwise maximization technique of the linear benchmark is no longer available. As a result, classical methods break down, and the analysis of optimal auctions becomes substantially more complex.

This paper develops a unified framework for studying optimal auctions when the auctioneer's expected revenue can be written in reduced form as an integral of bidder-specific revenue terms. These terms may be nonlinear in interim winning probabilities. We allow for full asymmetry across bidders and do not impose any ad hoc restrictions on the auctions available to the auctioneer. Our framework accommodates a wide range of economic environments, including the linear IPV model, settings with endogenous valuations, and settings with CRA preferences cited above.

The core difficulty in nonlinear environments is the loss of separability: the auctioneer's problem no longer decomposes across types, and allocations cannot be chosen type by type. Making progress therefore requires resolving three intertwined challenges.

Feasibility. Bidders' interim winning probabilities must be consistent with some ex post allocation, and characterizing this consistency requires checking Border's condition. This condition provides a set of necessary and sufficient inequalities, but this set is high-dimensional—one inequality for every possible cutoff vector of types. While tractable for symmetric auctions or with two bidders, it quickly becomes unmanageable as we move to asymmetric auctions with a larger number of bidders. This leaves little hope of solving the fully constrained problem directly.

Extremality. The set of feasible reduced forms—tuples of bidders' interim winning probabilities—is convex. Consequently, when the auctioneer's objective is convex—as in many endogenous-valuation environments—an optimum is attained at an extreme point of the feasible set. A tractable description of extremal reduced forms could therefore sharply restrict the set of candidates and reduce the subsequent maximization problem to a much lower-dimensional search. Yet, existing characterizations are largely confined to symmetric settings or to two bidders, leaving the structure of extreme points poorly understood in richer settings.

Optimality. Ultimately, the core obstacle is that nonlinearities destroy the pointwise structure of the linear benchmark. Feasibility and extremality help identify the relevant constraint set and restrict attention to plausible candidates, but the remaining task is still to optimize a nonlinear objective over that set. Doing so requires a methodology that both solves the auctioneer's problem and delivers structural insights about the optimal auction.

In this paper, we address these challenges by developing new characterizations for feasibility and extremality and then leveraging them to solve the auctioneer's problem.

Our first contribution is a new representation of feasibility that collapses Border's high-dimensional family of constraints to a unidimensional test. Its building block is a novel transformation of CDFs, the *δ-transform*, which reparametrizes each bidder's interim winning probability as a curve in the type-probability square along which the product of type and probability decays at a constant rate. Applying it coordinatewise to a reduced form yields a path inside the unit cube—its *principal curve*—along which all bidders' type-probability products are equalized. We show that Border's inequalities hold everywhere if and only if they hold along this single curve, so that feasibility reduces to checking the sign of one scalar function of one variable. This makes feasibility straightforward to verify even in fully asymmetric environments with many bidders.

Our second contribution is a sharp characterization of extremality. Building on our feasibility representation, we show that a reduced form is extremal precisely when Border's constraints bind along its principal curve. This criterion dramatically simplifies the auctioneer's problem in environments where the optimum is known to be extremal. It also clarifies how to induce reduced forms through ex post allocations. Specifically, we show that every extremal reduced form can be implemented by a simple *score allocation*. In a score allocation, each bidder's type is mapped into a unidimensional score, and the good is awarded to the bidder with the highest positive score, with ties occurring with probability zero. We prove that score allocations are exactly the allocations that induce extremal reduced forms. Moreover, the allocation inducing a given extremal reduced form is unique and is obtained from its *canonical scores*—each type times its winning probability—so that extremality can be tested without any search over scores. This yields a sharper form of the classical equivalence between Bayesian and dominant-strategy implementation originally established by Manelli and Vincent [2010]. In particular, we explain that any feasible reduced form is implementable as a mixture of deterministic, dominant-strategy incentive compatible, and nonbossy allocations, in the sense of Satterthwaite and Sonnenschein [1981].

Our third contribution concerns optimality in nonlinear environments. We first show that, in regular environments where it is optimal to always sell the good, the optimum is a score allocation whose scores are constructed to equalize bidders' marginal revenues along the principal curve. Unlike in linear settings, where the marginal revenue equalization directly yields the solution (Myerson [1981] and Bulow and Roberts [1989]), it is a fixed-point problem in nonlinear settings. This is because each bidder's marginal revenue depends on own interim winning probability, and bidders' interim winning probabilities are interconnected through feasibility. The principal curve untangles this fixed point: expressed in terms of δ-transforms, the equalization becomes a system of equations that can be solved at each point of the curve separately, directly from the primitives. To certify optimality, we develop a new methodology that recasts the auctioneer's problem as a continuous-time optimal control problem. In this problem, the principal curve is built over "time" moving from high to low types, the auctioneer chooses how quickly to allocate winning probability to each bidder, and Border's condition enters as a dynamic feasibility constraint along the curve. The system describing the marginal revenue equalization reappears there as the first-order condition, and regularity—which extends Myerson's regularity to nonlinear settings so that no ironing is needed—guarantees its sufficiency.

We show how this characterization applies to the endogenous-valuation model, where bidders choose ex ante investments that affect their valuations. In this environment, interim winning probabilities matter not only for rent extraction but also because they determine the return to investing. Any bidder with a positive interim winning probability has an incentive to invest, even though only one ultimately wins, so investment costs are largely duplicated. The optimal auction departs from Myerson's auction by shifting winning probability toward the strongest bidder and away from the weakest one until bidders' marginal revenues are equalized along the optimal principal curve. Relative to Myerson's auction, this reallocation reduces excessive duplication of investment costs, thereby increasing the auctioneer's expected revenue.

When exclusion is optimal, its form depends on the curvature of bidders' revenue functions in interim winning probabilities. In concave environments, as in settings with risk-averse bidders, exclusion operates at the intensive margin. At low types, the marginal revenue from awarding the full unit can be negative, and, rather than excluding such types from trade altogether, the auctioneer awards them the fraction at which marginal revenue is exactly zero. We illustrate this case in the CRA model with two groups of bidders, risk-neutral and risk-averse. Letting risk-averse bidders trade at scaled-down probabilities at the bottom of the type distribution is profitable, since it lowers information rents at a small efficiency cost, and the optimal auction favors the risk-averse bidders at low types and the risk-neutral bidders at high types. This echoes Myerson's logic that the optimal auction should tilt toward ex ante "weak" bidders to intensify competition, with the twist that, in the nonlinear CRA model, who is "weak" depends on interim winning probabilities. At high interim winning probabilities, risk-averse bidders are effectively stronger, since they are willing to pay a premium to secure the good with near certainty; at low interim winning probabilities, the ranking is reversed.

In convex environments, as in the endogenous-valuation model, exclusion operates at the extensive margin: the auctioneer either awards the good with certainty or cuts the type off. Where to cut is determined by a trade-off that is absent in the linear model. Excluding a bidder's low types relaxes feasibility for the other bidders, whose winning probabilities can then be raised. As a result, bidders are excluded at positive and bidder-specific marginal revenues, as opposed to the linear model, where every bidder is excluded at the common threshold of zero. We illustrate this point in the endogenous-valuation model with shipping costs, where the cheapest-to-serve bidder is excluded last, once marginal revenue reaches own cost, whereas every other bidder is excluded while marginal revenue still exceeds own cost.

1.1 Relation to literature

A substantial body of literature has studied feasibility, extremality, and optimality of reduced forms. Below, we describe the key existing contributions and relate our results to them.

Feasibility. In a classic result, Border [1991], building on Maskin and Riley [1984] and Matthews [1984], characterized the set of feasible symmetric reduced forms—that is, symmetric interim winning probabilities inducible by some ex post allocation—via a unidimensional family of linear inequalities. Border [2007] and Mierendorff [2011] generalized this result to asymmetric reduced forms, where feasibility is described by a family of inequalities whose dimension grows with the number of bidders. Che et al. [2013], in turn, extended these results to multiunit auction settings with capacity constraints. Hart and Reny [2015] related Border's conditions to second-order stochastic dominance, and He et al. [2025] showed that Border-type inequalities can also characterize the set of feasible belief distributions generated by private-private information structures.

The closest paper to ours regarding feasibility in auction settings is Cai et al. [2018], who showed that, in finite (potentially only partially ordered) type spaces, one can reduce the family of asymmetric Border inequalities to one dimension. Our result is complementary to theirs: it relies on additional structure—continuous, unidimensional, linearly ordered types and monotone interim winning probabilities—but delivers a sharper characterization. In this setting, we characterize feasibility via a unidimensional monotone curve, which we describe in closed form through an explicit transformation of reduced forms.

Extremality. Since a quasiconvex functional attains its maximum at an extreme point of a convex compact set,[1] characterizing extreme points has become an increasingly powerful approach. It allows one to solve many optimization problems of economic interest at once. In a landmark contribution, Kleiner et al. [2021] characterized extreme points of the set of monotone functions that majorize, or are majorized by, a given monotone function. Aside from applications to contests, optimal persuasion, and delegation, this yielded an immediate description of extreme points of the set of feasible symmetric reduced forms, since the symmetric Border constraint could be viewed as a majorization constraint.[2]

Extreme points of asymmetric reduced forms in auction settings are still not fully understood. A notable contribution is Manelli and Vincent [2010], who characterized some extremal asymmetric reduced forms—those that are step functions—and used this description to establish the BIC–DIC equivalence. More recently, Yang and Yang [2025] provided an equivalent characterization of extremal asymmetric reduced forms in the two-bidder case. Their approach differs from ours: it characterizes extreme points of the set of rationalizable tuples of monotone functions. While the connection between rationalizability and majorization is transparent with two bidders, it is unclear how to extend it to characterize extremal reduced forms when the number of bidders exceeds two. In contrast, our approach delivers an explicit description of all

[1]This holds under continuity by a generalization of Bauer's maximum principle to quasiconvex functionals; see Ball [2023].

[2]Kleiner et al. [2021] also covered multiunit symmetric environments.

extreme points irrespective of the number of bidders.

A few other papers used the extreme-points approach to solve problems of economic significance. Among them, Manelli and Vincent [2007] employed it to study multidimensional screening. Nikzad [2023] and Candogan and Strack [2023] characterized extreme points of the same set as in Kleiner et al. [2021] but with additional linear constraints and applied this approach to study optimal information disclosure. Arieli et al. [2023] described extreme points of the set of mean-preserving contractions of a given prior. Kleiner et al. [2024] studied extreme points of the set of fusions, the multidimensional analog of mean-preserving contractions, and Lahr and Niemeyer [2024] studied extreme points in multidimensional monopolistic screening settings. Finally, Yang and Zentefis [2024] characterized extreme points of the set of monotone functions bounded by two given monotone functions and applied this to political economy, persuasion, the psychology of judgment, and security design.

Optimality. As noted above, once the objective is nonlinear in interim winning probabilities, the optimization problem becomes hard because pointwise optimization is no longer available. Despite these difficulties, several papers have made substantial progress. Gershkov et al. [2021] leveraged the Fan–Lorentz inequality to solve for the optimal symmetric allocation with many goods when valuations are endogenous and provided conditions under which the globally optimal reduced form is symmetric in two-bidder environments; Andreyanov et al. [2024] provided analysis for a related setting with contractible actions. Zhang [2017] characterized the optimum in the symmetric two-bidder endogenous-valuations model with quadratic costs, whereas Andreyanov et al. [2026] solved for the optimum in general symmetric two-bidder settings with endogenous valuations. Finally, Gershkov et al. [2022] characterized the optimal auction in symmetric environments with non-expected-utility bidders exhibiting constant risk aversion. Our approach covers these applications and unifies and extends the analysis to general asymmetric settings.

1.2 Structure of the paper

The paper contains several distinct results, and, to help the reader navigate, we indicate key theorems and propositions in brackets throughout this outline. Because the main objects and transformations we introduce are nonstandard, and later arguments rely on this shared language, we recommend reading the paper linearly.

Section 2 introduces the model and explains how several leading auction environments fit into our framework. Section 3 presents an illustrative example that previews the economics of the optimal auction, and Section 4 collects the simplifications used throughout.

Sections 5–7 develop the main methodological results. Specifically, Section 5 reduces Border's high-dimensional feasibility constraints to a unidimensional test along the principal curve (Theorem 1). Section 6 characterizes extremal reduced forms and links them to implementation via score allocations (Theorems 2 and 3). Section 7 characterizes the optimal auction in regular environments when it is optimal to always sell the good (Theorem 4) and applies it to discrimination in the endogenous-valuation model. It also outlines the methodology for certifying optimality via optimal control (Proposition 2).

Section 8 extends the characterization to environments in which exclusion is optimal, distinguishing between concave and convex environments (Theorems 5 and 6), and applies the concave case to discrimination in the CRA-preferences model.

Section 9 concludes, and the appendix collects proofs and supplementary results.

2 Setting

We consider an auction environment with n bidders, indexed by $i = 1, \ldots, n$, and a single auctioneer who sells one indivisible unit of a good. Each bidder i has a private unidimensional type $\theta_i \in \Theta_i \subset \mathbf{R}_+$, where Θ_i is a compact interval. The type θ_i is drawn from a distribution F_i on Θ_i that admits a strictly positive and continuous density denoted by f_i. Types are independently distributed across bidders, although their marginal distributions may differ. We use boldface for vectors of individual variables, e.g., $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$ stands for a type profile.

Let $\boldsymbol{\Theta} = \Theta_1 \times \cdots \times \Theta_n$ denote the set of all possible type profiles. An **allocation** is a function $\boldsymbol{z} = (z_1, \ldots, z_n) : \boldsymbol{\Theta} \to [0, 1]^n$ satisfying

$$\sum_{i=1}^{n} z_i(\boldsymbol{\theta}) \leq 1 \tag{1}$$

for all type profiles $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. Here, $z_i(\boldsymbol{\theta})$ represents the probability that bidder i receives the good when the type profile is $\boldsymbol{\theta}$, and the constraint simply means that the auctioneer can sell at most one unit. For bidder i, the **interim winning probability** given their type θ_i is defined by

$$x_i(\theta_i) = \mathbf{E}[z_i \mid \theta_i], \tag{2}$$

where the expectation is taken with respect to the independent draws of the other bidders' types.

The setting described thus far corresponds to the standard IPV auction model with linear utilities. It is well known that, in this setting, the auctioneer's expected revenue can be expressed as the following function of interim winning probabilities:

$$\sum_{i=1}^{n} \int_{\Theta_i} \underbrace{\left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right)}_{=:\mathrm{MVV}_i(\theta_i)} x_i(\theta_i) dF_i(\theta_i),$$

where the term in parentheses is often referred to as **Myerson's virtual value** (henceforth **MVV**).

In contrast to the standard model, we remain agnostic about bidders' preferences and instead represent the auctioneer's implied revenue in reduced form. Specifically, let $J_i(x_i, \theta_i)$ denote the revenue obtained from selling the good to type θ_i with probability x_i, so that the expected revenue can be written as

$$\sum_{i=1}^{n} \int_{\Theta_i} J_i(x_i(\theta_i), \theta_i) dF_i(\theta_i). \tag{3}$$

This specification subsumes the standard model studied by Myerson, in which bidder i's **marginal revenue** $\frac{\partial J_i}{\partial x_i}$ (henceforth **MR**) coincides with that bidder's MVV, and it is far more general. For example, this framework accommodates environments with endogenous valuations and bidders with CRA preferences discussed below.

We are interested in finding the revenue-maximizing auction when the auctioneer's revenue is given by (3) subject to bidders' interim winning probabilities, defined by (2), being weakly increasing. The monotonicity requirement means that the auctioneer can implement any allocation in which bidders' interim winning probabilities are weakly increasing in their types, and is restricted to search among such allocations. This condition is the only substantive assumption we impose, and, as we explain below, it is satisfied in the standard model as well as in settings with endogenous valuations and non-expected-utility CRA preferences. In all these cases, monotonicity is both necessary and sufficient for implementability, which is a direct consequence of bidders' incentive compatibility.

To streamline the exposition, we assume that J_i is twice continuously differentiable, and we normalize $J_i(0, \theta_i) = 0$ for all θ_i. This normalization is without loss of generality, since the term $\int_{\Theta_i} J_i(0, \theta_i) dF_i(\theta_i)$ does not depend on the allocation. In the rest of this section, we show how our general framework encompasses several important auction environments studied in the literature.

2.1 IPV setting with linear utilities

This is the classical and most extensively studied case, already mentioned above. Bidder i's utility takes the form $\theta_i z_i - t_i$, where t_i denotes bidder i's payment to the auctioneer, and the auctioneer seeks to maximize the expected sum of bidders' payments. It is well known that only allocations in which bidders' interim winning probabilities are monotone can be implemented in Bayesian Nash equilibrium. In this environment, the auctioneer's revenue coincides with (3) with

$$J_i(x_i, \theta_i) = \mathrm{MVV}_i(\theta_i) x_i$$

up to a constant. This constant is pinned down by the participation constraint of the lowest types, which binds at zero, so the auctioneer's problem is exactly of the form described above.

The solution has been well understood since Myerson [1981]. In the regular case—where bidders' MVVs are strictly increasing—the good is allocated to the bidder with the highest positive MVV.[3] To see this, note that, due to linearity in interim winning probabilities, we can

[3] More generally, the optimal allocation is determined by ironed MVVs and may involve randomization among tied bidders.

rewrite the objective in (3) as

$$\sum_{i=1}^{n} \int_{\Theta_i} J_i(x_i(\theta_i), \theta_i) dF_i(\theta_i) = \mathbf{E}\left[\sum_{i=1}^{n} \mathrm{MVV}_i \cdot z_i\right], \tag{4}$$

which can be easily maximized pointwise in the space of allocations satisfying (1).

2.2 IPV setting with CRA preferences

In an important contribution, Gershkov et al. [2022] (henceforth GMSZ-CRA) extended optimal auction theory to environments in which bidders' preferences over interim winning probabilities depart from expected utility.[4] Under constant risk aversion, they showed that the auctioneer's revenue can still be expressed as in (3), with

$$J_i(x_i, \theta_i) = \theta_i x_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} g_i(x_i)$$

up to a constant. Here, $g_i(x_i)$ is the certainty equivalent of a lottery paying \$1 with probability x_i and \$0 otherwise. If $g_i'' \equiv 0$, bidder i has expected-utility preferences; if $g_i'' \geq 0$, bidder i is risk-averse.

Monotonicity of $\boldsymbol{x}$ remains both necessary and sufficient for implementability. Thus, the auctioneer's problem retains the same structure as before, but it is considerably harder because J_i is generally nonlinear in x_i. As a result, one cannot express the auctioneer's expected revenue as a function of the allocation similarly to (4) and maximize it pointwise.

To gain tractability, GMSZ-CRA focused on symmetric settings in which J_i and F_i are identical across bidders. Under risk aversion, J_i is concave in x_i, which was their main case of interest. They showed that symmetry of the environment, combined with concavity of J_i, implies that the optimal allocation is symmetric, and they derived a closed-form characterization of the corresponding interim winning probabilities.[5] Their analysis, however, does not cover asymmetric environments, which we take up in this manuscript.

2.3 IPV setting with endogenous valuations

In another influential paper, Gershkov et al. [2021] (henceforth GMSZ-EV) analyzed models with ex ante investments and showed that they naturally generate preferences that are convex in interim winning probabilities. In their framework, before the auction, each bidder may take a private action a_i that increases their valuation from winning. Specifically, bidder i's utility is

$$v_i(a_i, \theta_i) z_i - C_i(a_i) - t_i,$$

where C_i is a convex cost incurred regardless of the auction outcome.

[4] Their standing assumption is quasilinearity in money.

[5] GMSZ-CRA also described an approach to the optimal symmetric auction when bidders are symmetric but not necessarily risk-averse.

GMSZ-EV demonstrated that, when bidders are expected-utility maximizers, the auctioneer's expected revenue again takes the form in (3), with

$$J_i(x_i,\theta_i) = \omega_i(x_i,\theta_i) - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}\frac{\partial}{\partial\theta_i}\omega_i(x_i,\theta_i),$$

up to a constant, where

$$\omega_i(x_i,\theta_i) := \max_{a_i} v_i(a_i,\theta_i)x_i - C_i(a_i)$$

is the maximal expected utility of type θ_i given their interim winning probability x_i. Since ω_i is a pointwise maximum of functions linear in x_i, it is convex in x_i, and typically J_i inherits this convexity.[6]

Finding the optimal auction therefore requires maximizing the convex functional in (3) over monotone interim winning probabilities $\boldsymbol{x}$ that are consistent with an allocation, i.e., that satisfy (1) and (2). Very little is known about problems of this form. Notable exceptions include Zhang [2017], who found the optimum in an additively separable example with $n = 2$; GMSZ-EV, who analyzed optimal symmetric allocations (and their global optimality) in symmetric environments; and Andreyanov et al. [2026], who characterized the asymmetric optimum with two bidders. To the best of our knowledge, no results exist for asymmetric environments with more than two bidders. To address this gap, we develop a novel approach to the study of optimal auctions in general nonlinear settings without imposing symmetry or restricting the auctioneer to symmetric auctions.

3 Illustrative example

Since our approach requires several steps to develop, it is useful to first preview the main ideas in a concrete economic example. The purpose of this example is threefold. First, it illustrates how nonlinearities in J_i can be microfounded through bidders' ex ante investments. Second, it explains why Myerson's auction can be suboptimal once valuations are endogenous. Third, it previews the structure of the optimal auction that reallocates probability so as to equalize bidders' MRs. Throughout, we work with n bidders—one strong and $n-1$ symmetric weak ones. Already for $n = 3$, this environment lies beyond the reach of existing results, which are confined either to symmetric settings or to two bidders.

Consider an endogenous-valuation model in which each bidder i can take a private action a_i at cost a_i^2, which increases their own valuation from winning the good to $v_i(a_i,\theta_i) = 2\sqrt{\theta_i}a_i$. The maximal utility of bidder i is then $\omega_i(x_i,\theta_i) = \theta_i x_i^2$. Relative to the standard linear model, the key difference is the quadratic dependence of utility on the interim winning probability, which arises endogenously from bidders' optimal ex ante investment decisions. It is easy to see that bidders' revenue functions inherit the quadratic dependence on their interim winning

[6] It may happen that J_i itself is nonconvex in x_i, whereas its positive part $\max\{0, J_i\}$ is convex. As explained in GMSZ-EV, the difference is inconsequential for the structure of the optimum under suitable regularity; convexity of the positive part is also the condition under which we characterize the optimum with exclusion in Section 8.3.

probability, i.e.,

$$J_i(x_i, \theta_i) = \mathrm{MVV}_i(\theta_i) x_i^2.$$

In this environment, Myerson's auction is generally suboptimal. The reason is that such an allocation discriminates against stronger bidders so as to reduce their information rents (e.g., see Carroll and Segal [2019]), thereby inducing excessive duplication of investment costs at the ex ante stage. By reallocating winning probability away from weaker bidders, the auctioneer can mitigate over-investment and increase expected revenue.

To see this more concretely, suppose that bidders' type distributions satisfy[7]

$$\mathrm{MVV}_i(\theta_i) = (F_i(\theta_i))^{1/\beta_i}$$

for some $\beta_i \in (0, 1)$, where a larger value of β_i is indicative of a higher bidder i's strength in the sense of hazard ratio ordering.[8] Specifically, suppose that bidder 1 is strong with $\beta_1 = \frac{7}{8}$, whereas the remaining $n-1$ weak bidders have $\beta_2 = \cdots = \beta_n = \frac{1}{2}$. Then, bidders' interim winning probabilities in Myerson's auction $\boldsymbol{x}^{\mathrm{M}}$ are given by

$$x_1^{\mathrm{M}}(\theta_1) = \mathbf{Pr}\Big(\mathrm{MVV}_1 \geq \max_{i \neq 1} \mathrm{MVV}_i | \theta_1\Big) = (F_1(\theta_1))^{\frac{4(n-1)}{7}}, \qquad x_i^{\mathrm{M}}(\theta_i) = (F_i(\theta_i))^{\frac{4n-1}{4}} \;\; \forall i \neq 1.$$

Direct calculations show that Myerson's auction yields revenue of

$$\int_{\Theta_1} J_1(x_1^{\mathrm{M}}(\theta_1), \theta_1) dF_1(\theta_1) + \sum_{i \neq 1} \int_{\Theta_i} J_i(x_i^{\mathrm{M}}(\theta_i), \theta_i) dF_i(\theta_i) = \frac{7}{8n+7} + (n-1) \cdot \frac{2}{4n+5},$$

which equals $\frac{243}{527} \approx 0.461$ for $n = 3$. In contrast, trading exclusively with the strongest bidder—bidder 1—gives

$$\int_{\Theta_1} J_1(1, \theta_1) dF_1(\theta_1) + \sum_{i \neq 1} \int_{\Theta_i} J_i(0, \theta_i) dF_i(\theta_i) = \frac{7}{15} + (n-1) \cdot 0 = \frac{7}{15} \approx 0.467$$

independently of n. So, for $n = 3$, Myerson's auction is strictly dominated even by the simple policy of always trading with the strong bidder.[9]

Trading only with bidder 1, however, is still suboptimal. Excluding the weak bidders entirely forgoes the surplus from letting them win at the top of their type distributions, where their marginal contributions to the auctioneer's revenue can exceed bidder 1's. Concretely, consider instead the allocation in which bidder 1 wins when $(F_1(\theta_1))^{\frac{1}{7}} \geq \max_{i \neq 1} F_i(\theta_i)$, and otherwise the good goes to the weak bidder with the highest type. The induced interim winning probabilities

[7] To obtain F_i, remark that its inverse v_i satisfies $v_i(u) - (1-u)v_i'(u) = u^{1/\beta_i}$. Solving this equation yields $v_i(u) = \frac{1}{1-u} \int_u^1 s^{1/\beta_i} ds$.

[8] For two cumulative distribution functions, F and G, with the same support, F is stronger in the sense of hazard ratio ordering than G if $\frac{1-F}{f} \geq \frac{1-G}{g}$. See Section 7.2 for the general definition and discussion of this order.

[9] It can be verified that this domination occurs if and only if $n \leq 4$. However, the optimal auction constructed below strictly dominates both Myerson's auction and trading exclusively with bidder 1 for every n.

$\boldsymbol{x}^*$ are

$$x_1^*(\theta_1) = (F_1(\theta_1))^{\frac{n-1}{7}}, \qquad x_i^*(\theta_i) = (F_i(\theta_i))^{n+5} \;\; \forall i \neq 1,$$

and the auctioneer's expected revenue is

$$\int_{\Theta_1} J_1(x_1^*(\theta_1), \theta_1) dF_1(\theta_1) + \sum_{i \neq 1} \int_{\Theta_i} J_i(x_i^*(\theta_i), \theta_i) dF_i(\theta_i) = \frac{7}{2n+13} + (n-1) \cdot \frac{1}{2n+13},$$

which exceeds the revenue from Myerson's auction and exclusive trade for every n, and equals $\frac{9}{19} \approx 0.474$ for $n = 3$. Relative to Myerson's auction, this allocation shifts winning probability toward the strongest bidder—$x_1^* > x_1^{\mathrm{M}}$ and $x_i^* < x_i^{\mathrm{M}}$ for all $i \neq 1$ and every n—but stops short of exclusion, so that the weak bidders continue to win at high types.

The distinguishing feature of this auction is that the bidder with the highest MR wins the good with certainty, that is,

$$x_i^*(\theta_i) = \mathbf{Pr}\left(\frac{\partial}{\partial x_i} J_i(x_i^*, \theta_i) \geq \frac{\partial}{\partial x_j} J_j(x_j^*, \theta_j) \;\; \forall j \neq i | \theta_i \right). \tag{5}$$

In particular, bidders' MRs are equalized whenever they are tied, $\frac{\partial J_i}{\partial x_i} = \frac{\partial J_j}{\partial x_j}$. As explained in Bulow and Roberts [1989], the equalization of MRs is in fact the organizing principle behind the optimality of Myerson's auction in linear settings. There, due to linearity, $\frac{\partial}{\partial x_i} J_i(x_i^{\mathrm{M}}(\theta_i), \theta_i) = \mathrm{MVV}_i(\theta_i)$ does not depend on the interim winning probability. In our nonlinear example, by contrast, bidders' MRs in (5) depend on their interim winning probabilities, so this equation is a genuine fixed-point problem. It is natural to conjecture that a fixed point $\boldsymbol{x}^*$ characterizes the revenue-maximizing auction. As we show later, this conjecture is correct in regular nonlinear environments when it is optimal to always sell the good (as happens in this example). The verification is nontrivial, however, since local perturbations do not certify global optimality due to convexity of J_i, and admissible perturbations must remain implementable by some ex post allocation.

Our approach overcomes these difficulties by changing variables. Rather than solving directly for interim winning probabilities, we recast the auctioneer's problem in terms of a reparametrization of them. This reparametrization builds implementability into the problem and separates the optimality conditions that are entangled in (5). As a result, we can both verify the fixed-point characterization and solve it explicitly.

4 Preliminaries and simplifications

Before moving on, it is useful to introduce several simplifications and auxiliary observations that streamline exposition and reduce the dimensionality of the problem.

First, we can absorb all asymmetries across bidders into their revenue functions by working with quantiles of bidders' types. Specifically, let $u_i = F_i(\theta_i)$ denote the quantile of θ_i, which is uniformly distributed on $[0, 1]$, and define

$$H_i(x_i, u_i) := J_i(x_i, F_i^{-1}(u_i)).$$

Then, our original problem is equivalent to one in which types, denoted by $\boldsymbol{u}$, are uniformly and identically distributed. In this problem, the auctioneer's revenue from selling the good to type u_i with probability x_i is given by $H_i(x_i, u_i)$. In this formulation, both allocations and interim winning probabilities can be viewed as functions of $\boldsymbol{u}$, so that (3) becomes[10]

$$\sum_{i=1}^{n} \int_0^1 H_i(x_i(u), u) du. \tag{6}$$

Since $\frac{\partial}{\partial x_i} H_i(x_i, u_i) = \frac{\partial}{\partial x_i} J_i(x_i, F_i^{-1}(u_i))$, we continue to refer to $\frac{\partial H_i}{\partial x_i}$ as bidder i's MR.

Second, observe that the expected revenue in (6) depends only on interim winning probabilities. It has been known since at least Border [1991] and Mierendorff [2011] that $\boldsymbol{x}$ is consistent with some allocation $\boldsymbol{z}$ satisfying analogues of (1) and (2) in the $\boldsymbol{u}$-space if and only if the following **Border constraint** holds:

$$B(\boldsymbol{u}) = \prod_{i=1}^{n} u_i + \sum_{i=1}^{n} \int_{u_i}^1 x_i(u) du \leq 1 \;\; \forall \boldsymbol{u} \in [0, 1]^n. \tag{7}$$

in which case we say that $\boldsymbol{z}$ **induces** $\boldsymbol{x}$. Intuitively, (7) requires that, for every cutoff vector $\boldsymbol{u}$, the total probability assigned to the bidders whose types lie in the upper tails $[u_i, 1]$ does not exceed the probability that at least one such bidder is present.

Finally, it is without loss of generality to require each x_i to be right-continuous with $x_i(1) = 1$. To see this, recall that x_i is weakly increasing and thus has at most countably many jump discontinuities. Redefining x_i at those points (and possibly at $u = 1$) affects neither the objective in (6) nor the Border constraint in (7). By this observation, we can restrict attention to $\boldsymbol{x}$ that are cumulative distribution functions (CDFs) on the unit interval. Let $\mathcal{X}$ denote the set of such CDFs.

We refer to a tuple $\boldsymbol{x} \in \mathcal{X}^n$, which captures interim winning probabilities, as a **reduced form**, and we use

$$\underline{u}_i := \inf\{u_i \in [0, 1] : x_i(u_i) > 0\}$$

to denote bidder i's **exclusion type**, below which the bidder never wins. A reduced form $\boldsymbol{x}$ is **feasible** if it satisfies the Border constraint. For a feasible reduced form $\boldsymbol{x}$, the probability that the good is sold, $\mathbf{E}[\sum_{i=1}^n z_i]$, equals $B(\mathbf{0})$, and we say that the reduced form has **no exclusion** if this probability is equal to one.[11] Finally, a reduced form $\boldsymbol{x}$ is **optimal** if it solves the auctioneer's

[10] To obtain allocations in the original type space, we can simply substitute back $\boldsymbol{u} = (F_1(\theta_1), \ldots, F_n(\theta_n))$, e.g., $x_i(F_i(\theta_i))$ stands for bidder i's interim winning probability.

[11] If $\boldsymbol{x}$ has no exclusion, then necessarily $\prod_{i=1}^n \underline{u}_i = 0$; the converse, however, need not hold.

problem, which can be succinctly restated as

$$\max_{\boldsymbol{x} \in \mathcal{X}^n} \sum_{i=1}^{n} \int_0^1 H_i(x_i(u), u) du \quad \text{s.t.} \quad B(\boldsymbol{u}) \leq 1 \quad \forall \boldsymbol{u} \in [0,1]^n. \tag{8}$$

While compact, the maximization problem in (8) remains challenging due to the nonlinearities in H_i, and since the feasible set is defined by a large and intricate family of inequalities. To make progress, we introduce a novel characterization of feasibility that dramatically simplifies the characterization of the feasible set. Furthermore, since that set is convex, additional tractability can be obtained by analyzing its extreme points.[12] For example, if each revenue function H_i is convex in its first argument—as often arises in settings with endogenous valuations—then an optimal reduced form can be found at an extreme point of the feasible set. We call a feasible reduced form $\boldsymbol{x}$ **extremal** if it is an extreme point of the set of CDFs satisfying the Border constraint.

5 Feasibility

The Border constraint tells us exactly which reduced forms are feasible. The drawback, however, is that it requires checking an entire n-dimensional family of inequalities—one for every cutoff vector $\boldsymbol{u} \in [0,1]^n$. When $n = 2$, this is still manageable, but, as soon as there are three or more bidders, the condition becomes challenging to verify directly. Our goal in this section is to show that this system of constraints can, in fact, be reduced to a simple unidimensional test.

Our key insight is that feasibility can be checked along a single continuous curve inside the unit cube. That is, for every reduced form $\boldsymbol{x}$, there exists a path

$$t \in \mathbf{R}_+ \mapsto \boldsymbol{\nu}(t) \in [0,1]^n$$

such that Border's inequality holds everywhere if and only if it holds along this path. Instead of monitoring a continuum of constraints in n dimensions, we can therefore follow just one curve. We refer to such a curve associated with a given reduced form as its **principal curve**.

5.1 δ-transform

To derive the principal curve, we first reparametrize CDFs as curves in the type-probability square. To fix ideas, remark that the uniform CDF on $[0,1]$ can be thought of as a straight line in the type-probability square. For each parameter $t \in \mathbf{R}_+$ (henceforth time), this line prescribes type $u = e^{-t/2}$ and probability $x = e^{-t/2}$. This curve starts from the top-right corner and is such that the product of type and probability moves toward the origin at rate 1. Clearly, we can recover the uniform distribution from this parametric description as $x(u) = e^{-t/2}\big|_{t=-2\ln u}$.

[12] Recall that a point in a convex set is extreme if it cannot be expressed as a convex combination of two distinct points in the set. As explained in Section 2 of Kleiner et al. [2021], by the Krein–Milman theorem, any convex and compact set in a locally convex space coincides with the closed convex hull of its extreme points.

We now seek to reparametrize an arbitrary CDF $x \in \mathscr{X}$ by a curve that starts at the top-right corner and whose coordinates multiply at each time t to e^{-t}. In general, such a curve can be expressed as[13]

$$t \in \mathbf{R}_+ \mapsto \left(e^{-\delta(t)}, e^{\delta(t)-t}\right), \tag{9}$$

where δ belongs to the set $\mathscr{D}$ of absolutely continuous functions with $\delta(0) = 0$ and $\delta' \in [0, 1]$. Indeed, the coordinates multiply to e^{-t}, and the initial condition places the curve at the corner $(1, 1)$. The bounds on δ' make both coordinates weakly decreasing in t, so that the curve traces the graph of a weakly increasing function from top to bottom.

The representing function δ can be explicitly constructed from the CDF x in two consecutive steps. First, multiply x by u to obtain the product ux and construct the inverse of that product, $(ux)^{-1}$. Note that ux itself is a CDF and is strictly increasing wherever x is strictly positive. Second, apply the log change of variables to obtain

$$\delta(t) = -\ln(ux)^{-1}(e^{-t}). \tag{10}$$

Conversely, we can reverse these steps starting from any $\delta \in \mathscr{D}$ by finding the time at which the curve reaches type u using the inverse of δ and evaluating the probability coordinate at that time, that is,

$$x(u) = e^{\delta(t)-t}\big|_{t=\delta^{-1}(-\ln u)}. \tag{11}$$

The following lemma records that nothing is lost in the process described above.

Lemma 1. *The assignments in (10) and (11) are mutually inverse and define a bijection between $\mathscr{X}$ and $\mathscr{D}$.*[14]

In view of this lemma, we can use CDFs and functions in $\mathscr{D}$ interchangeably. For each $x \in \mathscr{X}$, we shall refer to the corresponding $\delta \in \mathscr{D}$ as the **δ-transform of x**. A more detailed discussion of the geometry of the δ-transform can be found in Appendix A.1.2.

5.2 Feasibility through the principal curve

The δ-transform is the tool that allows us to collapse Border's n-dimensional condition into a unidimensional one. Given a reduced form $\boldsymbol{x}$, which is not necessarily feasible, let $(\delta_1, \ldots, \delta_n)$ denote the δ-transforms of its coordinates, and define the **principal curve** by applying (9)

[13] In logarithmic coordinates $(-\ln u, -\ln x)$, the parametrization in (9) is the classical parametrization of monotone graphs by the sum of their coordinates, which goes back to Minty [1962]; e.g., see Proposition 1.1 in Alberti and Ambrosio [1999].

[14] The inverses used in (10) and (11) are generalized right- and left-continuous, respectively. We provide their precise definitions in Appendix A.1.

coordinatewise:

$$t \in \mathbf{R}_+ \mapsto \boldsymbol{\nu}(t) := \left(e^{-\delta_1(t)}, \ldots, e^{-\delta_n(t)}\right). \tag{12}$$

The principal curve is continuous and weakly decreasing. It starts at the top corner $(1, \ldots, 1)$ and descends toward the profile of bidders' exclusion types $(\underline{u}_1, \ldots, \underline{u}_n)$, below which the bidders stop receiving the good. By construction, all bidders' types times their probabilities are equalized along the principal curve and equal e^{-t}. The main result of this section is the following characterization of feasibility through the principal curve.

Theorem 1. *A reduced form $\boldsymbol{x}$ is feasible if and only if*

$$B(\boldsymbol{\nu}(t)) \leq 1 \;\; \forall t \in \mathbf{R}_+.$$

Moreover, along the principal curve, the value of B can be expressed as the following function of bidders' δ-transforms alone:

$$B(\boldsymbol{\nu}(t)) = e^{-\sum_{i=1}^n \delta_i(t)} + \int_0^t e^{-\tau} \sum_{i=1}^n \delta_i'(\tau) d\tau. \tag{13}$$

Theorem 1 shows that the Border constraint needs to be checked only along the principal curve. Moreover, the test along this curve can be directly expressed in terms of the sum of bidders' δ-transforms, so verifying feasibility requires no optimization—only checking the sign of a single scalar function of one variable.

Example 1. *To illustrate, consider reduced forms of the power type $x_i(u) = u^{1/\alpha_i - 1}$, where $\alpha_i \in (0, 1)$. Direct computations reveal that $\delta_i(t) = \alpha_i t$, which gives*

$$B(\boldsymbol{\nu}(t)) = e^{-t \sum_{i=1}^n \alpha_i} + (1 - e^{-t}) \sum_{i=1}^n \alpha_i. \tag{14}$$

Applying the theorem, we conclude that $\boldsymbol{x}$ is feasible if and only if $\sum_{i=1}^n \alpha_i \leq 1$.

We now sketch the logic behind the proof of the theorem, leaving additional details to Appendix A.2. For each level $s \in [0, 1]$, consider the following auxiliary problem:

$$G(s) := \max_{\boldsymbol{u} \in [0,1]^n} s^n + \sum_{i=1}^n \int_{u_i}^1 x_i(u) du \;\; \text{s.t.} \;\; \prod_{i=1}^n u_i \geq s^n. \tag{15}$$

Its value dominates the Border constraint: evaluating G at $s^n = \prod_{i=1}^n u_i$ shows that $G(s) \geq B(\boldsymbol{u})$ for every cutoff vector, whereas, at any maximizer, the objective is weakly below B. Hence, $\boldsymbol{x}$ is feasible if and only if $G(s) \leq 1$ for all s. The gain from replacing $\prod_{i=1}^n u_i$ with the constant s^n in the objective is that the auxiliary problem becomes concave, since each tail integral is concave in u_i and the constraint set is convex. Using concavity and (generalized) first-order conditions, we show in the appendix that (15) always admits a maximizer at which the product

$u_i x_i(u_i)$ is common across bidders. But such points are precisely the points of the principal curve, which, by construction, equalizes bidders' type-probability products and—dwelling at jumps—traverses every common value of these products. Finally, evaluating B at $\boldsymbol{\nu}(t)$ and changing the variable of integration along the curve, $\int_{\nu_i(t)}^1 x_i(u)du = \int_0^t e^{-\tau}\delta_i'(\tau)d\tau$, yields the expression in (13).

To sum up, this unidimensional reduction provides a complete and tractable test for feasibility. Just as important, it prepares the ground for our next steps, namely, the characterization of extremality and optimality that we will be covering in the subsequent sections.

6 Extremality

In this section, we study the structure of extremal reduced forms and the structure of the corresponding allocations that induce them.

6.1 Extremal reduced forms

Whenever bidders' revenue functions H_i are convex in x_i—as they are in many applications with endogenous valuations—the objective in (8) is convex. Maximizing a continuous convex function over a convex compact set always yields at least one solution at an extreme point. In such cases, the optimal auction can be implemented by an extremal reduced form. Even when the objective is not convex, extremal reduced forms remain essential. They serve as the "building blocks" of the feasible set, since any feasible reduced form can be expressed as a mixture of them.[15]

This perspective of extremality was first emphasized in the influential paper Kleiner et al. [2021], which demonstrated the surprising tractability of extreme points of feasible symmetric reduced forms. That paper also showed how they illuminate problems in auctions, delegation, and persuasion. Here, we take the next natural step in this agenda by studying extreme points of (potentially asymmetric) feasible reduced forms. We show that this broader set also admits a tractable description, one that can be directly exploited in economic applications.

At an intuitive level, extremal reduced forms correspond to "corners" of the feasible set, where Border's condition is maximally tight. If the inequality were slack anywhere, then the reduced form could be perturbed in different directions and written as a convex combination of other feasible points. Our main result in this section establishes that extremality arises precisely when the feasibility constraints along the principal curve bind at every point.

Theorem 2. *A feasible reduced form $\boldsymbol{x}$ is extremal if and only if*

$$B(\boldsymbol{\nu}(t)) = 1 \;\; \forall t \in \mathbf{R}_+,$$

[15]For example, see Kleiner et al. [2021], who explained that every feasible reduced form can be written as an integral over extremal ones due to the Choquet theorem.

or equivalently:

$$\sum_{i=1}^{n} \delta_i(t) = \min\{t, T\} \quad \forall t \in \mathbf{R}_+ \;\; \textit{for some } T \in [0, \infty]. \tag{16}$$

The theorem provides a unidimensional test analogous to the feasibility test developed in the previous section. By construction, the principal curve starts at the northeast corner with $B(\boldsymbol{\nu}(0)) = 1$. Then, feasibility requires that $B(\boldsymbol{\nu}(\cdot))$ never climb back above one, whereas extremality requires that it remain equal to one throughout.

The second part of the theorem builds on the representation of the Border constraint along the principal curve in (13) and delivers an explicit characterization of extremal reduced forms in terms of their δ-transforms, that is, (16). To see it, differentiate the expression for the Border constraint along the principal curve to obtain

$$\frac{d}{dt} B(\boldsymbol{\nu}(t)) = \left(\sum_{i=1}^{n} \delta_i'(t) \right) \left(e^{-t} - e^{-\sum_{i=1}^{n} \delta_i(t)} \right). \tag{17}$$

Thus, $B(\boldsymbol{\nu}(\cdot))$ stays at one if and only if, at almost every t, the principal curve either does not move, $\sum_{i=1}^{n} \delta_i'(t) = 0$, or satisfies $\sum_{i=1}^{n} \delta_i(t) = t$. Since a motionless curve can never rejoin the set $\sum_{i=1}^{n} \delta_i(t) = t$ once time has passed it by, the curve must ride this set up to some time T and stand still forever after, which is exactly (16).

The stopping time T has a direct economic interpretation in terms of exclusion. At this time, the principal curve stops at the profile of bidders' exclusion types $\underline{\boldsymbol{u}} = \boldsymbol{\nu}(T)$.[16] Since bidders' interim winning probabilities for types below the exclusion types are zero, the ex ante probability of selling the good satisfies

$$B(\mathbf{0}) = B(\underline{\boldsymbol{u}}) - \prod_{i=1}^{n} \underline{u}_i = 1 - e^{-T}.$$

Consequently, an extremal reduced form has no exclusion precisely when its principal curve never terminates, i.e., when $T = \infty$.

Example 1 (continued). Since $\sum_{i=1}^{n} \delta_i(t) = t \sum_{i=1}^{n} \alpha_i$, (16) holds—that is, a reduced form of the power type is extremal—if and only if $\sum_{i=1}^{n} \alpha_i = 1$. In that case, $T = \infty$, so the reduced form also has no exclusion.

Although the extremality condition restricts only the sum of bidders' δ-transforms, it pins down their interim winning probabilities along the principal curve. By construction, for each bidder i, $\nu_i(t) x_i(\nu_i(t)) = e^{-t}$ at every time outside bidder i's dwell intervals,[17] whereas extremality states that $\prod_{j=1}^{n} \nu_j(t) = e^{-t}$. Dividing the latter by the former yields the following expression

[16] Individual coordinates, however, may reach them earlier. Specifically, δ_i becomes constant at the first time the principal curve hits bidder i's exclusion type, which is $-\ln\left(\underline{u}_i x_i\left(\underline{u}_i\right)\right)$, and stays constant from then on, while at least one of the coordinates keeps changing before T.

[17] Within a dwell interval, the curve traverses a jump of x_i, and the identity holds with $x_i(\nu_i(t))$ replaced by the probability coordinate $e^{\delta_i(t)-t}$.

for bidder i's interim winning probability along the principal curve:

$$x_i(v_i(t)) = \prod_{j \neq i} v_j(t). \tag{18}$$

In other words, bidder i with type $v_i(t)$ wins precisely against the type profiles in the box $\prod_{j \neq i} \left[0, v_j(t)\right)$ and loses otherwise. That is, bidder i beats every other bidder j whose type u_j lies below that bidder's position $v_j(t)$ on the principal curve. As we explain in the next subsection, this "rectangular" structure of winning regions is the defining geometric feature of the allocations that induce extremal reduced forms.

Figure 1 illustrates this point with a simple numerical example. In this figure, the black line corresponds to the principal curve with u_1 and u_2 on the horizontal and vertical axes, respectively, constructed from an extremal reduced form. For example, bidder 1 with type $u_1 = \frac{1}{2}$ wins with probability $\frac{1}{2}$, and with type $u_1 = \frac{3}{4}$ wins with probability $\frac{3}{4}$. Using (18), we can recover the reduced form consistent with the principal curve shown in the figure:

$$x_1(u_1) = u_1 \cdot \mathbf{1}_{[1/4, 1/2) \cup [3/4, 1]}(u_1) + 1/2 \cdot \mathbf{1}_{[1/2, 3/4)}(u_1), \tag{19}$$

$$x_2(u_2) = u_2 \cdot \mathbf{1}_{[1/4, 1/2) \cup [3/4, 1]}(u_2) + 3/4 \cdot \mathbf{1}_{[1/2, 3/4)}(u_2). \tag{20}$$

This representation makes clear that the allocation inducing this extremal reduced form is unique: bidder 1 wins in the gray region, bidder 2 wins in the red region, and otherwise the good remains unallocated.

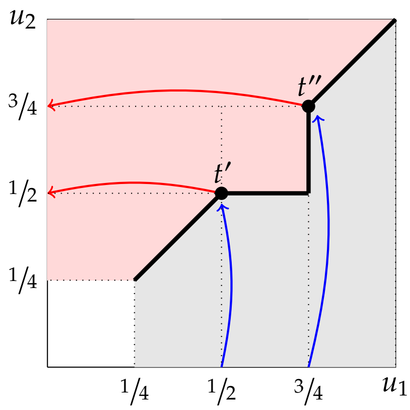

Figure 1: Principal curve (black) induced by an extremal reduced form. The points t', t'' correspond to bidder 1 types $u_1 = \frac{1}{2}, \frac{3}{4}$ and the associated bidder 2's cutoffs $v_2(t') = \frac{1}{2}, v_2(t'') = \frac{3}{4}$.

Though the logic outlined in (18) applies to any number of bidders, the two-bidder case admits a further geometric reading that can also be seen in Figure 1. With $n = 2$, extremality gives $v_2(t) = e^{-t}/v_1(t)$ for $t \leq T$, so that the principal curve is nothing but the completed graph of x_1 drawn in the (u_1, u_2)-square. Whenever bidder 1's coordinate moves, $v_2(t) = x_1(v_1(t))$, and whenever it dwells, the curve descends vertically, filling the corresponding jump of x_1—which is exactly the black line in the figure. It follows that the reduced form satisfies two simple inverse relations: $x_1 = x_2^{-1}$ on $\left[\underline{u}_1, 1\right]$ and $x_2 = x_1^{-1}$ on $\left[\underline{u}_2, 1\right]$. For the two-bidder case, Yang and

Yang [2025] established that this inverse-relation property is equivalent to extremality. With more than two bidders, no such unidimensional inverse relationship exists, and (18) can be viewed as its natural multidimensional generalization.

The proof of the "only if" direction of Theorem 2—asserting that slack anywhere on the principal curve destroys extremality—is rather involved. We show in Appendix A.3 that a feasible reduced form $\boldsymbol{x}$ with a stretch of slack can be perturbed by a function $\boldsymbol{\Delta} \neq \boldsymbol{0}$ so that both $\boldsymbol{x} + \varepsilon\boldsymbol{\Delta}$ and $\boldsymbol{x} - \varepsilon\boldsymbol{\Delta}$ are feasible reduced forms provided that $\varepsilon > 0$ is small enough. The construction is delicate because the principal curve itself moves with the perturbation in a nontrivial fashion, since the δ-transform depends nonlinearly on the reduced form. The "if" direction follows from the analysis of the next subsection by a different route. Specifically, as we explain there, reduced forms with $B(\boldsymbol{\nu}(t)) \equiv 1$ maximize suitable linear functionals uniquely, making them not only extreme but exposed points of the feasible set.

6.2 Allocations inducing extremal reduced forms

The analysis in the previous section showed that extremal reduced forms are entirely determined by their principal curves and that Border's condition must bind along this curve. This already indicates that such reduced forms can arise only from highly structured allocation rules. We now turn to these allocation rules. Understanding them clarifies the geometry behind extremality and allows us to revisit—and sharpen—the classical BIC–DIC equivalence.

The BIC–DIC equivalence, first established by Manelli and Vincent [2010], states that every extremal reduced form can be implemented in dominant strategies. That is, one can find an allocation $\boldsymbol{z}$ that induces the reduced form in question such that each z_i is weakly increasing in bidder i's own type.[18] Their original proof considers piecewise-constant reduced forms together with an approximation argument. Our characterization of extremality provides a more direct route without relying on limiting arguments. Specifically, we explicitly describe allocations that induce extremal reduced forms and establish that they satisfy the DIC property as well as two further properties that could not be seen with the earlier approach. Since any feasible reduced form can be expressed as an integral over extremal ones (see Footnote 15), the equivalence emerges almost automatically. To make this structure precise, we introduce the notion of a score allocation.

Definition 1. *An allocation $\boldsymbol{z}$ is a* ***score allocation*** *if there exists a tuple of weakly increasing, right-continuous functions $\boldsymbol{q} = (q_1, \ldots, q_n) : [0,1] \to \mathbf{R}^n$ such that, for a.e. type profile $\boldsymbol{u}$,*

$$\# \left\{ i : q_i(u_i) = \max_j q_j(u_j) > 0 \right\} \leq 1, \tag{21}$$

[18] See also Gershkov et al. [2013] and Goeree and Kushnir [2023].

and

$$z_i(\boldsymbol{u}) = \begin{cases} \mathbf{1}_{(0,\infty)}(q_i(u_i)) & \text{if } q_i(u_i) \geq \max_{j\neq i} q_j(u_j), \\ 0 & \text{otherwise}. \end{cases} \tag{22}$$

In this definition, $q_i(u_i)$ is the **score** assigned to bidder i with type u_i, and the good is awarded to the bidder with the highest positive score. Condition (21) ensures that the event of two or more bidders tied at the highest positive score is null, so that the score allocation is feasible.[19] Score allocations are familiar in classical settings; e.g., Myerson's optimal auction is a score allocation whenever bidders' MVVs are strictly increasing in their types, with the MVVs themselves serving as scores.

Every score allocation induces an extremal reduced form. To see it, take a tuple of scores $\boldsymbol{q}$ and remark that the associated score allocation $\boldsymbol{z}^{\boldsymbol{q}}$ is the a.e.-unique maximizer of

$$\sum_{i=1}^{n} \int_{[0,1]^n} \Big[q_i(u_i) - \mathbf{1}_{(-\infty,0]}(q_i(u_i))\Big] z_i(\boldsymbol{u}) d\boldsymbol{u}$$

in the space of all allocations. Here, uniqueness is a direct consequence of weak monotonicity of scores, (21), and the fact that the bracketed weights are negative wherever scores are nonpositive. As a result, the reduced form $\boldsymbol{x}^{\boldsymbol{q}}$ induced by that score allocation is the unique maximizer of the corresponding linear functional in the space of all feasible reduced forms. This shows that $\boldsymbol{x}^{\boldsymbol{q}}$ is not merely extremal but in fact exposed.[20] Strikingly, the converse is also true, as shown in the theorem below.

Theorem 3. *A feasible reduced form $\boldsymbol{x}$ is extremal if and only if it is induced by a score allocation. In that case, the allocation inducing $\boldsymbol{x}$ is unique, and it is the score allocation that can be obtained with* ***canonical scores*** *$q_i = u_i x_i$.*

Theorem 3 is a complete description of how extremal reduced forms are induced. First, any tuple of weakly increasing scores satisfying (21) induces an extreme point, and every extreme point arises this way. Second, the ex post allocation behind an extremal reduced form is unique, even though the scores representing it are not; i.e., applying a common strictly increasing, sign-preserving transformation to all scores leaves the allocation unchanged. Third, testing a candidate $\boldsymbol{x}$ requires no search over arbitrary scores, and it suffices to use its canonical scores $u_i x_i$, which are constructed directly from the reduced form itself.

Score allocations have appeared in the social choice literature, but their connection to extremality has not, to the best of our knowledge, been recognized. Mishra and Quadir [2014] referred to them as simple utility maximizers and showed that they can be axiomatized by three basic properties:

[19] To be clear, ties can occur with positive probability among those bidders that are not winning the good.

[20] Recall that a point in a convex subset of a topological vector space is exposed if some continuous linear functional attains its strict maximum at that point, and every exposed point is extremal, e.g., see Kleiner et al. [2021]. Here, the objective is a continuous linear functional when the set of CDFs on the unit interval is viewed as a subset of Lebesgue integrable functions endowed with the norm topology.

1. **Deterministic.** Either some bidder gets the good with certainty or the good is not traded.

2. **Monotone.** Each bidder's allocation is weakly increasing in their own type.

3. **Nonbossy.** A bidder cannot change the allocation of others without also changing their own. Formally,

$$z_i(u_i, \boldsymbol{u}_{-i}) = z_i(u_i', \boldsymbol{u}_{-i}) \implies \boldsymbol{z}(u_i, \boldsymbol{u}_{-i}) = \boldsymbol{z}(u_i', \boldsymbol{u}_{-i}) \ \forall i, \ \ \boldsymbol{u}_{-i},$$

in the sense of Satterthwaite and Sonnenschein [1981]. If bidder i's own outcome does not change when their type is varied, then the outcome for all other bidders must also remain unchanged. This rules out "bossy" manipulations where the bidder could leave their own allocation fixed while reshuffling who among their rivals receives the good.

The second property alone already delivers the BIC–DIC equivalence, whereas the full trio yields more: any BIC allocation is implementable as a lottery over allocations that are also deterministic and nonbossy. These two properties are invisible to the original approach of Manelli and Vincent [2010]. Taken together, these properties have a sharp geometric implication. The region of the type space where the auctioneer keeps the good, $\{\boldsymbol{u} \in [0,1]^n : \boldsymbol{z}(\boldsymbol{u}) = \boldsymbol{0}\}$, and the regions where each bidder i's type u_i wins, $\{\boldsymbol{u}_{-i} \in [0,1]^{n-1} : z_i(u_i, \boldsymbol{u}_{-i}) = 1\}$, are hyperrectangles aligned with the axes. This rectangular structure is exactly the winning-region geometry identified in (18), and it is what makes score allocations induce extremal reduced forms. Not every deterministic and monotone allocation admits this structure; e.g., suppose the auctioneer awards the good to bidder 1 when $u_1 > \frac{u_2+u_3}{2}$, and otherwise allocates it efficiently among bidders 2 and 3. Clearly, the set of (u_2, u_3) such that bidder 1's type u_1 wins is not a rectangle.

Turning to the proof, the "if" direction and uniqueness both follow from the argument preceding the theorem. That argument shows that a score allocation is the a.e.-unique maximizer of a linear functional whose value depends on the allocation only through the induced reduced form. Hence, the induced reduced form is exposed—in particular extremal.[21] The "only if" direction and the third point are direct consequences of (18) and the construction of the δ-transform in (10).

7 Optimality: the case of no exclusion

Having characterized feasibility and extremality, we now turn to the auctioneer's problem. We begin with the case in which the optimal reduced form, $\boldsymbol{x}^*$, has no exclusion, since there our approach to optimality through the principal curve is most transparent. Later, we provide

[21]With finitely many types, the set of feasible reduced forms is a polytope, its extreme points coincide with its exposed points, and they are known to be implemented by score-allocation-type mechanisms; e.g., see Vohra [2011]. With a continuum of types, exposed points are still implemented in this way; this is, in essence, the chain characterization of Delbaen [1974] for cores of convex games. The general theory, however, guarantees only that exposed points are dense in the extreme points, whose structure has otherwise remained unknown. Theorem 3 shows that, in our setting, the two notions in fact coincide.

conditions under which this case indeed describes the optimum and discuss how the analysis changes once exclusion arises.

Recall the conjecture of Section 3 based on Bulow and Roberts [1989], according to which the good should go to the bidder with the highest positive MR evaluated at the optimal interim winning probabilities themselves, that is,

$$x_i^*(u_i) = \mathbf{Pr}\left(\frac{\partial}{\partial x_i}H_i(x_i^*, u_i) \geq \frac{\partial}{\partial x_j}H_j(x_j^*, u_j)\ \ \forall j \neq i | u_i\right) \tag{23}$$

whenever bidder i's MR is positive, and $x_i^*(u_i) = 0$ otherwise. The difficulty is that this condition is self-referential, since the ranking of bidders depends on the very reduced form we are trying to determine. The only case where this circularity disappears is Myerson's linear model, in which bidders' MRs

$$\frac{\partial}{\partial x_i}H_i(x_i, u_i) = \zeta_i(u_i), \tag{24}$$

where ζ_i is bidder i's MVV in the quantile space, do not depend on interim winning probabilities.[22]

The results of the preceding sections allow us to remove this circularity and express the conjecture directly in terms of the principal curve. To streamline the exposition and highlight the main idea, we specialize to the case in which the optimal reduced form is continuous on $[0, 1)$ and bidders' MRs at the optimum, $\frac{\partial}{\partial x_i}H_i(x_i^*(u_i), u_i)$, are weakly increasing in their own types. We also require these MRs to be strictly increasing whenever positive. The general characterization in Theorem 4 shows that, in regular environments, the optimal reduced form indeed has the properties just listed, and that it satisfies the system we arrive at below.

Observe that (23) describes the score allocation in which bidder i's score is own MR. By Theorem 3, such an allocation induces an extremal reduced form and thus can also be induced by the canonical scores, i.e., for almost every $\boldsymbol{u}$,[23]

$$u_i x_i^*(u_i) \geq u_j x_j^*(u_j)\ \ \forall j \neq i \iff \frac{\partial}{\partial x_i}H_i(x_i^*(u_i), u_i) \geq \frac{\partial}{\partial x_j}H_j(x_j^*(u_j), u_j)\ \ \forall j \neq i. \tag{25}$$

In other words, bidders' MRs induce the same ranking of winners as their canonical scores. Since both sets of scores are continuous in types below the top, whenever canonical scores are tied, MRs are tied as well. But the canonical scores are tied precisely along the principal curve, and, absent exclusion, the curve never stops descending from the northeast corner. At each time t, every bidder whose coordinate has left the top sits at type $v_i^*(t) = e^{-\delta_i^*(t)} < 1$ with winning

[22] As explained in Section 2.1, bidder i's MVV as a function of their type θ_i equals $\theta_i - \frac{1-F_i}{f_i}$. In the quantile space, this object can be expressed as $\zeta_i = v_i - (1 - u_i)v_i'$, where v_i is the inverse of F_i.

[23] Equation (25) also holds for u_i below the exclusion type $\underline{u}_i$. To see it, remark that, in the auction with no exclusion, at least one bidder j is never excluded, i.e., $\underline{u}_j = 0$. By (23), that bidder's MR is strictly positive at all positive types, so each bidder i with $u_i < \underline{u}_i$ loses to bidder j on both sides of (25)—their canonical score is zero and their MR is not maximal.

probability $x_i^*(v_i^*(t)) = e^{\delta_i^*(t)-t}$ and canonical score

$$v_i^*(t)x_i^*(v_i^*(t)) = e^{-t}.$$

By (25), these bidders' MRs must share a common value denoted by $p^*(t)$.

A bidder whose coordinate is still at the top, $v_i^*(t) = e^{-\delta_i^*(t)} = 1$, is in a different position. The jump of x_i^* at $u_i = 1$ means that all of this bidder's types have canonical scores lower than e^{-t}, so, by (25), their MRs fall below $p^*(t)$. This, however, is a statement about MRs at winning probabilities up to $x_i^*(1-)$. Joining the principal curve at time t would give the bidder's top type the winning probability e^{-t}, and the conjecture is silent about the MR there. To complete the description, we require that no such bidder wants to join, that is, that the MR of its top type at the entry probability e^{-t} be at most the common MR,

$$\frac{\partial}{\partial x_i} H_i(e^{-t}, 1) \leq p^*(t).$$

As we explain in the sequel, it is exactly the condition that the auctioneer's problem delivers in regular environments.

Recalling from Theorem 2 that extremality with no exclusion means the δ-transforms sum to t, we are led to the following system:

$$\frac{\partial}{\partial x_i} H_i\left(e^{\delta_i^*(t)-t}, e^{-\delta_i^*(t)}\right) \leq p^*(t) \text{ with equality whenever } \delta_i^*(t) > 0 \ \ \forall i, \qquad \sum_{i=1}^{n} \delta_i^*(t) = t. \tag{26}$$

Thus far, (26) is a set of conditions we expect the optimum to satisfy, but it can also be read as a way to find it. To this end, treat (26) as a system of $n+1$ equations in $n+1$ unknowns parametrized by time t. The system involves only positions—no rates, nothing linking one time to another—so it can be solved at each time separately, directly from the primitives. The main result of the section that follows is that, in regular environments where the optimum has no exclusion, the path satisfying (26) characterizes the optimal auction. Section 8 studies environments in which it is optimal to exclude some bidders' types.

7.1 Main result

We begin with a notion of regularity that extends Myerson's regularity—MVVs are strictly increasing—to nonlinear environments so that no ironing is needed. We say that an environment is **regular** if the following two conditions hold.

(A) For each bidder i and all $x_i, u_i \in (0,1)$,

$$\frac{\partial}{\partial \ln x_i}\left(\frac{\partial}{\partial x_i} H_i(x_i, u_i)\right) < \frac{\partial}{\partial \ln u_i}\left(\frac{\partial}{\partial x_i} H_i(x_i, u_i)\right).$$

(B) The solution $(\boldsymbol{\delta}^{\mathrm{s}}(t), p^{\mathrm{s}}(t))$ to (26)—unique under Condition (A) at each time t, with the superscript "s" for (26)—is such that p^{s} is strictly decreasing and each δ_i^{s} is strictly increasing whenever positive.[24]

To understand the content of regularity, consider again Myerson's linear model in (24). Since $\frac{\partial H_i}{\partial x_i}$ does not depend on the winning probability, the left-hand side of Condition (A) vanishes, and the condition reads $u_i \zeta_i' > 0$, that is, bidders' MVVs are strictly increasing—precisely Myerson's regularity. Condition (B) then holds automatically, since bidders' MRs in (26) do not vary with time.[25] In nonlinear environments, bidders' MRs in that system vary with time through the winning probabilities. Condition (A) then means that bidder i's MR is strictly decreasing in δ_i at each time t, and Condition (B) becomes an additional restriction that rules out ironing. Appendix A.5 further discusses both conditions and unpacks Condition (B) in terms of primitives.

By construction, the path $\boldsymbol{\delta}^{\mathrm{s}}$ equalizes bidders' MRs at every time but never contemplates exclusion, since $\sum_{i=1}^{n} \delta_i^{\mathrm{s}}(t) = t$. We refer to $\boldsymbol{\delta}^{\mathrm{s}}$ as the **iso-MR path** and to p^{s} as the **common MR**, and we say that bidder i is **active** at time t if $\delta_i^{\mathrm{s}}(t) > 0$. Each coordinate of the iso-MR path is a weakly increasing and absolutely continuous function, and thus it can be inverted using (11) to obtain the following CDF:

$$x_i^{\mathrm{s}}(u_i) = e^{\delta_i^{\mathrm{s}}(t_i) - t_i}\Big|_{t_i = (\delta_i^{\mathrm{s}})^{-1}(-\ln u_i)}. \tag{27}$$

By Theorem 1, the tuple of CDFs $\boldsymbol{x}^{\mathrm{s}}$ is a feasible reduced form, which we term the **iso-MR reduced form**. Our main result explains that the iso-MR reduced form is optimal provided that the common MR stays positive.

Theorem 4. *Suppose the environment is regular and $p^{\mathrm{s}} \geq 0$. Then, the iso-MR reduced form is the unique optimal reduced form. Furthermore, it is extremal, has no exclusion, and satisfies (23).*

Theorem 4 explicitly constructs the optimal auction in regular nonlinear environments where the optimum has no exclusion. In such an auction, each bidder's type is mapped to the time at which the iso-MR path reaches it, $t_i = (\delta_i^{\mathrm{s}})^{-1}(-\ln u_i)$, and the bidder's canonical score takes the particularly simple form $u_i x_i^*(u_i) = e^{-t_i}$. The good then goes to the bidder whose type is reached by the path the earliest, so that the good is sold with certainty. At every type u_i in $[\underline{u}_i^*, 1)$, bidder i's MR in the optimal auction equals $p^{\mathrm{s}}(t_i)$—the common MR at the time the path reaches their type. Consequently, the ranking by canonical scores in Theorem 4 is exactly the MR ranking conjectured in (23).

The role of the condition $p^{\mathrm{s}} \geq 0$ is to guarantee that no exclusion is optimal. Intuitively, it says that, at every time, there is at least one bidder with a positive MR along the optimal principal curve. If instead p^{s} turned negative at some time, the auctioneer would prefer to stop selling to the types reached after that time, and the optimum would involve exclusion—the case

[24] At $t = 0$, we use the convention $p^{\mathrm{s}}(0) = \max_i \frac{\partial}{\partial x_i} H_i(1, 1)$ so that p^{s} is continuous.

[25] Using (26), we may express δ_i^{s} as a function of p^{s} alone and then find the value of p^{s} by solving $\prod_{i=1}^{n} e^{-\delta_i^{\mathrm{s}}} = e^{-t}$. Since MVVs do not depend on time, raising t only lowers p^{s}.

we take up in the sequel. It is worth mentioning that, since p^{s} is defined explicitly by (26), this condition is directly on primitives. In Lemma 9 in the appendix, we unpack it directly in terms of $\frac{\partial H_i}{\partial x_i}$.

In the linear model, bidders' MRs are pinned down by their types alone, and the optimal auction reduces to Myerson's auction, which awards the good to the bidder with the highest positive MVV. In nonlinear environments, MRs depend also on bidders' winning probabilities, so bidder i's MR in the optimal auction—unlike MVV—incorporates information about other bidders. For instance, whether bidder i's type u_i beats bidder j's type u_j depends on the entire environment, including the type distributions and revenue functions of the remaining bidders. In Myerson's auction, by contrast, this comparison involves the two bidders alone. Strikingly, despite the presence of nonlinearities that render standard approaches inapplicable, the characterization of the optimum coincides formally in the general and linear cases. In both, the optimal reduced form is obtained by solving (26) at each time and inverting the resulting δ-transforms via (27).

We defer the discussion of the proof to Section 7.3 below and first revisit our motivating example through the lens of Theorem 4.

Example 2 (Illustrative example, continued)**.** *To see Theorem 4 at work, consider again the parametric specification of Section 3, in which bidder i's revenue function in the quantile space is $H_i = u_i^{1/\beta_i} x_i^2$. Condition (A) holds since*

$$\frac{\partial}{\partial \ln x_i}\left(\frac{\partial H_i}{\partial x_i}\right) = \frac{\partial H_i}{\partial x_i} < \frac{1}{\beta_i}\frac{\partial H_i}{\partial x_i} = \frac{\partial}{\partial \ln u_i}\left(\frac{\partial H_i}{\partial x_i}\right).$$

Along the parametrization underlying the δ-transform, bidder i's MR equals

$$\frac{\partial}{\partial x_i} H_i\left(e^{\delta_i - t}, e^{-\delta_i}\right) = 2e^{-(1/\beta_i - 1)\delta_i} e^{-t}.$$

Equalizing these MRs across bidders and requiring that bidders' δ-transforms add up to t, we obtain

$$\delta_i^{\mathrm{s}}(t) = \frac{\kappa}{1/\beta_i - 1} t, \qquad p^{\mathrm{s}}(t) = 2e^{-(1+\kappa)t},$$

where κ is the normalization constant so that $\sum_{i=1}^n \delta_i^{\mathrm{s}}(t) = t$. This shows that Condition (B) holds as well, and so the environment is regular with $p^{\mathrm{s}} \geq 0$.

Inverting $\boldsymbol{\delta}^{\mathrm{s}}$ via (27) yields that the optimal reduced form is given by

$$x_i^*(u_i) = u_i^{(1/\beta_i - 1)/\kappa - 1}.$$

The canonical scores are $u_i x_i^*(u_i) = u_i^{(1/\beta_i - 1)/\kappa}$, so bidder i beats bidder j if and only if $u_i^{1/\beta_i - 1} \geq u_j^{1/\beta_j - 1}$. Returning to the original type space, this shows that the optimal auction awards the good to the bidder with the highest value of $(F_i(\theta_i))^{1/\beta_i - 1}$, as claimed in Section 3.

The example confirms the conjectured structure of the optimum, which, as explained in

Section 3, favors stronger bidders relative to Myerson's auction. Intuitively, in the endogenous-valuation model, the auctioneer internalizes the effect of changing interim winning probabilities on bidders' investment incentives. By equalizing MRs, the auctioneer shifts winning probability toward bidders whose investments generate the most revenue, reducing the duplication of investment costs. In the next section, we show, using Theorem 4, that this phenomenon holds beyond that single example.

7.2 Application to discrimination in the endogenous-valuation model

Consider a setting with endogenous valuations, in which the reduced-form utility of every bidder i is equal to $\omega_i(x_i, \theta_i) = \theta_i h(x_i)$, where $h', h'' > 0$. As explained in Section 2.3, bidder i's revenue function J_i takes the form

$$J_i(x_i, \theta_i) = \left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)}\right) h(x_i),$$

where the term in parentheses is that bidder's MVV. Bidders' MVVs are assumed to be strictly increasing to avoid dealing with ironing that does not add much economic insight. This setting generalizes that of Example 2 and is chosen because Myerson's virtual values arise naturally as scaling factors of bidders' revenue functions, even though the utilities are nonlinear.[26] As a result, we can naturally compare bidders' strengths within this framework in terms of their MVVs and contrast the optimal auction to Myerson's auction, which allocates the good to the bidder with the highest positive MVV.

One natural order that has been employed to rank bidders' strength (e.g., see Carroll and Segal [2019]) is the **hazard ratio order**. A cumulative distribution function F dominates another cumulative distribution function G in this order, denoted $F \geq_{hr} G$, when their survival functions satisfy

$$(1 - F(\theta'))(1 - G(\theta)) \geq (1 - F(\theta))(1 - G(\theta')) \;\; \forall \theta \leq \theta'.$$

See Chapter 1 in Shaked and Shanthikumar [2007]. As explained in this book, this ranking implies $\frac{1-F}{f} \geq \frac{1-G}{g}$ over the common support of these two distributions. Intuitively, high types are more likely under F than under G conditional on any interval of the form $[\theta, \infty)$. Hence, a bidder whose types are drawn from F should receive a higher information rent than the one whose types are drawn from G. As a result, if there are two distinct bidders i, j with the same value $\theta_i = \theta_j$ and $F_i \geq_{hr} F_j$, then Myerson's auction allocates the good more often to the weaker bidder j.

As shown in the previous section, to determine the optimum, one needs to work in the quantile space of bidders' types. The hazard ratio order, however, is silent about bidders' strength in the quantile space, where their MVVs can be expressed using $v_i = F_i^{-1}$ as in Footnote

[26] The primitive utility over action-type (a_i, θ_i) pairs that gives rise to such an ω_i can be recovered via $\min_x(\theta_i h(x) + a_i)/x$ when the investment costs are normalized to $C_i(a_i) = a_i$; see Andreyanov et al. [2026].

22, that is,

$$\zeta_i(u) = v_i(u)\left(1 - (1-u)\frac{v_i'(u)}{v_i(u)}\right). \tag{28}$$

This happens because the second term in (28) can be lower for bidder i than for bidder j when F_i is much more unequal than F_j. This is so even if the two bidders have the same quantile $u_i = u_j$ and $F_i \geq_{hr} F_j$, so that $v_i \geq v_j$. In our comparison of the optimal and Myerson's auctions, we use one more standard stochastic order allowing us to rank bidders' strengths in the quantile space as well.

Following Chapter 4 in Shaked and Shanthikumar [2007], we say that a cumulative distribution function F dominates another cumulative distribution function G in the **star order**, denoted $F \geq_* G$, when their quantile functions satisfy

$$F^{-1}(u')G^{-1}(u) \geq F^{-1}(u)G^{-1}(u') \ \ \forall u \leq u'.$$

The star order means that G is more equal relative to its mean than F. There are many parametric families of distributions for which $F \geq_{hr} G$ automatically implies $G \geq_* F$. For example, this is true for (1) Normal, Logistic, Gumbel distributions ordered by a location parameter; (2) Exponential distributions ordered by a scale parameter; (3) Power distributions ordered by a shape parameter; as well as (4) distributions of the form in Example 2. A notable family for which this is not true is Pareto distributions, where $F \geq_{hr} G$ implies $F \geq_* G$ due to a thicker tail.

Going back to our auction framework, consider two bidders i and j with $F_i \geq_{hr} F_j$ and, in addition, $F_j \geq_* F_i$. Then, not only are bidder i's values, v_i, stochastically higher, but also per-unit information rents, $(1-u)\frac{v_i'}{v_i}$, are lower. This means that the stronger bidder i is allocated the good more often in the quantile space than bidder j. Equipped with these two approaches to ranking bidders' strengths, we can now state our result formally.

Proposition 1. *Consider the model as above, and assume that the environment is regular with $p^{\mathrm{s}} \geq 0$ and that*

$$F_1 \geq_{hr} F_i \geq_{hr} F_n, \quad F_n \geq_* F_i \geq_* F_1 \ \ \forall i.$$

Then, the optimal reduced form $\boldsymbol{x}^*$ and the reduced form induced by Myerson's auction $\boldsymbol{x}^{\mathrm{M}}$ satisfy:

1. *The strongest bidder 1 is always favored by the optimum, $x_1^* \geq x_1^{\mathrm{M}}$;*

2. *The weakest bidder n is always discriminated against by the optimum, $x_n^* \leq x_n^{\mathrm{M}}$.*

Proposition 1 compares the interim winning probabilities of the strongest and weakest bidders across two auctions—optimal and Myerson's. It shows that the optimum allocates the good more often to the strongest bidder at the expense of the weakest one. It is instructive to unpack the argument behind this proposition to see the economics at play as well as the role of its assumptions. Let $\boldsymbol{\delta}^*$ and $\boldsymbol{\delta}^{\mathrm{M}}$ denote the δ-transforms of the two reduced forms $\boldsymbol{x}^*$ and $\boldsymbol{x}^{\mathrm{M}}$. Below, we assume that all bidders are active in both auctions in the sense that $\delta_i^*(t), \delta_i^{\mathrm{M}}(t) > 0$ for all $t > 0$. In Appendix A.4.3, we complete the argument by handling bidders that are dwelling at the top of either principal curve. We also show that, under the stated assumptions, Myerson's auction has no exclusion, so both principal curves indeed satisfy $\sum_{i=1}^{n} \delta_i(t) = t$ at all times.

As explained in the previous section, Myerson's auction equates bidders' MVVs along its principal curve, that is,

$$\zeta_1\left(e^{-\delta_1^{\mathrm{M}}(t)}\right) = \zeta_i\left(e^{-\delta_i^{\mathrm{M}}(t)}\right) \quad \forall i \neq 1.$$

Since the strongest bidder's distribution F_1 is more equal in the sense of the star order than any other bidder i's distribution F_i, their MVVs satisfy $\zeta_1 \geq \zeta_i$ whenever $\zeta_i \geq 0$. As a result, bidder 1's δ-transform corresponding to Myerson's auction is the highest one.

By equating MVVs, Myerson's auction completely ignores the curvature of bidders' revenue functions—the channel through which interim winning probabilities affect ex ante investments—leaving bidder 1 with the highest MR along its principal curve:

$$\zeta_1\left(e^{-\delta_1^{\mathrm{M}}(t)}\right) h'\left(e^{\delta_1^{\mathrm{M}}(t)-t}\right) \geq \zeta_i\left(e^{-\delta_i^{\mathrm{M}}(t)}\right) h'\left(e^{\delta_i^{\mathrm{M}}(t)-t}\right) \quad \forall i \neq 1. \tag{29}$$

This suggests that the auctioneer can raise revenue by reallocating more winning probability to bidder 1. Since the cumulative height along the principal curve is fixed at t, i.e., $\sum_{i=1}^{n} \delta_i^{\mathrm{M}}(t) = t$, the only such reallocation must be at the expense of other weaker bidders. But this is precisely what the optimum does—it equates MRs along its principal curve, internalizing the curvature in h, so that

$$\zeta_1\left(e^{-\delta_1^{*}(t)}\right) h'\left(e^{\delta_1^{*}(t)-t}\right) = \zeta_i\left(e^{-\delta_i^{*}(t)}\right) h'\left(e^{\delta_i^{*}(t)-t}\right) \quad \forall i \neq 1, \tag{30}$$

and the cumulative height stays the same, $\sum_{i=1}^{n} \delta_i^{*}(t) = t$. Since bidders' MRs are strictly decreasing in own δ-transforms (due to regularity), we should have $\delta_1^{*}(t) \geq \delta_1^{\mathrm{M}}(t)$, which is equivalent to the first point in Proposition 1.

As can be seen from our argument, the only place where the star order is used is to conclude that bidder 1 is favored in the quantile space in Myerson's auction, which is the same as $\delta_1^{\mathrm{M}}(t) \geq \delta_i^{\mathrm{M}}(t)$. So, the conclusion of this proposition holds under much weaker conditions; in fact, $x_1^{*}(u) \geq x_1^{\mathrm{M}}(u)$ for all high enough u without any extra structure. On the other hand, not much can be said at this level of generality about the bidders other than the strongest and the weakest one. The reason is that, relative to Myerson's auction, the optimum simultaneously reallocates interim winning probability from them (to the strongest bidder) and to them (from the weakest bidder). Which force dominates depends on fine details of the model.

7.3 Methodology behind Theorem 4

The derivation of (26) in Section 7 was heuristic—it presumed the conjecture (23), extracted its implications for bidders on the principal curve, and simply required that no bidder at the top want to join. The proof of Theorem 4 must instead establish that the iso-MR reduced form solves the auctioneer's problem among all feasible reduced forms, conjectured or not. This is where the δ-transforms fully pay off—they turn the auctioneer's problem into a textbook optimal control problem in which (26) reappears as the first-order condition, and regularity guarantees its sufficiency.

Given a reduced form $\boldsymbol{x}$ (not necessarily feasible), the auctioneer's revenue can be written as

$$\sum_{i=1}^{n}\int_0^\infty e^{-t}R_i(\delta_i(t),t)dt, \quad \text{where} \quad R_i(\delta_i,t) := \int_0^{\delta_i}\frac{\partial}{\partial x_i}H_i(e^{\tau-t},e^{-\tau})d\tau. \tag{31}$$

The flow payoff R_i aggregates the MRs of bidder i's types along the parametrization underlying the δ-transform up to position δ_i, so, by construction,

$$\frac{\partial}{\partial\delta_i}R_i(\delta_i,t) = \frac{\partial}{\partial x_i}H_i\left(e^{\delta_i-t},e^{-\delta_i}\right) \tag{32}$$

is exactly the MR that (26) equalizes.

Putting the revenue and the feasibility characterization in Theorem 1, we obtain the following representation of the auctioneer's problem.

Proposition 2. *A reduced form is optimal if and only if its δ-transforms solve the following program:*

$$\max_{\boldsymbol{\delta}\in\mathscr{D}^n}\int_0^\infty e^{-t}\sum_{i=1}^{n}R_i(\delta_i(t),t)dt \;\; \text{s.t.} \;\; \int_0^t e^{-\tau}\sum_{i=1}^{n}\delta_i'(\tau)d\tau \le 1-e^{-\sum_{i=1}^{n}\delta_i(t)} \;\; \forall t\in\mathbf{R}_+. \tag{33}$$

Moreover, there exists an optimal reduced form that is extremal if and only if there exists a solution to (33) *satisfying*

$$\sum_{i=1}^{n}\delta_i'(t) = \mathbf{1}_{[0,T]}(t) \;\; \textit{for some } T\in[0,\infty]. \tag{34}$$

Proposition 2 restates the optimal auction problem as maximization over principal curves parametrized by $\boldsymbol{\delta} = (\delta_1,\dots,\delta_n)\in\mathscr{D}^n$. Specifically, the program in (33) is an optimal control problem in which $\boldsymbol{\delta}$ serves as the state variable and the rates $\boldsymbol{\delta}'$ are the controls. The feasibility constraint is then a pure state constraint in the sense of Chapter 6 in Seierstad and Sydsaeter [1986]. Indeed, it can be written as $\mu(t)\le 1-e^{-\sum_{i=1}^n\delta_i(t)}$, where μ is an auxiliary state variable with dynamics $\mu'(t) = e^{-t}\sum_{i=1}^n\delta_i'(t)$ and $\mu(0)=0$. The second part of the proposition, which follows from Theorem 2, sharpens the search—any extremal reduced form corresponds to a front-loaded path that assigns probability at the maximal feasible rate up to the cutoff time T and ceases thereafter. Since Theorems 1 and 2 do not rely on regularity, Proposition 2 is valid in irregular environments as well.

To give (26) a variational meaning, apply the maximum principle to (33) along the trajectories satisfying (34) with $\delta_i'>0$ whenever δ_i is positive. Since the Hamiltonian is linear in the controls, the unit rate of growth is assigned to the coordinates with the highest MRs, and p^{s} equals that maximum. The growing coordinates thus share the same MR, whereas those staying at zero have weakly lower ones—which is exactly (26). Hence, (26) is a necessary optimality condition along such trajectories, and the iso-MR path is the trajectory that satisfies it at each time.

The proof itself does not rely on the maximum principle. In regular environments, bidders' flow payoffs are strictly concave in own δ-transforms, so the control problem in (33) is concave. For such problems, a feasible path is the unique global optimum whenever it satisfies the first-order conditions together with appropriate multipliers on the state constraint. In Appendix A.4.2, we construct these multipliers explicitly from p^{s} and verify these conditions for the iso-MR path.

8 Extensions: the case of exclusion

In the previous section, we developed our methodology for characterizing the optimal auction when keeping the good is suboptimal. In some cases, however, exclusion can be beneficial. For instance, in the regular linear model, this happens when bidders' MVVs are negative at the bottom of the type distribution. In the language of the previous section, the common MR eventually turns negative, so the iso-MR reduced form assigns positive winning probabilities to types with negative MRs and is no longer optimal. The question then becomes whom to exclude and how.

Below, we answer this question in nonlinear environments, focusing on two economically important cases—concave environments, covering auctions with risk-averse bidders, as in Section 2.2; and convex environments, covering the endogenous-valuation model, as in Section 2.3. It turns out that the structure of exclusion is very different in the two cases. In concave environments, the auctioneer excludes at the intensive margin, scaling winning probabilities down. By contrast, in convex environments, the auctioneer excludes at the extensive margin, cutting low types off.

8.1 Concave settings

We start with the concave case, in which each H_i is concave in x_i. In regular environments, the common MR strictly decreases and reaches zero at the time

$$T^{\mathrm{s}} = \inf\{t \geq 0 : p^{\mathrm{s}}(t) \leq 0\}.$$

On the iso-MR path at time T^{s}, bidder i's coordinate sits at the type

$$\overline{u}_i := e^{-\delta_i^{\mathrm{s}}(T^{\mathrm{s}})},$$

and, for every bidder active at T^{s},

$$\frac{\partial}{\partial x_i} H_i(x_i^{\mathrm{s}}(u_i), u_i) < 0 = \frac{\partial}{\partial x_i} H_i(x_i^{\mathrm{s}}(\overline{u}_i), \overline{u}_i) \ \ \forall u_i < \overline{u}_i.$$

This suggests that the auctioneer can do better by reducing bidders' winning probabilities relative to the iso-MR reduced form for bidder i's types below $\overline{u}_i$. In what follows, we formalize this idea using fractional score allocations.

Definition 2. *An allocation $\boldsymbol{z}$ is a **fractional score allocation** if there exist scores $\boldsymbol{q}$ as in Definition 1 and right-continuous functions $\boldsymbol{r} = (r_1, \ldots, r_n) : [0,1] \to [0,1]^n$ such that*

$$z_i(\boldsymbol{u}) = \begin{cases} r_i(u_i) & \text{if } q_i(u_i) > \max_{j \neq i} q_j(u_j), \\ 0 & \text{otherwise.} \end{cases} \tag{35}$$

Fractional score allocations naturally generalize score allocations by letting the winner receive the good with own **winning fraction** $r_i(u_i) \in [0,1]$ rather than with certainty, while the auctioneer keeps the good with the remaining probability. Equipped with this definition, we can now characterize the optimum.

Theorem 5. *Suppose the environment is regular with each H_i concave in x_i and supermodular. Then, the auctioneer's problem admits a unique optimal reduced form given by*

$$x_i^*(u_i) = \begin{cases} x_i^{\mathrm{s}}(u_i) & u_i \geq \overline{u}_i, \\ \max \operatorname{argmax}_{x_i \in [0,1]} H_i(x_i, u_i) & u_i < \overline{u}_i. \end{cases} \tag{36}$$

Moreover, $\boldsymbol{x}^*$ is induced by the fractional score allocation whose scores are the canonical scores of the iso-MR reduced form and whose winning fractions are $r_i^* = x_i^*/x_i^{\mathrm{s}}$.[27]

Theorem 5 extends the characterization of optimal auctions to regular concave environments. Here, the good is withheld with positive probability, and yet low types may continue to receive it—the auctioneer excludes only the types whose MR is negative already at zero winning probability.[28] Economically, withholding at the intensive margin dominates exclusion at the extensive margin because the units of winning probability below each type's unconstrained optimum are precisely the profitable ones. Cutting a type off destroys these units, whereas scaling the winning fraction down removes only the rest. We elaborate on this point in the next section, where we examine the optimal pattern of discrimination in the CRA-preferences model.

The proof extends the multipliers constructed for Theorem 4 by setting them to zero after T^{s}, which is consistent precisely because $p^{\mathrm{s}}(T^{\mathrm{s}}) = 0$. The added supermodularity assumption

[27] As a convention, we set $0/0 = 0$.

[28] The optimal allocation is not unique; implementing it through fractional score allocations is one convenient way to induce the optimal reduced form.

ensures that ironing is not needed when bidders' winning probabilities are scaled down. Appendix **??** verifies in addition that the continuation is feasible—the unconstrained optima never exceed the winning probabilities along the principal curve, so the good is never overallocated.

8.2 Application to discrimination in the CRA-preferences model

Before going to the convex case, we use Theorem 5 to examine how the optimal auction favors bidders as a function of their risk attitudes in the CRA-preferences model. It turns out that the optimal allocation discriminates against more risk-averse bidders at high types and against less risk-averse bidders at low types. To make the comparison across risk attitudes as transparent as possible, we assume that all bidders share the same type distribution F with inverse $v = F^{-1}$ and belong to one of two groups, risk-neutral and risk-averse.

A risk-neutral bidder has a linear certainty equivalent equal to x, whereas a risk-averse bidder has a strictly convex certainty equivalent $g(x) \leq x$ satisfying $g(0) = 0$, $g(1) = 1$, $g' > 0$, and $g'' > 0$. There are k risk-neutral bidders, indexed by $i = 1, \ldots, k$, and $n - k$ risk-averse bidders, indexed by $i = k + 1, \ldots, n$. As explained in Section 2.2, bidders' revenue functions in the quantile space are

$$H_i(x_i, u_i) = v(u_i)x_i - (1 - u_i)v'(u_i)g_i(x_i),$$

where $g_i(x) = x$ for the risk-neutral group and $g_i = g$ for the risk-averse one, so each H_i is concave in x_i. Supermodularity, which we assume throughout, completes the requirements of Theorem 5.

Because the environment is concave and symmetric within each group, the optimal auction is group-symmetric as well. Hence, it suffices to characterize the interim winning probabilities and their δ-transforms for a representative risk-neutral bidder, denoted $i = 1$, and a representative risk-averse bidder, denoted $i = n$. The system (26) specializes to

$$\begin{aligned} p &= v\left(e^{-\delta_1}\right) - v'\left(e^{-\delta_1}\right)(1 - e^{-\delta_1}) = v\left(e^{-\delta_n}\right) - v'\left(e^{-\delta_n}\right)(1 - e^{-\delta_n})g'\left(e^{\delta_n - t}\right), \\ t &= k\delta_1 + (n - k)\delta_n, \end{aligned} \tag{37}$$

whose solution gives δ_1^{s}, δ_n^{s}, and the common MR p^{s}. By Theorem 5, the optimum coincides with the iso-MR reduced form at the types reached before T^{s}. Once p^{s} reaches zero at T^{s}, the risk-neutral bidders' types below $\overline{u}_1$—the type at which the common MVV vanishes—are excluded outright, whereas the risk-averse bidders' low types continue to trade at the unconstrained optima. Intuitively, awarding a fractional unit to risk-averse bidders at the bottom of the type distribution lowers information rents while generating only small efficiency losses.

We are interested in comparing x_1^* and x_n^*. Since they are determined through the MR equalization in (37), the optimal pattern of discrimination depends on comparing $\frac{\partial H_1}{\partial x_1}$ and $\frac{\partial H_n}{\partial x_n}$, and thus on comparing g' with 1. In principle, the difference $1 - g'(x)$ may change sign arbitrarily many times. To gain traction and deliver a global result, we henceforth assume that $x - g(x)$ is unimodal, attaining its maximum at some $m \in (0, 1)$. This assumption holds, for instance, in the quadratic specification $g(x) = \alpha x^2 + (1-\alpha)x$, where $m = {}^1\!/\!{}_2$ regardless of $\alpha \in (0,1]$. As shown by GMSZ-CRA, the quadratic model is consistent with many important models of non-expected utility, such as versions of the disappointment-aversion theories of Loomes and Sugden [1986] and Jia et al. [2001], **?**'s loss-averse utility, and modified mean-variance preferences [Blavatskyy, 2010].[29] Under this assumption, the direction of discrimination becomes monotone in types, as shown in the proposition below.

Proposition 3. *Consider the model as above, and assume that the environment is regular and that* $\frac{\partial}{\partial x_n} H_n(m, m^{1/(n-1)}) > 0$. *Then,* $x_1^*(u) \gtreqless x_n^*(u)$ *if and only if* $u \gtreqless m^{1/(n-1)}$ *for all* $u \geq \overline{u}_n$, *whereas* $x_1^*(u) = 0 \leq x_n^*(u)$ *for all* $u < \overline{u}_n$.

Proposition 3 shows that the optimal auction tilts in favor of the more risk-averse bidder at low quantiles and in favor of the less risk-averse bidder at high quantiles. Intuitively, when types are low, allocating probability to risk-averse bidders generates relatively smaller information rents. For high types, the concern shifts toward allocative efficiency, and the auction favors bidders whose valuations respond more strongly to additional probability—namely, the less risk-averse ones. This reversal of priorities across the type distribution is a direct consequence of MR equalization along the optimal principal curve.

Example 3. *Consider the quadratic specification with* $\alpha = 1$ *and the uniform type distribution, so that* $g(x) = x^2$ *and* $v(u) = u$. *To illustrate, fix* $n = 3$.[30] *When* $k = 2$ *(two risk-neutral bidders and one risk-averse bidder), the optimal reduced form satisfies*

$$x_1^*(u) = \frac{2u^2 + 2u - 1}{2u + {}^1\!/\!{}_u} \cdot \mathbf{1}_{[{}^1\!/\!{}_2,1]}(u), \quad x_n^*(u) = \frac{u}{2(1-u)} \cdot \mathbf{1}_{[0,{}^1\!/\!{}_3)}(u) + \frac{2 - u^2 - \sqrt{3 - 2u^2}}{2(1-u)^2} \cdot \mathbf{1}_{[{}^1\!/\!{}_3,1]}(u).$$

Here, T^{s} *is finite, and low types of the risk-averse bidder receive their unconstrained optima,* $x_n^*(u) = \frac{u}{2(1-u)}$, *as predicted by Theorem 5. When* $k = 1$ *(one risk-neutral bidder and two risk-averse bidders), the optimal reduced form satisfies*

$$x_1^*(u) = \frac{12u^2 - 8u + (2u-1)^{3/2}\sqrt{10u - 1} + 1}{8u^2} \cdot \mathbf{1}_{[{}^1\!/\!{}_2,1]}(u), \quad x_n^*(u) = \frac{u+1}{2(1-u) + {}^2\!/\!{}_u}.$$

In contrast to the previous case, T^{s} *is infinite, and the optimum coincides with the iso-MR reduced form everywhere. In both cases, the tipping point in Proposition 3 equals* $m^{1/2} = \frac{1}{\sqrt{2}} \approx 0.707$, *and risk-averse bidders are favored for* $u < \frac{1}{\sqrt{2}}$ *and disadvantaged for* $u > \frac{1}{\sqrt{2}}$.

8.3 Convex settings

We now complete the characterization of optimal auctions with exclusion by taking up the convex case, in which each H_i is supermodular and $\max\{0, H_i\}$ is convex in x_i, as in the

[29] The unimodality condition also holds for $g(x) = \frac{x}{1+\alpha(1-x)}$, which corresponds to Gul's disappointment-averse preferences [Gul, 1991] with linear utility over outcomes.

[30] It is routine to verify that the environment is regular for all n, k (see Appendix A.5 for the discussion of regularity in terms of primitives) and that the added assumption in Proposition 3 is satisfied. The derivations for this example are collected in Appendix **??**.

endogenous-valuation model of Section 2.3. Since $H_i(\cdot, u_i)$ is then nonpositive up to some winning probability and convex beyond it, its maximum over $[0,1]$ is attained at a corner for every type u_i. Exclusion therefore operates at the extensive margin: the optimal reduced form assigns zero winning probability to each bidder's types below own exclusion type, which is now potentially positive.

A natural conjecture, in view of Theorem 5, is to truncate the iso-MR reduced form at the types $\overline{\boldsymbol{u}}$. That is, one sets each bidder's winning probabilities to zero below own type $\overline{u}_i$, which is reached at the time T^{s} when the common MR hits zero. In nonlinear settings, however, the auctioneer can typically do better by excluding more types. To see why, recall from (31) that the revenue of the truncated reduced form can be expressed as

$$\sum_{i=1}^{n} \int_{\overline{u}_i}^{1} H_i(x_i^{\mathrm{s}}(u_i), u_i) du_i = \sum_{i=1}^{n} \int_0^{T^{\mathrm{s}}} e^{-t} R_i(\delta_i^{\mathrm{s}}(t), t) dt + \sum_{i=1}^{n} \int_{T^{\mathrm{s}}}^{\infty} e^{-t} R_i(\delta_i^{\mathrm{s}}(T^{\mathrm{s}}), t) dt.$$

Truncating instead at the slightly earlier time $T^{\mathrm{s}} - \varepsilon$ changes the revenue at the rate

$$-\sum_{i=1}^{n} \delta_i^{\mathrm{s}\prime}(T^{\mathrm{s}}) \int_{T^{\mathrm{s}}}^{\infty} e^{-t} \frac{\partial}{\partial x_i} H_i\left(e^{\delta_i^{\mathrm{s}}(T^{\mathrm{s}})-t}, e^{-\delta_i^{\mathrm{s}}(T^{\mathrm{s}})}\right) dt = -\sum_{i=1}^{n} \delta_i^{\mathrm{s}\prime}(T^{\mathrm{s}}) \cdot \overline{u}_i H_i(x_i^{\mathrm{s}}(\overline{u}_i), \overline{u}_i),$$

where the second expression follows from changing the variable of integration from t to bidder i's winning probability. For every bidder active at T^{s}, $\frac{\partial}{\partial x_i} H_i(x_i^{\mathrm{s}}(\overline{u}_i), \overline{u}_i) = 0$, which implies $H_i(x_i^{\mathrm{s}}(\overline{u}_i), \overline{u}_i) \leq 0$ by convexity of the positive part of H_i. The inequality is strict outside of degenerate cases such as Myerson's linear model. In fact, since these terms differ across bidders, truncating different bidders' winning probabilities at different types can do even better.

To account for this, we let the auctioneer choose bidders' exclusion types directly. Fix a profile of exclusion types $\underline{\boldsymbol{u}} = (\underline{u}_1, \ldots, \underline{u}_n) \in [0,1]^n$, and let the auctioneer equalize MRs subject to each bidder's types below own exclusion type receiving zero winning probability. This leads to the following modification of (26):

$$\frac{\partial}{\partial x_i} H_i\left(e^{\delta_i(t)-t}, e^{-\delta_i(t)}\right) \begin{cases} \leq p(t) & \text{if } \delta_i(t) = 0 < -\ln \underline{u}_i, \\ = p(t) & \text{if } \delta_i(t) \in (0, -\ln \underline{u}_i), \\ \geq p(t) & \text{if } \delta_i(t) = -\ln \underline{u}_i > 0, \end{cases} \quad \sum_{i=1}^{n} \delta_i(t) = \min\left\{t, -\sum_{j=1}^{n} \ln \underline{u}_j\right\}. \quad (38)$$

In words, (38) is a system of $n+1$ equations in $n+1$ unknowns $(\boldsymbol{\delta}(t), p(t))$ at each time t. Akin to (26), it asks that MRs be equalized among the bidders who have left the top, $e^{-\delta_i(t)} < 1$, and are not excluded, $e^{-\delta_i(t)} > \underline{u}_i$, and that the bidders at the top have weakly lower MRs. However, the excluded bidders may have weakly higher MRs than the remaining ones; as a result, the solution to (38) in general differs from the iso-MR path of Section 7.1.

To ensure that this construction is well-behaved for every profile of exclusion types, we strengthen regularity accordingly. We say that an environment is **strongly regular** if Condition (A) holds together with the following strengthening of Condition (B).

(B*) For each profile $\underline{\boldsymbol{u}}$, the solution $\left(\boldsymbol{\delta}^{\mathrm{s}}(t|\underline{\boldsymbol{u}}), p^{\mathrm{s}}(t|\underline{\boldsymbol{u}})\right)$ to (38)—unique under Condition (A) at each $t < -\sum_{j=1}^{n} \ln \underline{u}_j$—is such that $p^{\mathrm{s}}(\,\cdot\,|\underline{\boldsymbol{u}})$ is strictly decreasing on that time interval and each $\delta_i^{\mathrm{s}}(\,\cdot\,|\underline{\boldsymbol{u}})$ is weakly increasing, and strictly so whenever bidder i is active, i.e., $0 < \delta_i(t) < -\ln \underline{u}_i$.

Condition (B) is Condition (B*) evaluated at $\underline{\boldsymbol{u}} = \mathbf{0}$, so a strongly regular environment is regular and $\boldsymbol{\delta}^{\mathrm{s}} = \boldsymbol{\delta}^{\mathrm{s}}(\,\cdot\,|\mathbf{0})$. We refer to $\boldsymbol{\delta}^{\mathrm{s}}(\,\cdot\,|\underline{\boldsymbol{u}})$ as the **$\underline{\boldsymbol{u}}$-capped iso-MR path**. We denote the reduced form obtained from it via (11) by $\boldsymbol{x}^{\mathrm{s}}(\,\cdot\,|\underline{\boldsymbol{u}})$, the **$\underline{\boldsymbol{u}}$-capped iso-MR reduced form**. By Theorem 1, it is a feasible reduced form, and its exclusion types are $\underline{\boldsymbol{u}}$.[31] For later reference, we let $T_i(\underline{\boldsymbol{u}}) := -\ln\left(\underline{u}_i x_i^{\mathrm{s}}(\underline{u}_i|\underline{\boldsymbol{u}})\right)$ denote the **exclusion time** at which bidder i's coordinate reaches own exclusion type, which is infinite if it never does. We also set $T(\underline{\boldsymbol{u}}) := \max_i T_i(\underline{\boldsymbol{u}})$ for the time at which the $\underline{\boldsymbol{u}}$-capped iso-MR path stops.

The revenue that a profile of exclusion types generates can be expressed as

$$\sum_{i=1}^{n} \int_{\underline{u}_i}^{1} H_i\left(x_i^{\mathrm{s}}(u_i|\underline{\boldsymbol{u}}), u_i\right) du_i. \tag{39}$$

The exclusion types enter (39) through two channels—directly, through the lower limits of integration; and indirectly, through the $\underline{\boldsymbol{u}}$-capped iso-MR reduced form, which depends on the entire profile $\underline{\boldsymbol{u}}$. Because of the latter, the exclusion types cannot be optimized bidder by bidder. Yet, maximizing (39) is a finite-dimensional problem, and it yields the following characterization.

Theorem 6. *Suppose the environment is strongly regular with each H_i supermodular and $\max\{0, H_i\}$ convex in x_i. Then, the auctioneer's problem admits a unique optimal reduced form $\boldsymbol{x}^*$, namely, the $\underline{\boldsymbol{u}}^*$-capped iso-MR reduced form, where $\underline{\boldsymbol{u}}^*$ is a maximizer of (39). Moreover, for every bidder i,*

$$\underline{u}_i^* H_i\left(x_i^*(\underline{u}_i^*), \underline{u}_i^*\right) \le \int_{T_i(\underline{\boldsymbol{u}}^*)}^{T(\underline{\boldsymbol{u}}^*)} e^{-t} p^{\mathrm{s}}(t|\underline{\boldsymbol{u}}^*)\, dt \quad \text{with equality whenever } \underline{u}_i^* < 1. \tag{40}$$

Finally, $\boldsymbol{x}^*$ is extremal and is induced by the score allocation with its canonical scores.

Theorem 6 shows that the optimal auction can be found by choosing bidders' exclusion types and then equalizing MRs above them. In contrast to Theorem 5, the optimum remains extremal and is still induced by the canonical scores.

The optimality conditions (40) pin down the exclusion types. To interpret them, it is convenient to think about the optimum through its principal curve, which descends from the northeast corner as time progresses. Bidder i's coordinate reaches own exclusion type at the time $T_i(\underline{\boldsymbol{u}}^*)$ and stays there forever after, and we say that the bidder drops out at that time. Of course, the auction itself is static, and an earlier drop-out means nothing more than a higher exclusion type. In this language, (40) says that the revenue generated by bidder i's exclusion

[31] Strictly speaking, this is true when $\prod_{i=1}^{n} \underline{u}_i > 0$. However, if $\prod_{i=1}^{n} \underline{u}_i = 0$, then bidder i's exclusion type $\lim_{t\to\infty} e^{-\delta_i^{\mathrm{s}}(t|\underline{\boldsymbol{u}})}$ may be larger than $\underline{u}_i$. This distinction is inconsequential for our analysis.

type—the left-hand side—equals the flow of the common MR collected between own drop-out time and the last one—the right-hand side. In particular, the bidder that drops out last does so when own exclusion type generates zero MR, whereas every other bidder drops out while own exclusion type still generates positive MR.

To see where this trade-off comes from, differentiate (39) with respect to $\underline{u}_i$ to obtain

$$-H_i\big(x_i^{\mathrm{s}}(\underline{u}_i|\underline{\boldsymbol{u}}),\underline{u}_i\big)+\sum_{j=1}^{n}\int_{\underline{u}_j}^{1}\frac{\partial}{\partial x_j}H_j\big(x_j^{\mathrm{s}}(u_j|\underline{\boldsymbol{u}}),u_j\big)\frac{\partial}{\partial \underline{u}_i}x_j^{\mathrm{s}}(u_j|\underline{\boldsymbol{u}})du_j.$$

The first term is the revenue lost by excluding bidder i's marginal type $\underline{u}_i$. The second term is the change in the revenue from all other types, whose winning probabilities are rearranged along the $\underline{\boldsymbol{u}}$-capped iso-MR path. We show in Appendix A.4.4 that it equals

$$\frac{1}{\underline{u}_i}\int_{T_i(\underline{\boldsymbol{u}})}^{T(\underline{\boldsymbol{u}})}e^{-t}p^{\mathrm{s}}(t|\underline{\boldsymbol{u}})dt. \tag{41}$$

To see why, note that only the types reached after $T_i(\underline{\boldsymbol{u}})$ are affected, since the path up to that time does not depend on $\underline{u}_i$, and that their MRs equal the common MR by (38). Furthermore, differentiating the last equation of (38) and using the definition of the principal curve, one can obtain

$$\sum_{j=1}^{n}\frac{\partial}{\partial \underline{u}_i}x_j^{\mathrm{s}}(u_j|\underline{\boldsymbol{u}})du_j=\frac{e^{-t}}{\underline{u}_i}dt$$

for the types reached at time $t\in\big(T_i(\underline{\boldsymbol{u}}),T(\underline{\boldsymbol{u}})\big)$. Weighting this by the common MR and integrating over t gives (41).

The optimality condition (40) also explains why all bidders drop out at once in Myerson's linear model but at different times in general. In the linear model, $H_i(x_i,u_i)=\zeta_i(u_i)x_i$, so, at any positive winning probability, $H_i=0$ if and only if $\frac{\partial H_i}{\partial x_i}=0$. As a result, both sides of (40) vanish together when every bidder drops out at T^{s}. In nonlinear settings, this equivalence breaks down because MRs depend on winning probabilities. When the auctioneer excludes bidder i's low types, the other bidders' winning probabilities go up. The auctioneer thus gives up the revenue of the excluded types in exchange for the revenue that the other bidders generate at higher winning probabilities, and (40) balances these two at bidder-specific exclusion types. The following example illustrates.

Example 4. *Consider the setting of Section 3 with symmetric bidders, $\beta_i=\beta$, but suppose that the auctioneer incurs a shipping cost $c_i\in[0,1)$ when selling the good to bidder i, with $c_1\le c_2\le\cdots\le c_n$. Then, bidder i's revenue function in the quantile space is $H_i=u_i^{1/\beta}x_i^2-c_ix_i$, which is convex in x_i. Along the parametrization underlying the δ-transform, bidder i's MR equals $2e^{-(1/\beta-1)\delta_i-t}-c_i$. Since H_i is quadratic in own winning probability, the revenue of bidder i's exclusion type can be written in* *closed form using own MR at the exclusion time, which equals the common MR there:*

$$\underline{u}_i H_i\left(x_i^{\mathrm{s}}\left(\underline{u}_i|\underline{\boldsymbol{u}}\right), \underline{u}_i\right) = \frac{e^{-T_i\left(\underline{\boldsymbol{u}}\right)}}{2}\left(p^{\mathrm{s}}\left(T_i\left(\underline{\boldsymbol{u}}\right)|\underline{\boldsymbol{u}}\right) - c_i\right).$$

This revenue is decreasing in own cost, and one can verify that bidders are excluded in the order of their costs, the costliest first, $T_1\left(\underline{\boldsymbol{u}}^*\right) \geq T_2\left(\underline{\boldsymbol{u}}^*\right) \geq \cdots \geq T_n\left(\underline{\boldsymbol{u}}^*\right)$. Between the exclusions of bidders $i+1$ and i, bidders $1, \ldots, i$ are active, and the MR equalization gives

$$({}^1/_\beta - 1)\delta_j = \ln 2 - t - \ln\left(p + c_j\right) \tag{42}$$

for each of them. The coordinates of bidders $i+1, \ldots, n$ stay at their exclusion types. Summing up, we obtain that the common MR on this interval solves

$$\sum_{j=1}^{i} \ln\left(p^{\mathrm{s}}\left(t|\underline{\boldsymbol{u}}^*\right) + c_j\right) = i \ln 2 - \left(i + {}^1/_\beta - 1\right) t - ({}^1/_\beta - 1) \sum_{j=i+1}^{n} \ln \underline{u}_j^*. \tag{43}$$

Finally, the optimality conditions (40) give $p^{\mathrm{s}}\left(T_1\left(\underline{\boldsymbol{u}}^*\right)|\underline{\boldsymbol{u}}^*\right) = c_1$ for the cheapest bidder, who is excluded last, and

$$p^{\mathrm{s}}\left(T_i\left(\underline{\boldsymbol{u}}^*\right)|\underline{\boldsymbol{u}}^*\right) - c_i = 2e^{T_i\left(\underline{\boldsymbol{u}}^*\right)} \int_{T_i\left(\underline{\boldsymbol{u}}^*\right)}^{T_1\left(\underline{\boldsymbol{u}}^*\right)} e^{-t} p^{\mathrm{s}}\left(t|\underline{\boldsymbol{u}}^*\right) dt \tag{44}$$

for every other bidder, who is excluded while the common MR still exceeds own cost, with the gap proportional to the flow of the common MR collected until the last exclusion. Hence, every bidder but the cheapest is excluded strictly before the common MR reaches own cost, and all are excluded before it reaches zero. Moreover, bidder i's exclusion type depends on the costs of all cheaper bidders.

For example, take $\beta = {}^1/_2$ with a costless bidder 1, $c_1 = 0$, and $n-1$ costly bidders, $c_2 = \cdots = c_n$. Since $c_1 = 0$, bidder 1 never drops out, i.e., $T_1\left(\underline{\boldsymbol{u}}^*\right) = \infty$ and $\underline{u}_1^* = 0$. The costly bidders are symmetric and hence drop out together at some $T_2\left(\underline{\boldsymbol{u}}^*\right) = \cdots = T_n\left(\underline{\boldsymbol{u}}^*\right)$ that satisfies

$$-\ln \underline{u}_n^* = \ln 2 - T_n\left(\underline{\boldsymbol{u}}^*\right) - \ln\left(p^{\mathrm{s}}\left(T_n\left(\underline{\boldsymbol{u}}^*\right)|\underline{\boldsymbol{u}}^*\right) + c_n\right) \tag{45}$$

due to (42). After that time, only bidder 1 is active, so that (43) gives

$$p^{\mathrm{s}}\left(t|\underline{\boldsymbol{u}}^*\right) = \frac{2e^{-2t}}{\left(\underline{u}_n^*\right)^{n-1}}.$$

Substituting this expression into (44), we obtain

$$p^{\mathrm{s}}\left(T_n\left(\underline{\boldsymbol{u}}^*\right)|\underline{\boldsymbol{u}}^*\right) - c_n = 2e^{T_n\left(\underline{\boldsymbol{u}}^*\right)} \int_{T_n\left(\underline{\boldsymbol{u}}^*\right)}^{\infty} e^{-t} p^{\mathrm{s}}\left(t|\underline{\boldsymbol{u}}^*\right) dt = \frac{2}{3} p^{\mathrm{s}}\left(T_n\left(\underline{\boldsymbol{u}}^*\right)|\underline{\boldsymbol{u}}^*\right),$$

which gives $p^{\mathrm{s}}\left(T_n\left(\underline{\boldsymbol{u}}^*\right)|\underline{\boldsymbol{u}}^*\right) = 3c_n$. Combining this value with (45) and the common MR above, we *finally obtain*

$$\underline{u}_2^* = \cdots = \underline{u}_n^* = \sqrt[n+1]{\frac{8c_n}{3}}, \qquad T_n\left(\underline{\boldsymbol{u}}^*\right) = \ln\left(\frac{\underline{u}_n^*}{2c_n}\right).$$

The proof of Theorem 6 starts by using convexity of $\max\{0, H_i\}$ to conclude that the optimal reduced form is extremal. By Theorem 2, this means that its δ-transforms sum to $\min\{t, T\}$ for some $T \in [0, \infty]$. We then relax the auctioneer's optimal control problem of Section 7.3 at the exclusion types of the optimum itself. Specifically, we keep only the constraints of (38)—each coordinate between zero and own exclusion level, and their sum equal to $\min\{t, -\sum_j \ln \underline{u}_j^*\}$—so that the relaxed problem can be solved time by time. Strong regularity ensures that the time-by-time solution is the δ-transform of a reduced form, so the relaxation is tight and the optimum is the $\underline{\boldsymbol{u}}^*$-capped iso-MR path. Maximality of (39) at $\underline{\boldsymbol{u}}^*$ and the optimality conditions (40) then follow from the envelope theorem applied to the exclusion types.

9 Final remarks

In this paper, we develop new tools to study asymmetric auction settings with unidimensional types, unit demands, and a single good for sale. We believe that these tools extend well beyond our specific framework and can help analyze other economically significant problems. One immediate application is information design under privacy constraints in the spirit of He et al. [2025], who established a connection between such information structures and Border-type feasibility in auction settings. Other potential applications include persuasion, delegation, and related problems studied in Kleiner et al. [2021] through the lens of symmetric Border-type feasibility.

Extending our approach to feasibility and extremality to more general auction settings is another fruitful direction for future research. Concretely, our method for identifying the principal curve of tightest feasibility constraints delivers a sharp characterization of feasibility and an explicit description of extremality. It may extend to settings with multiunit demands or even multidimensional types. Border-type feasibility characterizations are known in some of these environments (e.g., see Che et al. [2013]), and our approach may help make them more tractable and operational.

Our approach to optimality may also prove valuable in broader principal–agent models and in persuasion. Even within our auction framework, we focus on regular environments in which ironing is not needed, and we illustrate how the results apply to settings with endogenous valuations and non-EU preferences. This focus is largely due to space constraints. Further developing the economics of optimal auctions in these applications and extending the analysis beyond regular environments are natural next steps.

A Appendix

A.1 The δ-transform

In this section, we discuss the reparametrization of CDFs on the unit interval through the δ-transform. Throughout, every CDF $x \in \mathcal{X}$ is equipped with its *right-continuous inverse*, i.e.,

$$x^{-1}(\iota) := \sup\{u \in [0,1] : x(u) \le \iota\}.$$

For a function $\delta \in \mathcal{D}$, we use δ^{-1} to denote its *left-continuous inverse* defined by

$$\delta^{-1}(\iota) := \inf\{t \in \mathbf{R}_+ : \delta(t) \ge \iota\}$$

with the convention $\inf \emptyset = \infty$.

As explained in Section 5.1, the δ-transform is constructed in two steps, and it is convenient to give the inverse in (10) a name. For a CDF $x \in \mathcal{X}$, let

$$\psi(\iota) := (ux)^{-1}(\iota) \tag{46}$$

denote the inverse of the product $u \mapsto ux(u)$, termed the **ψ-transform of** x. Since $ux(u) \le u$, we have $\psi(\iota) \in [\iota, 1]$, so that $\iota/\psi(\iota)$ is well-defined for $\iota > 0$. Furthermore, since ux is strictly increasing whenever positive, ψ is continuous with $\psi(0) = \underline{u}$, which is the exclusion type of x defined in Section 4. We abbreviate by Ψ the subset of $\mathcal{X}$ that consists of CDFs satisfying these properties.

Figure 2 illustrates the construction: panel (a) shows a step-shaped CDF x, panel (b) its product with u, and panel (c) the resulting ψ. The picture highlights how ψ inverts the joint map ux—jumps of x become flats of ψ, and flats of x become linear stretches of ψ.

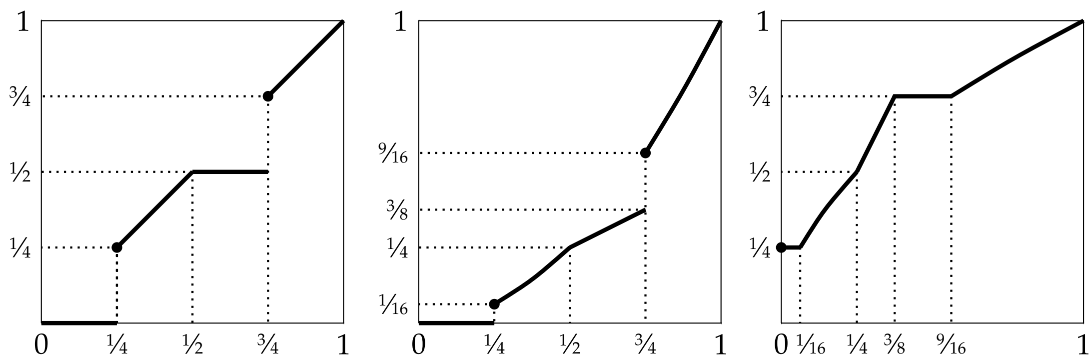

(a) x (b) ux (c) ψ

Figure 2: Illustration of ψ for $x(u) = u \cdot \mathbf{1}_{[1/4,1/2)\cup[3/4,1]}(u) + 1/2 \cdot \mathbf{1}_{[1/2,3/4)}(u)$.

Then, the δ-transform of x at time t equals $-\ln \psi(\iota)$ evaluated at $\iota = e^{-t}$. Since $\psi \in \Psi$, it is easy to see that $\delta \in \mathcal{D}$, i.e., δ is absolutely continuous with $\delta' \in [0,1]$ and $\delta(0) = 0$. Indeed, evaluating at $t = 0$, we have $\delta(0) = -\ln \psi(1) = 0$. Monotonicity of ψ makes δ weakly increasing, and monotonicity of $\iota/\psi(\iota) = e^{\delta(t)-t}\big|_{t=-\ln\iota}$ makes $\delta(t)-t$ weakly decreasing. Such a function is necessarily Lipschitz with constant 1, thus it is an element of $\mathscr{D}$.

Below, we first discuss the opposite direction and complete the proof of Lemma 1. Then, we elaborate more on the geometry of the δ-transform.

A.1.1 Proof of Lemma 1

Proof. We have already argued that the assignment in (10) maps $\mathscr{X}$ into $\mathscr{D}$. By construction, for any u with $ux(u) > 0$, we have $\delta^{-1}(-\ln u) = -\ln\psi^{-1}(u) < \infty$, where ψ^{-1} denotes the right-continuous inverse of ψ and satisfies $\psi^{-1}(u) = ux(u)$. Hence,

$$e^{\delta(t)-t}\Big|_{t=\delta^{-1}(-\ln u)} = \frac{e^{-t}}{\psi(e^{-t})}\Big|_{e^{-t}=ux(u)} = \frac{ux(u)}{\psi(\psi^{-1}(u))} = x(u).$$

On the other hand, if $u > 0$ and $ux(u) = 0$, then $\delta^{-1}(-\ln u) = \infty$ and $e^{\delta(t)-t}\big|_{t=\delta^{-1}(-\ln u)} = 0 = x(u)$, since $\psi(0) \geq u > 0$. Taking the two cases together, we conclude that the output of (11) applied to the δ-transform of the CDF x is x itself.

We now turn to the converse assignment in (11). Given a function $\delta \in \mathscr{D}$, we set $\psi(\iota) := e^{-\delta(-\ln\iota)}$ for $\iota \in (0,1]$ and $\psi(0) := \lim_{\iota\to 0}\psi(\iota)$, which exists by monotonicity. Due to the properties of δ, ψ is an element of Ψ. Hence, x, defined as $x(u) := \psi^{-1}(u)/u$, extended to $u = 0$ by right continuity, is an element of $\mathscr{X}$. By construction, $ux(u) = \psi^{-1}(u)$, meaning that the δ-transform of the CDF x is δ itself. □

A.1.2 Geometry of the δ-transform

The δ-transform yields an explicit geometric decomposition of a CDF $x \in \mathscr{X}$ into (a) vertical segments (jumps), corresponding to $\delta' = 0$; (b) horizontal segments with no mass, corresponding to $\delta' = 1$; and (c) strictly increasing continuous segments, corresponding to $\delta' \in (0,1)$. To see it, recall from (9) that the curve $t \mapsto \left(e^{-\delta(t)}, e^{\delta(t)-t}\right)$ traces the completed graph of x from the top-right corner downward. When $\delta' = 0$, the type coordinate dwells while the probability coordinate falls—the curve descends through a jump of x. When $\delta' = 1$, the probability coordinate dwells while the type coordinate falls—the curve moves left along a flat of x. When $\delta' \in (0,1)$, both coordinates fall, and the curve follows a strictly increasing stretch of x.

For example, Figure 2 plots the ψ-transform of

$$x(u) = u \cdot \mathbf{1}_{[1/4,1/2)\cup[3/4,1]}(u) + 1/2 \cdot \mathbf{1}_{[1/2,3/4)}(u),$$

which (a) jumps at $u = 1/4$ and $u = 3/4$, (b) is horizontal on $[1/2, 3/4)$, and (c) coincides with the diagonal on $[1/4, 1/2)$ and $[3/4, 1]$; below the exclusion type $\underline{u} = 1/4$, it vanishes. The resulting

δ-transform satisfies

$$\delta'(t) = \begin{cases} 0, & t \in [2\ln(4/3), \ln(8/3)) \cup [2\ln 4, \infty), \\ 1, & t \in [\ln(8/3), 2\ln 2), \\ 1/2, & t \in [0, 2\ln(4/3)) \cup [2\ln 2, 2\ln 4), \end{cases}$$

and is illustrated in Figure 3. Because the δ-transform proceeds from $u = 1$ downward, the first jump of x, at the exclusion type $u = 1/4$, appears as the last interval where $\delta' = 0$. This interval is unbounded, since the curve dwells at this jump forever. The other pieces of the decomposition are read off similarly. The stretch $[0, 1/4)$ on which x vanishes is never traversed by the curve.

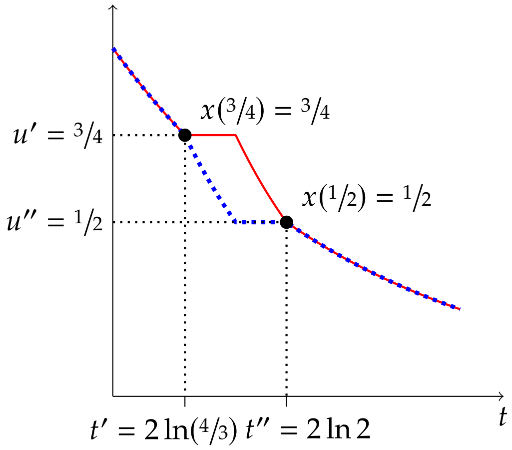

Figure 3: Reading of the CDF x in Figure 2 at $u = u', u''$ from the parametric curves $e^{-\delta(t)}$ and $e^{\delta(t)-t}$, shown in red (solid) and blue (dotted).

A.2 Feasibility

This section proves Theorem 1. The proof proceeds in two steps. The first key step is Lemma 2, which analyzes (15) and expresses its solution in terms of the ψ-transforms $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_n)$ of a reduced form $\boldsymbol{x}$, defined coordinatewise in (46), and the map $\boldsymbol{\rho} = (\rho_1, \ldots, \rho_n)$ given by

$$\rho_i(s) := \psi_i\left(\bar{\psi}^{-1}(s)\right), \quad \text{where } \bar{\psi} := \sqrt[n]{\prod_{i=1}^{n} \psi_i}. \tag{47}$$

We use these objects extensively in the proofs of Theorems 1 and 2.

A.2.1 The auxiliary problem

Lemma 2. *A reduced form $\boldsymbol{x}$ is feasible if and only if $G \leq 1$, where G is the optimal value in (15). Furthermore, letting $\boldsymbol{\psi}$ denote the vector of ψ-transforms of $\boldsymbol{x}$ and $\boldsymbol{\rho}$ be defined by (47), we have the following.*

(1) The point $\boldsymbol{\rho}(s)$ is optimal in (15) and uniquely optimal whenever $s \geq \bar{\psi}(0)$ and $s > 0$.

(2) *The map $s \in [0,1] \mapsto \boldsymbol{\rho}(s)$ is continuous and weakly increasing, and it is strictly increasing in at least one coordinate on $[\bar{\psi}(0), 1]$.*

(3) *The optimal value is given by*

$$G(s) = s^n + \sum_{i=1}^{n} \int_{\bar{\psi}^{-1}(s)}^{1} \iota d \ln \psi_i(\iota). \tag{48}$$

Proof. Note that, in (15), s^n is fixed in the objective, so decreasing u_i increases $\int_{u_i}^1 x_i(u)du$ due to monotonicity of x_i. As a result, a maximizer can be taken with the product constraint binding, which gives

$$G(s) = \max_{\boldsymbol{u} \in [0,1]^n} B(\boldsymbol{u}) \text{ s.t. } \prod_{i=1}^{n} u_i = s^n$$

for each $s \in [0,1]$. Since every cutoff vector $\boldsymbol{u} \in [0,1]^n$ satisfies the constraint at $s^n = \prod_{i=1}^n u_i$, the reduced form $\boldsymbol{x}$ is feasible if and only if $G(s) \leq 1$ for all $s \in [0,1]$. Below, we prove three remaining claims.

Claim (1). To begin, remark that, for each $s \in (0,1)$, the auxiliary program is concave, since each integral $\int_{u_i}^1 x_i(u)du$ is concave in u_i and the constraint set is convex. It follows that the auxiliary program can be solved through the following necessary (and sufficient due to concavity) first-order conditions:

$$\lambda \prod_{j=1}^{n} u_j \in [x_i(u_i-)u_i, x_i(u_i)u_i] \;\; \forall i,$$

where $\lambda \geq 0$ is the Lagrange multiplier on the constraint, and we use the convention $x_i(0-) = 0$.

If $\lambda > 0$, then the constraint binds, so $s^n = \prod_{i=1}^n u_i$. It follows from the definition of the right-continuous inverse that the first-order conditions can be rewritten as

$$u_i = \psi_i(\lambda s^n) \;\; \forall i.$$

Taking the geometric mean across bidders gives $s = \bar{\psi}(\lambda s^n)$, and therefore $s \geq \bar{\psi}(0)$ due to monotonicity of $\bar{\psi}$. On the other hand, if $\lambda = 0$, the first-order conditions require $x_i(u_i-) = 0$, that is,

$$u_i \in [0, \psi_i(0)] \;\; \forall i,$$

which implies $s \leq \bar{\psi}(0)$.

Equipped with this description of the first-order conditions, we distinguish between various values of the parameter s.

- *Case $s \in \left(\bar{\psi}(0), 1\right)$.* As shown above, we must have $\lambda > 0$, so that $s = \bar{\psi}(\lambda s^n)$. The largest Lagrange multiplier with this property can be obtained by taking the right-continuous inverse of $\bar{\psi}$, that is, $\lambda s^n = \bar{\psi}^{-1}(s)$, which is strictly positive by continuity of $\bar{\psi}$. By the definition of $\boldsymbol{\psi}$, the point $\boldsymbol{\rho}(s)$ satisfies the first-order conditions with this multiplier and the constraint as equality, which implies that it is optimal in (15). Since each ψ_i is weakly increasing, it is constant on every interval on which $\bar{\psi}$ is constant. Hence, if $\tilde{\lambda} \in (0, \lambda)$ is such that $s = \bar{\psi}(\tilde{\lambda} s^n)$, then $\boldsymbol{\psi}(\tilde{\lambda} s^n) = \boldsymbol{\rho}(s)$, which shows that $\boldsymbol{\rho}(s)$ is the unique maximizer.

- *Case* $s = \bar{\psi}(0) > 0$. By construction, $\boldsymbol{\psi}(0)$ satisfies the first-order conditions with $\lambda = 0$; thus, it is optimal. Any maximizer with $\lambda = 0$ satisfies $\prod_{i=1}^n u_i \geq s^n = \prod_{i=1}^n \psi_i(0)$ and $u_i \leq \psi_i(0)$ for all i, which forces $\boldsymbol{u} = \boldsymbol{\psi}(0)$. Any maximizer with $\lambda > 0$ equals $\boldsymbol{\psi}(\lambda s^n) = \boldsymbol{\psi}(0)$ because each ψ_i is constant on $\left[0, \bar{\psi}^{-1}(\bar{\psi}(0))\right]$. Hence, $\boldsymbol{\psi}(0)$ is the unique maximizer, and it coincides with $\boldsymbol{\rho}(s)$ at $s = \bar{\psi}(0)$ by the same constancy.

- *Case* $s \in \left(0, \bar{\psi}(0)\right)$. By the same argument, $\boldsymbol{\psi}(0)$ satisfies the first-order conditions with $\lambda = 0$ and is optimal, and it equals $\boldsymbol{\rho}(s)$, since $\bar{\psi}^{-1}(s) = 0$.

- *Case* $s = 1$. Remark that the only feasible point in the auxiliary program is $\boldsymbol{\psi}(1) = (1, \ldots, 1)$, and thus it is uniquely optimal.

- *Case* $s = 0$. Note that the constraint in the auxiliary program has no bite. By construction, $x_i = 0$ on $\left[0, \psi_i(0)\right)$; as a result, $\boldsymbol{\psi}(0) = \left(\underline{u}_1, \ldots, \underline{u}_n\right)$ maximizes the objective in (15) when the constraint is ignored.

Claim (2). The map $s \mapsto \boldsymbol{\rho}(s) = \boldsymbol{\psi}\left(\bar{\psi}^{-1}(s)\right)$ is weakly increasing as a composition of weakly increasing maps. It is continuous because $\bar{\psi}^{-1}$ can jump only across an interval on which $\bar{\psi}$ is constant, where each ψ_i is constant as well. To see strict monotonicity in at least one coordinate, take $\underline{s} < \bar{s}$ in $\left[\bar{\psi}(0), 1\right]$. By continuity of $\bar{\psi}$, we have $\prod_{i=1}^n \rho_i(s) = s^n$ for $s \in \left[\bar{\psi}(0), 1\right]$; hence,

$$\prod_{i=1}^n \rho_i(\underline{s}) = \underline{s}^n < \bar{s}^n = \prod_{i=1}^n \rho_i(\bar{s}),$$

which shows that at least one coordinate strictly increases.

Claim (3). By Claim (1), $\boldsymbol{\rho}(s)$ is optimal in (15) for every $s \in [0, 1]$, so $G(s) = s^n + \sum_{i=1}^n \int_{\rho_i(s)}^1 x_i(u) du$, and it suffices to verify that

$$\int_{\psi_i(\iota)}^1 x_i(u) du = \int_{\iota}^1 \sigma \, d \ln \psi_i(\sigma)$$

for each i and $\iota = \bar{\psi}^{-1}(s)$. Indeed, since $\psi_i^{-1}(u) = u x_i(u)$ on $[\psi_i(0), 1]$, we obtain

$$\begin{aligned}
\int_{\psi_i(\iota)}^1 x_i(u) du &= \int_{\psi_i(0)}^1 \mathbf{1}_{(\iota,1]}\left(\psi_i^{-1}(u)\right) \psi_i^{-1}(u) d \ln u \\
&= \int_0^1 \mathbf{1}_{(\iota,1]}(\sigma) \sigma d \ln \psi_i(\sigma) \\
&= \int_{\iota}^1 \sigma d \ln \psi_i(\sigma).
\end{aligned}$$

The first line uses $x_i(u)du = \psi_i^{-1}(u)d\ln u$, and the second line is due to the change of variables of integration, i.e., $u = \psi_i(\sigma)$, which maps $d\ln u$ on $(\psi_i(0), 1]$ to $d\ln\psi_i(\sigma)$ on $(0, 1]$. $\square$

A.2.2 Proof of Theorem 1

Proof. In view of Lemma 2, since G is strictly increasing on $\left[0, \bar{\psi}(0)\right]$, $\boldsymbol{x}$ is feasible if and only if $G(s) \le 1$ for all $s \ge \bar{\psi}(0)$. On this interval, the maximizers of the auxiliary program described in that lemma trace the principal curve. Indeed, recalling that $\nu_i(t) = e^{-\delta_i(t)} = \psi_i(e^{-t})$ by the construction of the δ-transform, we have $\bar{\psi}(e^{-t}) = e^{-\frac{1}{n}\sum_{i=1}^n \delta_i(t)}$. Moreover, setting $s = \bar{\psi}(e^{-t})$, we get $\boldsymbol{\rho}(s) = \boldsymbol{\psi}(e^{-t}) = \boldsymbol{\nu}(t)$, because each ψ_i is constant wherever $\bar{\psi}$ is constant. Since $t \mapsto \bar{\psi}(e^{-t})$ is continuous and takes every value above $\bar{\psi}(0)$, we conclude that $G(s) \le 1$ on $\left[\bar{\psi}(0), 1\right]$ if and only if $B(\boldsymbol{\nu}(t)) \le 1$ for all $t \in \mathbf{R}_+$. This establishes the first part of the theorem.

To see the expression for the Border constraint along the principal curve, evaluate (48) at $s = \bar{\psi}(e^{-t})$ to obtain

$$B(\boldsymbol{\nu}(t)) = e^{-\sum_{i=1}^n \delta_i(t)} + \sum_{i=1}^n \int_{e^{-t}}^1 \iota d\ln\psi_i(\iota) = e^{-\sum_{i=1}^n \delta_i(t)} + \int_0^t e^{-\tau} \sum_{i=1}^n \delta_i'(\tau)d\tau,$$

where the second equality follows from the change of variables $\iota = e^{-\tau}$ and $\ln\psi_i(e^{-\tau}) = -\delta_i(\tau)$. This is exactly (13). $\square$

A.3 Extremality

In this section, we study extremality of reduced forms and prove Theorems 2 and 3. We show the two theorems in conjunction with each other, establishing three statements in the following order:

> *Statement I.* If a feasible reduced form $\boldsymbol{x}$ is extremal, then Border's constraint binds along the principal curve, i.e., (16) holds, as explained in Section 6.1.
>
> *Statement II.* If a feasible reduced form $\boldsymbol{x}$ satisfies (16), then it is induced by the score allocation with its canonical scores.
>
> *Statement III.* If a feasible reduced form $\boldsymbol{x}$ is induced by a score allocation, then it is extremal and the allocation inducing it is unique.

Statement I corresponds to the "only if" direction of Theorem 2, and the other two statements establish Theorem 3 provided that Statement I has been shown. Taken together, I–III imply the "if" direction of Theorem 2.

The most challenging part of the analysis is the first statement, which says that Border's constraint binds along the principal curve for extremal reduced forms. In view of Lemma 2 and the proof of Theorem 1, this is equivalent to requiring that the function G, which is defined in (15) and then explicitly solved for in (48), satisfy

$$G(s) = 1 \;\; \forall s \in \left[\bar{\psi}(0), 1\right].$$

This condition is equivalent to nonexistence of a **sag**, that is, an open interval $(\underline{s}, \overline{s})$ in $[\bar{\psi}(0), 1]$ on which G is strictly less than 1.

For each feasible (potentially nonextremal) reduced form $\boldsymbol{x}$, if the set of sags is nonempty, then, since $G(1) = 1$ and $G < 1$ on $[0, \bar{\psi}(0))$, there are two mutually exclusive and exhaustive cases.

1. There exists a sag so that G equals 1 at its endpoints.
2. There exists a sag that starts at $\bar{\psi}(0)$ and ends at $\overline{s}$ so that G equals 1 on $[\overline{s}, 1]$.

We shall show that both cases are inconsistent with extremality of $\boldsymbol{x}$, thereby establishing that no extremal reduced form admits a sag. To this end, we first state and prove six auxiliary lemmas that record various properties of elements of the feasible set and its extreme points. With these tools in hand, we then prove the two theorems stated in Section 6.

A.3.1 Preliminaries

Lemmas 3 and 4 show that certain modifications of a feasible reduced form leave (some of) Border's constraints unchanged. For brevity, when we need to emphasize the dependence on the reduced form, we write $B_{\boldsymbol{x}}$ for Border's constraint in (7) corresponding to $\boldsymbol{x}$.

The other four lemmas in this section study sags of extremal reduced forms. Specifically, Lemma 5 establishes that x_i is piecewise-constant with finitely many jumps along a sag of the principal curve. Lemmas 6 and 7 discuss jumps at the endpoints of the sag. Finally, Lemma 8 shows that any sag corresponds to a nontrivial portion of the principal curve. Throughout this analysis, we use the following notation. Given a reduced form $\boldsymbol{x}$ and two cutoffs $\underline{v} < \overline{v}$, define the set of jumps of x_i in $[\underline{v}, \overline{v}]$ by

$$D_i(\underline{v}, \overline{v}) := \{u \in [\underline{v}, \overline{v}] | x_i(u-) \neq x_i(u)\}$$

with the convention $x_i(0-) = 0$. We extend the definition of D_i to degenerate intervals by setting

$$D_i(v, v) := \{v\}.$$

Lemma 3 (Perturbations at the bottom of the type space)**.** *Let $\boldsymbol{x}$ be a feasible reduced form that satisfies $B_{\boldsymbol{x}}(\overline{\boldsymbol{v}}) = 1$ for some $\overline{\boldsymbol{v}} \in [0,1]^n$, $\prod_{i=1}^n \overline{v}_i > 0$. Consider $\boldsymbol{\Delta} = (\Delta_1, \ldots, \Delta_n)$ with $\Delta_i : [0,1] \to \mathbf{R}$ such that $\boldsymbol{x} + \boldsymbol{\Delta}$ is a reduced form, $\boldsymbol{\Delta} \neq 0$, and $\Delta_i(u) = 0$ on $(\overline{v}_i, 1]$ for each i. Then, $\boldsymbol{x} + \boldsymbol{\Delta}$ is feasible if and only if*

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) \leq 1 \;\; \forall \boldsymbol{u} \;\; s.t. \;\; \mathbf{0} \leq \boldsymbol{u} \leq \overline{\boldsymbol{v}}. \tag{49}$$

Proof. The "only if" direction is immediate, so we focus on the converse. Specifically, suppose that (49) holds. We now show that $\boldsymbol{x} + \boldsymbol{\Delta}$ is feasible, meaning that the inequality in (49) holds for all $\boldsymbol{u}$ other than $\mathbf{0} \leq \boldsymbol{u} \leq \overline{\boldsymbol{v}}$. To this end, we distinguish between several cases.

Case $\prod_{i=1}^n u_i = 0$. In this case,

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) \le \sum_{i=1}^n \int_0^1 (x_i(u) + \Delta_i(u))du = B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{0}) \le 1,$$

where the first inequality holds because $\prod_{i=1}^n u_i = 0$ and $x_i + \Delta_i \ge 0$, and the last inequality follows from (49).

Conversely, suppose $\boldsymbol{u} \in [0,1]^n$ is such that $\prod_{i=1}^n u_i > 0$ and $u_i > \overline{v}_i$ precisely for the first $k \ge 1$ coordinates, i.e., $i = 1, \dots, k$—after relabeling if necessary.

Case $k = n$. Since $\boldsymbol{u} \ge \overline{\boldsymbol{v}}$, we have $\int_{u_i}^1 \Delta_i(u)du = 0$ for each i, and hence

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) = B_{\boldsymbol{x}}(\boldsymbol{u}) \le 1.$$

Case $k < n$. Set $a_1 = \prod_{i=1}^k u_i$, $\overline{a}_1 = \prod_{i=1}^k \overline{v}_i$, $a_2 = \prod_{i=k+1}^n u_i$, and $\overline{a}_2 = \prod_{i=k+1}^n \overline{v}_i$.

Since $\boldsymbol{x}$ is feasible and $B_{\boldsymbol{x}}(\overline{\boldsymbol{v}}) = 1$, we have

$$B_{\boldsymbol{x}}\big(u_1, \dots, u_k, \overline{v}_{k+1}, \dots, \overline{v}_n\big) \le B_{\boldsymbol{x}}\big(\overline{\boldsymbol{v}}\big),$$

which expands to

$$a_1 \cdot \overline{a}_2 + \sum_{i=1}^k \int_{u_i}^1 x_i(u)du \le \overline{a}_1 \cdot \overline{a}_2 + \sum_{i=1}^k \int_{\overline{v}_i}^1 x_i(u)du. \tag{50}$$

Next, consider evaluating $B_{\boldsymbol{x}+\boldsymbol{\Delta}}$ at $\big(\overline{v}_1, \dots, \overline{v}_k, u_{k+1}, \dots, u_n\big)$. Using $\Delta_i(u) = 0$ for $u > \overline{v}_i$, we obtain

$$\overline{a}_1 \cdot a_2 + \sum_{i=1}^k \int_{\overline{v}_i}^1 x_i(u)du + \sum_{i=k+1}^n \int_{u_i}^1 x_i(u)du + \sum_{i=1}^n \int_{u_i}^1 \Delta_i(u)du \le 1, \tag{51}$$

due to (49).

Since $a_1 \ge \overline{a}_1$ and $a_2 \le \overline{a}_2$, we have $(a_1 - \overline{a}_1) \cdot (a_2 - \overline{a}_2) \le 0$, i.e.,

$$a_1 \cdot a_2 \le \overline{a}_1 \cdot a_2 + a_1 \cdot \overline{a}_2 - \overline{a}_1 \cdot \overline{a}_2.$$

Combining this inequality with (50) and (51), we conclude that

$$
\begin{aligned}
B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) &= a_1 \cdot a_2 + \sum_{i=1}^{n} \int_{u_i}^{1} x_i(u)du + \sum_{i=1}^{n} \int_{u_i}^{1} \Delta_i(u)du \\
&\leq \overline{a}_1 \cdot a_2 + a_1 \cdot \overline{a}_2 - \overline{a}_1 \cdot \overline{a}_2 + \sum_{i=1}^{n} \int_{u_i}^{1} x_i(u)du + \sum_{i=1}^{n} \int_{u_i}^{1} \Delta_i(u)du \\
&\leq \overline{a}_1 \cdot a_2 + \sum_{i=1}^{k} \int_{\overline{v}_i}^{1} x_i(u)du + \sum_{i=k+1}^{n} \int_{u_i}^{1} x_i(u)du + \sum_{i=1}^{n} \int_{u_i}^{1} \Delta_i(u)du \\
&\leq 1,
\end{aligned}
$$

as claimed. ◻

Lemma 4 (Perturbations in the middle of the type space). *Let $\boldsymbol{x}$ be a feasible reduced form that satisfies $B_{\boldsymbol{x}}(\underline{\boldsymbol{v}}) = B_{\boldsymbol{x}}(\overline{\boldsymbol{v}}) = 1$ for some $\underline{\boldsymbol{v}}, \overline{\boldsymbol{v}} \in [0,1]^n$ such that $\prod_{i=1}^n \overline{v}_i > 0$, $\underline{\boldsymbol{v}} \neq \overline{\boldsymbol{v}}$, and $\underline{\boldsymbol{v}} \leq \overline{\boldsymbol{v}}$. Consider $\boldsymbol{\Delta} = (\Delta_1, \ldots, \Delta_n)$ with $\Delta_i : [0,1] \to \mathbf{R}$ such that $\boldsymbol{x} + \boldsymbol{\Delta}$ is a reduced form, $\boldsymbol{\Delta} \neq 0$, $\Delta_i(u) = 0$ on $\left[0, \underline{v}_i\right) \cup \left(\overline{v}_i, 1\right]$ for each i, and*

$$\sum_{i=1}^{n} \int_{\underline{v}_i}^{\overline{v}_i} \Delta_i(u)du = 0.$$

Then, $\boldsymbol{x} + \boldsymbol{\Delta}$ is feasible if and only if

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) \leq 1 \ \ \forall \boldsymbol{u} \ \text{ s.t. } \ \underline{\boldsymbol{v}} \leq \boldsymbol{u} \leq \overline{\boldsymbol{v}}. \tag{52}$$

Proof. In view of Lemma 3, the reduced form $\boldsymbol{x} + \boldsymbol{\Delta}$ is feasible if and only if (49) holds. We thus need to show that (49) is satisfied if and only if (52) holds. The "only if" direction is immediate, so we focus on the converse, again distinguishing between several cases.

Case $\prod_{i=1}^n u_i = 0$. In this case,

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) \leq \sum_{i=1}^{n} \int_{0}^{1} (x_i(u) + \Delta_i(u))du = B_{\boldsymbol{x}}(\mathbf{0}) \leq 1,$$

where the first inequality holds because $\prod_{i=1}^n u_i = 0$ and $x_i + \Delta_i \geq 0$, the equality follows from $\sum_{i=1}^n \int_0^1 \Delta_i(u)du = 0$, and the last inequality follows from feasibility of $\boldsymbol{x}$.

Now, suppose $\boldsymbol{u} \leq \overline{\boldsymbol{v}}$ is such that $\prod_{i=1}^n u_i > 0$ and $u_i < \underline{v}_i$ precisely for the first $k \geq 1$ coordinates, i.e., $i = 1, \ldots, k$—after relabeling if necessary.

Case $k = n$. Since $\boldsymbol{u} \leq \underline{\boldsymbol{v}}$, it follows that $\sum_{i=1}^k \int_{u_i}^1 \Delta_i(u)du = 0$, and hence

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) = B_{\boldsymbol{x}}(\boldsymbol{u}) \leq 1.$$

Case $k < n$. The argument is identical to the one in the proof of Lemma 3 but replaces $\overline{\boldsymbol{v}}$ with $\underline{\boldsymbol{v}}$. Setting $a_1 = \prod_{i=1}^k u_i$, $\underline{a}_1 = \prod_{i=1}^k \underline{v}_i$, $a_2 = \prod_{i=k+1}^n u_i$, $\underline{a}_2 = \prod_{i=k+1}^n \underline{v}_i$, and repeating the same steps as before, we can obtain

$$a_1 \cdot \underline{a}_2 + \sum_{i=1}^{k} \int_{u_i}^{1} x_i(u)du \leq \underline{a}_1 \cdot \underline{a}_2 + \sum_{i=1}^{k} \int_{\underline{v}_i}^{1} x_i(u)du,$$

$$\underline{a}_1 \cdot a_2 + \sum_{i=1}^{k} \int_{\underline{v}_i}^{1} x_i(u)du + \sum_{i=k+1}^{n} \int_{u_i}^{1} x_i(u)du + \sum_{i=1}^{n} \int_{u_i}^{1} \Delta_i(u)du \leq 1,$$

$$a_1 \cdot a_2 \leq a_1 \cdot \underline{a}_2 + \underline{a}_1 \cdot a_2 - \underline{a}_1 \cdot \underline{a}_2.$$

Combining these inequalities yields the desired result:

$$B_{\boldsymbol{x}+\boldsymbol{\Delta}}(\boldsymbol{u}) \leq 1.$$

□

Lemma 5 (Structure along a sag)**.** *Let $\boldsymbol{x}$ be an extremal reduced form. Consider a sag $(\underline{s}, \overline{s})$. Then, for each i such that $\rho_i(\underline{s}) < \rho_i(\overline{s})$, the function x_i is piecewise-constant on $[\rho_i(\underline{s}), \rho_i(\overline{s})]$ with finitely many elements in $D_i(\rho_i(\underline{s}), \rho_i(\overline{s}))$, i.e., x_i takes finitely many values.*

Proof. Let i be such that $\rho_i(\underline{s}) < \rho_i(\overline{s})$. There are two mutually exclusive possibilities: either all (weak) lower contour sets of x_i are open in $(\rho_i(\underline{s}), \rho_i(\overline{s}))$, i.e.,

$$\{u \in (\rho_i(\underline{s}), \rho_i(\overline{s})) | x_i(u) \leq x_i(\underline{v})\} \text{ is open } \forall \underline{v} \in (\rho_i(\underline{s}), \rho_i(\overline{s})), \tag{53}$$

or there exists at least one such lower contour set that is not open. In what follows, we first prove, in two logical steps, that the second case cannot occur given that $\boldsymbol{x}$ is extremal, and then complete the proof in one additional step.

Step 1. Towards a contradiction, suppose that (53) is violated at $\underline{v}$, i.e., the associated lower contour set equals $(\rho_i(\underline{s}), \underline{v}]$. This means that there exists a closed interval $[\underline{v}, \overline{v}]$ compactly contained in $(\rho_i(\underline{s}), \rho_i(\overline{s}))$ such that $x_i(u) > x_i(\underline{v})$ for all $u \in (\underline{v}, \overline{v}]$. Denote the value of

$$\max_{\boldsymbol{u} \in [0,1]^n} B_{\boldsymbol{x}}(\boldsymbol{u}) \text{ s.t. } u_i \in [\underline{v}, \overline{v}]$$

by ω. Since $\boldsymbol{x}$ is feasible, we have $\omega \leq 1$. We now show that ω is strictly less than one.

Suppose not, i.e., $\omega = 1$. Let $\boldsymbol{u}$ be a maximizer that attains ω and set $s \in [0,1]$ so that $s^n = \prod_{j=1}^{n} u_j$. Note that $s > 0$; otherwise, $B_{\boldsymbol{x}}(\boldsymbol{u}) = 1$ would require $x_j = 0$ on $[0, u_j)$ for all j, i.e., $u_j \leq \psi_j(0)$, contradicting $u_i \geq \underline{v} > \rho_i(\underline{s}) \geq \psi_i(0)$. By definition of the auxiliary problem in Lemma 2, $G(s)$ is an upper bound on ω. This bound is tight because $\boldsymbol{x}$ is assumed to be feasible, i.e., $G \leq 1$, and $\omega = 1$. Since G is strictly less than one on $[0, \bar{\psi}(0))$, we must have $s \geq \bar{\psi}(0)$; thus,

$$B_{\boldsymbol{x}}(\boldsymbol{u}) = G(s) = B_{\boldsymbol{x}}(\boldsymbol{\rho}(s))$$

as established in the proof of Theorem 1.

Since G is strictly less than one on $(\underline{s}, \overline{s})$, we cannot have $s \in (\underline{s}, \overline{s})$. By Claim (2) of Lemma 2, $\rho_i(s) \le \rho_i(\underline{s}) < \underline{v}$ whenever $s \le \underline{s}$, and $\rho_i(s) \ge \rho_i(\overline{s}) > \overline{v}$ whenever $s \ge \overline{s}$. In either case, we obtain $\rho_i(s) \notin [\underline{v}, \overline{v}]$. As a result, $\boldsymbol{u} \neq \boldsymbol{\rho}(s)$, contradicting the uniqueness of the optimizer stated in Lemma 2. Therefore, we conclude that $\omega < 1$, as claimed.

Step 2. We now construct a perturbation $\boldsymbol{\Delta} : [0,1] \longrightarrow \mathbf{R}^n$ so that $\boldsymbol{x} \pm \varepsilon\boldsymbol{\Delta}$ is a feasible reduced form provided that $\varepsilon > 0$ is small enough. The existence of such a perturbation contradicts extremality of $\boldsymbol{x}$, thereby ruling out the existence of a non-open lower contour set.

First of all, for each $j \neq i$, set $\Delta_j = 0$. As shown in the proof of Theorem 1 in Kleiner et al. [2021], there exist a number $x^* \in [x_i(\underline{v}), x_i(\overline{v})]$ and an interval $[\underline{v}^*, \overline{v}^*]$ contained in $[\underline{v}, \overline{v}]$ so that Δ_i defined by

$$\Delta_i(u) = \begin{cases} x_i(u) - x_i(\underline{v}) & u \in [\underline{v}, \underline{v}^*), \\ x^* - x_i(u) & u \in [\underline{v}^*, \overline{v}^*), \\ x_i(u) - x_i(\overline{v}) & u \in [\overline{v}^*, \overline{v}], \\ 0 & u \notin [\underline{v}, \overline{v}] \end{cases}$$

is such that $\Delta_i \neq 0$, $x_i \pm \Delta_i$ is in $\mathcal{X}$, and $\int_{\underline{v}}^{\overline{v}} \Delta_i(u)du = 0$. In fact, Kleiner et al. [2021] provided a closed-form expression for the endpoints of the subinterval used above:

$$\begin{aligned} \underline{v}^* &= \inf_{u \in [\underline{v}, \overline{v}]} u \;\; \text{s.t.} \;\; 2x_i(u) \ge x_i(\underline{v}) + x^*, \\ \overline{v}^* &= \inf_{u \in [\underline{v}, \overline{v}]} u \;\; \text{s.t.} \;\; 2x_i(u) \ge x_i(\overline{v}) + x^* \end{aligned}$$

and proved that the value of x^* making $\int_{\underline{v}}^{\overline{v}} \Delta_i(u)du = 0$ (provided the endpoints are chosen as above) is well-defined.

Set

$$\overline{\varepsilon} = \frac{1 - \omega}{\max_{u_i \in [\underline{v}, \overline{v}]} \left| \int_{u_i}^{\overline{v}} \Delta_i(u)du \right|} > 0$$

and choose $\varepsilon \in (0,1)$ such that $\varepsilon \le \overline{\varepsilon}$. Then, for every $\boldsymbol{u} \in [0,1]^n$ with $u_i \in [\underline{v}, \overline{v}]$, we have

$$B_{\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}}(\boldsymbol{u}) \le B_{\boldsymbol{x}}(\boldsymbol{u}) + \varepsilon \left| \int_{u_i}^{1} \Delta_i(u)du \right| \le B_{\boldsymbol{x}}(\boldsymbol{u}) + (1 - \omega).$$

Since $B_{\boldsymbol{x}}(\boldsymbol{u}) \le \omega$ for every $\boldsymbol{u} \in [0,1]^n$ with $u_i \in [\underline{v}, \overline{v}]$, Border's constraint holds for $\boldsymbol{x} \pm \varepsilon\boldsymbol{\Delta}$ at all such $\boldsymbol{u}$. For every $\boldsymbol{u} \in [0,1]^n$ with $u_i \notin [\underline{v}, \overline{v}]$, we have $B_{\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}}(\boldsymbol{u}) = B_{\boldsymbol{x}}(\boldsymbol{u}) \le 1$. Hence, the reduced form $\boldsymbol{x} \pm \varepsilon\boldsymbol{\Delta}$ is feasible. However, since $\boldsymbol{x}$ is assumed to be an extreme point of the feasible set, this contradicts extremality. Therefore, no such interval $[\underline{v}, \overline{v}]$ can exist.

To sum up, we have shown that all (weak) lower contour sets of x_i are open in $(\rho_i(\underline{s}), \rho_i(\overline{s}))$. This means that x_i is piecewise-constant and every jump in that interval other than the last one has an immediate successor. We can thus meaningfully talk about consecutive jumps.

Step 3. To conclude the proof, we need to show that the number of jumps is finite. Towards a contradiction, suppose that x_i admits three consecutive jump points in $\big(\rho_i(\underline{s}), \rho_i(\overline{s})\big)$, denoted $\underline{v} < v^* < \overline{v}$. Similarly to the previous step, we rule out this possibility by explicitly constructing a perturbation $\boldsymbol{\Delta} = (\Delta_1, \dots, \Delta_n) : [0,1] \to \mathbf{R}^n$ so that $\boldsymbol{x} \pm \varepsilon\boldsymbol{\Delta}$ is a feasible reduced form provided that $\varepsilon > 0$ is small enough.

First of all, for each $j \neq i$, set $\Delta_j = 0$. For coordinate i, set

$$\Delta_i(u) = \kappa \cdot \left(\frac{\mathbf{1}_{[\underline{v}, v^*)}(u)}{v^* - \underline{v}} - \frac{\mathbf{1}_{[v^*, \overline{v})}(u)}{\overline{v} - v^*} \right) \tag{54}$$

for some $\kappa > 0$ small enough to ensure $\Delta_i \neq 0$ and $x_i \pm \Delta_i \in \mathscr{X}$, e.g., any choice such that

$$\frac{\kappa}{v^* - \underline{v}} \leq x_i(\underline{v}) - x_i(\underline{v}-), \quad \frac{\kappa}{v^* - \underline{v}} + \frac{\kappa}{\overline{v} - v^*} \leq x_i(v^*) - x_i(v^*-), \quad \frac{\kappa}{\overline{v} - v^*} \leq x_i(\overline{v}) - x_i(\overline{v}-)$$

will do.

Then, applying the same argument as in the previous construction, we conclude that the perturbed reduced form $\boldsymbol{x} \pm \varepsilon\boldsymbol{\Delta}$ is feasible for sufficiently small $\varepsilon > 0$. This contradicts extremality of $\boldsymbol{x}$, and hence no such triple jump structure can exist. □

Lemma 6 (Discontinuity at the top of a sag)**.** *Let $\boldsymbol{x}$ be an extremal reduced form. Consider a sag $(\underline{s}, \overline{s})$ satisfying $G(\overline{s}) = 1$. Then, $\rho_i(\overline{s}) \in D_i\big(\rho_i(\underline{s}), \rho_i(\overline{s})\big)$ for all i, i.e., x_i is discontinuous at $\rho_i(\overline{s})$ whenever $\rho_i(\underline{s}) < \rho_i(\overline{s})$.*

Proof. Fix an index i with $\rho_i(\underline{s}) < \rho_i(\overline{s})$, and suppose, towards a contradiction, that x_i is continuous at $\rho_i(\overline{s})$. We distinguish between two cases.

Case $\rho_i(\overline{s}) < 1$. Clearly, we have $\overline{s} < 1$. For s near $\overline{s}$, define $\boldsymbol{u}(s)$ by keeping all coordinates except i fixed at their values on $\rho(\cdot)$ at $\overline{s}$, that is, $u_j(s) = \rho_j(\overline{s})$ for $j \neq i$, and adjusting coordinate i to satisfy the product constraint:

$$u_i(s) = \frac{\rho_i(\overline{s})}{\overline{s}^n} s^n.$$

By construction, $\prod_{j=1}^n u_j(s) = s^n$, so $\boldsymbol{u}(s)$ is feasible in the optimization problem that defines $G(s)$ in Lemma 2. Hence, $G(s) \geq \underline{G}(s)$, where $\underline{G}(s)$ is the objective evaluated at $\boldsymbol{u}(s)$:

$$\underline{G}(s) = s^n + \sum_{j=1}^n \int_{u_j(s)}^1 x_j(u)du = s^n + \int_{\frac{\rho_i(\overline{s})}{\overline{s}^n} s^n}^1 x_i(u)du + \sum_{j \neq i} \int_{\rho_j(\overline{s})}^1 x_j(u)du. \tag{55}$$

On the one hand, $\underline{G}(\overline{s}) = G(\overline{s}) = 1$. On the other hand, since $(\underline{s}, \overline{s})$ is a sag, we have $\underline{G}(s) \leq G(s) < 1$ for $s < \overline{s}$.

By Lemma 5, the function x_i is piecewise-constant on $\big[\rho_i(\underline{s}), \rho_i(\overline{s})\big]$ with finitely many jumps. Since x_i is assumed to be continuous at the right endpoint, it is constant on a left neighborhood of this point. Consequently, $\underline{G}$ is an affine function of s^n for $s < \overline{s}$. But $\underline{G}(s) < \underline{G}(\overline{s})$ for $s < \overline{s}$, meaning that the (left) slope of this affine function is strictly positive. Since x_i is continuous at $\rho_i(\bar{s})$, the slope of $\underline{G}$ is continuous at $\bar{s}$; therefore, $\underline{G}(s) > 1$ for $s > \bar{s}$ close to $\bar{s}$. This implies that $G(s)$ strictly exceeds 1, contradicting $\underline{G} \leq G \leq 1$ on a right neighborhood of $\bar{s}$, as required by feasibility. Thus, we conclude that x_i cannot be continuous at $\rho_i(\bar{s})$.

Case $\rho_i(\bar{s}) = 1$. The argument is virtually identical to the previous case, with the only difference that now $x_i(1) = 1$ and continuity at $\rho_i(\bar{s}) = 1$ implies $x_i(u) = 1$ on a left neighborhood of 1. Hence, for $s < \bar{s}$ close enough to $\bar{s}$, the same construction of $\boldsymbol{u}(s)$ and the same lower bound $\underline{G}$ in (55) apply, and $\underline{G}$ is affine in s^n on a left neighborhood of $\bar{s}$. Its (left) slope equals

$$\frac{\partial}{\partial s^n}\underline{G}(\bar{s}-) = 1 - \frac{\rho_i(\bar{s})}{\bar{s}^n} x_i(\rho_i(\bar{s})) = 1 - \frac{1}{\bar{s}^n} \leq 0.$$

Since $\underline{G}(\bar{s}) = G(\bar{s}) = 1$, linearity and the nonpositive slope imply that $\underline{G}(s) \geq 1$ for all $s < \bar{s}$ close enough to $\bar{s}$, contradicting that $(\underline{s}, \bar{s})$ is a sag. Therefore, x_i must be discontinuous at $\rho_i(\bar{s}) = 1$. □

Lemma 7 (Discontinuity at the bottom of a sag)**.** *Let $\boldsymbol{x}$ be an extremal reduced form. Consider a sag $(\underline{s}, \bar{s})$ satisfying $G(\underline{s}) = 1$ with $\underline{s} > 0$. Then, $\rho_i(\underline{s}) \in D_i(\rho_i(\underline{s}), \rho_i(\bar{s}))$ for all i, i.e., x_i is discontinuous at $\rho_i(\underline{s})$ whenever $\rho_i(\underline{s}) < \rho_i(\bar{s})$.*

Proof. Let i be such that $\rho_i(\underline{s}) < \rho_i(\bar{s})$. Since $\underline{s} > 0$, it must be the case that $\rho_i(\underline{s}) > 0$. The argument is identical to the first case in the proof of Lemma 6. Specifically, suppose, towards a contradiction, that x_i is continuous at $\rho_i(\underline{s})$. Then, the lower bound $\underline{G}$ is well-defined in a neighborhood of $\underline{s}$, and we have

$$\frac{\partial}{\partial s^n}\underline{G}(\underline{s}-) = \frac{\partial}{\partial s^n}\underline{G}(\underline{s}+) = 1 - \frac{\rho_i(\underline{s})}{\underline{s}^n} x_i(\rho_i(\underline{s})) < 0,$$

where the equality follows from the continuity of x_i at $\rho_i(\underline{s})$, and the inequality follows from the fact that x_i is constant in a right neighborhood of $\rho_i(\underline{s})$, together with the assumption that $\underline{G} \leq G < 1$ on $(\underline{s}, \bar{s})$. Since the slope of $\underline{G}$ is continuous at $\underline{s}$, this implies that $G \geq \underline{G} > 1$ on a left neighborhood of $\underline{s}$, contradicting feasibility. Therefore, x_i is discontinuous at $\rho_i(\underline{s})$. □

Lemma 8 (Nontriviality along a sag)**.** *Let $\boldsymbol{x}$ be a feasible reduced form. Consider a sag $(\underline{s}, \bar{s})$. Then, we have (1) $\rho_i(\underline{s}) < \rho_i(\bar{s})$ for at least one index i; and (2) it is true for at least two such indices whenever G equals 1 at the endpoints of the sag.*

Proof. The first part of this claim has been established in Lemma 2, and so we focus on the second part.

Towards a contradiction, suppose that $\rho_i(\underline{s}) < \rho_i(\bar{s})$ but $\rho_j(\underline{s}) = \rho_j(\bar{s})$ for all $j \neq i$. For each such $j \neq i$, since ρ_j is weakly increasing by Claim (2) of Lemma 2, it equals $\rho_j(\bar{s})$ on $[\underline{s}, \bar{s}]$. Therefore, for all $s \in [\underline{s}, \bar{s}]$, we have

$$\rho_i(s) = \frac{s^n}{\prod_{j \neq i} \rho_j(\bar{s})} = \frac{\rho_i(\bar{s})}{\bar{s}^n} s^n.$$

It follows that $G(s)$ equals $\underline{G}(s)$ defined in (55) for all $s \in [\underline{s}, \overline{s}]$. Clearly, G is a concave function of s^n on the respective interval. Since $G(\underline{s}) = G(\overline{s}) = 1$, it must be that $G = 1$ on $(\underline{s}, \overline{s})$ due to concavity. This contradicts the premise of the claim. $\square$

A.3.2 Proofs of Theorems 2 and 3: Statement I

Proof. Let $\boldsymbol{x}$ be an extremal reduced form. As explained above, there are two mutually exclusive cases. We now analyze each in detail, starting with the latter.

Case 2 with $\overline{s} > \bar{\psi}(0)$. Assume $G < 1$ on $[\bar{\psi}(0), \overline{s})$ and $G = 1$ on $[\overline{s}, 1]$ for some $\overline{s} > \bar{\psi}(0)$. As explained in Lemma 2, we also know that G is strictly increasing on $[0, \bar{\psi}(0))$, and thus $G < 1$ on that interval.

To simplify notation, set $\overline{\boldsymbol{v}} = \boldsymbol{\rho}(\overline{s})$ and $\underline{\boldsymbol{v}} = \boldsymbol{\psi}(0)$. Since $\overline{s} > 0$, we have $\prod_{i=1}^n \overline{v}_i > 0$. Lemmas 5 and 6 jointly imply that, for each i with $\underline{v}_i < \overline{v}_i$, the function x_i is piecewise-constant on this interval with finitely many jumps. Furthermore, the upper endpoint is a discontinuity point, i.e., $\overline{v}_i \in D_i(\underline{v}_i, \overline{v}_i)$. Since $x_i = 0$ on $[0, \underline{v}_i)$ by construction, it follows that the discontinuity points up to $\overline{v}_i$ are the same whether one starts at $\underline{v}_i$ or 0, that is

$$D_i(\underline{v}_i, \overline{v}_i) = D_i(0, \overline{v}_i).$$

Step 1. To begin, we show that there exists an index i for which the discontinuity set $D_i(0, \overline{v}_i)$ contains at least two elements. Towards a contradiction, suppose that

$$D_i(0, \overline{v}_i) \subseteq \{\overline{v}_i\} \ \ \forall i.$$

This means that $x_i = 0$ on $[0, \overline{v}_i)$. Since, by the definition of the transformation ψ_i, we have $\psi_i(0) = \sup\{u \in [0,1] : u x_i(u) = 0\}$, the upper endpoint $\overline{v}_i$ is not higher than $\psi_i(0)$ for all i. As a result,

$$\overline{s} = \sqrt[n]{\prod_{i=1}^n \overline{v}_i} \leq \bar{\psi}(0),$$

contradicting the assumption that $\overline{s} > \bar{\psi}(0)$.

Step 2. Next, fix i such that $D_i(0, \overline{v}_i)$ contains at least two elements, which exists due to the argument above. To rule out the second case with $\overline{s} > \bar{\psi}(0)$, we construct a perturbation $\boldsymbol{\Delta} : [0,1] \longrightarrow \mathbf{R}^n$ so that $\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}$ is a feasible reduced form provided $\varepsilon > 0$ is small enough. As in the proof of Lemma 5, the existence of such $\boldsymbol{\Delta}$ contradicts extremality.

To this end, let $\underline{v}^*$, $\overline{v}^*$ be the last two points in $D_i(0, \overline{v}_i)$, where $\overline{v}^* = \overline{v}_i$ as explained above. Denote the value of

$$\max_{\boldsymbol{u} \in [0,1]^n} B_{\boldsymbol{x}}(\boldsymbol{u}) \ \text{ s.t. } \ \boldsymbol{u} \in \prod_{j=1}^n \left(D_j(0, \overline{v}_j) \cup \{0, \overline{v}_j\}\right), \ \ \boldsymbol{u} \neq \overline{\boldsymbol{v}}$$

by ω. Due to finiteness of bidders' discontinuity sets and $G < 1$ on $[0, \overline{s})$, ω is well-defined and strictly below 1.

First of all, for each $j \neq i$, set $\Delta_j = 0$. For coordinate i, set

$$\Delta_i(u) = \kappa \cdot \frac{\mathbf{1}_{\left[\underline{v}^*, \overline{v}^*\right)}(u)}{\overline{v}^* - \underline{v}^*}$$

for some $\kappa > 0$ small enough to ensure $x_i \pm \Delta_i \in \mathscr{X}$. For example, the choice analogous to (54) will do.

Next, pick any $\varepsilon \in (0,1)$ such that $\varepsilon \leq \overline{\varepsilon}$, where

$$\overline{\varepsilon} = \frac{1-\omega}{\max_{u_i \in \left[\underline{v}^*, \overline{v}^*\right]} \left| \int_{u_i}^{\overline{v}^*} \Delta_i(u) du \right|} > 0.$$

Since $\boldsymbol{x}$ is piecewise-constant on $\left[0, \overline{\boldsymbol{v}}\right]$, the mapping

$$u_j \in \left[0, \overline{v}_j\right] \mapsto B_{\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}}(\boldsymbol{u})$$

is piecewise linear, holding other coordinates $(u_1, \ldots, u_{j-1}, u_{j+1}, \ldots, u_n)$ fixed in $\prod_{k \neq j} \left[0, \overline{v}_k\right]$. As a result, (49) is equivalent to

$$B_{\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}}(\boldsymbol{u}) \leq 1 \quad \forall \boldsymbol{u} \in \prod_{j=1}^{n} \left(D_j\left(0, \overline{v}_j\right) \cup \left\{0, \overline{v}_j\right\}\right). \tag{56}$$

By construction, our perturbation $\boldsymbol{\Delta}$ verifies (56) at all grid points, including $\overline{\boldsymbol{v}}$ (due to the definition of $\boldsymbol{\Delta}$), and hence (49) holds as well. But then Lemma 3 implies that $\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}$ is feasible, which contradicts extremality of $\boldsymbol{x}$.

To sum up, we have shown that Case 2 with $\overline{s} > \bar{\psi}(0)$ is inconsistent with extremality. We now turn to Case 1, where we consider a sag $\left(\underline{s}, \overline{s}\right)$ so that G equals 1 at its endpoints. As before, we denote $\overline{\boldsymbol{v}} = \boldsymbol{\rho}\left(\overline{s}\right)$ and $\underline{\boldsymbol{v}} = \boldsymbol{\rho}\left(\underline{s}\right)$, which satisfy $\prod_{i=1}^{n} \overline{v}_i > 0$ due to $\overline{s} > 0$ and $\prod_{i=1}^{n} \underline{v}_i > 0$ whenever $\underline{s} > 0$. Below, we study $\underline{s} > 0$ and $\underline{s} = 0$ separately, as Lemma 7 does not cover the latter.

Case 1 with $\underline{s} > 0$. By Lemmas 5, 6, and 7, for each i with $\underline{v}_i < \overline{v}_i$, the function x_i is piecewise-constant on this interval with finitely many jumps, including $\underline{v}_i, \overline{v}_i$.

By Lemma 8, we can find two distinct coordinates $i \neq j$ such that $\underline{v}_i < \overline{v}_i$ and $\underline{v}_j < \overline{v}_j$. Following the logic of Step 2, we shall construct a perturbation to contradict extremality of $\boldsymbol{x}$ assumed in the premise.

Let $\underline{v}_i^*$, $\overline{v}_i^*$ and $\underline{v}_j^*$, $\overline{v}_j^*$ be the last two points in $D_i\left(\underline{v}_i, \overline{v}_i\right)$ and $D_j\left(\underline{v}_j, \overline{v}_j\right)$, respectively, where $\overline{v}_i^* = \overline{v}_i$ and $\overline{v}_j^* = \overline{v}_j$. As before, denote the value of

$$\max_{\boldsymbol{u} \in [0,1]^n} B_{\boldsymbol{x}}(\boldsymbol{u}) \ \text{ s.t. } \ \boldsymbol{u} \in \prod_{k=1}^{n} D_k\left(\underline{v}_k, \overline{v}_k\right), \ \boldsymbol{u} \neq \underline{\boldsymbol{v}}, \overline{\boldsymbol{v}}$$

by ω, which is again well-defined and strictly less than 1 due to $G < 1$ on $(\underline{s}, \overline{s})$.

Now, define a perturbation $\boldsymbol{\Delta} : [0,1] \to \mathbf{R}^n$ by setting $\Delta_k = 0$ for all indices k other than i, j. For coordinates i and j, set

$$\Delta_i(u) = \kappa \cdot \frac{\mathbf{1}_{[\underline{v}_i^*, \overline{v}_i^*)}(u)}{\overline{v}_i^* - \underline{v}_i^*}, \quad \Delta_j(u) = -\kappa \cdot \frac{\mathbf{1}_{[\underline{v}_j^*, \overline{v}_j^*)}(u)}{\overline{v}_j^* - \underline{v}_j^*}$$

for some $\kappa > 0$ small enough to ensure $x_i \pm \Delta_i \in \mathscr{X}$ and $x_j \pm \Delta_j \in \mathscr{X}$. For example, the minimum of the two values defined analogously to (54) for coordinates i and j will suffice.

Note that

$$\int_{\underline{v}_i}^{\overline{v}_i} \Delta_i(u)du + \int_{\underline{v}_j}^{\overline{v}_j} \Delta_j(u)du = 0,$$

which implies

$$B_{\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}}(\boldsymbol{u}) \leq 1 \quad \forall \boldsymbol{u} \in \prod_{k=1}^{n} D_k(\underline{v}_k, \overline{v}_k), \tag{57}$$

for all $\varepsilon \in (0,1)$ such that $\varepsilon \leq \overline{\varepsilon}$, where

$$\overline{\varepsilon} = \frac{1}{2} \frac{1 - \omega}{\max\left\{\max_{u_i \in [\underline{v}_i^*, \overline{v}_i^*]} \left|\int_{u_i}^{\overline{v}_i^*} \Delta_i(u)du\right|, \max_{u_j \in [\underline{v}_j^*, \overline{v}_j^*]} \left|\int_{u_j}^{\overline{v}_j^*} \Delta_j(u)du\right|\right\}} > 0.$$

Since each function x_k is piecewise-constant on $[\underline{v}_k, \overline{v}_k]$, the argument in Case 2 implies that (57) is equivalent to (52). We conclude that $\boldsymbol{x} \pm \varepsilon \boldsymbol{\Delta}$ is feasible due to Lemma 4, which contradicts extremality of $\boldsymbol{x}$.

Case 1 with $\underline{s} = 0$*.* Since $\underline{s} = 0$, we have $G(0) = 1$, which means $\bar{\psi}(0) = 0$ because $G < 1$ on $[0, \bar{\psi}(0))$ due to Lemma 2 and the proof of Theorem 1. Since $x_i = 0$ on $[0, \underline{v}_i)$, the discontinuity points up to $\overline{v}_i$ are the same whether one starts at $\underline{v}_i$ or 0, exactly as happens in Case 2. By the first step in the proof of that case, the discontinuity set $D_i(0, \overline{v}_i)$ contains at least two points for some index i. Note also that, for any $\boldsymbol{u} \leq \overline{\boldsymbol{v}}$ with $\prod_{j=1}^n u_j = 0$ and $\boldsymbol{u} \nleq \underline{\boldsymbol{v}}$, we have $u_j > \underline{v}_j$ for some j and $x_j > 0$ above $\underline{v}_j$, so that

$$B_{\boldsymbol{x}}(\boldsymbol{u}) = \sum_{j=1}^{n} \int_{u_j}^{1} x_j(u)du < \sum_{j=1}^{n} \int_{\underline{v}_j}^{1} x_j(u)du = 1.$$

Furthermore, by the argument of Case 1 with $\underline{s} > 0$, there cannot be two such indices; hence, we have

$$D_j(0, \overline{v}_j) \subseteq \{\overline{v}_j\} \quad \forall j \neq i.$$

We proceed in two logical steps. First, we argue that $D_i(0, \overline{v}_i)$ must contain at least three elements, and then we complete the proof by ruling out this possibility as well.

Step 1. Towards a contradiction, suppose that $D_i(0, \overline{v}_i)$ contains exactly two discontinuity points, where one of them is necessarily $\overline{v}_i$ due to Lemma 6. Let $\underline{v}_i^* < \overline{v}_i$ be the other discontinuity of x_i.

Suppose first that $\underline{v}_i^* > 0$ is strictly positive. Then, for each sufficiently small $s > 0$ (any value strictly less than $\sqrt[n]{\underline{v}_i^* \prod_{j \neq i} \overline{v}_j}$ will do), we have

$$G(s) = s^n + \sum_{j=1}^{n} \int_0^1 x_j(u) du$$

because $x_j = 0$ on $[0, \overline{v}_j)$ for all j other than i and $x_i = 0$ on $[0, \underline{v}_i^*)$. But this implies that

$$\sum_{j=1}^{n} \int_0^1 x_j(u) du < 1,$$

contradicting the fact that $G(\underline{s}) = 1$, since $\bar{\psi}(0) = \underline{s} = 0$.

It follows that we must have $\underline{v}_i^* = 0$, meaning that x_i jumps at $u = 0$ and $u = \overline{v}_i$. Let χ denote the size of the first jump, that is, the constant value of x_i on the whole interval $[0, \overline{v}_i)$. Recall that G equals 1 at the endpoints of the sag $(\underline{s}, \overline{s})$, which can be expressed as

$$\overline{v}_i \cdot \chi + \sum_{j=1}^{n} \int_{\overline{v}_j}^1 x_j(u) du = \overline{v}_i \cdot \prod_{j \neq i} \overline{v}_j + \sum_{j=1}^{n} \int_{\overline{v}_j}^1 x_j(u) du = 1. \tag{58}$$

One immediate implication of (58) is that the jump size of x_i satisfies

$$\chi = \prod_{j \neq i} \overline{v}_j.$$

Take any $s \in (0, \overline{s})$ and consider $\boldsymbol{u}$ that equals $\overline{\boldsymbol{v}}$ for all coordinates but

$$u_i = \frac{\overline{v}_i}{\overline{s}^n} s^n.$$

Evaluate Border's constraint at this point to obtain

$$B_{\boldsymbol{x}}(\boldsymbol{u}) = s^n + (\overline{v}_i - u_i) \cdot \chi + \sum_{j=1}^{n} \int_{\overline{v}_j}^1 x_j(u) du = 1,$$

where the last equality follows from the definition of χ, u_i, and (58). Since $\boldsymbol{u}$ is feasible in the problem defining $G(s)$ in Lemma 2, we must have $G(s) = 1$, which contradicts the assumption that $G < 1$ on $(\underline{s}, \overline{s})$, as $\underline{s} = 0$.

Step 2. We have established that there are at least three elements in $D_i(0, \overline{v}_i)$. To conclude the proof, we shall again invoke a perturbational argument that has been extensively used earlier in this proof as well as to establish the auxiliary lemmas.

Let $\underline{v}_i^* < \overline{v}_i^* < \overline{v}_i$ be the last three such points, where the fact that x_i jumps at the right endpoint is due to Lemma 6. Denote the value of the following program

$$\max_{\boldsymbol{u}\in[0,1]^n} B_{\boldsymbol{x}}(\boldsymbol{u}) \text{ s.t. } \boldsymbol{u} \in \prod_{j=1}^{n} \left(D_j\left(0,\overline{v}_j\right) \cup \left\{0,\overline{v}_j\right\}\right),\ \boldsymbol{u} \nleq \underline{\boldsymbol{v}},\ \boldsymbol{u} \neq \overline{\boldsymbol{v}}$$

by ω. It is well-defined because each discontinuity set is finite, and it is strictly less than one because $G < 1$ on the interior of the sag and, as noted above, $B_{\boldsymbol{x}} < 1$ at grid points with a zero coordinate. Grid points $\boldsymbol{u} \leq \underline{\boldsymbol{v}}$ are excluded, since $B_{\boldsymbol{x}}$ equals 1 there.

Consider a perturbation $\boldsymbol{\Delta} : [0,1] \to \mathbf{R}^n$ used in the proof of the second part of Lemma 5, that is, $\Delta_j = 0$ for all $j \neq i$, whereas, for coordinate i, we have

$$\Delta_i(u) = \kappa \cdot \left(\frac{\mathbf{1}_{\left[\underline{v}_i^*,\overline{v}_i^*\right)}(u)}{\overline{v}_i^* - \underline{v}_i^*} - \frac{\mathbf{1}_{\left[\overline{v}_i^*,\overline{v}_i\right)}(u)}{\overline{v}_i - \overline{v}_i^*} \right)$$

for some $\kappa > 0$ small enough to ensure $\Delta_i \neq 0$ and $x_i \pm \Delta_i \in \mathcal{X}$. Again, the choice analogous to (54) will do.

By construction,

$$\int_{\underline{v}_i^*}^{\overline{v}_i} \Delta_i(u)du = 0.$$

Hence, the perturbation preserves Border's constraint at $\overline{\boldsymbol{v}}$ and at any grid point $\boldsymbol{u} \leq \underline{\boldsymbol{v}}$, where $u_i \leq \underline{v}_i \leq \underline{v}_i^*$. In particular, we obtain

$$B_{\boldsymbol{x}\pm\varepsilon\boldsymbol{\Delta}}(\boldsymbol{u}) \leq 1 \quad \forall \boldsymbol{u} \in \prod_{j=1}^{n} \left(D_j\left(0,\overline{v}_j\right) \cup \left\{0,\overline{v}_j\right\}\right) \tag{59}$$

for every $\varepsilon \in (0,1)$ such that $\varepsilon \leq \overline{\varepsilon}$, where

$$\overline{\varepsilon} = \frac{1-\omega}{\max_{u_i \in \left[\underline{v}_i^*,\overline{v}_i\right]} \left| \int_{u_i}^{\overline{v}_i} \Delta_i(u)du \right|} > 0.$$

Similarly to Case 2, since each function x_j is piecewise-constant on $\left[0,\overline{v}_j\right]$, (59) is equivalent to (49). Thus, by Lemma 3, both $\boldsymbol{x} \pm \varepsilon\boldsymbol{\Delta}$ are feasible, contradicting extremality of $\boldsymbol{x}$.

Putting all pieces together. To sum up, we have established that, if $\boldsymbol{x}$ is an extremal reduced form, then we cannot have any sag $\left(\underline{s},\overline{s}\right)$ with $\overline{s} > \bar{\psi}(0)$. But this means that

$$G(s) = 1 \quad \forall s \in \left[\bar{\psi}(0), 1\right], \tag{60}$$

which, as explained in the proof of Theorem 1, is equivalent to Border's constraint being identically 1 along the whole principal curve.

As explained in Section A.1, the ψ-transform can be inverted to obtain the underlying CDF in $\mathcal{X}$. It is easy to see that $\overline{x} = \bar{\psi}^{-1}(u)/u$ is the CDF corresponding to $\bar{\psi}$. Using $\overline{x}$, we can unpack

(60) for any $s \geq \bar{\psi}(0)$ as

$$\begin{aligned} \frac{1-s^n}{n} &= \int_{\bar{\psi}^{-1}(s)}^{1} \iota d \ln \bar{\psi}(\iota) \\ &= \int_{\bar{\psi}\left(\bar{\psi}^{-1}(s)\right)}^{1} \overline{x}(u) du \\ &= \int_{s}^{1} \overline{x}(u) du, \end{aligned}$$

where the first line uses (48) in Lemma 2, the second line is due to the integral identity in the proof of that lemma, and the third line follows from continuity of $\bar{\psi}$. But this means that $\overline{x}(u)$ equals u^{n-1} above $\bar{\psi}(0)$ and is zero below that threshold. This is equivalent to $\bar{\psi}(\iota) = \max\left\{\bar{\psi}(0), \sqrt[n]{\iota}\right\}$, or (16) in the theorem, since $\bar{\psi}(e^{-t}) = e^{-\frac{1}{n}\sum_{i=1}^{n}\delta_i(t)}$. □

A.3.3 Proofs of Theorems 2 and 3: Statement II

Proof. Take a feasible reduced form $\boldsymbol{x}$ that satisfies (16), which, as shown at the end of the proof of Statement I, is equivalent to $\bar{\psi}(\iota) = \max\left\{\bar{\psi}(0), \sqrt[n]{\iota}\right\}$ in terms of the ψ-transforms $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_n)$. By the definition of the ψ-transform and its bijectivity established in Section A.1, we have $u x_i(u) = \psi_i^{-1}(u)$ for all bidders i, that is, the canonical scores of $\boldsymbol{x}$ are $q_i(u) = \psi_i^{-1}(u)$. These scores are weakly increasing and right-continuous, and they satisfy (21) because each ψ_i^{-1} is strictly increasing wherever positive, making a tie at a positive score a null event. Thus, the score allocation with the canonical scores is well-defined.

Fix an index i and let $\tilde{x}_i$ be the interim winning probability of that bidder induced by the score allocation with these scores, that is,

$$\widetilde{x}_i(u) = \begin{cases} \prod_{j \neq i} \mathbf{E}\left[\mathbf{1}_{\left[0, \psi_i^{-1}(u)\right)}\left(\psi_j^{-1}(u_j)\right)\right] & \text{if } u \geq \psi_i(0), \\ 0 & \text{if } u < \psi_i(0). \end{cases}$$

We shall show that x_i and $\tilde{x}_i$ coincide. Clearly, both these functions equal 0 on $[0, \psi_i(0))$, and so it suffices to consider their values on the complement of that set.

Let $u \geq \psi_i(0)$. Since, for each j, by definition, ψ_j^{-1} is zero on $[0, \psi_j(0))$ and strictly increasing on its complement, we have

$$\tilde{x}_i(u) = \prod_{j \neq i} \psi_j\left(\psi_i^{-1}(u)\right).$$

Multiply each side by $u = \psi_i\left(\psi_i^{-1}(u)\right)$ to obtain

$$\begin{aligned}
u\widetilde{x}_i(u) &= \prod_{j=1}^{n} \psi_j\left(\psi_i^{-1}(u)\right) \\
&= \left[\bar{\psi}\left(\psi_i^{-1}(u)\right)\right]^n \\
&= \max\left\{\left[\bar{\psi}(0)\right]^n, \psi_i^{-1}(u)\right\} \\
&= ux_i(u),
\end{aligned}$$

where the second line follows from the definition of $\bar{\psi}$, and the third line is due to $\bar{\psi}(\iota) = \max\left\{\bar{\psi}(0), \sqrt[n]{\iota}\right\}$. Finally, the last line is due to $ux_i(u) = \psi_i^{-1}(u)$ and

$$\psi_i^{-1}(u) \geq \psi_i^{-1}(\psi_i(0)) \geq \left[\bar{\psi}(0)\right]^n.$$

To see the last inequality, note that $\psi_i(\iota) = \psi_i(0)$ for $\iota \leq \left[\bar{\psi}(0)\right]^n$ due to monotonicity of $(\psi_1, \ldots, \psi_n)$. $\square$

A.3.4 Proofs of Theorems 2 and 3: Statement III

Proof. This directly follows from the argument in Section 6.2 establishing that every score allocation induces an extremal reduced form and is the a.e.-unique allocation inducing it. $\square$

A.4 Optimality

A.4.1 Proof of Proposition 2

Proof. Using $H_i(0, u) = 0$ and the ψ-transform of x_i denoted ψ_i, we can obtain

$$\begin{aligned}
\int_0^1 H_i(x_i(u), u)du &= \int_0^1 \int_0^{ux_i(u)} \frac{\partial}{\partial x_i} H_i\left(\frac{\iota}{u}, u\right) d\iota\, d\ln u \\
&= \int_0^1 \int_{\psi_i(\iota)}^1 \frac{\partial}{\partial x_i} H_i\left(\frac{\iota}{u}, u\right) d\ln u\, d\iota \\
&= \int_0^\infty e^{-t} R_i(\delta_i(t), t)dt,
\end{aligned}$$

where the first line uses the change of variables $x_i = \iota/u$, and the second one follows from the change of the order of integration. The last line is due to the change of variables $\iota = e^{-t}, u = e^{-\tau}$, and the definition of R_i in (31).

Next, by Theorem 1 and (13), a reduced form is feasible if and only if its δ-transforms satisfy the constraint in (33) at all times. Since the δ-transform is a bijection between $\mathscr{X}$ and $\mathscr{D}$ by Lemma 1, the first part of the proposition follows. For the second part, Theorem 2 says that a feasible reduced form is extremal if and only if $\sum_{i=1}^n \delta_i(t) = \min\{t, T\}$ for some $T \in [0, \infty]$. As $\boldsymbol{\delta}(0) = \mathbf{0}$ and $\boldsymbol{\delta}$ is absolutely continuous, this is (34) almost everywhere. $\square$

A.4.2 Proof of Theorem 4

Proof. The proof proceeds in two steps. First, we show that the iso-MR path $\boldsymbol{\delta}^{\mathrm{s}}$ is the unique solution to the optimal control problem in Proposition 2. Second, we establish that the corresponding reduced form is extremal, has no exclusion, and satisfies (23).

Step 1. By the implicit function theorem applied to (26), p^{s} is continuously differentiable on each of the at most n time intervals on which the set of active bidders is constant. At a time at which a bidder becomes active, the system of equations changes, and p^{s} may have a kink. Since p^{s} is strictly decreasing and nonnegative, $\lambda := -(p^{\mathrm{s}})'$ is positive, piecewise continuous with at most n jumps, and integrable with

$$\int_t^{\infty} \lambda(\tau)d\tau = p^{\mathrm{s}}(t) - p^{\mathrm{s}}(\infty),$$

where $p^{\mathrm{s}}(\infty) \geq 0$ denotes the limit of p^{s} as $t \to \infty$. For any $\boldsymbol{\delta}$ that is feasible in the optimal control problem of Proposition 2, define

$$\mathscr{C}(\boldsymbol{\delta}) := \int_0^{\infty} \left[1 - e^{-\sum_{i=1}^n \delta_i(t)} - \int_0^t e^{-\tau} \sum_{i=1}^n \delta_i'(\tau)d\tau\right] \lambda(t)dt + p^{\mathrm{s}}(\infty)\left[1 - \int_0^{\infty} e^{-\tau} \sum_{i=1}^n \delta_i'(\tau)d\tau\right], \tag{61}$$

and remark that

$$\mathscr{C}(\boldsymbol{\delta}) \geq 0,$$

because each term in the square brackets is nonnegative—the first by the constraint of the control problem at time t, and the second by the same constraint in the limit as $t \to \infty$. To proceed, we rewrite $\mathscr{C}(\boldsymbol{\delta})$ solely in terms of $\boldsymbol{\delta}$ rather than its derivatives. First, integrating by parts and using $\boldsymbol{\delta}(0) = \mathbf{0}$ together with $e^{-t} \sum_{i=1}^n \delta_i(t) \to 0$, we can obtain

$$\begin{aligned}\mathscr{C}(\boldsymbol{\delta}) = \int_0^{\infty} &\left(1 - e^{-\sum_{i=1}^n \delta_i(t)} - e^{-t}\sum_{i=1}^n \delta_i(t) - \int_0^t e^{-\tau} \sum_{i=1}^n \delta_i(\tau)d\tau\right)\lambda(t)dt \\ &+ p^{\mathrm{s}}(\infty)\left(1 - \int_0^{\infty} e^{-\tau} \sum_{i=1}^n \delta_i(\tau)d\tau\right).\end{aligned} \tag{62}$$

Second, changing the order of integration in the terms involving $\int_0^t e^{-\tau} \sum_{i=1}^n \delta_i(\tau)d\tau$ and using the relationship between λ and p^{s}, we can further rewrite (62) as

$$\mathscr{C}(\boldsymbol{\delta}) = \int_0^{\infty} \left(\left(1 - e^{-\sum_{i=1}^n \delta_i(t)} - e^{-t}\sum_{i=1}^n \delta_i(t)\right)\lambda(t) - e^{-t}\sum_{i=1}^n \delta_i(t)p^{\mathrm{s}}(t)\right)dt + p^{\mathrm{s}}(\infty). \tag{63}$$

By construction, the expression in (63) is nonnegative for all feasible policies irrespective of their extremality or optimality; therefore,

$$\begin{aligned}\int_0^\infty e^{-t}\sum_{i=1}^n R_i(\delta_i(t),t)dt &\le \int_0^\infty e^{-t}\sum_{i=1}^n R_i(\delta_i(t),t)dt + \mathscr{C}(\boldsymbol{\delta})\\ &= p^{\mathrm{s}}(0) + \int_0^\infty e^{-t}\left[\sum_{i=1}^n R_i(\delta_i(t),t) - \left(e^{t-\sum_{i=1}^n \delta_i(t)} + \sum_{i=1}^n \delta_i(t)\right)\lambda(t) - \sum_{i=1}^n \delta_i(t)p^{\mathrm{s}}(t)\right]dt. \quad (64)\end{aligned}$$

It is easy to see that $\mathscr{C}(\boldsymbol{\delta}^{\mathrm{s}}) = 0$; therefore, the bound in (64) is actually tight for the iso-MR path. We now make the bound more permissive while keeping its tightness when evaluated at $\boldsymbol{\delta}^{\mathrm{s}}$. Specifically, since the term in the square brackets in (64) is strictly concave in $\boldsymbol{\delta}(t)$, this concave function lies below its tangent at $\boldsymbol{\delta}^{\mathrm{s}}(t)$. Combining this gradient inequality with $\left(-e^{t-\sum_{i=1}^n \delta_i^{\mathrm{s}}(t)} + 1\right)\lambda(t) = 0$, which is due to our construction of $\boldsymbol{\delta}^{\mathrm{s}}$ and λ, we obtain

$$\begin{aligned}\int_0^\infty e^{-t}\sum_{i=1}^n R_i(\delta_i(t),t)dt + \mathscr{C}(\boldsymbol{\delta}) &\le \int_0^\infty e^{-t}\sum_{i=1}^n R_i(\delta_i^{\mathrm{s}}(t),t)dt\\ &+ \int_0^\infty e^{-t}\sum_{i=1}^n\left[\frac{\partial}{\partial\delta_i}R_i(\delta_i^{\mathrm{s}}(t),t) - p^{\mathrm{s}}(t)\right]\left(\delta_i(t)-\delta_i^{\mathrm{s}}(t)\right)dt. \quad (65)\end{aligned}$$

The term in the square brackets is zero for every active bidder due to the definition of the iso-MR path in (26), whereas, for a bidder with $\delta_i^{\mathrm{s}}(t) = 0$, it is nonpositive and multiplies $\delta_i(t) - \delta_i^{\mathrm{s}}(t) = \delta_i(t) \ge 0$. Hence, the second line of (65) is nonpositive for every feasible policy and vanishes at $\boldsymbol{\delta}^{\mathrm{s}}$, thereby establishing optimality of this policy in the optimal control problem.

As for uniqueness, any optimal policy attains equality throughout (64) and (65). Since the function in the second line of (64) is strictly concave in $\boldsymbol{\delta}$ under Condition (A), equality in the gradient inequality forces $\boldsymbol{\delta}(t) = \boldsymbol{\delta}^{\mathrm{s}}(t)$ for almost every t, hence for every t by continuity.

Step 2. Let $\boldsymbol{x}^{\mathrm{s}}$ be the reduced form that corresponds to $\boldsymbol{\delta}^{\mathrm{s}}$, i.e., by (11), we have

$$ux_i^{\mathrm{s}}(u) = e^{-(\delta_i^{\mathrm{s}})^{-1}(-\ln u)}. \quad (66)$$

By construction, $\boldsymbol{\delta}^{\mathrm{s}}$ satisfies $\sum_{i=1}^n \delta_i^{\mathrm{s}}(t) = t$ at all times, meaning that $\boldsymbol{x}^{\mathrm{s}}$ is extremal due to Theorem 2. Moreover, since some bidder's coordinate grows without bound, the score allocation with the canonical scores (66) awards the good with probability one at almost every type profile, so the auction has no exclusion. In view of Theorem 3, $\boldsymbol{x}^{\mathrm{s}}$ is induced by awarding the good to bidder i whenever

$$u_i x_i^{\mathrm{s}}(u_i) \ge u_j x_j^{\mathrm{s}}(u_j) \;\; \forall j \ne i.$$

Equivalently, using (66) and applying p^{s}, which is strictly decreasing, to each side, we can rewrite this condition as the requirement that bidder i's MR

$$p^{\mathrm{s}}\left((\delta_i^{\mathrm{s}})^{-1}(-\ln u_i)\right) = \frac{\partial}{\partial x_i}H_i(x_i^{\mathrm{s}}(u_i), u_i)$$

be the highest one, where the equality is due to (26). Since $p^{\mathrm{s}} \geq 0$, this is exactly (23), with the types below the exclusion types covered by the argument in Footnote 23 of the main text. $\square$

A.4.3 Proof of Proposition 1

Proof. Let $\boldsymbol{\delta}^{\mathrm{M}}$ and $\boldsymbol{\delta}^{*}$ be the δ-transforms of Myerson's reduced form and the optimum. We begin with two details deferred in the main text. First, Myerson's auction has no exclusion, since bidders' MRs are their MVVs times $h' > 0$, MRs and MVVs have the same sign, so Lemma 9 applies to both auctions alike. Second, (29) holds for all bidders, including those dwelling at the top of Myerson's principal curve. For a bidder i that has left the top, it holds by the MVV equalization, $\zeta_1\big(e^{-\delta_1^{\mathrm{M}}}\big) = \zeta_i\big(e^{-\delta_i^{\mathrm{M}}}\big) \geq 0$, together with $\delta_1^{\mathrm{M}} \geq \delta_i^{\mathrm{M}}$ and monotonicity of h'. For a bidder i dwelling at the top, $\delta_i^{\mathrm{M}}(t) = 0$, Myerson's auction gives $\zeta_i(1) \leq \zeta_1\big(e^{-\delta_1^{\mathrm{M}}(t)}\big)$, and again $h'\big(e^{\delta_1^{\mathrm{M}}(t)-t}\big) \geq h'(e^{-t})$, so (29) follows.

Equipped with these two facts, we are ready to complete the argument in the main text. Suppose, towards a contradiction, that there exists some time t so that $\delta_1^{\mathrm{M}}(t) > \delta_1^{*}(t)$. Then, there must exist a bidder i for whom the inequality is reversed at that time, that is, $\delta_i^{\mathrm{M}}(t) < \delta_i^{*}(t)$, because $\sum_{i=1}^{n} \delta_i^{\mathrm{M}}(t) = \sum_{i=1}^{n} \delta_i^{*}(t) = t$. In particular, $\delta_i^{*}(t) > 0$, so bidder i is active at the optimum and, by (26) and strict monotonicity of MRs in own coordinate under Condition (A),

$$p^{\mathrm{s}}(t) = \frac{\partial}{\partial \delta_i} R_i(\delta_i^{*}(t), t) < \frac{\partial}{\partial \delta_i} R_i(\delta_i^{\mathrm{M}}(t), t) \leq \frac{\partial}{\partial \delta_1} R_1(\delta_1^{\mathrm{M}}(t), t) < \frac{\partial}{\partial \delta_1} R_1(\delta_1^{*}(t), t),$$

where the middle inequality is (29). But (26) requires $\frac{\partial}{\partial \delta_1} R_1(\delta_1^{*}(t), t) \leq p^{\mathrm{s}}(t)$ whether bidder 1 is active or dwelling at the top, a contradiction. We conclude that $\delta_1^{\mathrm{M}}(t) \leq \delta_1^{*}(t)$ for all t.

Finally, note that δ_1^{*} being larger than δ_1^{M} implies that bidder 1's interim winning probability at the optimum x_1^{*} is larger than in Myerson's auction due to (11). The argument for the weakest bidder is identical and hence omitted. $\square$

Proof. The proof of Theorem 4 extends almost as is, and we focus on the differences that arise with $T^{\mathrm{s}} < \infty$. Set $\lambda := -(p^{\mathrm{s}})'$ on $[0, T^{\mathrm{s}})$ and $\lambda := 0$ afterward, so that $\int_t^{\infty} \lambda(\tau) d\tau = \max\{p^{\mathrm{s}}(t), 0\}$, and define $\mathscr{C}(\boldsymbol{\delta})$ by the first term of (61) alone. The chain (61)–(65) goes through verbatim with p^{s} replaced by $\max\{p^{\mathrm{s}}, 0\}$, and $\mathscr{C}(\boldsymbol{\delta}^{*}) = 0$, since $\sum_{i=1}^{n} \delta_i^{*}(t) = t$ on $[0, T^{\mathrm{s}})$, where λ is supported. The term in the square brackets in (65) is now

$$\frac{\partial}{\partial \delta_i} R_i(\delta_i^{*}(t), t) - \max\{p^{\mathrm{s}}(t), 0\}.$$

This term is zero for active bidders on $[0, T^{\mathrm{s}})$ and nonpositive at $\delta_i^{*}(t) = 0$. On $[T^{\mathrm{s}}, \infty)$, it is bidder i's MR at the unconstrained optimum in (36), which is zero whenever $0 < \delta_i^{*}(t) < t$, nonpositive at $\delta_i^{*}(t) = 0$, where it multiplies $\delta_i(t) \geq 0$, and nonnegative at $\delta_i^{*}(t) = t$, where it multiplies $\delta_i(t) - t \leq 0$. Hence, the second line of (65) is again nonpositive for every feasible policy and vanishes at $\boldsymbol{\delta}^{*}$, and uniqueness follows as before. To sum up, we have established that

$$\int_0^\infty e^{-t}\sum_{i=1}^n R_i(\delta_i^*(t),t)dt \geq \int_0^\infty e^{-t}\sum_{i=1}^n R_i(\delta_i(t),t)dt$$

for every $\boldsymbol{\delta}$ feasible in (33).

It remains to show that $\boldsymbol{\delta}^*$, and hence the reduced form $\boldsymbol{x}^*$, is itself feasible. On $[0,T^{\mathrm{s}})$, $\boldsymbol{\delta}^*$ coincides with the iso-MR path, which is in $\mathscr{D}^n$ by Condition (B) and satisfies the constraint in (33) with equality, since $\sum_{i=1}^n \delta_i^{\mathrm{s}}(t)=t$. On $[T^{\mathrm{s}},\infty)$, $\delta_i^*(t)$ is the maximizer of $R_i(\cdot,t)$ on $[0,t]$, that is, the unconstrained optimum in (36) along the level set $u_ix_i=e^{-t}$; it solves $\frac{\partial}{\partial x_i}H_i\big(e^{\delta_i^*(t)-t},e^{-\delta_i^*(t)}\big)=0$ whenever this equation has a root in $(0,t)$, and equals 0 or t otherwise, with $(\delta_i^*)'\in\{0,1\}$ there. In the former case, differentiating with respect to t gives

$$(\delta_i^*)' = \frac{\frac{\partial}{\partial \ln x_i}\left(\frac{\partial H_i}{\partial x_i}\right)}{\frac{\partial}{\partial \ln x_i}\left(\frac{\partial H_i}{\partial x_i}\right)-\frac{\partial}{\partial \ln u_i}\left(\frac{\partial H_i}{\partial x_i}\right)} \in [0,1)$$

at $(x_i,u_i)=\big(e^{\delta_i^*(t)-t},e^{-\delta_i^*(t)}\big)$ whenever $0<\delta_i^*(t)<t$, as the numerator is nonpositive by concavity of H_i in x_i and the denominator is negative by Condition (A). The two pieces agree at T^{s}, where $p^{\mathrm{s}}(T^{\mathrm{s}})=0$ makes the iso-MR condition and the unconstrained optimum coincide. Hence, $\boldsymbol{\delta}^*\in\mathscr{D}^n$. Furthermore, since $\delta_i^*\leq\delta_i^{\mathrm{s}}$ on $[T^{\mathrm{s}},\infty)$ due to $0\geq p^{\mathrm{s}}$ there and strict monotonicity of MRs in own coordinate, we have $x_i^*\leq x_i^{\mathrm{s}}$. Hence, $\boldsymbol{x}^*$ is feasible along with $\boldsymbol{x}^{\mathrm{s}}$, and $r_i^*=x_i^*/x_i^{\mathrm{s}}\in[0,1]$ under the convention $0/0=0$. □

A.4.4 Proof of Theorem 6

Proof. The proof proceeds in three steps. First, we show that the optimal reduced form is unique and extremal. Second, we show that it is the $\underline{\boldsymbol{u}}^*$-capped iso-MR reduced form for a maximizer $\underline{\boldsymbol{u}}^*$ of (39). Third, we derive (40).

Step 1. We first record two consequences of our assumptions. Since $H_i(0,\cdot)=0$ and $\frac{\partial H_i}{\partial x_i}$ is weakly increasing in u_i by supermodularity,

$$H_i(x_i,u_i)=\int_0^{x_i}\frac{\partial}{\partial x_i}H_i(s,u_i)ds$$

is weakly increasing in u_i. Since $\max\{0,H_i\}$ is convex in x_i, nonnegative, and vanishes at $x_i=0$, for $x_i<y_i$ we have

$$\max\{0,H_i(x_i,u_i)\}\leq\frac{x_i}{y_i}\max\{0,H_i(y_i,u_i)\}\leq\max\{0,H_i(y_i,u_i)\},$$

so $\max\{0,H_i\}$ is weakly increasing in x_i.

For a feasible reduced form $\boldsymbol{x}$, let $\tilde{\boldsymbol{x}}$ withhold the good whenever $H_i\leq 0$, that is, $\tilde{x}_i$ is the right-continuous modification of $x_i\cdot\mathbf{1}_{(0,\infty)}(H_i(x_i(u),u))$. Due to the two properties of H_i derived above and monotonicity of each x_i, $\tilde{\boldsymbol{x}}$ is a reduced form, and it is feasible, as it is dominated by

$\boldsymbol{x}$. Its revenue satisfies

$$\sum_{i=1}^{n} \int_0^1 H_i(\tilde{x}_i(u), u) du = \sum_{i=1}^{n} \int_0^1 \max\{0, H_i(x_i(u), u)\} du \geq \sum_{i=1}^{n} \int_0^1 H_i(x_i(u), u) du.$$

Next, the optimal reduced form $\boldsymbol{x}^*$ is unique, since the objective in (33) is strictly concave in $\boldsymbol{\delta}$ under Condition (A), its feasible set is convex, and $\boldsymbol{\delta} \mapsto \boldsymbol{x}$ is a bijection by Lemma 1. It follows that $\boldsymbol{x}^*$ maximizes the convex functional $\sum_{i=1}^{n} \int_0^1 \max\{0, H_i(x_i(u), u)\} du$, and every other maximizer of this functional is at least as large as $\boldsymbol{x}^*$, because it truncates to $\boldsymbol{x}^*$ by uniqueness. Since a maximizer of a convex functional that is dominated by every other maximizer is extremal, we conclude that $\boldsymbol{x}^*$ is extremal.

Step 2. Let $\underline{\boldsymbol{u}}^*$ be the exclusion types of $\boldsymbol{x}^*$. By Step 1 and Theorem 2, the δ-transforms of $\boldsymbol{x}^*$ satisfy $\delta_i^*(t) \in \left[0, -\ln \underline{u}_i^*\right]$ for all i and $\sum_{i=1}^{n} \delta_i^*(t) = \min\left\{t, -\sum_{j=1}^{n} \ln \underline{u}_j^*\right\}$ for all t, and the constraint in (33) then holds automatically. Hence, by Proposition 2, the auctioneer's problem can be expressed as

$$\max_{\underline{\boldsymbol{u}} \in [0,1]^n} \max_{\boldsymbol{\delta} \in \mathscr{D}^n} \int_0^\infty e^{-t} \sum_{i=1}^{n} R_i(\delta_i(t), t) dt \ \text{ s.t. } \ \delta_i(t) \in \left[0, -\ln \underline{u}_i\right] \ \forall i,$$
$$\sum_{i=1}^{n} \delta_i(t) = \min\left\{t, -\sum_{j=1}^{n} \ln \underline{u}_j\right\} \ \forall t \in \mathbf{R}_+. \quad (67)$$

Now, consider the relaxation of the inner problem in which the requirement $\boldsymbol{\delta} \in \mathscr{D}^n$ is dropped. This relaxed problem separates across t, and its value at t is

$$W(t, \underline{\boldsymbol{u}}) := \max_{\boldsymbol{\delta} \in \mathbf{R}^n} \sum_{i=1}^{n} R_i(\delta_i, t) \ \text{ s.t. } \ \delta_i \in \left[0, -\ln \underline{u}_i\right] \ \forall i, \ \sum_{i=1}^{n} \delta_i = \min\left\{t, -\sum_{j=1}^{n} \ln \underline{u}_j\right\}.$$

Let p, $\lambda_i \geq 0$, and $\mu_i \geq 0$ denote the multipliers on the equality constraint, the lower bound $\delta_i \geq 0$, and the cap $\delta_i \leq -\ln \underline{u}_i$, respectively. The first-order conditions read

$$\frac{\partial}{\partial \delta_i} R_i(\delta_i, t) - p + \lambda_i - \mu_i = 0, \qquad \lambda_i \delta_i = 0, \qquad \mu_i\left(\delta_i + \ln \underline{u}_i\right) = 0 \ \forall i.$$

By (32), these are exactly (38), and, since the objective is strictly concave under Condition (A), they are also sufficient. Hence, the unique solution is $\boldsymbol{\delta}^{\mathrm{s}}(t|\underline{\boldsymbol{u}})$, the multiplier on the equality constraint is $p^{\mathrm{s}}(t|\underline{\boldsymbol{u}})$, and the multiplier on bidder i's cap is

$$\mu_i = \frac{\partial}{\partial \delta_i} R_i\left(-\ln \underline{u}_i, t\right) - p^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) \ \text{ if } \ t > T_i(\underline{\boldsymbol{u}}), \qquad \mu_i = 0 \ \text{ otherwise.}$$

Condition (B*) is exactly what makes the relaxation tight. Under this condition, each coordinate of $\boldsymbol{\delta}^{\mathrm{s}}(\cdot\,|\underline{\boldsymbol{u}})$ is weakly increasing, and, since the coordinates sum to $\min\left\{t, -\sum_{j=1}^{n} \ln \underline{u}_j\right\}$, each is then 1-Lipschitz, so $\boldsymbol{\delta}^{\mathrm{s}}(\cdot\,|\underline{\boldsymbol{u}}) \in \mathscr{D}^n$. Hence, the inner problem in (67) is uniquely solved by the $\underline{\boldsymbol{u}}$-capped iso-MR path, and its value $V(\underline{\boldsymbol{u}})$ equals (39) by (31). Thus, $\underline{\boldsymbol{u}}^*$ maximizes (39),

$\boldsymbol{x}^*$ is the $\underline{\boldsymbol{u}}^*$-capped iso-MR reduced form by uniqueness of the inner solution, and every other maximizer yields the same reduced form by uniqueness of $\boldsymbol{x}^*$. The last claim of the theorem follows from Theorem 3.

Step 3. It remains to derive (40). Differentiating (39) with respect to $\underline{u}_i$ gives

$$\frac{\partial}{\partial \underline{u}_i} V(\underline{\boldsymbol{u}}) = -H_i\big(x_i^{\mathrm{s}}(\underline{u}_i|\underline{\boldsymbol{u}}), \underline{u}_i\big) + \sum_{j=1}^{n} \int_{\underline{u}_j}^{1} \frac{\partial}{\partial x_j} H_j\big(x_j^{\mathrm{s}}(u_j|\underline{\boldsymbol{u}}), u_j\big) \frac{\partial}{\partial \underline{u}_i} x_j^{\mathrm{s}}(u_j|\underline{\boldsymbol{u}}) du_j,$$

and (40) is the first-order condition for $\underline{\boldsymbol{u}}^*$ multiplied by $\underline{u}_i^*$ once we show that the second term equals

$$\frac{1}{\underline{u}_i} \int_{T_i(\underline{\boldsymbol{u}})}^{T(\underline{\boldsymbol{u}})} e^{-t} p^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) dt. \tag{68}$$

To this end, we compute $\frac{\partial V}{\partial \underline{u}_i}$ a second time, now from (31), which represents $V(\underline{\boldsymbol{u}})$ as $\int_0^\infty e^{-t} \sum_{j=1}^n R_j\big(\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}), t\big) dt$. By the implicit function theorem applied to (38), $\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}})$ is differentiable in $\underline{u}_i$ at every t at which the set of bidders with $\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) \in (0, -\ln \underline{u}_j)$ is locally constant, which excludes at most finitely many t under Condition (B*). Hence,

$$\frac{\partial}{\partial \underline{u}_i} V(\underline{\boldsymbol{u}}) = \int_0^\infty e^{-t} \sum_{j=1}^n \frac{\partial}{\partial \delta_j} R_j\big(\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}), t\big) \frac{\partial}{\partial \underline{u}_i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) dt. \tag{69}$$

Below, we evaluate the derivatives of $\boldsymbol{\delta}^{\mathrm{s}}(\,\cdot\,|\underline{\boldsymbol{u}})$ on various time intervals.

- *Case $t < T_i(\underline{\boldsymbol{u}})$.* By construction, $\underline{u}_i$ does not enter (38), and thus $\frac{\partial}{\partial \underline{u}_i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = 0$ for all j.
- *Case $t > T(\underline{\boldsymbol{u}})$.* Here, $\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = -\ln \underline{u}_j$ for all j, and thus $\frac{\partial}{\partial \underline{u}_i} \delta_i^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = -\frac{1}{\underline{u}_i}$ and $\frac{\partial}{\partial \underline{u}_i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = 0$ for $j \neq i$.
- *Case $t \in \big(T_i(\underline{\boldsymbol{u}}), T(\underline{\boldsymbol{u}})\big)$.* As in the previous case, $\delta_i^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = -\ln \underline{u}_i$, and thus $\frac{\partial}{\partial \underline{u}_i} \delta_i^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = -\frac{1}{\underline{u}_i}$. For $j \neq i$, we necessarily have

  $$\left(\frac{\partial}{\partial \delta_j} R_j\big(\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}), t\big) - p^{\mathrm{s}}(t|\underline{\boldsymbol{u}})\right) \frac{\partial}{\partial \underline{u}_i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = 0$$

  due to the first-order conditions. Since $\sum_{j\neq i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = t - \delta_i^{\mathrm{s}}(t|\underline{\boldsymbol{u}})$, we obtain

  $$\sum_{j \neq i} \frac{\partial}{\partial \delta_j} R_j\big(\delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}), t\big) \frac{\partial}{\partial \underline{u}_i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = p^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) \sum_{j \neq i} \frac{\partial}{\partial \underline{u}_i} \delta_j^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) = \frac{p^{\mathrm{s}}(t|\underline{\boldsymbol{u}})}{\underline{u}_i}.$$

Substituting the three cases into (69), we obtain

$$\frac{\partial}{\partial \underline{u}_i} V(\underline{\boldsymbol{u}}) = -\frac{1}{\underline{u}_i} \int_{T_i(\underline{\boldsymbol{u}})}^{\infty} e^{-t} \frac{\partial}{\partial \delta_i} R_i\big(-\ln \underline{u}_i, t\big) dt + \frac{1}{\underline{u}_i} \int_{T_i(\underline{\boldsymbol{u}})}^{T(\underline{\boldsymbol{u}})} e^{-t} p^{\mathrm{s}}(t|\underline{\boldsymbol{u}}) dt.$$

By (32), the first integrand equals $e^{-t}\frac{\partial}{\partial x_i}H_i\big(\frac{e^{-t}}{\underline{u}_i},\underline{u}_i\big)$, so the change of variables $x = \frac{e^{-t}}{\underline{u}_i}$, which maps $T_i(\underline{\boldsymbol{u}})$ to $x_i^{\mathrm{s}}(\underline{u}_i|\underline{\boldsymbol{u}})$ by the definition of the exclusion time, turns the first term into $-H_i\big(x_i^{\mathrm{s}}(\underline{u}_i|\underline{\boldsymbol{u}}),\underline{u}_i\big)$. Comparing with the first display of this step establishes (68). □

A.4.5 Proof of Proposition 3

Proof. As explained in the text, it is without loss of generality to look at group-symmetric policies, i.e., $\delta_1^{\mathrm{s}}, \delta_n^{\mathrm{s}}$ and δ_1^*, δ_n^*, corresponding to the risk-neutral and risk-averse bidders. By Theorem 5, δ_i^* coincides with δ_i^{s} before T^{s} and is given by the unconstrained optimum in (36) after that threshold. As shown in the proof of that theorem, δ_n^* is weakly increasing on $[T^{\mathrm{s}},\infty)$, whereas δ_1^* stays at $\delta_1^{\mathrm{s}}(T^{\mathrm{s}})$ there, since H_1 is linear in x_1. We prove the proposition in two steps. First, we shall show that both $\delta_1^{\mathrm{s}}(t), \delta_n^{\mathrm{s}}(t) \lesseqgtr t + \ln m$ if and only if $t \gtreqless t^o$, where $t^o = -\frac{n}{n-1}\ln m < T^{\mathrm{s}}$. Then, we shall rank $\delta_1^{\mathrm{s}}(t), \delta_n^{\mathrm{s}}(t)$ (as well as $\delta_1^*(t), \delta_n^*(t)$) before t^o and after that time, and use this ranking to complete the proof.

Step 1. Note first that $\delta_1^{\mathrm{s}}(t)$ and $\delta_n^{\mathrm{s}}(t)$ cannot lie on opposite sides of the line $t+\ln m$ with at least one of them strictly. Indeed, if $\delta_1^{\mathrm{s}}(t) \le t+\ln m < \delta_n^{\mathrm{s}}(t)$, then $g'\big(e^{\delta_n^{\mathrm{s}}(t)-t}\big) > 1$ due to unimodality, and hence

$$\frac{\partial}{\partial \delta_n}R_n(\delta_n^{\mathrm{s}}(t),t) < v\big(e^{-\delta_n^{\mathrm{s}}(t)}\big) - v'\big(e^{-\delta_n^{\mathrm{s}}(t)}\big)\big(1-e^{-\delta_n^{\mathrm{s}}(t)}\big) < \frac{\partial}{\partial \delta_1}R_1(\delta_1^{\mathrm{s}}(t),t),$$

where the second inequality is due to regularity and $\delta_n^{\mathrm{s}}(t) > \delta_1^{\mathrm{s}}(t)$, contradicting the MR equalization in (37). The case $\delta_n^{\mathrm{s}}(t) \le t+\ln m < \delta_1^{\mathrm{s}}(t)$ is ruled out symmetrically, since then $g'\big(e^{\delta_n^{\mathrm{s}}(t)-t}\big) \le 1$ and both inequalities are reversed.

Next, both $\delta_1^{\mathrm{s}}(t), \delta_n^{\mathrm{s}}(t)$ are strictly above $t+\ln m$ for t sufficiently close to 0, as $t+\ln m \to \ln m < 0$ as $t \to 0$. On the other hand, at least one of them is strictly below $t+\ln m$ for t sufficiently large, as

$$\min\{\delta_1^{\mathrm{s}}(t), \delta_n^{\mathrm{s}}(t)\} \le \frac{k}{n}\delta_1^{\mathrm{s}}(t) + \left(1-\frac{k}{n}\right)\delta_n^{\mathrm{s}}(t) = \frac{t}{n}.$$

By continuity, the minimum of $\delta_1^{\mathrm{s}}, \delta_n^{\mathrm{s}}$ crosses the line $t+\ln m$ at some time, say t^o. By the observation above, $\delta_1^{\mathrm{s}}(t^o) = \delta_n^{\mathrm{s}}(t^o) = t^o + \ln m$, so that $g'\big(e^{\delta_n^{\mathrm{s}}(t^o)-t^o}\big) = g'(m) = 1$ and, by the second line of (37), $\delta_1^{\mathrm{s}}(t^o) = \delta_n^{\mathrm{s}}(t^o) = \frac{t^o}{n}$. Clearly, this solves to the unique crossing time $t^o = -\frac{n}{n-1}\ln m$. Consequently, both $\delta_1^{\mathrm{s}}(t), \delta_n^{\mathrm{s}}(t)$ are strictly above $t+\ln m$ for $t < t^o$ and strictly below it for $t > t^o$. This crossing time is below T^{s}, since

$$p^{\mathrm{s}}(t^o) = v\left(m^{1/(n-1)}\right) - v'\left(m^{1/(n-1)}\right)\left(1-m^{1/(n-1)}\right) = \frac{\partial}{\partial x_n}H_n\left(m, m^{1/(n-1)}\right) > 0,$$

by assumption.

Step 2. Let $t < t^o$. Since $\delta_n^{\mathrm{s}}(t) > t+\ln m$, unimodality gives $g'\big(e^{\delta_n^{\mathrm{s}}(t)-t}\big) > 1$, and the MR equalization yields

$$\frac{\partial}{\partial \delta_1} R_1(\delta_1^{\mathrm{s}}(t), t) = \frac{\partial}{\partial \delta_n} R_n(\delta_n^{\mathrm{s}}(t), t) < \frac{\partial}{\partial \delta_1} R_1(\delta_n^{\mathrm{s}}(t), t),$$

which implies $\delta_1^{\mathrm{s}}(t) > \delta_n^{\mathrm{s}}(t)$ due to regularity. For $t \in (t^o, T^{\mathrm{s}})$, we have $\delta_n^{\mathrm{s}}(t) < t + \ln m$, so $g'\left(e^{\delta_n^{\mathrm{s}}(t)-t}\right) < 1$, both inequalities are reversed, and $\delta_1^{\mathrm{s}}(t) < \delta_n^{\mathrm{s}}(t)$. This ranking persists on $[T^{\mathrm{s}}, \infty)$, where δ_1^* is constant and δ_n^* is weakly increasing. Putting all these pieces together, we conclude: $\delta_1^*(t) > \delta_n^*(t)$ for $t < t^o$, $\delta_1^*(t) = \delta_n^*(t)$ for $t = t^o$, and $\delta_1^*(t) < \delta_n^*(t)$ for $t > t^o$. Finally, using (11), we can obtain $x_1^*(u) > x_n^*(u)$ for $u > e^{-\delta_i^*(t^o)} = m^{1/(n-1)}$, $x_1^*(u) < x_n^*(u)$ for $\overline{u}_n \le u < m^{1/(n-1)}$, and $x_1^*(u) = x_n^*(u)(= m)$ for $u = m^{1/(n-1)}$, as claimed. Below $\overline{u}_n < \overline{u}_1$, x_1^* vanishes, as δ_1^* never exceeds $\delta_1^{\mathrm{s}}(T^{\mathrm{s}})$. □

A.5 Regularity

This section unpacks the regularity conditions of Sections 7.1 and 8.3 in terms of primitives. Recall from (32) that bidder i's MR along the parametrization underlying the δ-transform is $\frac{\partial}{\partial \delta_i} R_i(\delta_i, t) = \frac{\partial}{\partial x_i} H_i(e^{\delta_i - t}, e^{-\delta_i})$. We track its curvature through two semi-elasticities

$$\xi_i^x = \frac{\partial}{\partial \ln x_i}\left(\frac{\partial H_i}{\partial x_i}\right) = x_i \frac{\partial^2 H_i}{\partial x_i^2}, \qquad \xi_i^u = \frac{\partial}{\partial \ln u_i}\left(\frac{\partial H_i}{\partial x_i}\right) = u_i \frac{\partial^2 H_i}{\partial x_i \partial u_i},$$

evaluated at $(x_i, u_i) = (e^{\delta_i - t}, e^{-\delta_i})$; differentiating R_i twice then gives

$$\frac{\partial^2 R_i}{\partial \delta_i^2} = \xi_i^x - \xi_i^u, \qquad \frac{\partial^2 R_i}{\partial \delta_i \partial t} = -\xi_i^x. \tag{70}$$

A.5.1 Condition (A)

By (70), Condition (A) says exactly that each R_i is strictly concave in its first argument. This holds automatically when H_i is strictly supermodular and concave in x_i, since then $\xi_i^x \le 0 < \xi_i^u$. More generally, Condition (A) requires that supermodularity is sufficiently strong relative to convexity. In Myerson's linear model, $\xi_i^x \equiv 0$, and (A) is equivalent to bidders' MVVs being strictly increasing.

A.5.2 Condition (B)

We now assume Condition (A) and examine the added restrictions imposed by Condition (B). Under strict concavity of R_i, the iso-MR path is well-defined. Recall that bidder i is active at time t if $\delta_i^{\mathrm{s}}(t) > 0$. When (B) holds, a bidder that becomes active stays active forever, so the set of active bidders changes at most n times as time progresses. By the implicit function theorem, on a time interval on which the set of active bidders is constant, say equal to I, we have

$$(\xi_i^x - \xi_i^u)(\delta_i^{\mathrm{s}})' - \xi_i^x = (p^{\mathrm{s}})' \;\; \forall i \in I, \qquad \sum_{i \in I} (\delta_i^{\mathrm{s}})' = 1,$$

where the semi-elasticities are evaluated along the iso-MR path, whereas $(\delta_i^{\mathrm{s}})' = 0$ for $i \notin I$. Solving this linear system yields

$$\underbrace{\left[\sum_{j\in I} \frac{1}{\xi_j^x - \xi_j^u}\right]}_{<0} \cdot (p^{\mathrm{s}})' = 1 - \sum_{j\in I} \frac{\xi_j^x}{\xi_j^x - \xi_j^u}, \quad \underbrace{\left[\sum_{j\in I} \frac{\xi_i^x - \xi_i^u}{\xi_j^x - \xi_j^u}\right]}_{>0} \cdot (\delta_i^{\mathrm{s}})' = 1 - \sum_{j\in I} \frac{\xi_j^x - \xi_i^x}{\xi_j^x - \xi_j^u} \quad \forall i \in I. \tag{71}$$

Under Condition (A), the first square bracket in (71) is negative, whereas the second one is positive. Hence, the positivity of the right-hand sides in (71) is sufficient for Condition (B). This can be expressed as the following requirement:

$$\sum_{j\in I} \frac{\xi_j^x - \max\{0, \max_{i\in I} \xi_i^x\}}{\xi_j^x - \xi_j^u} < 1, \tag{72}$$

which trivially holds in Myerson's linear model, as the x-semi-elasticity is identically zero.

More generally, (72) demands that the semi-elasticities be not too strong and not differ across bidders too much. For example, consider the following multiplicatively separable class in the spirit of our illustrative example and Section 7.2:

$$H_i(x_i, u_i) = u_i^{1/\beta_i} h_i(x_i)$$

with $\beta_i > 0$ and strictly increasing h_i. Along the iso-MR path, active bidders share $u_i^{1/\beta_i} h_i'(x_i) = p^{\mathrm{s}}$, so $\xi_i^u = p^{\mathrm{s}}/\beta_i$ and $\xi_i^x = \left(\frac{x_i h_i''}{h_i'}\right) p^{\mathrm{s}}$. Thus, (A) is $\left(\frac{x_i h_i''}{h_i'}\right) < 1/\beta_i$, and, for convex h_i, (72) becomes

$$\sum_{i\in I} \frac{\max_{j\in I}\left(\frac{x_j h_j''}{h_j'}\right) - \left(\frac{x_i h_i''}{h_i'}\right)}{1/\beta_i - \left(\frac{x_i h_i''}{h_i'}\right)} < 1.$$

Here, $\frac{x_i h_i''}{h_i'}$ is the elasticity of bidder i's MR with respect to own winning probability. Condition (B) adds nothing to Condition (A) when those elasticities coincide across bidders, as happens in the illustrative example.

When each H_i is strictly supermodular and concave in x_i, so that (A) holds automatically, all ξ_j^x are nonpositive, and (72) reads

$$\sum_{j\in I} \frac{\xi_j^x}{\xi_j^x - \xi_j^u} < 1.$$

If, in addition, the environment is symmetric, then the iso-MR path is symmetric, that is, $\delta_i^{\mathrm{s}}(t) = {}^t/_n$ and $x_i^{\mathrm{s}}(u) = u^{n-1}$, and the condition transforms into

$$n \frac{\xi^x}{\xi^x - \xi^u} < 1$$

evaluated at $(x_i, u_i) = (u^{n-1}, u)$, nesting the regularity condition in Theorem 1 of Gershkov et al. [2022] for the symmetric CRA specification $H(x, u) = v(u)x - (1-u)v'(u)g(x)$.

A.5.3 Condition (B*)

We keep assuming Condition (A), so that the $\underline{\boldsymbol{u}}$-capped iso-MR path is well-defined for each profile of exclusion types, $\underline{\boldsymbol{u}}$. Condition (B*) extends Condition (B) to such capped paths. Following the same argument as in the previous section, we can obtain that it is verified if, along every $\underline{\boldsymbol{u}}$-capped iso-MR path,

$$\sum_{j \in I} \frac{\xi_j^x - \max\left\{0, \max_{i \in I} \xi_i^x, \max_{i:\delta_i^{\mathrm{s}}\left(t|\underline{\boldsymbol{u}}\right) = -\ln \underline{u}_i} \xi_i^x\right\}}{\xi_j^x - \xi_j^u} < 1,$$

that is, (72) with I the set of active bidders and the inner maximum extended over the bidders at their exclusion levels. This trivially holds in Myerson's linear model, as the x-semi-elasticity is identically zero. More generally, the added terms require that the common MR decline at least as fast as the MR of each bidder at own exclusion level.

A.6 The condition $p^{\mathrm{s}} \geq 0$ in primitives

Lemma 9. *In a regular environment, the common MR p^{s} is nonnegative if and only if*

$$\prod_{i=1}^{n} \inf\left\{u_i \in [e^{-t}, 1] : \frac{\partial}{\partial x_i} H_i\left(\frac{e^{-t}}{u_i}, u_i\right) \geq 0\right\} \leq e^{-t}, \tag{73}$$

with the convention $\inf \emptyset = 1$.

In words, the ith factor in (73) is the lowest type at which bidder i's MR is still nonnegative on the level set $u_i x_i = e^{-t}$ of the canonical score. The condition says that the surface $\prod_{i=1}^{n} u_i = e^{-t}$, on which the principal curve sits at time t, contains a point at which every bidder's MR is nonnegative.

Proof. Fix each $t > 0$, and, for $p \in \mathbf{R}$, let

$$\delta_i(p) := \sup\left\{\delta_i \in [0, t] : \frac{\partial}{\partial x_i} H_i\left(e^{\delta_i - t}, e^{-\delta_i}\right) \geq p\right\},$$

so that, with $u_i = e^{-\delta_i}$, (73) reads $\sum_{i=1}^{n} \delta_i(0) \geq t$. By Condition (A), bidder i's MR is continuous and strictly decreasing in δ_i, so $\delta_i(p)$ is weakly decreasing in p, and (26) gives $\delta_i^{\mathrm{s}}(t) = \delta_i(p^{\mathrm{s}}(t))$ for all i. If $p^{\mathrm{s}}(t) \geq 0$, then

$$\sum_{i=1}^{n} \delta_i(0) \geq \sum_{i=1}^{n} \delta_i(p^{\mathrm{s}}(t)) = \sum_{i=1}^{n} \delta_i^{\mathrm{s}}(t) = t.$$

Conversely, if $p^s(t) < 0$, then $\delta_j(0) \leq \delta_j(p^s(t)) = \delta_j^s(t)$ for every bidder j. For an active bidder i (which exists, since $\sum_{j=1}^n \delta_j^s(t) = t > 0$), the inequality is strict. Hence,

$$\sum_{j=1}^{n} \delta_j(0) < \sum_{j=1}^{n} \delta_j^s(t) = t,$$

so (73) fails. □

References

G. Alberti and L. Ambrosio. A geometrical approach to monotone functions in. *Mathematische Zeitschrift*, 230(2):259–316, 1999.

P. Andreyanov, I. Krasikov, and A. Suzdaltsev. Scoring and favoritism in optimal procurement design. 2024.

P. Andreyanov, I. Krasikov, and A. Suzdaltsev. Auctions after actions. 2026.

I. Arieli, Y. Babichenko, R. Smorodinsky, and T. Yamashita. Optimal persuasion via bi-pooling. *Theoretical Economics*, 18(1):15–36, 2023.

I. Ball. Bauer's maximum principle for quasiconvex functions. *arXiv preprint arXiv:2305.04893*, 2023.

P. R. Blavatskyy. Modifying the mean-variance approach to avoid violations of stochastic dominance. *Management Science*, 56(11):2050–2057, 2010.

K. C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica: Journal of the Econometric Society*, pages 1175–1187, 1991.

K. C. Border. Reduced form auctions revisited. *Economic Theory*, 31(1):167–181, 2007.

J. Bulow and J. Roberts. The simple economics of optimal auctions. *Journal of political economy*, 97(5):1060–1090, 1989.

Y. Cai, C. Daskalakis, and S. M. Weinberg. A constructive approach to reduced-form auctions with applications to multi-item mechanism design. *arXiv preprint arXiv:1112.4572*, 2018.

O. Candogan and P. Strack. Optimal disclosure of information to privately informed agents. *Theoretical Economics*, 18(3):1225–1269, 2023.

G. Carroll and I. Segal. Robustly optimal auctions with unknown resale opportunities. *The Review of Economic Studies*, 86(4):1527–1555, 2019.

Y.-K. Che, J. Kim, and K. Mierendorff. Generalized reduced-form auctions: A network-flow approach. *Econometrica*, 81(6):2487–2520, 2013.

F. Delbaen. Convex games and extreme points. *Journal of Mathematical Analysis and Applications*, 45(1):210–233, 1974.

A. Gershkov, J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi. On the equivalence of bayesian and dominant strategy implementation. *Econometrica*, 81(1):197–220, 2013.

A. Gershkov, B. Moldovanu, P. Strack, and M. Zhang. A theory of auctions with endogenous valuations. *Journal of Political Economy*, 129(4):1011–1051, 2021.

A. Gershkov, B. Moldovanu, P. Strack, and M. Zhang. Optimal auctions: Non-expected utility and constant risk aversion. *The review of economic studies*, 89(5):2630–2662, 2022.

J. K. Goeree and A. Kushnir. A geometric approach to mechanism design. *Journal of Political Economy Microeconomics*, 1(2):321–347, 2023.

F. Gul. A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, pages 667–686, 1991.

S. Hart and P. J. Reny. Implementation of reduced form mechanisms: a simple approach and a new characterization. *Economic Theory Bulletin*, 3(1):1–8, 2015.

K. He, F. Sandomirskiy, and O. Tamuz. Private private information. *arXiv preprint arXiv:2112.14356*, 2025.

J. Jia, J. S. Dyer, and J. C. Butler. Generalized disappointment models. *Journal of Risk and Uncertainty*, 22(1):59–78, 2001.

A. Kleiner, B. Moldovanu, and P. Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021.

A. Kleiner, B. Moldovanu, P. Strack, and M. Whitmeyer. The extreme points of fusions. *arXiv preprint arXiv:2409.10779*, 2024.

P. Lahr and A. Niemeyer. Extreme points in multi-dimensional screening. *arXiv preprint arXiv:2412.00649*, 2024.

G. Loomes and R. Sugden. Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, 53(2):271–282, 1986.

A. M. Manelli and D. R. Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic theory*, 137(1):153–185, 2007.

A. M. Manelli and D. R. Vincent. Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica*, 78(6):1905–1938, 2010.

E. Maskin and J. Riley. Optimal auctions with risk averse buyers. *Econometrica: Journal of the Econometric Society*, pages 1473–1518, 1984.

S. A. Matthews. On the implementability of reduced form auctions. *Econometrica: Journal of the Econometric Society*, pages 1519–1522, 1984.

K. Mierendorff. Asymmetric reduced form auctions. *Economics Letters*, 110(1):41–44, 2011.

G. J. Minty. Monotone (nonlinear) operators in hilbert space. 1962.

D. Mishra and A. Quadir. Non-bossy single object auctions. *Economic Theory Bulletin*, 2:93–110, 2014.

R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

A. Nikzad. Constrained majorization: Applications in mechanism design. 2023.

M. A. Satterthwaite and H. Sonnenschein. Strategy-proof allocation mechanisms at differentiable points. *The Review of Economic Studies*, 48(4):587–597, 1981.

A. Seierstad and K. Sydsaeter. *Optimal control theory with economic applications*. Elsevier North-Holland, Inc., 1986.

M. Shaked and J. G. Shanthikumar. *Stochastic orders*. Springer, 2007.

R. V. Vohra. *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs. Cambridge University Press, Cambridge, 2011.

F. Yang and K. H. Yang. Multidimensional monotonicity and economic applications. *arXiv preprint arXiv:2502.18876*, 2025.

K. H. Yang and A. K. Zentefis. Monotone function intervals: Theory and applications. *American Economic Review*, 114(8):2239–2270, 2024.

M. Zhang. Auctions vs. negotiations: Optimal selling mechanism with endogenous bidder values. *Negotiations: Optimal Selling Mechanism with Endogenous Bidder Values (August 30, 2017)*, 2017.