The Cost of Simple Pricing

Christoph Schlom*

March 15, 2025

I derive novel upper bounds on the revenue loss from mechanism simplicity in two related economic selling problems. First, in the Bulow and Roberts (1989) capacity-constrained selling problem, I derive a tight upper bound on the revenue ratio between the optimal mechanism and the best posted-price mechanism. This bound has value $2 - c$, where c is the seller's capacity. Second, I extend this result to give an upper bound on the revenue ratio between the optimal auction and the best posted-price mechanism in the (symmetric, potentially irregular) Myersonian multi-item auction. This bound is tight in the large auction limit, where it has limiting value $2 - m/n$, for an m-item, n-bidder auction. My derivations make novel use of a concavification procedure; the technique appears portable to other approximation questions in economic theory.

1 Introduction

Simplicity is a desirable feature in economic design. Simpler mechanisms and policies are easier for participants to play (Pathak and Sönmez (2008), Li (2017)), easier for designers to statistically optimize (Kitagawa and Tetenov (2018)), and easier for institutions to bureaucratically administer. Of course, simple mechanisms also carry performance costs, relative to optimal mechanisms. To paint a full picture of the tradeoffs that come with simplicity, we would like to quantify these performance costs.

In this paper, I perform such a quantification exercise, in two related settings which have attracted attention in the economic theory literature. First, I study a capacity-constrained selling problem, in which a seller sells a good to a single buyer, subject to an

*University of California, Davis.

I thank Maria Betto, Piotr Dworczak, Andrew Ferdowsian, Jason Hartline, Ellen Muir, Burkhard Schipper, Lars Stole, Takuya Ura, Mark Whitmeyer, Yu Fu Wong; and seminar audiences at UC Davis, University of Nevada Reno, and Arizona State University for helpful feedback.

ex-ante capacity constraint. This problem was suggested by Bulow and Roberts (1989), and studied more explicitly by Loertscher and Muir (2022). Loertscher and Muir (2022) show that the revenue-maximizing selling mechanism may feature two prices: a high price, which guarantees receipt of the good, and a low price with rationing.

Such a two-price mechanism is more complex than a posted-price mechanism; I therefore ask how much revenue a seller could stand to lose, by using a posted-price mechanism instead of the optimal mechanism. More precisely, I ask: what is the highest possible *revenue ratio* (over all buyer valuation distributions) between the optimal mechanism and the best posted-price mechanism? I give a sharp answer to this question: $2 - c$, where c is the seller's capacity. My argument makes novel use of a geometrical concavification procedure. I extend this argument to give a sharp bound in the case where the valuation distribution's support is known at the time the bound is assessed.

Second, I use this result to study a symmetric m-item, n-buyer Myersonian auction setting. In this setting, a posted-price mechanism offers some price, p, to all buyers; any who accept the price receive a good, with uniform rationing if more than m accept. As in the capacity-constrained selling setting, I ask: what is the highest possible revenue ratio (again over buyer valuation distributions) between the optimal auction and the best posted-price mechanism?

Since an $m+1$st-price auction with reserve price p generates at least as much revenue as a posted price of p, the answer to this question gives an upper-bound for a second interesting question: what is the highest possible revenue ratio between the optimal auction and the optimal $m + 1$st-price auction with a reserve price? Or, more succinctly, since the optimal auction is an $m + 1$st-price auction with a reserve price and ironing: how important is ironing?

The key to my analysis here is to draw a parallel between the m-item, n-buyer setting, and the capacity-constrained single-buyer setting with capacity m/n. Given this parallel, one might expect the maximum revenue ratio in the auction setting to be $2 - m/n$. I establish this result, in a limiting sense, for large auctions; for small auctions, I show that the maximum revenue ratio lies in the interval $[2 - m/n, (2 - m/n)\kappa(m, n)]$, where κ is a statistical object called the *correlation gap*, which satisfies $\kappa(m, n) \in [1, 1.6]$, for all m, n.

My results contribute to the literature in two ways. First, relative to the literature on approximation in mechanism design (see Hartline (2020)), I establish, to my knowledge, the first non-uniform upper-bound on the worst-case revenue ratio between the optimal auction and the best posted-price mechanism. Compared to the (tight) uniform upper-bound of 2 given by Chawla, Hartline, Malec, and Sivan (2010), my upper-bound is tighter for many m, n, including for all large auctions. See Section 3.3 for further discussion.

Second, the concavification argument I develop in the capacity-constrained selling setting appears to be portable to other linear optimization problems in economic theory, and could deliver similar tight approximation results there.[1] Such linear models (whose "calling card" is an optimal policy described by a small, finite menu) have proliferated in economic theory in recent years, and are particularly relevant to the inequality-aware market design literature, starting with Dworczak, Kominers, and Akbarpour (2021). In a recent article surveying that literature, Dworczak (2024) writes:

> In [a simplified version of the Dworczak et al. (2021) framework], the optimal mechanism [...] involves at most two prices [...]. One can ask: what fraction of the optimal welfare can the designer achieve by relying only on posted-price mechanisms? [...] In more complex settings, the optimal mechanism can become very complicated, and finding simpler mechanisms that achieve good welfare guarantees is of first-order importance for practical purposes.

My analysis in Section 2 takes up that call.

2 A tight bound in the single-buyer model

The first model I analyze will feature an (ex-ante) capacity-constrained seller, selling optimally to a single buyer. This model was originally suggested by Bulow and Roberts (1989), and studied in detail by Loertscher and Muir (2022).

2.1 Model

A monopolist seller, who faces an ex-ante capacity constraint, sells a good to a single buyer. The buyer's private valuation, v, is drawn from a commonly-known distribution, F, which has support lying in $V := [0, 1]$.[2] The buyer has quasi-linear utility in money, and receives utility v whenever she gets the good (otherwise 0), and utility $-t$ whenever she pays a monetary transfer of t. All agents are risk-neutral. We will consider two different selling mechanisms for the seller: the optimal mechanism, and the best posted-price mechanism.

[1]This is because this paper's arguments essentially rely only on linearity, rather than on details specific to the capacity-constrained selling problem.

[2]The substantive assumption we are making is that F has bounded support; once we have assumed this, we can choose units for money so that the support of F lies in $[0, 1]$. I assume throughout that F does not put unit mass on 0.

Optimal mechanism. The seller seeks to maximize his expected revenue, and designs an arbitrary extensive-form game (a mechanism) between himself and the buyer, each outcome of which is a (possibly negative) monetary transfer from the buyer to the seller, and a probability of the buyer getting the good. The mechanism must respect the buyer's ex-interim participation constraint, that each buyer type receive non-negative utility in expectation. Further, the mechanism must transfer the good to the buyer with ex-ante probability no greater than a given "capacity" for the seller, $c \in (0, 1)$. This capacity could represent, for example, the fraction of the total (relevant) population that a stadium or concert venue could hold.[3]

By a standard application of the revelation principle, it is without loss for the seller to use a direct mechanism, satisfying the capacity constraint, incentive compatibility (IC) and individual rationality (IR). For our purposes, since the support of F may be a strict subset of V (the unit interval), it will be more convenient to work with "semi-direct" mechanisms, which are maps from V to outcomes. In particular, a semi-direct mechanism consists of two functions, $q : V \to [0, 1]$ and $t : V \to \mathbb{R}$, which map a reported "type" (now in V) to an assignment probability and a transfer. One can prove a revelation principle for semi-direct mechanisms, and show that, as expected, there exists an optimal mechanism which is semi-direct, and solves optimization problems (1) and (3). This is a relatively standard argument; I will not make it explicitly in this version of the paper.

Having put aside the distinction between direct and semi-direct mechanisms, let us state the seller's basic optimization problem, which is otherwise standard. The seller chooses q and t to solve:

$$\max_{q,t} \int t(v) dF(v) \tag{1}$$

$$\text{s.t. } vq(v) - t(v) \geq vq(\hat{v}) - t(\hat{v}) \text{ for all } v, \hat{v} \in V \tag{IC}$$

$$vq(v) - t(v) \geq 0 \text{ for all } v \in V \tag{IR}$$

$$\int q(v) dF(v) \leq c. \tag{Cap}$$

That is, he chooses an IC and IR mechanism, which satisfies the ex-ante capacity constraint.

Next, the standard implementability result tells us that a pair of functions, (q, t), satisfies (IC) and (IR) if and only if q is non-decreasing, and t satisfies the envelope formula,

$$t(v) = vq(v) - \int_0^v q(x)dx - k, \text{ for some } k \geq 0 \text{ and for all } v \in V. \tag{2}$$

[3]The model with a single buyer and an ex-ante capacity constraint should be thought of as a stand-in for a model with a continuum of buyers and an ex-post capacity constraint (and a law of large numbers).

Following Börgers (2015)'s presentation of Manelli and Vincent (2007), let $\mathcal{F}$ be the vector space of bounded functions on $[0,1]$, with the $\mathcal{L}^1$ norm, and let $\mathcal{A} \subset \mathcal{F}$ be the subset of $\mathcal{F}$ consisting of non-decreasing functions from $[0,1] \to [0,1]$. We will call any element of $\mathcal{A}$ an **allocation function**. Since the seller should optimally set $k = 0$ in (2), the seller's problem reduces to choosing an allocation function, q, to solve:

$$\max_{q \in \mathcal{A}} \int \left[vq(v) - \int_0^v q(x)dx \right] dF(v) \tag{3}$$

$$\text{s.t.} \quad \int q(v)dF(v) \leq c. \tag{Cap}$$

Denote the seller's value from problem (3), against buyer valuation distribution F, by $rev1^*(F; c)$.

Best posted-price mechanism. Now, the seller is constrained to sell the good via a posted price. If he sets a low price, where the "demand," $1 - F_-(p)$, exceeds the capacity, c, the good is rationed so that it is sold with probability exactly c.[4] Denote the seller's value from using a posted-price mechanism with price p, against buyer valuation distribution F, by:

$$rev1_p(F; c) := p \times \min\{1 - F_-(p), c\}.$$

Why should we be interested in posted-price mechanisms specifically? The perspective this paper takes is that posted-price mechanims are *simple*, in the sense that they offer the buyer a choice between a small number of options. To be precise, we can define the class of **finite menus with opt-out** to be those mechanisms which offer the buyer a choice between a finite number of menu items (consisting of an allocation probability, q, and a transfer, t), one of which is opt-out ($q = 0, t = 0$). Let us then define the **complexity** of such a mechanism to be the number of menu items it contains, aside from opt-out. Clearly, posted-price mechanisms have complexity 1; it is also easy to show that the best posted-price mechanism is optimal within the class of complexity 1 mechanisms. We will show shortly that the optimal *mechanism* has complexity at most 2; thus, the comparison we will make between the optimal mechanism and the best posted-price mechanism is in fact a comparison between the optimal complexity-2 and the optimal complexity-1 mechanisms.

[4] We can think of the seller using the direct mechanism, (q_p, t_p), with

$$q_p(v) := \begin{cases} 0 & v < p \\ a & v \geq p \end{cases} \qquad t_p(v) := \begin{cases} 0 & v < p \\ ap & v \geq p, \end{cases}$$

where $a = \max\{1, c/(1 - F_-(p))\}$.

2.2 The concavification bound

We will be interested in how large the ratio between the revenue of the optimal mechanism and of the best posted-price mechanism can be. First, define the **revenue ratio** at some valuation distribution, F, to be:

$$R(F;c) := \frac{rev1^*(F;c)}{\max_p rev1_p(F;c)}. \tag{4}$$

Given an F, the revenue ratio measures the proportional gain from using the optimal mechanism rather than the best posted-price mechanism. Notice that the maximum in the denominator of (4) is well-defined, since $rev1_p$ is upper semi-continuous in p, and so attains a maximum on $[0,1]$.

Then, define the **concavification bound**, $B1(c)$, to be:

$$B1(c) := \sup_{F\in\Delta([0,1])} R(F;c) = \sup_{F\in\Delta([0,1])} \frac{rev1^*(F;c)}{\max_p rev1_p(F;c)}. \tag{5}$$

The concavification bound gives the worst-case (i.e., largest possible) revenue ratio over all distributions, F. It therefore measures the largest possible proportional gain from using the optimal mechanism rather than the best posted-price mechanism. Since the supremum in (5) is taken only over F, the concavification bound depends on the capacity, c. This reflects the idea that the seller knows his capacity, but is (at the time the bound is assessed) uncertain about the distribution of buyers he will face.

In this section, we will explicitly compute $B1(c)$. Our argument will proceed in two steps. First, we will compute the worst-case revenue ratio over a *restricted set* of valuation distributions – namely, those which are supported on exactly two points. Second, we will argue that this is without loss, so that the revenue ratio we have found this way is in fact $B1(c)$. To this end, define the **two-point concavification bound**, $\hat{B}1(c)$, to be:

$$\hat{B}1(c) := \sup_{F\in\Delta([0,1]),|\operatorname{supp}(F)|=2} R(F;c) = \sup_{F\in\Delta([0,1]),|\operatorname{supp}(F)|=2} \frac{rev1^*(F;c)}{\max_p rev1_p(F;c)}; \tag{6}$$

we will begin by computing $\hat{B}1(c)$.

Since we are, for the moment, focusing on two-point valuation distributions, F, it will be convenient to use a different normalization of F going forward. We will assume $F = ((1-\gamma)\circ 1, \gamma\circ v)$, for some $\gamma\in(0,1)$, $v>1$ – that is, the lower valuation is 1.[5]

[5]If the lower valuation is 0 and the higher valuation is v, then a posted-price mechanism with price v is optimal, and so $R(F;c)=1$ for all c. Since $R(F;c)$ can never fall below 1, we need not consider these

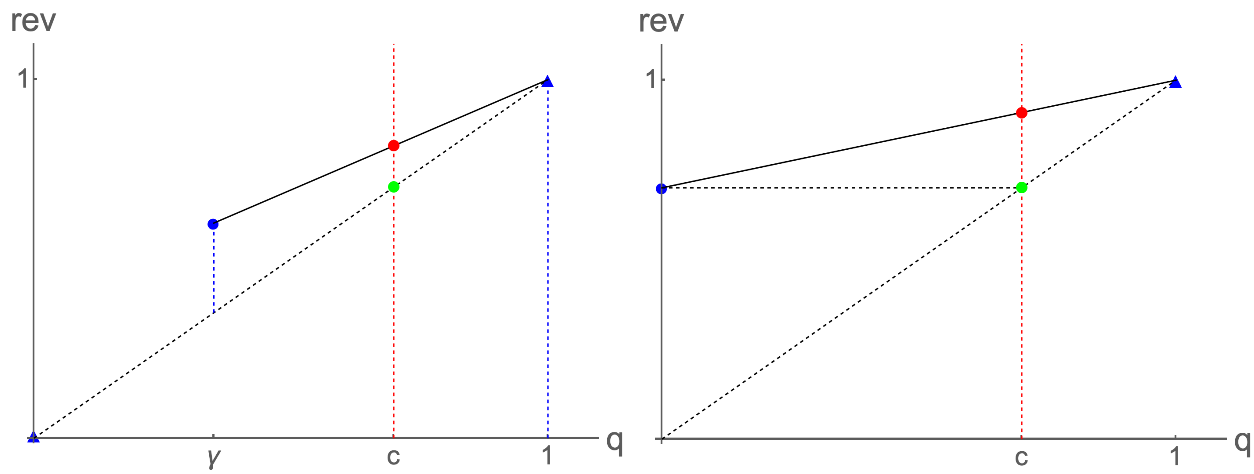

(a) The optimal mechanism, against a *generic two-point* valuation distribution.

(b) The optimal mechanism, against the *revenue-ratio maximizing* valuation distribution.

Figure 1: The concavification approach, in quantity-revenue space.

With this normalization, our problem lends itself to a geometrical analysis. Let us first find the seller's optimal mechanism, against some F of the aforementioned form. This exercise is shown in Figure 1a, and is inspired by a similar analysis in Loertscher and Muir (2022).

To begin, observe that, to assess a mechanism's performance, only its quantity (meaning ex-ante probability of trade) and revenue are relevant – quantity to assess feasibility, and revenue since it is the seller's objective. Examining problem (3), both of these objects are linear in the allocation function, q. This motivates us to use an extreme point approach.

First, following Manelli and Vincent (2007), recall that the extreme points of the set, $\mathcal{A}$, of allocation functions, are the 0-1 step functions, which correspond to posted-price mechanisms (now, without rationing). We will compute the quantity and revenue for each of these extremal allocation functions, and plot them in quantity-revenue space. These are the blue points and lines in Figure 1a. The blue circle at γ represents a posted price equal to the high valuation; the blue line at γ represents posted prices between the low and the high valuation. The blue triangle at 1 represents a posted price equal to the low valuation; the blue line at 1 represents posted prices below the low valuation. The blue triangle at 0 represents posted prices above the high valuation.

Now, *any* allocation function, q, can be expressed as a convex combination of extremal

distributions when maximizing $R(F; c)$. If F is any other two-point distribution, say $((1-\gamma)\circ v_1, \gamma\circ v_2)$, with $v_2 > v_1$, then $((1-\gamma)\circ 1, \gamma\circ v_2/v_1)$ is a two-point distribution whose lower point is 1, and which has the same revenue ratio as F.

allocation functions.[6] Since both quantity and revenue are linear in q, this means that the point representing q in quantity-revenue space is that same convex combination of the points representing its component extremal allocation functions. Thus, any IC (and hence non-decreasing) allocation function lies, in quantity-revenue space, in the convex hull of the blue points and lines. Further, any feasible IC allocation function lies to the left of the capacity constraint, c. Together, these facts select the red circle as the quantity and revenue of the optimal mechanism. Notice that this approach is reminiscent of the concavification technique used to study Bayesian persuasion, starting with Kamenica and Gentzkow (2011).

We can also visualize the revenue of the best posted-price mechanism (now with rationing) in Figure 1a. The two candidate posted prices are the high posted price (the blue circle at γ) and the low posted price, with rationing (the green circle at (c, c)). In this case, the green circle is higher, and so the revenue ratio at F is the ratio between the revenue of the red and the green circles. Finally, notice that if the arrangement of the blue points were qualitatively different than in Figure 1a (i.e., if both were to the right of c, or if the circle at γ were higher than the star at 1), then one of the posted-price mechanisms would be optimal, and so the revenue ratio would be 1.

The quantitative implications of this discussion are recorded in Lemma 1, which is proved in the appendix using an algebraic argument.

Lemma 1 (Concavification approach)**.** *Fix $c \in (0, 1)$, and suppose F is the distribution $((1-\gamma) \circ 1, \gamma \circ \frac{r}{\gamma})$, with $\frac{r}{\gamma} > 1$. Suppose further that $0 < \gamma < c$, and that $r < 1$. Then, the optimal mechanism gives revenue:*

$$rev1^*(F; c) = \alpha + (1 - \alpha)r,$$

where $\alpha := \frac{c-\gamma}{1-\gamma}$; and the best posted-price mechanism gives revenue:

$$\max_p rev1_p(F; c) = \max\{c, r\}.$$

Moreover, if $\gamma \geq c$ or $r \geq 1$, then $R(F; c) = 1$.

Understanding this concavification approach for finding the seller's optimal mechanism will guide our search for the revenue-ratio maximizing (worst-case) two-point valuation distribution. Imagine that, by choosing the valuation distribution, F, we can arbitrarily position the blue circle (the blue triangle is frozen at $(1, 1)$ due to our normalization of F).

[6] Or arbitrarily well approximated by such a convex combination – this is the Krein-Milman theorem in our setting.

Where should we place the blue circle, to maximize the ratio between the heights of the red circle and the higher of the blue and the green circles?

The answer is intuitive. First, whatever the height of the blue circle, we can increase the revenue ratio by sliding it all the way to the left. This raises the red circle, without affecting the height of the blue or the green circles. Second, if the blue circle is below the green circle, we should move it as high as the green circle: again, this raises the red circle, without affecting the height of either the blue or green circles. Finally, a bit of algebra shows that we should move the blue circle no higher than the green circle. Taking the height of the blue circle to be $r \geq c$, we have:

$$R(F;c) = \frac{c \times 1 + (1-c) \times r}{r} = \frac{c}{r} + (1-c),$$

which is evidently maximized for r as small as possible – i.e., $r = c$. The result of these optimizations is shown in Figure 1b, and results in a worst-case revenue ratio, $\hat{B}1(c) = \frac{c \times 1 + (1-c) \times c}{c} = 2 - c$. This result is recorded in Lemma 2.

Lemma 2. *The two-point concavification bound is given by* $\hat{B}1(c) = 2 - c$.

To complete our argument for Lemma 2, we should show that there is, in fact, some two-point valuation distribution (or sequence of such distributions), which achieves this desired position for the blue circle. I show in the proof of Lemma 2 that the sequence of distributions, indexed by H, given by $F_H = ((1 - \frac{1}{H}) \circ 1, \frac{1}{H} \circ cH)$, achieves this revenue ratio, as $H \to \infty$. This is intuitive if we examine Figure 1b. The blue circle corresponds to the high posted-price mechanism. It generates quantity 0 and revenue c, consistent with mass $1/H$ of buyers having valuation cH, for large H.

Next, we will show that, given *any* valuation distribution $F \in \Delta([0,1])$, and any $c \in [0,1]$, we can find some two-point $\tilde{F}$, such that $R(\tilde{F};c) \geq R(F;c)$. This will enable us to conclude that it is without loss to take the worst-case over two-point distributions, and hence that $B1(F;c) = \hat{B}1(F;c)$.

We will begin by stating a lemma which shows that, for any F and c, the seller's optimal mechanism is a finite menu with complexity no more than 2. This result is well-known, and a version of it appears as Proposition 1 in Loertscher and Muir (2022). I prove a slightly stronger version, which applies for arbitrary (non-smooth) valuation distributions, F. In particular, some additional work is required to ensure that an upper semi-continuous q is always optimal when F has mass points.

Lemma 3 (2-price mechanism)**.** *Given any* $F \in \Delta([0,1])$*, and any* $c \in (0,1)$*, there exists a semi-direct mechanism,* (q,t)*, which is optimal among all mechanisms, and such that* *either:*

$$q(v) = \begin{cases} 0 & v < v_1 \\ a & v \geq v_1, \end{cases}$$

for some $v_1 \in [0, 1]$ *and* $a \in [0, 1]$, *or:*

$$q(v) = \begin{cases} 0 & v < v_1 \\ a & v \in [v_1, v_2) \\ 1 & v \geq v_2, \end{cases}$$

for some $v_1, v_2 \in [0, 1]$ *with* $v_1 \leq v_2$, *and* $a \in [0, 1]$.

Lemma 3 is again a consequence of an extreme-point approach. The seller's problem, (3), is to choose a non-decreasing allocation function, q, to maximize a linear objective, subject to a single linear constraint. By Bauer's maximum principle, a maximum will be attained at some extreme point of the set of feasible allocation functions.

Two results are relevant. The first is the aforementioned result that the extreme points of the set of allocation functions are simply the 0-1 step functions. The second is a result by Winkler (1988), which relates the extreme points of a linearly constrained set to the extreme points of the unconstrained set. It says that the extreme points of the constrained set are convex combinations of at most $n+1$ extreme points of the unconstrained set, where n is the number of linear constraints.[7] In our case, $n = 1$, and so the extreme points of the set of *feasible* allocation functions are convex combinations of at most two 0-1 step functions; this is the form of the optimal q indicated in the Lemma.

Next, we will use Lemma 3 to prove our desired result, which allows us to consider only two-point distributions in our worst-case analysis.

Lemma 4 (2-point distribution)**.** *Fix* $c \in (0, 1)$. *Given any* $F \in \Delta([0, 1])$, *there exists a distribution* $\tilde{F} \in \Delta([0, 1])$, *such that* $|\operatorname{supp}(\tilde{F})| = 2$ *and* $R(\tilde{F}; c) \geq R(F; c)$.

The proof of Lemma 4 proceeds in three steps. First, given some $F \in \Delta([0, 1])$, we use Lemma 3 to find an optimal assignment function, q_F, which is piecewise constant, with jumps at v_1 and v_2. We then construct the distribution $\tilde{F}_0$ by shifting all the probability mass F puts above v_2 to v_2; all the mass it puts between v_1 and v_2 to v_1; and all the mass it puts below v_1 to 0. By construction, this shift does not affect the performance of q_F, and so can not hurt the optimal mechanism; however, since it is a FOSD downward

[7] Winkler's result is a consequence of the well-known Caratheodory theorem.

shift of valuations, it (weakly) hurts any posted-price mechanism. Therefore, we find that $R(\tilde{F}_0; c) \geq R(F; c)$.

Second, we form $\tilde{F}_1$ by proportionally redistributing the mass $\tilde{F}_0$ places at 0 to the other two points in its support. That is, if $\tilde{F}_0 = ((1-\alpha) \circ 0, m_1 \circ v_1, m_2 \circ v_2)$, where $\alpha := m_1 + m_2$, then $\tilde{F}_1 = (m_1/\alpha \circ v_1, m_2/\alpha \circ v_2)$. I claim that $R(\tilde{F}_0; c) = R(\tilde{F}_1; c/\alpha)$.[8] The reason is that the RHS effectively grows the relevant population, and expands the capacity, both by factor $1/\alpha$. This (1) preserves the set of feasible allocation functions and (2) scales the revenue of any feasible allocation function by $1/\alpha$. Thus, the optimal mechanism's and the best posted-price mechanism's revenues both grow by factor $1/\alpha$, leaving the revenue ratio unchanged.

Third, since $\tilde{F}_1$ is a two-point distribution, we have that $R(\tilde{F}_1; c/\alpha) \leq \hat{B}1(c/\alpha) = 2 - c/\alpha$. But, since $\hat{B}1(c) = 2 - c \geq 2 - c/\alpha$, there must be some two-point $\tilde{F}$ such that $R(\tilde{F}; c) \geq 2 - c/\alpha$. Putting these arguments together, we obtain:

$$R(F; c) \leq R(\tilde{F}_0; c) = R(\tilde{F}_1; c/\alpha) \leq 2 - c/\alpha \leq R(\tilde{F}; c),$$

as desired.

Lemma 4 allows us to consider only two-point distributions when assessing the worst-case revenue ratio; we therefore immediately conclude that $B1(c) = \hat{B}1(c) = 2 - c$, which is this section's main result.

Proposition 1. *The concavification bound is given by* $B1(c) = 2 - c$.

2.3 Support restrictions

Now, we will extend the model by supposing that, at the time the concavification bound is assessed, the seller has additional knowledge about the *support* of the valuation distribution. Such additional knowledge seems realistic. For example, if the seller represents a venue selling tickets, the venue may know that fans' ticket valuations lie between \$50 and \$500, in which case the worst-case distribution found in Section 2.2 would be impossible. Naturally, such additional knowledge restricts the worst-case distribution, and hence decreases the concavification bound. We will be able to build on the techniques introduced in Section 2.2, to again obtain a sharp value for the concavification bound in this case.

Model. The model remains unchanged, except that, now, the valuation distribution, F, must have support in $[1, \bar{v}]$, with $\bar{v} > 1$. (As before, this is a normalization; we would obtain

[8] Where the RHS is defined to be 1 for $\alpha \leq c$.

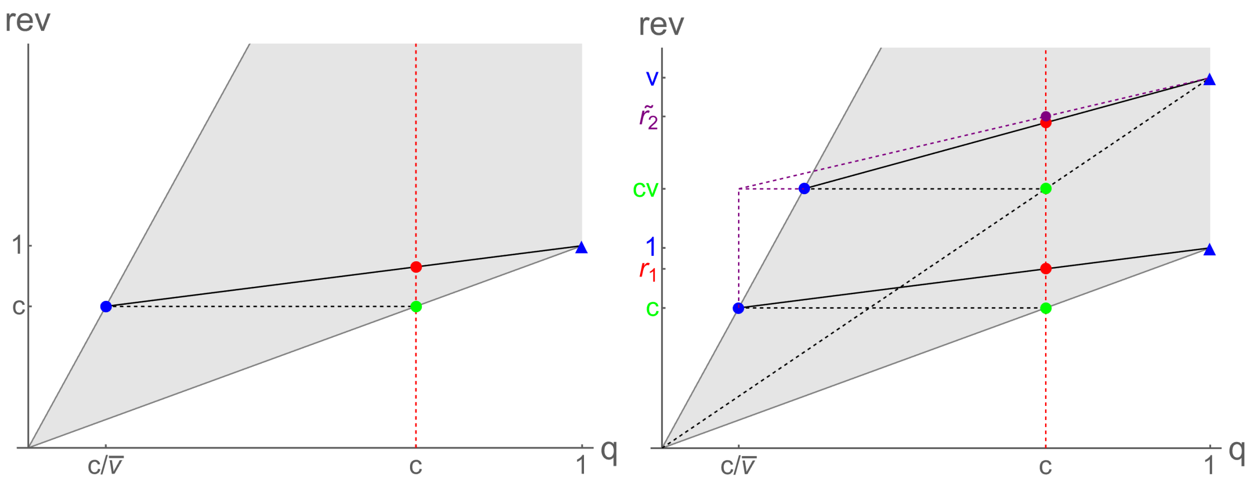

(a) The worst-case arrangement, in quantity-revenue space.

(b) A geometrical proof that any other candidate worst-case arrangement yields lower revenue ratio than the arrangement in (a).

Figure 2: The concavification approach for the restricted-support problem.

the same results by working with support $[v_1, v_2]$ as by working with support $[1, v_1/v_2]$.) We will define the **restricted-support concavification bound**, to be:

$$B1^{RS}(c, \bar{v}) := \sup_{F \in \Delta([1,\bar{v}])} \frac{rev1^*(F;c)}{\max_p rev1_p(F;c)}.$$

Analysis. We would again like to compute $B1^{RS}(c; \bar{v})$, and will take the same approach as in Section 2.2: we will first compute the worst-case revenue ratio over two-point distributions, and then argue that this was without loss. The second part of the argument will be relegated to the appendix; it is similar to the argument in Section 2.2. In the main text, we will take as given that it is without loss to consider two-point distributions.

To find the worst-case two-point distribution, as before, we would like to plot the revenue and quantity of the candidate posted-price mechanisms and of the optimal mechanism, against such a distribution. This is undertaken in Figure 2. Recall that our approach starts by plotting blue points corresponding to posted-price mechanisms, ignoring the capacity constraint. Now, however, the revenue and quantity that these blue points can generate are restricted: since any buyer must have $v \in [1, \bar{v}]$, it must be that any posted-price mechanism has $rev \in [q, q\bar{v}]$. Thus, in Figure 2, the blue points must lie in the grey cone, which consists of all points below the line of slope $\bar{v}$ through the origin, and above the line of slope 1 through the origin.

Therefore, our geometrical problem becomes: where, in the grey cone, should we place the two blue points to maximize the ratio between the height of the red circle (the optimal mechanism's revenue) and the higher of the blue and green circles (the best posted-price mechanism's revenue)? The same arguments used in Section 2.2 show that, wherever we place the blue triangle (i.e., whatever the lower buyer valuation), the blue circle should be (1) at the left edge of the grey cone and (2) the same height as the green circle. Let us call a pair of blue points satisfying these conditions a **candidate worst-case arrangement**.

Since the height of the blue triangle fully determines the candidate worst-case arrangement, it remains only to find the optimal value for this height. I claim that the blue triangle should be as low as possible, at $(1, 1)$, as shown in Figure 2a – call this pair of blue points "arrangement 1."

Figure 2b provides a geometrical proof of this claim. First, notice that the height of the green circle in arrangement 1 is c. Next, consider some other candidate worst-case arrangement, "arrangement 2," in which the blue triangle has height v. It follows that the green circle in arrangement 2 has height cv. Let r_1 and r_2 denote the heights of the red circles in arrangements 1 and 2, respectively. Construct the purple circle, with height $\tilde{r}_2$, by intersecting the line between $(c/\bar{v}, cv)$ and $(1, v)$ with the vertical line through c. We have:

$$r_1 = \alpha + (1-\alpha)c \qquad \tilde{r}_2 = \alpha v + (1-\alpha)cv,$$

where $\alpha = \frac{c-c/\bar{v}}{1-c/\bar{v}}$. Thus, $\tilde{r}_2 = r_1 \times v$. But, as Figure 2b shows, we have $r_2 < \tilde{r}_2$. Thus, we have:

$$\frac{r_2}{cv} < \frac{\tilde{r}_2}{cv} = \frac{r_1}{c},$$

so that the red-to-green ratio is higher under arrangement 1 than arrangement 2, as desired.

As before, we can read off the restricted-support concavification bound and the worst-case distribution, from Figure 2a.

Proposition 2. *The restricted-support concavification bound is given by:*

$$B1^{RS}(c, \bar{v}) = \frac{(2-c)\bar{v} - 1}{\bar{v} - c}.$$

A worst-case distribution which attains this bound is:

$$F^{WC} = \left(\frac{c}{\bar{v}} \circ \bar{v}, (1 - \frac{c}{\bar{v}}) \circ 1\right).$$

The worst-case distribution is straightforward: the low posted price gives revenue 1, while the high posted price gives probability of trade $c/\bar{v}$ and expected revenue c. This is consistent with the worst-case distribution given in Proposition 2. The expression for $B1^{RS}$

then follows algebraically. Notice that $\lim_{\bar{v}\to\infty} B1^{RS}(c,\bar{v}) = 2 - c = B1(c)$, as one would expect, since in this limit the valuation distribution becomes unrestricted.

2.4 General outperformance measures

Let us now generalize our analysis from Section 2.3 along a second dimension. We would like to be able to discuss measures, besides the revenue ratio, which capture the degree to which the optimal mechanism outperforms the best posted-price mechanism. For instance, if the seller derives a fixed monetary benefit from using a simple mechanism, the *revenue difference* between the optimal mechanism and the best posted-price mechanism may be the economically appropriate outperformance measure.[9][10]

Model. The model remains unchanged, except that the valuation distribution, F, now has support in $[\underline{v}, \bar{v}]$, with $0 \le \underline{v} < \bar{v}$. (That is, the normalization in Section 2.3 is not valid for general outperformance measures.)

Outperformance measures. Our new object of interest will be a general revenue comparison, of the following form.

Definition 1. *An* ***outperformance measure*** *is a map* $M : \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}$*, which is (weakly) increasing in the first argument, and (weakly) decreasing in the second argument.*

We will apply outperformance measures by taking the first argument to be the optimal mechanism's revenue, and the second argument to be the best posted-price mechanism's revenue.

It turns out that the key property of an outperformance measure, which will determine whether we can execute the arguments developed in Section 2.2, is its *responsiveness to scale.*

Definition 2. *An outperformance measure,* M*, is:*

- ***Increasing in multiplicative scale*** *(IMS) if* $M(\alpha r_1, \alpha r_2) \ge M(r_1, r_2)$*, for all* $\alpha \ge 1$ *and* $r_1 \ge r_2$*.*

[9] If, instead, using a simple mechanism saves the seller from paying a *variable* cost of complexity – e.g., if buyer aversion to complex mechanisms results in 5% less sales – then the revenue ratio may be appropriate.

[10] Notice that the revenue difference would not be sensible to discuss under the unrestricted-support model in Section 2.2, since there we could scale the valuation distribution to make the revenue difference arbitrarily large.

- ***Decreasing in additive scale*** *(DAS) if $M(r_1+\delta, r_2+\delta) \leq M(r_1, r_2)$, for all $\delta \geq 0$ and $r_1 \geq r_2$.*

- ***Intermediate scale responsive*** *(ISR) if it is IMS and DAS.*

Intuitively, IMS says that M is not too scale-hating, while DAS says that M is not too scale-loving. Together, these conditions constitute ISR, which will be our main notion of "appropriate" scale-responsiveness. ISR is satisfied by both the revenue ratio measure and the revenue difference measure. In a sense, the revenue ratio is the most scale-hating ISR measure (since it is only weakly IMS), while the revenue difference is the most scale-loving ISR measure (since it is only weakly DAS).

As before, we will be interested in how the outperformance varies with the valuation distribution. Let us define the **M-outperformance**, at valuation distribution F, to be:

$$R_M(F; c) := M(rev1^*(F; c), \max_p rev1_p(F; c)).$$

Finally, define the **M-worst-case outperformance** to be:

$$B1_M(c, \underline{v}, \bar{v}) := \sup_{F \in \Delta([\underline{v}, \bar{v}])} R_M(F; c).$$

Main result. We can now state this section's main result, which is developed in more detail in Appendix B.

Proposition 3. *Suppose M is an ISR outperformance measure. Then, given any valuation distribution, $F \in \Delta([\underline{v}, \bar{v}])$, there exists some $v \in [\underline{v}, \bar{v}]$, such that:*

$$R_M\left(\frac{cv}{\bar{v}} \circ \bar{v}, (1 - \frac{cv}{\bar{v}}) \circ v; c\right) \geq R_M(F; c).$$

Proposition 3 says that, as long as the outperformance measure is ISR, the worst-case distribution will yield a *candidate* worst-case arrangement, in quantity-revenue space. In other words, ISR is sufficient to make the arguments in Section 2.2, but not to make the additional argument in Section 2.3. Indeed, the latter should not be possible – one can visually see in Figure 2b that the alternative candidate worst-case arrangement we considered yields higher *revenue difference* than the arrangement in Figure 2a.

Proposition 3 follows by recreating the arguments we made in Section 2.2. Many of these involved modifying a valuation distribution to increase the optimal mechanism's revenue, without raising the best posted-price mechanism's; these modifications continue to raise M, by monotonicity. In three instances, monotonicity will not suffice for our arguments, and we must make use of M's stipulated scale responsiveness. First, in the second step of our proof of Lemma 4, we would like to redistribute the mass $\tilde{F}_0$ places on $\underline{v}$; this will weakly raise M by IMS. Second, we would like to argue that, given some 2-point distribution, $F = ((1-\alpha) \circ v_1, \alpha \circ v_2)$, we must have $R_M((1-\alpha) \circ \frac{\bar{v}}{v_2} v_1, \alpha \circ \bar{v}; c) \geq R_M(F; c)$, so that, in quantity-revenue space, the blue circle should be on the left edge of the grey cone. Again, this is implied by IMS. Finally, we would like to show that the blue circle should be no higher than the green circle; this is implied by DAS.

2.4.1 Example: The revenue difference

Given any ISR outperformance measure, Proposition 3 transforms our search for a worst-case distribution into a single-parameter optimization problem (i.e., we must only find the worst-case v). Let us demonstrate this approach by analyzing the worst-case outperfomance under the **revenue difference** measure,

$$M^{-}(r_1, r_2) := r_1 - r_2.$$

Fix some $v \in [\underline{v}, \bar{v}]$. Against the candidate binary distribution from Proposition 3,

$$F = \left(\frac{cv}{\bar{v}} \circ \bar{v}, (1 - \frac{cv}{\bar{v}}) \circ v \right),$$

the optimal mechanism generates revenue $\alpha(v)v + (1-\alpha(v))cv$, where $\alpha(v) := \frac{c\bar{v} - cv}{\bar{v} - cv}$. The best posted-price mechanism generates revenue cv. Thus, to obtain the worst-case revenue difference, we should choose v to maximize the quantity:

$$\overbrace{\alpha(v)v + (1-\alpha(v))cv}^{\text{optimal revenue}} - \overbrace{cv}^{\text{best posted-price revenue}} = \alpha(v)v(1-c);$$

this is evidently done by maximizing $\alpha(v)v$, or, equivalently, solving:

$$\max_{v \in [\underline{v}, \bar{v}]} \frac{\bar{v} - v}{\bar{v} - cv} v. \tag{7}$$

It can be shown that this objective is concave in v,[11] and hence attains a maximum at the critical point, $v^* = \frac{\bar{v}}{c}(1 - \sqrt{1-c})$, or at $\underline{v}$, if $v^* < \underline{v}$. (Note that $v^* < \bar{v}$ for $c < 1$.) We

[11] The objective's second derivative is $-\frac{2\bar{v}^2(1-c)}{(\bar{v}-cv)^3}$, which is negative for $v \in [\underline{v}, \bar{v}]$.

therefore find that the worst-case distribution is given by:

$$\left(\frac{cv}{\bar{v}} \circ \bar{v}, (1 - \frac{cv}{\bar{v}}) \circ v\right), \text{ where } v = \max\{\underline{v}, v^*\},$$

and the corresponding worst-case outperformance by:

$$B1_{M^-}(c, \underline{v}, \bar{v}) = \begin{cases} \frac{\bar{v}}{c}(\sqrt{1-c} - (1-c))^2 & v^* \geq \underline{v} \\ \frac{c(1-c)\underline{v}(\bar{v}-\underline{v})}{\bar{v}-c\underline{v}} & v^* < \underline{v}. \end{cases}$$

3 Myersonian auctions and the value of ironing

Let us now extend our results from the single-buyer model to a symmetric auction setting, with n ex-ante identical buyers and m identical items, following the seminal work of Myerson (1981). We will be interested in the maximum (over buyer valuation distributions) revenue ratio between the optimal auction and the best posted-price mechanism – in which the seller posts a price, and sells to as many buyers (up to m) as are willing to pay that price. Roughly, the idea of our analysis will be to establish a connection between the single-buyer setting with capacity constraint m/n and the auction setting.

Since the posted-price mechanism with price p has lower revenue than the $m+1$st-price auction with reserve price p, the revenue ratio we find will be an upper bound for the revenue ratio between the optimal auction and the optimal $m+1$st-price auction with reserve price.[12] One might say that this "bounds the value of ironing," since the optimal auction differs from the optimal $m+1$st-price auction with reserve in that it features ironing.

3.1 Auction model

There are n ex-ante symmetric buyers, indexed by i. Each buyer has a private valuation, v_i, which is drawn iid from a distribution, F, with support $V \subseteq [0, 1]$.[13] Define, for any $k \leq n$, the k-joint CDF as $F^{(k)} : V^k \to \mathbb{R}$, $F^{(k)}(x_1, \dots, x_k) := \prod_{i=1}^k F(x_i)$. There is a seller, who has m identical items to sell. Each buyer has unit demand: she receives payoff v_i if she gets an item and 0 otherwise. All agents are risk-neutral. We will consider two different selling mechanisms for the seller: the optimal auction, and the best posted-price mechanism.

[12] An $m+1$st-price auction with reserve price is optimal when F is Myerson regular.

[13] As in Section 2, the restriction that F's support lie in $[0, 1]$ is a normalization.

Optimal auction. The seller designs an incentive compatible (IC) and individually rational (IR) direct selling mechanism with transfers, described by an allocation function, $Q : V^n \to \Delta$, where $\Delta := \{x \in [0,1]^n : \sum x_i \le m\}$, and a transfer function, $T : V^n \to \mathbb{R}^n$. For each i, define the interim expected allocation and transfer functions, $q_i : V \to [0,1]$ and $t_i : V \to \mathbb{R}$ by:

$$\begin{aligned} q_i(v_i) &:= \int Q_i(v_i, v_{-i}) dF^{(n-1)}(v_{-i}) \\ t_i(v_i) &:= \int T_i(v_i, v_{-i}) dF^{(n-1)}(v_{-i}). \end{aligned}$$

The seller's mechanism, (Q,T), must satisfy the IC and IR conditions,

$$v_i q_i(v_i) - t_i(v_i) \ge v_i q_i(\hat{v}_i) - t_i(\hat{v}_i) \text{ for all } v_i, \hat{v}_i \in V \tag{IC}$$

$$v_i q_i(v_i) - t_i(v_i) \ge 0 \text{ for all } v_i \in V \tag{IR}$$

Such a mechanism is without loss for the seller, by the revelation principle and the agents' risk-neutrality. The seller maximizes his expected revenue; his problem is therefore given by:

$$\max_{Q,T} \sum_i \int T_i(v) dF^{(n)}(v) \qquad \text{s.t. (IC), (IR).} \tag{8}$$

Denote the seller's value from problem (8), against buyer valuation distribution F, by $rev^*(F; m, n)$.

Best posted-price mechanism. The seller now must choose some posted-price, p, and use the direct mechanism (Q^p, T^p), where:

$$\begin{aligned} Q_i^p(v) &= \gamma(v) \times \mathbf{1}\{v_i \ge p\} \\ T_i^p(v) &= \gamma(v) \times p \times \mathbf{1}\{v_i \ge p\}, \end{aligned}$$

where $\gamma(v)$ is the uniform rationing function,

$$\gamma(v) := \frac{m}{\max\{m, \sum_i \mathbf{1}\{v_i \ge p\}\}}.$$

Equivalently, the seller sets a price, p, and sells to as many buyers (up to m, with ties broken randomly) as are willing to pay p. Denote the seller's value from using posted price p, against buyer valuation distribution F, by $rev_p(F; m, n)$.

As in Section 2, we will be interested in the maximum ratio, over distributions, F, between the revenue of the optimal mechanism and of the best posted-price mechanism. Define the **ironing bound**, $B(m, n)$, to be:

$$B(m, n) := \sup_{F \in \Delta([0,1])} \frac{rev^*(F; m, n)}{\max_p rev_p(F; m, n)}. \tag{9}$$

As before, notice that the maximum in the denominator of (9) is well-defined, since rev_p is upper semi-continuous in p, and so attains a maximum on $[0, 1]$.

3.2 Bounding the value of ironing

We will now work towards establishing an upper bound on $B(m, n)$, which we will do by connecting the auction problem to an ex-ante capacity-constrained selling problem, with capacity m/n ("the ex-ante problem"). Accordingly, $B(m, n)$ will be closely related to $B1(m/n)$.

The key ingredients in our argument will be Lemma 5 and Lemma 6. Given some buyer valuation distribution, F, Lemma 5 upper-bounds the payoff from the optimal auction, in terms of the payoff from the optimal mechanism in the ex-ante problem; while Lemma 6 lower-bounds the payoff from any posted-price mechanism, in terms of the payoff from that posted-price mechanism in the ex-ante problem.

Lemma 5. *For any buyer valuation distribution, F, and any m, n, we have:*

$$rev^*(F; m, n) \leq n \times rev1^*(F; m/n).$$

Lemma 5 says that, against any F, the revenue from the optimal auction is no more than n times the revenue of the ex-ante problem. This is true because, due to the symmetry of the model, there exists some optimal auction which is ex-ante symmetric, and therefore assigns an item to each buyer with ex-ante probability no more than m/n. Since each agent's ex-interim assignment function in such a mechanism is feasible for (1) with capacity constraint $c = m/n$, it must generate lower revenue than $rev1^*(F; m/n)$. Thus, in total, the optimal auction must generate lower revenue than $n \times rev1^*(F; m/n)$.

To state Lemma 6, which lower-bounds the payoff from the posted-price mechanism, we will use a mathematical object called the correlation gap.[14] The **correlation gap**, $\kappa(m,n)$, is defined by:

$$\kappa(m,n) := \frac{m}{\sum_{i=1}^{m-1} i\binom{n}{i}\left(\frac{m}{n}\right)^i\left(1-\frac{m}{n}\right)^{n-i} + m\sum_{i=m}^{n}\binom{n}{i}\left(\frac{m}{n}\right)^i\left(1-\frac{m}{n}\right)^{n-i}}.$$

One can think of the correlation gap as the ratio of expected winnings between two casino games. In the "fair" game, we flip n coins, each of which comes up heads with probability m/n; your winnings equal the number of heads (m on expectation). In the "capped" game, we again flip n of these coins, but your winnings equal the minimum of the number of heads and m. The correlation gap is the ratio between your expected winnings in the fair and the capped game. As I show in Lemma 7, $\kappa(m,n) \in [1, \frac{e}{e-1} \approx 1.6]$, for all m,n.

Lemma 6. *For any buyer valuation distribution, F, any m,n, and any price, p, we have:*

$$rev_p(F;m,n) \geq \frac{1}{\kappa(m,n)} \times n \times rev1_p(F;m/n).$$

Given any distribution, F, and price, p, Lemma 6 gives a lower-bound on the payoff from a posted price of p in the auction problem, in terms of the payoff from that same posted price in the ex-ante problem. One might expect the former to be n times the latter: in the auction problem, we have n times the buyers (n vs 1), and n times the items (m vs capacity m/n), so that the auction problem resembles n copies of the ex-ante problem.

What this analysis misses is that the ex-ante problem is over-optimistic (relative to the auction problem) about avoiding "collisions" between buyers. Consider, for example, a 2-buyer, 1-item auction. The ex-ante problem is a single-buyer selling problem with capacity-constraint 1/2. One feasible direct mechanism in that problem is to assign the item if the buyer's valuation exceeds the median buyer valuation, and not assign it otherwise (this can be accomplished by setting a posted price equal to the median valuation). However, in the auction setting, since both buyers sometimes have valuations exceeding the median, this assignment is no longer feasible. The correlation gap captures the severity of this "over-optimism" in the ex-ante analysis.

Combining Lemmas 5 and 6 gives this section's main result, a novel upper-bound on $B(m,n)$. Proposition 4 combines this upper-bound with a known lower-bound on $B(m,n)$ (which appears as Example 1.1 of Roughgarden and Schrijvers (2016)).[15]

[14] The correlation gap I work with is a particular instance of a general class of objects called correlation gaps, which were introduced by Agrawal, Ding, Saberi, and Ye (2012). Yan (2011) exposits the application of the theory of correlation gaps to worst-case bounds in auction theory; my $\kappa(m,n)$ is $k/\Phi(k,n)$ in his notation – see his Lemmas 4.2 and 4.3.

[15] Specifically, Roughgarden and Schrijvers (2016) give the lower bound for a single-unit auction; Propo-

Additionally, Proposition 4 gives analogous bounds for the maximum revenue ratio between the optimal auction and the optimal $m+1$*st-price auction* with reserve price. Formally, let $rev_p^{SP}(F;m,n)$ denote the seller's expected revenue from using an $m+1$st-price auction with reserve price p,[16] and define the $m+1$st-price ironing bound,

$$B^{SP}(m,n) := \sup_{F\in\Delta([0,1])} \frac{rev^*(F;m,n)}{\max_p rev_p^{SP}(F;m,n)}.$$

Notice that, for any p, we have $rev_p^{SP}(F;m,n) \geq rev_p(F;m,n)$;[17] therefore, we clearly have $B^{SP}(m,n) \leq B(m,n)$. In fact, Proposition 4 gives the same upper- and lower-bounds for the two objects.

Proposition 4. *The ironing bound, $B(m,n)$, and the $m+1$st-price ironing bound, $B^{SP}(m,n)$, satisfy:*

$$B(m,n), B^{SP}(m,n) \in [2-m/n, (2-m/n)\times\kappa(m,n)].$$

In order to better understand the content of Proposition 4, we would like to have a grasp on the behavior of κ. If $\kappa(m,n)$ is very large, then we know little about the ironing bound. If $\kappa(m,n)$ is close to 1, then we know the ironing bound very precisely. Lemma 7, which is based on Lemma 4.2 in Yan (2011), delivers such an understanding of κ.

Lemma 7 (Properties of κ)**.** *The following properties of $\kappa(m,n)$ hold:*

1. *$\kappa(m,n)$ is weakly increasing in n.*
2. $\lim_{n\to\infty}\kappa(m,n) = \frac{1}{1-\frac{m^m}{e^m m!}}$. *This is strictly decreasing in m, and satisfies:*
   a. $\lim_{n\to\infty}\kappa(1,n) = \frac{e}{e-1}$.
   b. $\lim_{m\to\infty}\lim_{n\to\infty}\kappa(m,n) = 1$.

Lemma 7 implies two important facts. First, for all m and n, we have $\kappa(m,n) \leq \frac{e}{e-1} \approx 1.6$. Second, in **large auctions** (meaning any auction with large m, even if n is much larger), $\kappa(m,n)\to 1$. Thus, the result in Proposition 4 is never looser than a factor of 1.6, and is tight in large auctions. The large auction result, that $B(m,n) = 2-m/n$, should not be too surprising, since the single-agent model, where $B1(m/n) = 2-m/n$, can be thought of as the large-market limit of the auction model.

sition 4 extends that result to a multi-unit auction. Roughgarden and Schrijvers (2016) attribute the single-unit lower-bound to Hartline (2020).

[16]Where buyers play the (weakly dominant) truthful bidding equilibrium.

[17]Since the $m+1$st-price auction yields higher revenues whenever the $m+1$ highest valuations exceed p and equal revenues otherwise; recall that each buyer bids her true value in equilibrium.

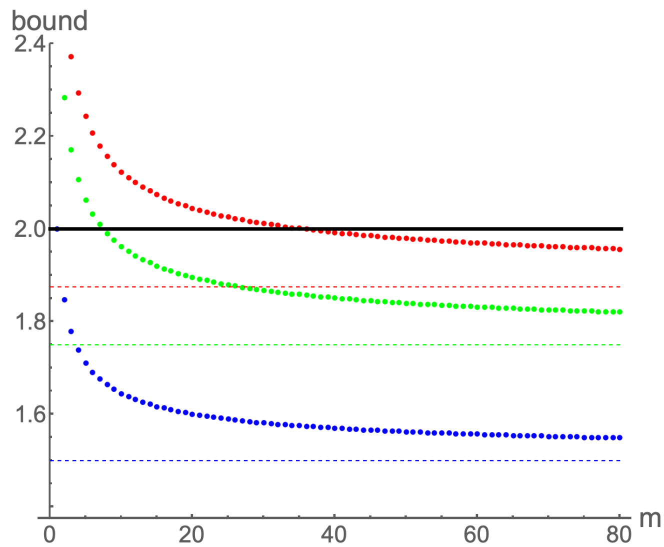

Figure 3: Upper bounds on $B(m,n)$. Black: The CHMS bound. Blue, green, red: The Proposition 4 bound for $m/n \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}\}$.

3.3 Comparison with existing upper-bounds

To my knowledge, the best existing upper-bound on $B(m,n)$ is due to Chawla, Hartline, Malec, and Sivan (2010), henceforth CHMS, who use an argument based on prophet inequalities[18] to give a (tight) uniform upper-bound of 2.

Result 1 (CHMS, Theorem 8; see also Daskalakis and Pierrakos (2011), Theorem 3)**.** *For all m, n, we have $B(m,n) \leq 2$.*

In large auctions, the current paper's result, $B(m,n) = 2 - m/n$, clearly tightens this uniform bound. However, in small auctions, the upper bound given in Proposition 4 may be looser than the CHMS bound, due to the imprecision introduced by the correlation gap factor, $\kappa(m,n)$.

I explore this relationship in Figure 3. The black line shows the CHMS bound. For an m-item, n-buyer auction, call m/n that auction's capacity. The blue, green and red series show the upper-bounds from Proposition 4, in constant-capacity auctions, as a function of m, for $m/n \in \{1/2, 1/4, 1/8\}$. (The dotted lines show these series' large-auction limits.) For capacity $1/2$ (or higher), the Proposition 4 bound is tighter than the CHMS bound for all m. For capacity $1/4$, it is tighter for $m \geq 8$; and for capacity $1/8$, it is tighter for $m \geq 36$.

[18] See Samuel-Cahn (1984) for the seminal paper, Lucier (2017) for a survey on prophet inequalities in economics, and Chapter 4 of Hartline (2020).

4 Conclusion and future work

This paper established novel upper-bounds on the cost of simplicity in two economic selling settings. First, in the capacity-constrained single-buyer selling model, with capacity c, it proved a tight upper-bound of $2 - c$, on the revenue ratio between the optimal mechanism and the best posted-price mechanism. Second, in the m-item, n-buyer Myersonian auction model, it proved an upper bound of $(2 - m/n)\kappa(m, n)$. In large auctions, this bound converges to $2 - m/n$, which is tight. This paper introduced a novel technique, based on concavification, to prove the former bound; it then used an approach based on correlation gaps to extend that bound to the auction setting.

Three avenues for future work seem particularly exciting. First, it may be possible to tighten the gap in Proposition 4. Intuitively, the sorts of distributions that make the single-buyer revenue ratio, $R(F; c)$, high, are different from the sorts of distributions that realize the full correlation gap. In fact, I conjecture that the lower bound in Proposition 4 is tight:

Conjecture 1. *The ironing bound is given by* $B(m, n) = 2 - m/n$.

It may be possible to prove this conjecture, or to tighten the gap in Proposition 4, by combining prophet inequality arguments with the concavification bound in Proposition 1.

Second, it may be possible to extend the analysis in Section 2.3 to consider settings where information on aspects of the distribution beyond its support – such as its mean or variance – is available. The work of Kang, Pernice, and Vondrák (2022) on distribution-informed bounds in the bilateral trade problem seems particularly relevant here.

Third, it would be interesting to apply the concavification arguments in Section 2 to other linear optimization problems in economic theory – particularly ones with two or more constraints. Preliminary evidence suggests that such an analysis could reveal rapidly diminishing marginal returns from mechanism complexity. In particular, I have analyzed a version of the capacity-constrained selling model, with an additional *price* constraint: no transaction may occur at a price exceeding k (compared to the maximum possible buyer valuation of 1). Proposition 1 shows that, when $c = 0.75$, the seller could gain as much as 25% by using the optimal mechanism (complexity 2) rather than a posted-price mechanism (complexity 1). In contrast, in the two-constraint setting, I find that, when $c = k = 0.75$, the seller could gain at most 2.6% by using the optimal mechanism (complexity 3) rather than the best complexity 2 mechanism.

References

Agrawal, S., Y. Ding, A. Saberi, and Y. Ye (2012). Price of Correlations in Stochastic Optimization. *Operations Research 60*(1), 150–162.

Börgers, T. (2015). *An Introduction to the Theory of Mechanism Design*. Oxford University Press, USA.

Bulow, J. and J. Roberts (1989). The Simple Economics of Optimal Auctions. *Journal of Political Economy 97*(5), 1060–1090.

Chawla, S., J. D. Hartline, D. L. Malec, and B. Sivan (2010). Multi-parameter Mechanism Design and Sequential Posted Pricing. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, pp. 311–320.

Daskalakis, C. and G. Pierrakos (2011). Simple, Optimal and Efficient Auctions. In *International Workshop on Internet and Network Economics*, Volume 7090, pp. 109–121. Springer.

Dworczak, P. (2024). Inequality and Market Design. *ACM SIGecom Exchanges 22*(1), 83–92.

Dworczak, P., S. D. Kominers, and M. Akbarpour (2021). Redistribution Through Markets. *Econometrica 89*(4), 1665–1698.

Hartline, J. D. (2020). Mechanism Design and Approximation. *Book draft. Available at https://jasonhartline.com/MDnA/.*

Kamenica, E. and M. Gentzkow (2011). Bayesian Persuasion. *American Economic Review 101*(6), 2590–2615.

Kang, Z. Y., F. Pernice, and J. Vondrák (2022). Fixed-price Approximations in Bilateral Trade. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2964–2985. SIAM.

Kitagawa, T. and A. Tetenov (2018). Who Should be Treated? Empirical Welfare Maximization Methods for Treatment Choice. *Econometrica 86*(2), 591–616.

Li, S. (2017). Obviously Strategy-proof Mechanisms. *American Economic Review 107*(11), 3257–3287.

Loertscher, S. and E. V. Muir (2022). Monopoly Pricing, Optimal Randomization, and Resale. *Journal of Political Economy 130*(3), 566–635.

Lucier, B. (2017). An Economic View of Prophet Inequalities. *ACM SIGecom Exchanges 16*(1), 24–47.

Manelli, A. M. and D. R. Vincent (2007). Multidimensional Mechanism Design: Revenue Maximization and the Multiple-good Monopoly. *Journal of Economic Theory 137*(1), 153–185.

Myerson, R. B. (1981). Optimal Auction Design. *Mathematics of Operations Research 6*(1), 58–73.

Pathak, P. A. and T. Sönmez (2008). Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *American Economic Review 98*(4), 1636–1652.

Riley, J. and R. Zeckhauser (1983). Optimal Selling Strategies: When to Haggle, When to Hold Firm. *The Quarterly Journal of Economics 98*(2), 267–289.

Roughgarden, T. and O. Schrijvers (2016). Ironing in the Dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 1–18.

Samuel-Cahn, E. (1984). Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *The Annals of Probability*, 1213–1216.

Winkler, G. (1988). Extreme Points of Moment Sets. *Mathematics of Operations Research 13*(4), 581–587.

Yan, Q. (2011). Mechanism Design via Correlation Gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 710–719. SIAM.

A Proofs and auxiliary results

Lemma 1.

Proof. By the (standard) revelation principle, since there are two buyer types, we can, without loss, focus on two-item menus $\langle(p_1, t_1), (p_2, t_2)\rangle$, where the low-valuation type selects the first menu item and the high-valuation type selects the second menu item. (Throughout, I will use angle-brace menus to denote allocation probabilities, and *per-unit probability* prices. That is, menu item (p_1, t_1) allocates the item with probability p_1 and at total price

$p_1 \times t_1$.) Then, standard zero surplus-at-the-bottom arguments imply that the menu must optimally take the form:

$$\langle (p_1, 1), (p_2, t) \rangle,$$

where $t = \frac{r}{\gamma} - \frac{p_1}{p_2}(\frac{r}{\gamma} - 1)$.

We will now argue that, optimally, $p_2 = 1$. Let q denote the allocation function induced by $\langle (p_1, 1), (p_2, t) \rangle$, and let $\tilde{q}$ denote the allocation function induced by $\langle (\tilde{p}_1, 1), (1, \tilde{t}) \rangle$, where $\tilde{p}_1$ satisfies:

$$\gamma(1 - p_2) = (1 - \gamma)(p_1 - \tilde{p}_1), \tag{10}$$

and $\tilde{t} = \frac{r}{\gamma} - \tilde{p}_1(\frac{r}{\gamma} - 1)$. Notice that $\tilde{p}_1 \geq 0$ because, by assumption, $\gamma < c$. We will use the *quantity* and *revenue* functions defined in the proof of Lemma 3. By construction, $quantity(\tilde{q}) = quantity(q)$. Further, $revenue(q) = p_1 + (p_2 - p_1)r$, and $revenue(\tilde{q}) = \tilde{p}_1 + (1 - \tilde{p}_1)r$, so that:

$$revenue(\tilde{q}) - revenue(q) = r(1 - p_2) - (1 - r)(p_1 - \tilde{p}_1). \tag{11}$$

Comparing equations (10) and (11), and recalling that $r > \gamma$, we conclude that $revenue(\tilde{q}) \geq revenue(q)$, so that $\tilde{q}$ is a feasible improvement over q.

Thus, we will consider menus, q, of the form:

$$\langle (p, 1), (1, t) \rangle,$$

where $t = \frac{r}{\gamma} - p(\frac{r}{\gamma} - 1)$. We have:

$$revenue(q) = p + (1 - p)r \qquad quantity(q) = p(1 - \gamma) + \gamma.$$

Since $r < 1$, *revenue* and *quantity* are both increasing in p, so the seller should optimally set p as high as possible – namely, $p = \alpha = \frac{c-\gamma}{1-\gamma}$. This yields the desired expression for the optimal mechanism's revenue.

Next, it is clear that the only posted prices the seller should consider are a price of 1 (with rationing), and a price of $\frac{r}{\gamma}$. These yield revenue c and r, respectively, so that the best posted-price mechanism gives revenue $\max\{c, r\}$.

Finally, we will show what happens when either of the inequalities is violated. By the high-type IR constraint, any sale must have a price no higher than $\frac{r}{\gamma}$. By the capacity constraint, sale must occur with probability no more than c. Together, these imply an upper bound on any mechanism's revenue,

$$rev1^*(F;c) \leq \frac{r}{\gamma}c.$$

If $\gamma \geq c$, then the high posted price (with rationing) achieves this upper bound.

Similarly, the revenue from the quantity-constrained problem can be no higher than the seller's revenue absent that constraint. By the classic Riley and Zeckhauser (1983) result, the seller's optimal mechanism absent the quantity constraint is a posted-price mechanism; such a mechanism generates $\max\{1, r\}$ revenue, and so we have:

$$rev1^*(F;c) \leq \max\{1, r\}.$$

If $\gamma < c$, but $r \geq 1$, then this upper bound is r, which is achieved by the (feasible) high-posted-price mechanism in the quantity-constrained problem. In both cases, $\max_p rev1_p(F;c) = rev1^*(F;c)$, so $R(F;c) = 1$. $\square$

Lemma 8. *Fix $c \in (0,1)$, and let $F := ((1-\gamma) \circ 1, \gamma \circ \frac{r}{\gamma})$, for some γ, r satisfying the conditions in Lemma 1. Then,*

$$rev1^*(F;c) \leq c + (1-c)r.$$

Proof. Define the function $\alpha_c : [0,c) \to \mathbb{R}, \alpha_c(\gamma) = \frac{c-\gamma}{1-\gamma}$.

Claim. The function α_c is continuous and non-increasing in γ on $[0,c)$.

Proof. Continuity is immediate. The function is non-increasing because α_c is differentiable on $[0,c)$ and $\alpha_c'(\gamma) = \frac{c-1}{(1-\gamma)^2}$, which is negative for all $\gamma \in [0,c)$.

By Lemma 1, $rev1^*(F;c) = \alpha + (1-\alpha)r$, where $\alpha = \frac{c-\gamma}{1-\gamma}$. Then, by *Claim*,

$$\alpha = \alpha_c(\gamma) \leq \alpha_c(0) = c. \tag{12}$$

(Notice that $\alpha_c(\gamma)$ is well-defined, since $\gamma < c$.) Since $r < 1$ by assumption, (12) gives:

$$rev1^*(F;c) = \alpha + (1-\alpha)r \leq c + (1-c)r,$$

as desired. $\square$

Lemma 2.

Proof. Upper bound: $\hat{B}1(c) \leq 2 - c$.

By Lemma 1 and the main text, it suffices to prove that $R(F;c) \leq 2-c$ for binary distributions $F = ((1-\gamma)\circ 1, \gamma \circ \frac{r}{\gamma})$, with r, γ satisfying the conditions in Lemma 1.

Case 1: $r \leq c$.

By Lemma 8,

$$rev1^*(F;c) \leq c + (1-c)r \leq c + (1-c)c, \tag{13}$$

where the second inequality follows from the case assumption. Then, by Lemma 1,

$$\max_p rev1_p(F;c) = c; \tag{14}$$

dividing inequality (13) by equation (14), we find:

$$R(F;c) = \frac{rev1^*(F;c)}{\max_p rev1_p(F;c)} \leq 2-c,$$

as desired.

Case 2: $r > c$.

By Lemma 8,

$$rev1^*(F;c) \leq c + (1-c)r \leq r + (1-c)r, \tag{15}$$

where the second inequality follows from the case assumption. Then, by Lemma 1,

$$\max_p rev1_p(F;c) = r; \tag{16}$$

dividing inequality (15) by equation (16), we find:

$$R(F;c) = \frac{rev1^*(F;c)}{\max_p rev1_p(F;c)} \leq 2-c,$$

as desired.

Lower bound: $\hat{B}1(c) \geq 2-c$.

We will propose a particular sequence of binary distributions, which achieves this lower bound. Fixing $c \in (0,1)$, let F_H be the distribution $((1-\frac{1}{H})\circ 1, \frac{1}{H}\circ cH)$, where H is a number which we will eventually make large. By Lemma 1, we have:

$$rev1^*(F_H;c) = \alpha_H + (1-\alpha_H)c, \tag{17}$$

where $\alpha_H = \frac{c-1/H}{1-1/H}$, and:

$$\max_p rev1_p(F; c) = c. \tag{18}$$

As we take $H \to \infty$, we have $\alpha_H \to c$, and so, dividing equation (17) by equation (18), we find:

$$R(F_H; c) \to 2 - c.$$

Since $R(F_H; c)$ can be arbitrarily close to $2-c$, and this holds for any c, the claim follows. $\square$

Lemma 3.

Proof. Following Chapter 2 of Börgers (2015), let $\mathcal{F}$ be the vector space of bounded functions on $[0,1]$, with the $\mathcal{L}^1$ norm, and let $\mathcal{A} \subset \mathcal{F}$ be the subset of $\mathcal{F}$ consisting of non-decreasing functions from $[0,1] \to [0,1]$. Finally, let $\overline{\mathcal{A}} \subset \mathcal{A}$ be the subset of $\mathcal{A}$ consisting of functions q such that $\int q(v)dF(v) \leq c$. We will be interested in the extreme points of $\overline{\mathcal{A}}$.

Börgers (2015) shows that the extreme points of $\mathcal{A}$ are the 0-1 step functions – that is, functions $q(v)$ where there exists some $x \in [0,1]$ such that $q(v) = 0$ for all $v < x$ and $q(v) = 1$ for all $v > x$. Proposition 2.1 of Winkler (1988) then implies that the extreme points of $\overline{\mathcal{A}}$ are contained within the set of convex combinations of two such extreme points of $\mathcal{A}$. This means that any extreme points, q, of $\overline{\mathcal{A}}$ satisfy:

$$q(v) = \begin{cases} 0 & v < v_1 \\ a & v \in (v_1, v_2) \\ 1 & v > v_2, \end{cases}$$

for some $v_1, v_2 \in [0,1]$ with $v_1 \leq v_2$. Notice that the behavior of q at v_1 and v_2 is left unspecified. By the Bauer maximum principle, the seller's problem, (3), is solved by some q which takes that form.

To fill in the gap between this result and the upper semi-continuous form for q claimed in the lemma, we will show that, if $q(v_1) < a$ or $q(v_2) < 1$, we can find a feasible $\tilde{q}$ which takes the form in the lemma, and which yields at least as much revenue as q. Define the quantities:

$$revenue(\hat{q}) := \int \left[v\hat{q}(v) - \int_0^v \hat{q}(x)dx \right] dF(v) \qquad quantity(\hat{q}) := \int \hat{q}(v)dF(v).$$

First, notice that if F does not have a mass point at the cutoff point in question, then it is trivial to "fix" q, by simply setting $\tilde{q}(v_1) = a$ or $\tilde{q}(v_2) = 1$, and leaving $\tilde{q}(v) = q(v)$ at all other v. This leaves $quantity(\tilde{q}) = quantity(q)$, so that $\tilde{q}$ remains feasible, and

$revenue(\tilde{q}) = revenue(q)$. We will therefore only show how to "fix" q when F has a mass point at v_1 or v_2 (if it has a mass point at both, we will peform both "fixes").

Suppose that $q(v_1) = e \leq a$, and $q(v_2) = f \leq 1$, and notice that the mechanism q can be described as follows: The buyer chooses an item from the five-item menu,

$$\langle (0,0), (e, v_1), (a, v_1), (f, t), (1, t) \rangle,$$

where $t = v_2 - a(v_2 - v_1)$; (indifferent) type v_1 buyers choose (e, v_1) and (indifferent) type v_2 buyers choose (f, t). The first coordinate of each menu item indicates the allocation probability, and the second indicates the per-unit-probability price (i.e., (e, v_1) means that the buyer receives the good with probability e and pays $e \times v_1$).

Possibility 1: $e < a$; $\mu_F(v_1) = m > 0$.

Our strategy will be to "merge" the menu items (e, v_1) and (a, v_1). This will not affect the mechanism's quantity, or the revenue from these low-price items. However, it will increase the price, and hence revenue, from the high-price items.

Let $\hat{m} := F_-(v_2) - F(v_1)$ be the mass of buyers who select the (a, v_1) item. Then, define $\tilde{a} := \frac{a\hat{m}+em}{\hat{m}+m}$, and let the mechanism $\tilde{q}$ be described by the alternative menu,

$$\langle (0,0), (\tilde{a}, v_1), (f, \tilde{t}), (1, \tilde{t}) \rangle,$$

where $\tilde{t} = v_2 - \tilde{a}(v_2 - v_1)$, with the tie-breaking rule where (indifferent) type v_1 buyers choose $(\tilde{a}, v_1)$ (type v_2 buyers still choose $(f, \tilde{t})$). In words, $\tilde{q}$ merges the low-price menu items, using the mean allocation probability. It then sets $\tilde{t}$ to preserve type v_2 indifference. Notice that $\tilde{q}$ is an upper semi-continuous mechanism, as indicated in the lemma statement.

By construction, $\tilde{q}$ does not change the mass of buyers who select either the low-price or the high-price menu items, and so, in particular, we have $quantity(q) = quantity(\tilde{q})$. Further, this maintains the revenue from low-price menu items. However, since $\tilde{a} < a$, we have $\tilde{t} > t$, and so $\tilde{q}$ has higher revenue from high-price menu items. Thus, $revenue(\tilde{q}) \geq revenue(q)$, as desired.

Having performed this first "fix," we are left with a menu of the form:

$$\langle (0,0), (a, v_1), (f, t), (1, t) \rangle,$$

and the tie-breaking rule that type v_1 buyers choose (a, v_1) while type v_2 buyers choose (f, t). We will again improve this mechanism in the relevant case.

Possibility 2: $f < 1$; $\mu_F(v_2) = m > 0$.

Our strategy will be to set $f = 1$, and then adjust other parts of the mechanism to make sure the quantity constraint continues to hold. As before, let $\hat{m} := F_{-}(v_2) - F_{-}(v_1)$ be the mass of buyers who select the (a, v_1) item.

Case 1: $m(1-f) \leq a\hat{m}$.

In this case, let $\tilde{q}$ be described by the alternative menu,

$$\langle (0,0), (\tilde{a}, v_1), (1, \tilde{t}) \rangle,$$

where $\tilde{a}$ satisfies $(1-f)m = (a - \tilde{a})\hat{m}$, and $\tilde{t} = v_2 - \tilde{a}(v_2 - v_1)$; and the tie-breaking rule that type v_1 buyers choose (a, v_1) and type v_2 buyers choose $(1, \tilde{t})$. (The case assumption guarantees that $\tilde{a} \geq 0$.) Notice again that $\tilde{q}$ is upper semi-continuous. By construction of $\tilde{a}$, we have that $quantity(\tilde{q}) = quantity(q)$, so that $\tilde{q}$ remains feasible. Further, it is straightforward to see that:

$$revenue(\tilde{q}) - revenue(q) = \overbrace{-(a-\tilde{a})\hat{m}v_1 + (1-f)mt}^{\text{reallocation to high-price buyers}} + \overbrace{(m + (1 - F(v_2)))(\tilde{t} - t)}^{\text{raising high price}}.$$

The "reallocation" term is ≥ 0 by the definition of $\tilde{a}$, and because $t \geq v_1$. The "raising high price" term is ≥ 0 because, as before, $\tilde{t} \geq t$. Thus, $revenue(\tilde{q}) \geq revenue(q)$.

Case 2: $m(1-f) > a\hat{m}$.

In this case, let $\hat{\tilde{q}}$ be described by the alternative menu,

$$\langle (0,0), (\hat{\tilde{f}}, v_2), (1, v_2) \rangle,$$

where $\hat{\tilde{f}}$ satisfies $(\hat{\tilde{f}} - f)m = a\hat{m}$, and the tie-breaking rule that type v_2 buyers choose $(\hat{\tilde{f}}, v_2)$. (The case assumption guarantees that $\hat{\tilde{f}} < 1$.) Notice that $\hat{\tilde{q}}$ is not upper semi-continuous. By construction of $\hat{\tilde{f}}$, we have that $quantity(\tilde{q}) = quantity(q)$, so that $\hat{\tilde{q}}$ remains feasible. Further, it is straightforward to see that:

$$revenue(\hat{\tilde{q}}) - revenue(q) = \overbrace{-a\hat{m}v_1 + (\hat{\tilde{f}} - f)mv_2}^{\text{reallocation to high-price buyers}} + \overbrace{(1 - F(v_2))(v_2 - t)}^{\text{raising high price}}.$$

The "reallocation" term is ≥ 0 by the definition of $\hat{\tilde{f}}$, and because $v_2 \geq v_1$. The "raising high price" term is ≥ 0 because $v_2 \geq t$. Thus, $revenue(\hat{\tilde{q}}) \geq revenue(q)$.

Finally, let $\tilde{q}$ be described by the menu $\langle (0,0), (\tilde{f}, v_2) \rangle$, where $\tilde{f} = \frac{\hat{\tilde{f}}m + (1-F(v_2))}{m + (1-F(v_2))}$, and the tie-breaking rule that type v_2 buyers choose $(\tilde{f}, v_2)$. By construction, $\tilde{q}$ simply merges the two top items in $\hat{\tilde{q}}$, and so generates the same quantity and revenue. Further, $\tilde{q}$ is upper semi-continuous (it takes the first form given by the lemma). □

Lemma 9. *For any $c \in (0,1)$, $F \in \Delta([0,1])$ and $\alpha \in [0,1]$, we have:*

$$R((1-\alpha)\delta_0 + \alpha F; c) = R(F; c/\alpha),$$

where δ_0 is the distribution which puts unit mass on 0, and $R(F;\hat{c})$ is defined to be 1 for any $\hat{c} \geq 1$.

Proof. Define $\hat{F} := (1-\alpha)\delta_0 + \alpha F$. We will refer to (3), under $\hat{F}$ and c, as the "$\hat{F}$ problem"; and, under F and c/α, as the "F problem."

Case 1: $c \geq \alpha$.

Notice that any allocation function is feasible for the $\hat{F}$ problem. Thus, by the Riley and Zeckhauser (1983) result, there is an optimal mechanism for the $\hat{F}$ problem which is a posted-price mechanism. We therefore have $R(\hat{F}; c) = 1$, as desired.

Case 2: $c < \alpha$.

For any allocation function, q, we have that (1) q is feasible for the $\hat{F}$ problem $\Leftrightarrow$ q is feasible for the F problem and (2) q generates α times as much revenue in the $\hat{F}$ problem as in the F problem. (Both claims follow from the linearity of *quantity* and *revenue* in F.)

It follows that any optimal mechanism in the $\hat{F}$ problem remains optimal in the F problem, and so $rev1^*(\hat{F}; c) = \alpha rev1^*(F; c/\alpha)$. It further follows that, for any price p, we have $rev1_p(\hat{F}; c) = \alpha rev1_p(F; c/\alpha)$, and so $\max_p rev1_p(\hat{F}; c) = \alpha \max_p rev1_p(F; c/\alpha)$. Thus, $R(\hat{F}; c) = R(F; c/\alpha)$. □

Lemma 4.

Proof. We will introduce a bit more notation for this proof. Let

$$rev1(q, F) := \int \left[vq(v) - \int_0^v q(x)dx \right] dF(v)$$

represent the seller's revenue from using an IC semi-direct mechanism with assignment function q, against distribution F.

Lemma 3 states that there is either a two-tiered or a three-tiered assignment function, q_F, which is optimal against F. Since a two-tiered assignment function is just a posted-price mechanism, in that case we have $R(F; c) = 1$, and so any two-point distribution, $\tilde{F}$, trivially has $R(\tilde{F}; c) \geq R(F; c)$.

In the interesting three-tiered case, Lemma 3 tells us there exists an F-optimal assignment function, q_F, which can be written as:

$$q_F(v) := \begin{cases} 0 & v < v_1 \\ a & v \in [v_1, v_2) \\ 1 & v \geq v_2, \end{cases}$$

for some a, v_1, v_2, with $v_1 \leq v_2$. Now, define $\tilde{F}_0 := (m_0 \circ 0, m_1 \circ v_1, m_2 \circ v_2)$, where $m_0 := P_{X\sim F}(X < v_1)$, $m_1 := P_{X\sim F}(v_1 \leq X < v_2)$ and $m_2 := P_{X\sim F}(X \geq v_2)$. That is, $\tilde{F}_0$ moves each unit of probability mass under F to the bottom of its q_F "bin." Notice that $rev1(q_F, \tilde{F}_0) = rev1(q_F, F) = rev1^*(F)$, and so we must have:

$$rev1^*(\tilde{F}_0) \geq rev1^*(F). \tag{19}$$

On the other hand, since $F \succeq_{FOSD} \tilde{F}_0$, any fixed posted-price assignment function performs weakly worse against $\tilde{F}_0$ than against F (weakly less buyers pay the posted price under $\tilde{F}_0$). Thus,

$$\max_p rev1_p(\tilde{F}_0) \leq \max_p rev1_p(F), \tag{20}$$

and so, dividing inequality (19) by inequality (20), we find that $R(\tilde{F}_0; c) \geq R(F; c)$.

Now, if $m_0 = 0$, the distribution $\tilde{F}_0$ is a two-point distribution, and so we are done. If not, define $\alpha := m_1 + m_2$, observe that $\alpha < 1$, and let $\tilde{F}_1 := (m_1/\alpha \circ v_1, m_2/\alpha \circ v_2)$. By Lemma 9, we have $R(\tilde{F}_0; c) = R(\tilde{F}_1; c/\alpha)$. Now, if $\alpha \leq c$, again by Lemma 9, $R(\tilde{F}_0; c) = 1$, and so $R(F; c) = 1$, and any two-point $\tilde{F}$ suffices to prove the lemma.

If instead $\alpha > c$, then since $\tilde{F}_1$ is a two-point distribution, we have, by Lemma 2, that $R(\tilde{F}_1; c/\alpha) \leq 2 - c/\alpha$. But, again using Lemma 2, we can find a two-point distribution $\tilde{F}$ with $R(\tilde{F}; c)$ arbitrarily close to $2 - c$, and, in particular, such that $R(\tilde{F}; c) \geq 2 - c/\alpha$. Putting together these observations, we obtain:

$$R(F; c) \leq R(\tilde{F}_0; c) = R(\tilde{F}_1; c/\alpha) \leq 2 - c/\alpha \leq R(\tilde{F}; c),$$

completing the proof. $\square$

Lemma 5.

Proof. Due to the symmetry of the problem, there exists an optimal mechanism, which is ex-interim symmetric,[19] and assigns the good to each agent with ex-ante probability at

[19]I.e., the interim expected allocation and transfer functions, q_i and t_i, are the same for each agent i.

most m/n. (Given some, possibly asymmetric, optimal mechanism, $\mathcal{M}$, we can construct such a mechanism, $\tilde{\mathcal{M}}$, by adding an ex-ante stage in which we choose a labeling of agents uniformly at random from all permutations of agents, and then run $\mathcal{M}$ on that labeling of agents. Clearly, $\tilde{\mathcal{M}}$ remains optimal – it is still IC and IR, and generates the same expected revenue as $\mathcal{M}$. Since $\tilde{\mathcal{M}}$ has equal ex-ante assignment probability for all agents, and always assigns at most m items to n agents, it must have ex-ante assignment probability at most m/n for each agent.)

Let (q,t) be the interim expected allocation and transfer functions of such an optimal mechanism, and observe that, by construction, (q,t) is feasible for (1), with capacity constraint $c = m/n$. Since $rev1^*(F; m/n)$ gives the optimal revenue for this problem, we have:

$$rev^*(F; m, n) = n \times \int t(v) dF(v) \leq n \times rev1^*(F; m/n),$$

as desired. □

Lemma 10. *Fix positive integers m, n with $m \leq n$, and define the function $\kappa_0 : [0,1] \to \mathbb{R}$ as:*

$$\kappa_0(\phi) := \frac{\min\{n\phi, m\}}{\sum_{i=1}^{m-1} i\binom{n}{i}\phi^i (1-\phi)^{n-i} + m\sum_{i=m}^{n}\binom{n}{i}\phi^i (1-\phi)^{n-i}}.$$

Then, κ_0 is maximized at $\phi = m/n$.

Proof. Notice that the denominator gives the expected number of 'heads' in the capped game (where n weighted coins, which come up heads with probability ϕ, are flipped, and the game is cut short once m have come up heads). Since this is clearly increasing in ϕ,[20] while the numerator is constant in ϕ for $\phi \geq m/n$, we must have $\kappa_0(m/n) \geq \kappa_0(\phi)$ for all $\phi \geq m/n$.

Suppose now that $\phi < m/n$, so that the numerator of κ_0 is simply $n\phi$. Write the denominator of κ_0 as:

$$n\phi - \sum_{i=m+1}^{n} (i-m)\binom{n}{i}\phi^i (1-\phi)^{n-i},$$

where the first term gives the expectation of the uncapped game, and the second gives the

[20]Consider $\phi, \hat{\phi}$ with $\phi < \hat{\phi}$, and choose sample spaces so that heads occurs in the ϕ game only when heads occurs in the $\hat{\phi}$ game. Then, in any outcome where k heads were realized in the ϕ game, at least k heads were realized in the $\hat{\phi}$ game.

loss due to the cap. Then, we have:

$$\frac{n}{\kappa_0(\phi)} = n - \sum_{i=m+1}^{n} (i-m) \binom{n}{i} \phi^{i-1} (1-\phi)^{n-i}$$

and so:

$$\frac{d}{d\phi} \left(\frac{n}{\kappa_0(\phi)} \right) = \sum_{i=m+1}^{n} (i-m) \binom{n}{i} \phi^{i-2} (1-\phi)^{n-i-1} \left[(n-i)\phi - (i-1)(1-\phi) \right].$$

I claim that the bracketed term is negative for all i. Notice that it is largest when i is as small as possible ($i = m+1$); in this case, its value is $(n-m-1)\phi - (m)(1-\phi) = (n-1)\phi - m < 0$, where the final inequality follows because $\phi < m/n$. Thus, $\kappa_0(\phi)$ is increasing for $\phi < m/n$, and so is maximized at $\phi = m/n$.

□

Lemma 6.

Proof. Let $\phi = 1 - F_{-}(p)$ be the probability that a buyer is willing to purchase at price p. Then:

$$rev1_p(F; m/n) = p \times \min\{\phi, m/n\}, \tag{21}$$

and:

$$rev_p(F; m, n) = p \times \left[\sum_{i=1}^{m-1} i \binom{n}{i} \phi^i (1-\phi)^{n-i} + m \sum_{i=m}^{n} \binom{n}{i} \phi^i (1-\phi)^{n-i} \right]. \tag{22}$$

Combining (21) and (22) gives:

$$\frac{rev_p(F; m, n)}{n \times rev1_p(F; m/n)} = \frac{1}{\kappa_0(\phi)} \geq \frac{1}{\kappa_0(m/n)} = \frac{1}{\kappa(m, n)},$$

where κ_0 is defined as in Lemma 10, and the second inequality is the content of Lemma 10. This proves the result. □

Proposition 4.

Proof. Since $B^{SP}(m, n) \leq B(m, n)$, it will suffice to prove the upper-bound for $B(m, n)$ and the lower-bound for $B^{SP}(m, n)$.

Upper bound: $B(m, n) \leq (2 - m/n) \times \kappa(m, n)$.

This follows immediately from Lemmas 5 and 6. Fix some distribution, $F \in \Delta([0,1])$. Taking p^* to be any maximizer of $rev1_p(F; m/n)$ (such a p^* exists since $rev1_p$ is upper semi-continuous), Lemma 6 gives us:

$$\begin{aligned} \max_p rev_p(F;m,n) &\geq rev_{p^*}(F;m,n) \\ &\geq \frac{1}{\kappa(m,n)} \times n \times rev1_{p^*}(F;m/n) \\ &= \frac{1}{\kappa(m,n)} \times n \times \max_p rev1_p(F;m/n). \end{aligned} \tag{23}$$

Then, from Lemma 5, we have:

$$rev^*(F;m,n) \leq n \times rev1^*(F;m/n). \tag{24}$$

Dividing inequality (24) by (23), we obtain:

$$\frac{rev^*(F;m,n)}{\max_p rev_p(F;m,n)} \leq \kappa(m,n) \frac{rev1^*(F;m/n)}{\max_p rev1_p(F;m/n)}. \tag{25}$$

Then, since (25) holds for any F, we also have:

$$\sup_{F\in\Delta([0,1])} \frac{rev^*(F;m,n)}{\max_p rev_p(F;m,n)} \leq \kappa(m,n) \sup_{F\in\Delta([0,1])} \frac{rev1^*(F;m/n)}{\max_p rev1_p(F;m/n)},$$

which relates $B(m,n)$ to $B1(m/n)$. Our result then follows since, by Proposition **??**, $B1(m/n) = 2 - m/n$.

Lower bound: $B^{SP}(m,n) \geq 2 - m/n$.

We will propose a particular sequence of distributions, which achieves this lower bound. Let F_H be the distribution $(\frac{1}{H} \circ m/n, (1-\frac{1}{H}) \circ 1/H)$, where H is a number which we will eventually make large.

The $m+1$st-price auction seller should only consider two candidate reserve prices against F_H: namely, m/n and $1/H$. If he sets the high reserve price, m/n, then, since the expected number of buyers who bid above the reserve price is $n \times \frac{1}{H}$, and since sale will occur at price m/n, his expected revenue is no greater than m/H.

If he sets the low reserve price, $1/H$, then all buyers will bid at least the reserve. The seller's revenue is then $m \times 1/H$, unless there are $m+1$ or more high-valuation buyers, in which case his revenue is $m \times m/n$. Thus, his expected revenue is given by:

$$m \times \frac{1}{H} + m \times \left(\frac{m}{n} - \frac{1}{H}\right) \times \sum_{i=m+1}^{n} \binom{n}{i} \left(\frac{1}{H}\right)^{i} \left(1 - \frac{1}{H}\right)^{n-i} = \frac{m}{H} + O((1/H)^{m+1}), \quad (26)$$

which is optimal, since it exceeds m/H.

Now, let us consider a different auction format, the "*l*-*h* auction," which we will use to lower-bound the revenue from the optimal auction. In this auction, all buyers will be constrained to report l (for low valuation) or h (for high valuation). All m goods will be assigned to buyers, with all h buyers served before any l buyers. (Ties within buyer reports will be broken randomly.) Any l buyer who receives a good will pay $1/H$, while any h buyer (whether or not she receives a good) will pay t, where t is chosen so that high-valuation buyers are indifferent between reporting l and h, if all other high-valuation buyers report h and all low-valuation buyers report l. By construction, it is an equilibrium for all buyers to report truthfully; the mechanism will select this equilibrium.

In the *l*-*h* auction, consider the probability, p_h, with which a buyer who reports h receives a good. Clearly, a particular h buyer receives a good whenever *all* h buyers receive a good, and so:

$$p_h \geq 1 - \sum_{i=m+1}^{n} \binom{n}{i} \left(\frac{1}{H}\right)^{i} \left(1 - \frac{1}{H}\right)^{n-i} = 1 - O((1/H)^{m+1}). \quad (27)$$

On the other hand, let p_l be the probability with which a buyer who reports l receives a good. In the best case, all buyers report l, in which case an individual buyer receives a good with probability m/n. We therefore have:

$$p_l \leq m/n. \quad (28)$$

For a high-valutaion buyer to be indifferent between reporting l and h, we must have:

$$\left(p_h \times \frac{m}{n}\right) - t = p_l \times \left(\frac{m}{n} - \frac{1}{H}\right). \quad (29)$$

Combining equation (29) with inequalities (27) and (28), we find:

$$t \geq \frac{m}{n} - \left(\frac{m}{n}\right)^2 + \frac{m}{n}\frac{1}{H} - O((1/H)^{m+1}) = \frac{m}{n} - \left(\frac{m}{n}\right)^2 + O(1/H). \quad (30)$$

Then, since the *l*-*h* auction always assigns m goods, we must have $n \times \left(\frac{1}{H} \times p_h + (1 - \frac{1}{H}) \times p_l\right) =$

m. Since $p_h \leq 1$, we find:

$$p_l \geq m/n - \frac{1}{H}. \tag{31}$$

Using inequalities (30) and (31), we find that the l-h auction generates expected revenue rev^{lh} satisfying:

$$rev^{lh} = n \times \left(\frac{1}{H} \times t + p_l \times \frac{1}{H}\right) \geq \frac{1}{H} \times n \times \left(2\frac{m}{n} - \left(\frac{m}{n}\right)^2\right) + O((1/H)^2). \tag{32}$$

Finally, since $rev^*(F_H; m, n) \geq rev^{lh}$, dividing inequality (32) by equation (26), we have:

$$\frac{rev^*(F_H; m, n)}{\max_p rev_p^{SP}(F_H; m, n)} \geq \frac{rev^{lh}}{\max_p rev_p^{SP}(F_H; m, n)} \geq 2 - \frac{m}{n} + O(1/H).$$

Taking $H \to \infty$ then establishes the result.

□

Lemma 7.

Proof. For Assertion 1, and the expression for $\lim_{n\to\infty} \kappa(m, n)$, see Yan (2011), Lemma 4.2 (b) and (c), respectively.

To see that $\lim_{n\to\infty} \kappa(m, n)$ is strictly decreasing in m, we will show that $\frac{m^m}{e^m m!}$ is strictly decreasing in m. We have:

$$\frac{(m+1)^{m+1}}{e^{m+1}(m+1)!} \Big/ \frac{m^m}{e^m m!} = \frac{1}{e}\left(1 + \frac{1}{m}\right)^m < 1,$$

where the inequality holds because $(1 + 1/m)^m$ is strictly increasing in m, with limit e as $m \to \infty$.

Finally, Assertion 2a is immediate, while Assertion 2b follows from the asymptotic approximation, $1 - \frac{m^m}{e^m m^m} \sim 1 - 1/\sqrt{2\pi m}$, given by Yan (2011) in Lemma 4.2 (c). □

B Support restrictions; general outperformance measures

In progress.