Settling the Competition Complexity of Additive Buyers over Independent Items

MAHSA DERAKHSHAN, Northeastern University, USA
EMILY RYU, Cornell University, USA
S. MATTHEW WEINBERG, Princeton University, USA
ERIC XUE, Princeton University, USA

The competition complexity of an auction setting is the number of additional bidders needed such that the simple mechanism of selling items separately (with additional bidders) achieves greater revenue than the optimal but complex (randomized, prior-dependent, Bayesian-truthful) optimal mechanism without the additional bidders. Our main result settles the competition complexity of n bidders with additive values over $m < n$ independent items at $\Theta(\sqrt{nm})$. The $O(\sqrt{nm})$ upper bound is due to [Beyhaghi and Weinberg, 2019], and our main result improves the prior lower bound of $\Omega(\ln n)$ to $\Omega(\sqrt{nm})$.

Our main result follows from an explicit construction of a Bayesian IC auction for n bidders with additive values over $m < n$ independent items drawn from the Equal Revenue curve truncated at $\sqrt{nm}$ ($\mathcal{ER}_{\leq\sqrt{nm}}$), which achieves revenue that exceeds $\text{SRev}_{n+\sqrt{nm}}\left(\mathcal{ER}^m_{\leq\sqrt{nm}}\right)$. Along the way, we show that the competition complexity of n bidders with additive values over m independent items is *exactly equal to* the minimum c such that $\text{SRev}_{n+c}(\mathcal{ER}^m_{\leq p}) \geq \text{Rev}_n(\mathcal{ER}^m_{\leq p})$ for all p (that is, *some* truncated Equal Revenue witnesses the worst-case competition complexity). Interestingly, we also show that the untruncated Equal Revenue curve does *not* witness the worst-case competition complexity when $n > m$: $\text{SRev}_n(\mathcal{ER}^m) = nm + O_m(\ln(n)) \leq \text{SRev}_{n+O_m(\ln(n))}(\mathcal{ER}^m)$, and therefore our result can only follow by considering all possible truncations.

CCS Concepts: • **Theory of computation → Algorithmic game theory and mechanism design**; **Algorithmic mechanism design**.

Additional Key Words and Phrases: Mechanism Design, Revenue Maximization, Competition Complexity, Bayesian Incentive Compatibility, Auction Design

ACM Reference Format:
Mahsa Derakhshan, Emily Ryu, S. Matthew Weinberg, and Eric Xue. 2024. Settling the Competition Complexity of Additive Buyers over Independent Items. In *The 25th ACM Conference on Economics and Computation (EC '24), July 8–11, 2024, New Haven, CT, USA.* ACM, New York, NY, USA, 27 pages. https://doi.org/10.1145/3670865.3673627

1 Introduction

Multi-dimensional mechanism design has become a core subdomain of TCS following the seminal work of Chawla, Hartline, and Kleinberg, introducing its study to Computer Science [Chawla et al., 2007]. In particular, while Myerson's seminal work in the *single*-dimensional setting elegantly

Authors' Contact Information: Mahsa Derakhshan, m.derakhshan@northeastern.edu, Northeastern University, Boston, MA, USA; Emily Ryu, eryu@cs.cornell.edu, Cornell University, Ithaca, NY, USA; S. Matthew Weinberg, smweinberg@princeton.edu, Princeton University, Princeton, NJ, USA; Eric Xue, ex3782@princeton.edu, Princeton University, Princeton, NJ, USA.

EC '24, July 8–11, 2024, New Haven, CT, USA

ACM ISBN 979-8-4007-0704-9/24/07
https://doi.org/10.1145/3670865.3673627

characterizes the optimal single-item auction in quite broad settings, Economists and Computer Scientists alike soon realized that optimal mechanisms in the multi-dimensional setting, even in restricted two-item instances, can be horribly intractable [Briest et al., 2015, Daskalakis et al., 2017, Hart and Nisan, 2013, Hart and Reny, 2015, Pavlov, 2011, Psomas et al., 2019, 2022, Rochet and Chone, 1998, Thanassoulis, 2004, Weinberg and Zhou, 2022]. In response, [Chawla et al., 2007] initiates a vast series of works establishing that simple mechanisms, while rarely optimal, achieve constant-factor approximations in quite rich settings [Babaioff et al., 2020, Cai and Zhao, 2017, Chawla et al., 2007, 2010, 2015, Chawla and Miller, 2016, Eden et al., 2021, Hart and Nisan, 2017, Li and Yao, 2013, Rubinstein and Weinberg, 2015, Yao, 2015]. These works help explain the prevalence of simple auctions in practice.

Still, constant-factor approximations do not tell the whole story – sticking with something simple that guarantees *something* is a reasonable starting point, but why not shoot for more? The Resource Augmentation paradigm offers a different perspective: running a complex auction is costly – is it perhaps more cost-effective instead to recruit extra bidders (the "resources") to participate in a simple auction? That is, prior-dependent (versus prior-independent) mechanisms are costly because you must learn the prior. Bayesian IC, BIC (versus Dominant Strategy IC, DSIC) mechanisms are costly because you must, at minimum, teach bidders the concept of Bayes-Nash equilibria (or set up auto-bidding infrastructure and convince them to trust it, etc.). Randomized mechanisms are costly because you must further ensure the risk-neutrality of your bidders. Computationally intractable mechanisms are costly simply because computation is expensive. What if recruiting extra bidders for a prior-independent, DSIC, deterministic, computationally tractable mechanism could outperform the complex optimum (without additional bidders) – might that be more cost-effective?

The mathematical question at hand, then, is to nail down *how many additional bidders are necessary for a simple auction to outperform the (intractable) optimum?* The seminal work of Bulow and Klemperer [Bulow and Klemperer, 1996] is the first to ask such a question and establish that the (prior-independent, DSIC, deterministic, computationally efficient) second-price auction with one additional bidder outperforms Myerson's (prior-dependent, DSIC, deterministic, computationally-efficient) revenue-optimal auction in single-item settings with regular[1] bidders.[2] [Roughgarden et al., 2012] are the first to ask this question in multi-dimensional settings, and [Eden et al., 2017] term the minimum number of bidders needed the *competition complexity*. Specifically, for a class C of distributions over valuation functions for a single bidder, the competition complexity $\textsc{Comp}_C(n) := \inf_{c\in\mathbb{N}_{\geq 0}}\{c \mid \text{VCG}_{n+c}(D) \geq \textsc{Rev}_n(D)\ \forall\, D \in C\}$.[3]

In the canonical domain of n additive bidders over m independent items, (the same domain studied in [Babaioff et al., 2020, Beyhaghi and Weinberg, 2019, Cai et al., 2016, Eden et al., 2017, Feldman et al., 2018, Hart and Nisan, 2017, Li and Yao, 2013, Yao, 2015]), [Eden et al., 2017] first establish a competition complexity bound of at most $n + 2(m - 1)$. That is, if $\mathcal{A}_m^R$ denotes the class of all valuation distributions which are additive across items, and each item valuation is drawn independently from a regular distribution, then $\textsc{Comp}_{\mathcal{A}_m^R}(n) \leq n + 2(m - 1)$. In other words, the VCG mechanism with $n+2(m-1)$ additional bidders outperforms the optimum (without additional bidders) for any distribution $D \in \mathcal{A}_m^R$.

[1] A single-variate distribution F is regular if the function $\varphi_F(x) := x - \frac{1-F(x)}{f(x)}$ is monotone non-decreasing.

[2] In this setting, Myerson's optimal auction is exceptionally simple: it is just a second-price auction with reserve. So [Bulow and Klemperer, 1996] essentially argues that one additional bidder removes the need for prior dependence and does not provide commentary on BIC vs. DSIC, randomized vs. deterministic, or computational tractability.

[3] C is a class of distributions such as "additive over m independent items." D is a distribution such as "the value for item j is drawn independently from $U([0, j])$." $\textsc{Rev}_n(D)$ denotes the optimal revenue of any BIC auction for n bidders drawn iid from D, and $\text{VCG}_n(D)$ denotes the expected revenue of the welfare-maximizing VCG mechanism [Clarke, 1971, Groves, 1973, Vickrey, 1961] for n bidders drawn iid from D. See Section 2 for further clarity.

In the "little n regime" ($n = O(m)$), their bound was later improved to $\text{Comp}_{\mathcal{A}_m^R}(n) = \Theta(n \ln(2 + m/n))$, which is tight (up to constant factors) [Beyhaghi and Weinberg, 2019, Feldman et al., 2018]. In the "big n regime" ($n = \Omega(m)$), [Beyhaghi and Weinberg, 2019] establish that $\text{Comp}_{\mathcal{A}_m^R}(n) \in [\Omega(\ln n), 9\sqrt{nm}]$, leaving open an exponential gap. Our main result closes the final gap in the "Big n" regime: the competition complexity is indeed $\Theta(\sqrt{nm})$. That is, when $m \geq 2$:

Main Result: $\text{Comp}_{\mathcal{A}_m^R}(n) = \Omega(\sqrt{nm})$: The competition complexity of n additive bidders over m independent items is $\Omega(\sqrt{nm})$ in the "Big n" regime. Combined with [Beyhaghi and Weinberg, 2019], this settles $\text{Comp}_{\mathcal{A}_m^R}(n) = \Theta(\sqrt{nm})$ in this regime (the "little n" regime is previously settled [Beyhaghi and Weinberg, 2019, Feldman et al., 2018]).

While our main result ultimately follows by designing a single BIC auction with high revenue, we highlight areas of technical interest briefly below. Similarly, while much of the journey towards our main result is not "necessary" for its final proof, several aspects of the journey are likely of independent interest, and we highlight these subsequently.

1.1 Main Result: Technical Highlights

Our main result ultimately follows by designing a BIC auction for n bidders whose values for m items are drawn from the Equal Revenue curve truncated at $T = \Theta(\sqrt{nm})$ ($\mathcal{ER}_{\leq T}$).[4] A priori, it is unclear what might be technically engaging about designing a BIC auction for a particular distribution. We briefly overview three technical highlights:

- When $n \gg m$ (the regime we study), selling each of m items separately to n bidders whose value for each item is drawn from $\mathcal{ER}_{\leq T}$ *already achieves expected revenue extremely close to the expected welfare*. To see this, observe that the expected welfare is clearly at most mT, and selling items separately using a second-price auction with reserve T achieves revenue $\geq mT \cdot (1 - e^{-n/T}) = mT(1 - e^{-\Omega(\sqrt{n/m})})$. This means there is little room for a more sophisticated auction to outperform selling separately without additional bidders, let alone with additional bidders.[5]
- The auction we design is not DSIC – we explicitly design an interim allocation rule together with interim payments and prove that the mechanism is implementable and BIC. To the best of our knowledge, there are not many prior instances of useful explicit designs of BIC-but-not-DSIC mechanisms – the only notable example is that of [Yao, 2017] for two bidders whose values for each of two items are drawn independently from the uniform distribution on $\{1, 2\}$. We design such an auction for any $m \geq 2$ and $n \geq m$.
- The method we use to design our auction is likely of use to future designs of BIC-but-not-DSIC auctions. We start by picking an allocation rule we would like to implement and prices we would like to charge that result in a clean analysis. Unfortunately, the prices we'd like to charge are not BIC, but the interim allocation rule (of our desired allocation rule) is well-structured, so the prices can be massaged to get fairly close to BIC. Further unfortunately, there does not appear to be an exactly-BIC implementation of this allocation rule at all, so we further slightly massage the allocation rule. This last step, in particular, is reminiscent of a specialized (for our mechanism and distribution) instantiation of an ε-BIC to BIC reduction [Bei and Huang, 2011, Cai et al., 2021, Daskalakis and Weinberg, 2012, Dughmi et al., 2017, Hartline et al., 2011, Rubinstein and Weinberg, 2015].

[4] $\mathcal{ER}_{\leq T}$ has CDF $F(x) = 1 - 1/x$ for $x \in [1, T)$, $F(T) = 1$, and $F(x) = 0$ for $x < 1$.

[5] But, it does also mean that we may not need to outperform selling separately by much in order to also outperform selling separately with additional bidders, because the additional bidders cannot possibly help much either.

1.2 Results of Independent Interest along the Journey

The journey towards our main result yields two additional results – while these are ultimately "unnecessary" steps along the journey, they help provide context for our results and approach.

Independent Result I: Let $\mathcal{ERC}_m$ denote the subclass of $\mathcal{A}_m^R$ containing only distributions of the form $\mathcal{ER}_{\leq T}^m$ for some truncation T. Then $\text{Comp}_{\mathcal{A}_m^R}(n) = \text{Comp}_{\mathcal{ERC}_m}(n)$ for all n. That is, for all n, the worst-case competition complexity of any distribution in $\mathcal{A}_m^R$ is witnessed by iid truncated Equal Revenue curves.

One direction of this equality is trivial, as $\mathcal{ERC}_m \subset \mathcal{A}_m^R$. The non-trivial direction does not at all follow from identifying an a priori worst-case distribution (indeed, the worst-case distribution for $\text{Comp}_{\mathcal{A}_m^R}(n)$ is $\mathcal{ER}_{\leq\Theta(\sqrt{mn})}$, which has no apparent a priori justification). It is natural to guess that an untruncated equal revenue curve may be the worst-case distribution, as it stochastically dominates all other distributions with the same single-bidder revenue. However, this intuition breaks rather quickly: (a) just because every marginal of D stochastically dominates those of D' does *not* imply that the optimal revenue for D exceeds that of D' due to the phenomenon of revenue non-monotonicity [Hart and Reny, 2015, Rubinstein and Weinberg, 2015, Yao, 2018], (b) even if moving from D' to D were guaranteed to make the revenue benchmark larger, it also improves the revenue of selling separately, so both sides of the desired inequality increase. Instead, we show that a modification of [Beyhaghi and Weinberg, 2019]'s approach upper bounds $\text{Comp}_{\mathcal{A}_m^R}(n)$ *if and only if it upper bounds* $\text{Comp}_{\mathcal{ERC}_m}(n)$. See Section 5 for further details.

In addition, Independent Result I provides further context for our main result. A practically-minded reader might wonder why it matters that $\Omega(\sqrt{mn})$ bidders are necessary for VCG to outperform the optimum in an instance where selling separately is already extremely close to optimal.[6] Independent Result I highlights that analyzing this instance is in some sense a necessary step to analyze instances where the gap might be larger.

Independent Result II: For some absolute constant c, $\text{VCG}_{n+cm\ln(n)}(\mathcal{ER}^m) \geq \text{Rev}_n(\mathcal{ER}^m)$. That is, $O(m\ln(n))$ additional bidders suffice for selling separately n bidders with values for m items drawn from iid untruncated Equal Revenue curves to exceed the optimal revenue. This implies that the untruncated Equal Revenue curve is *not* the worst-case instance for *any* $n \geq m$ (and it witnesses a lower bound that is exponentially suboptimal in n). This result follows by establishing that the optimal revenue for n bidders with values for each of m items drawn iid from the Equal Revenue curve is $nm + O(m^2\ln(n))$.

Previous disclaimers that an untruncated revenue curve is not *obviously* the worst-case distribution notwithstanding, it is still a tempting conjecture that the equal revenue curve may indeed be the worst-case (or at least, asymptotically close) – if one had hoped to improve [Beyhaghi and Weinberg, 2019]'s bounds, iid untruncated equal revenue curves is a natural first step. So it is interesting that the untruncated Equal Revenue curve witnesses an exponentially-suboptimal bound compared to a properly-truncated Equal Revenue curve. This provides further motivation for Independent Result I, as it highlights that there is indeed no a priori worst-case distribution. See Section 6 for further details.

Beyond the result itself, our analysis should be of independent interest. In particular, a first step towards our upper bound on the revenue is a flow in the [Cai et al., 2016] framework. To the best of our knowledge, prior works that approximate the optimal revenue all use a "region-separated" flow (see Section 6.4 for a formal definition) [Beyhaghi and Weinberg, 2019, Cai et al., 2016, 2022,

[6] See Section 1.3 for a very brief note on results such as [Cai and Saxena, 2021, Feldman et al., 2018] that explicitly consider resource augmentation to target a $(1-\varepsilon)$-approximation rather than truly exceeding the optimum.

Cai and Zhao, 2017, Eden et al., 2017, 2021].[7] We prove that such a flow cannot possibly witness an upper bound better than $nm + \Omega(m\sqrt{nm})$ on $\textsc{Rev}_n(\mathcal{ER}^m)$ by designing an auction that satisfies all "within-region" BIC constraints (but not the cross-region constraints). To cope with this, our analysis still begins from a region-separated flow (in fact, the same canonical flow used in [Eden et al., 2017]), but adds a novel second step to (necessarily) leverage cross-region constraints.

Finally, in Section 6.3 we also establish that bundling items together (to n bidders with m items from the Equal Revenue curve) with a second-price auction achieves expected revenue $nm + \Omega(m \ln(n))$. Our analysis follows primarily from a coupling argument as opposed to raw calculations, and also slightly improves the analysis of [Beyhaghi and Weinberg, 2019] from $nm + \Omega(\ln(n))$.

1.3 Related Work

We have already overviewed the most directly related work. [Bulow and Klemperer, 1996] first consider resource augmentation for Bayesian mechanism design, and establish that a single additional bidder suffices for the second-price auction with no reserve to outperform the revenue-optimal auction with any number of i.i.d. regular bidders and a single item. [Roughgarden et al., 2012] first consider resource augmentation for multi-dimensional mechanism design, and compare the VCG mechanism with additional bidders to the optimal deterministic DSIC auction for unit-demand bidders over independent items. [Eden et al., 2017] are the first to target outperformance of the optimal BIC randomized auction, and study the now-canonical setting of additive bidders over independent items. Their bounds have since been tightened by [Beyhaghi and Weinberg, 2019, Feldman et al., 2018], and our main result tightens the last remaining gap. Moreover, if $\mathcal{A}^R_{m,\mathcal{I}}$ denotes the class of valuation functions that are additive over m independent regular items subject to downwards-closed constraints $\mathcal{I}$, [Eden et al., 2017] also establish that $\textsc{Comp}_{\mathcal{A}^R_{m,\mathcal{I}}}(n) \leq \textsc{Comp}_{\mathcal{A}^R_m}(n) + m - 1$. Therefore, the study of additive buyers has implications for significantly more general settings as well. Other works, such as [Brustle et al., 2022, Fu et al., 2019, Liu and Psomas, 2018] consider the competition complexity of Bayesian mechanism design in other settings (such as dynamic auctions, non-iid single-dimensional bidders, or posted-price mechanisms). [Cai and Saxena, 2021, Feldman et al., 2018] further consider how many additional bidders are needed to recover a $(1 - \varepsilon)$-fraction of the optimal revenue, rather than truly exceeding the optimal revenue. Both results require strictly fewer bidders than would otherwise be necessary.

The concept of resource augmentation is well-represented within TCS broadly [Barman et al., 2012, Sleator and Tarjan, 1985], Economics broadly [Akbarpour et al., 2022, 2018], and also their intersection [Chawla et al., 2013, Roughgarden and Tardos, 2002].

We have also previously noted a vast literature justifying simple auctions in multi-dimensional settings, despite their suboptimality [Babaioff et al., 2020, Cai and Zhao, 2017, Chawla et al., 2007, 2010, 2015, Chawla and Miller, 2016, Eden et al., 2021, Hart and Nisan, 2017, Li and Yao, 2013, Rubinstein and Weinberg, 2015, Yao, 2015].[8] Our independent results use similar technical tools (such as the benchmark induced by [Cai et al., 2016]'s "canonical flow"), but deviates from these in seeking a $(1 - o(1))$-approximation to the optimal revenue, rather than a constant-factor approximation.

At a technical level, our results are similar-to-yet-distinct-from several themes in the literature on multi-dimensional mechanism design. As previously noted, we design an explicit BIC auction

[7] There are certainly works, such as [Daskalakis et al., 2017, Haghpanah and Hartline, 2015] that use more complex flows to derive *optimal* mechanisms in single-bidder settings.

[8] Note that [Hartline and Roughgarden, 2009] initiate a conceptually-similar line of work justifying exceptionally simple auctions in single-dimensional settings via constant-factor approximation guarantees.

that is not DSIC, which is also done in [Yao, 2017]. A difference is that [Yao, 2017] considers two bidders and two items and proves that the BIC auction strictly outperforms the optimal DSIC auction, whereas we consider arbitrarily-many bidders and items but do not explicitly compare to a DSIC auction. We have also mentioned that one step of our auction design bears similarity to ε-BIC to BIC reductions, which are developed in [Cai et al., 2021, Daskalakis and Weinberg, 2012, Rubinstein and Weinberg, 2015] based on techniques introduced in [Bei and Huang, 2011, Dughmi et al., 2017, Hartline et al., 2011]. Their results apply generally and are technically quite involved, whereas we directly massage a specific nearly-BIC auction for a specific distribution. There is also a line of works deriving *optimal* mechanisms for specific distributions [Daskalakis et al., 2017, Giannakopoulos and Koutsoupias, 2014, 2015, Haghpanah and Hartline, 2015]. These works consider single bidder settings, and most use some form of duality to establish optimality. Other works derive optimal mechanisms for simple classes of single-bidder distributions to establish computational hardness [Chen et al., 2022, Daskalakis et al., 2014]. In comparison, our work considers multi-bidder settings, and in some sense lies between these works and constant-factor approximations in terms of complexity: our upper bounds are slightly more involved than those sufficient for constant-factor approximations, but not as involved as those necessary for precise optimality. At the same time, we do not nail precisely the revenue-optimal auctions, but do derive bounds strictly better than what can be achieved by the simple duals sufficient for constant-factors.

2 Preliminaries

In this section, we provide the minimal preliminaries necessary to state and prove our main result. Section 4 provides additional preliminaries specific to our independent results.

The setting we study consists of n bidders with additive valuations over m items. Formally, the values of the bidders are drawn from an $n \times m$ dimensional joint distribution $\mathcal{D}$ where v_{ij} denotes bidder i's value for item j. Bidder i's value for a subset S of items is $\sum_{j \in S} v_{ij}$.

A mechanism is given by ex-post allocation and payment rules that specify the probabilities with which each bidder gets each item and how much each bidder pays for each item, respectively. We will also consider interim allocation rules of auctions, which suffice to understand whether mechanisms are Bayesian IC (see Definition 2.2 below).

Definition 2.1. The ex-post probabilities with which each bidder receives each item are given by a function $x : \text{supp}(\mathcal{D}) \to \Delta^{n \times m}$ where $x_{ij}(v_1, \ldots, v_n)$ denotes the probability with which bidder i receives item j given the bid profile v, and the ex-post payments $q : \text{supp}(\mathcal{D}) \to \Delta^{n \times m}$ has $q_i(v)$ denote the payment that bidder i makes given the bid profile v. Given an ex-post allocation rule x and price rule q, the interim probability with which bidder i receives item j when she bids v_i and the interim price paid are defined as $\pi_{ij}(v_i) := \mathbb{E}_{v \sim \mathcal{D}}\left[x_{ij}(v) \,\middle|\, v_i\right], p_i(v_i) := \mathbb{E}_{v \sim \mathcal{D}}\left[q_i(v) \mid v_i\right]$. That is, the interim probability is the expected probability with which bidder i receives item j when she bids v_i and the remaining bidders bid truthfully, and the interim price is the expected price when bidding v_i and the remaining bidders bid truthfully.

Definition 2.2. Let π denote an interim allocation rule and let p denote an interim payment rule. The mechanism (π, p) is Bayesian Incentive Compatible (BIC) if for all i, v_i, v_i', $\sum_j \pi_{ij}(v_i)v_{ij} - p_i(v_i) \geq \sum_j \pi_{ij}(v_i')v_{ij} - p_i(v_i')$. That is, each bidder's best response to her peers if they report their true values is also to report her true values.

In addition, we use the following terminology:

- $\mathcal{ER}$: the single-variate distribution with CDF $F(x) = 1 - \frac{1}{x}$, for $x \geq 1$.
- $\mathcal{ER}^m$: the multi-variate distribution that draws m values i.i.d. from $\mathcal{ER}$.
- $\mathcal{ER}^{n \times m}$: the multi-variate distribution drawing n bidders' values for m items i.i.d. from $\mathcal{ER}$.

- $\mathcal{ER}_{\leq T}$: the equal revenue distribution truncated at T; i.e. the single-variate distribution with CDF $F(x) = 1 - \frac{1}{x}$ for $x \in [1, T)$ and $F(T) = 1$.
- $\textsc{Rev}^M(\mathcal{D})$: the expected revenue of an auction M when played by bidders drawn from the joint distribution $\mathcal{D}$ over values of n bidders for m items.
- $\textsc{Rev}(\mathcal{D})$: the supremum over all BIC auctions M of $\textsc{Rev}^M(\mathcal{D})$.
- $\textsc{SRev}(\mathcal{D})$: the expected revenue of selling separately (using Myerson's optimal auction [Myerson, 1981]) to bidders drawn from the joint distribution $\mathcal{D}$ over values of n bidders for m items.
- $\text{VCG}(\mathcal{D})$: the expected revenue of the welfare-maximizing Vickrey-Clarke-Groves (VCG) auction when played by bidders drawn from the joint distribution $\mathcal{D}$ over values of n bidders for m items.
- All $\mathcal{D}$ considered in this paper are i.i.d. across bidders, and of the form D^n for some distribution over a single bidder's valuation function for m items. To simplify notation throughout, in these cases we notate $\textsc{Rev}_n(D) := \textsc{Rev}(\mathcal{D}), \textsc{SRev}_n(D) := \textsc{SRev}(\mathcal{D}), \text{VCG}_n(D) := \text{VCG}(\mathcal{D})$.

Below is a formal (re-)statement of the competition complexity. Our work considers additive bidders, where the VCG auction sells items separately using a second-price auction (so the distinction between VCG and selling separately is simply whether or not there is a reserve, when item values are regular, or whether items are sold using a second-price auction vs. Myerson's optimal auction in the general case).

Definition 2.3 (Competition Complexity)**.** Let C be a class of distributions over valuation functions for a single bidder. The *Competition Complexity* of C is the function $\textsc{Comp}_C(\cdot) : \mathbb{N}_+ \to \mathbb{N}_{\geq 0}$ where $\textsc{Comp}_C(n) := \inf_{c \in \mathbb{N}_{\geq 0}} \{c \mid \text{VCG}_{n+c}(D) \geq \textsc{Rev}_n(D)\ \forall\ D \in C\}$. The *Selling Separately Competition Complexity* is instead $\textsc{SSComp}_C(n) := \inf_{c \in \mathbb{N}_{\geq 0}} \{c \mid \textsc{SRev}_{n+c}(D) \geq \textsc{Rev}_n(D)\ \forall\ D \in C\}$.

3 Main Result: $\textsc{Comp}_{\mathcal{A}_m^R}(n) = \Omega(\sqrt{nm})$

For the class of truncated equal revenue distributions $\mathcal{ER}_{\leq T}$ with $T = \lambda\sqrt{nm}$ for some absolute constant $\lambda > 1$ (and $T < n$), we now provide an explicit construction of a BIC auction M with $\textsc{Rev}^M(\mathcal{ER}_{\leq T}^{n \times m}) > \textsc{SRev}_{n+c\sqrt{nm}}(\mathcal{ER}_{\leq T}^m)$ for some absolute constant c. This witnesses that $\textsc{Comp}_{\mathcal{A}_m^R}(n) = \Omega(\sqrt{nm})$, and $\textsc{SSComp}_{\mathcal{A}_m}(n) = \Omega(\sqrt{nm})$. For ease of readability, several calculation-based proofs are deferred to Appendix A.1.

3.1 Step One: Intuition & a Not-at-all BIC Auction

First, we explicitly compute $\textsc{SRev}_{n'}(\mathcal{ER}^m)$.

Lemma 3.1. *$\textsc{SRev}_{n'}(\mathcal{ER}_{\leq T}^m) = m \cdot T \cdot \left(1 - (1 - 1/T)^{n'}\right)$. One mechanism achieving this sells each item separately with a second-price auction at reserve T.*[9]

Note that selling separately is already nearly optimal when $T \ll n$. In particular, $\textsc{Rev}_n(\mathcal{ER}_{\leq T}^m) \leq mT$, and $\textsc{SRev}_n(\mathcal{ER}_{\leq T}^m) \approx mT$ when $T \ll n$. Moreover, for each item j, selling separately achieves the maximum value whenever it is T, and so the only possible room for improvement over selling separately is in the exponentially-unlikely cases that all bidders have value $< T$ for item j (where selling separately gets 0 revenue from item j, yet there is strictly positive value).

So, in order to possibly have an auction whose revenue exceeds $\textsc{SRev}_n(\mathcal{ER}_{\leq T}^m)$ (let alone $\textsc{SRev}_{n+c\sqrt{nm}}(\mathcal{ER}_{\leq T}^m)$), we must somehow get nonzero revenue from item j in cases when $v_{ij} < T$ for all i. One naive way to accomplish this is simply to remove the reserve, and sell each item with a second-price auction instead. Of course, this is still selling separately and thus offers no

[9] In fact, a second-price auction with any reserve $\leq T$ achieves this.

improvement over $\text{SRev}_n(\mathcal{ER}^m_{\leq T})$. But, it highlights the tradeoff that any BIC mechanism must face: allocating item j to a bidder i with $v_{ij} < T$ provides incentive for bidder i to misreport that $v_{ij} < T$ when in fact $v_{ij} = T$, and therefore risks revenue $< T$ in cases where selling separately achieves T.

So, the first idea in designing our BIC auction is to find opportunities to allocate item j to a bidder i with $v_{ij} < T$ without risking too much in cases where $v_{ij} = T$ instead. Below is an allocation rule that accomplishes this first step, but is not yet BIC. In particular, we only ever consider allocating an item j to a bidder i with $v_{ij} < T$ *if* $v_{ij'} = T$ *for at least one other item* j'.

Definition 3.2 (The Naive Auction)**.** The Naive Auction allocates each item j separately as follows.

(1) If there exists a bidder i with $v_{ij} = T$, then allocate item j uniformly at random to such a bidder and charge a price of T.
(2) If $v_{ij} < T$ for all i, but there exists a bidder i with both $v_{ij'} = T$ for some $j' \neq j$ and $v_{ij} \geq mn/T$, then allocate item j uniformly at random to such a bidder and charge a price of mn/T.
(3) Otherwise, do not allocate or elicit payments for item j.

The Naive Auction is certainly not BIC: a bidder whose values are T for every item achieves utility of 0 for reporting the truth, but > 0 for instead lowering one value to mn/T. Still, it clearly achieves revenue greater than $\text{SRev}_n(\mathcal{ER}^m_{\leq T})$. We first establish that the revenue of the Naive Auction further exceeds $\text{SRev}_{n+c\sqrt{nm}}(\mathcal{ER}^m_{\leq T})$ – the remainder of this section is then devoted to massaging the Naive Auction into a BIC auction without losing much of this additional revenue.

Lemma 3.3. *$\text{SRev}_{n+x}(\mathcal{ER}^m_{\leq T}) \leq \text{SRev}_n(\mathcal{ER}^m_{\leq T}) + mx(1-1/T)^n$.*

Proof. Couple draws from $\mathcal{ER}^{m\times(n+x)}_{\leq T}$ and $\mathcal{ER}^{m\times n}_{\leq T}$ so that the first n bidders' values are identical. For each item j, selling separately to $n+x$ bidders outperforms selling separately to n bidders (by exactly T) iff $v_{ij} < T$ for all $i \in [n]$ and $v_{ij} = T$ for some $i > n$. Therefore, the additional revenue is exactly: $m \cdot T \cdot (1-1/T)^n \cdot (1-(1-1/T)^x) \leq mT(1-1/T)^n \cdot x/T = mx(1-1/T)^n$. □

Lemma 3.4. *The Naive Auction satisfies $\text{Rev}^{NA}_n(\mathcal{ER}^m_{\leq T}) \geq \text{SRev}_n(\mathcal{ER}^m_{\leq T}) + \Omega(m\sqrt{mn}(1-1/T)^n)$.*[10]

Proof. For each item j, the Naive Auction achieves revenue T whenever selling separately achieves revenue T (whenever some bidder i has $v_{ij} = T$). The Naive Auction achieves additional revenue in cases where no bidder has value T.

For a fixed item j, the probability that all n bidders have value $< T$ is $(1-1/T)^n$. Conditioned on this, we want to find the probability that some bidder both has value at least mn/T for item j, and also T for some other item. These are independent events across both bidders and items.

For a fixed bidder i, the probability that $v_{ij} \geq mn/T$ conditioned on $v_{ij} < T$ is exactly $\frac{T/mn-1/T}{1-1/T} = \frac{\lambda/\sqrt{mn}-1/(\lambda\sqrt{mn})}{1-1/(\lambda\sqrt{mn})} = \Omega(1/\sqrt{mn}) = \Omega(1/T)$ (as $\lambda > 1$ is an absolute constant). The probability that a fixed bidder i has value T for some item $\neq j$ is simply $1-(1-1/T)^{m-1} = \Omega(m/T)$.[11] Therefore, the probability that a fixed bidder i has $v_{ij} > mn/T$ *and* $v_{ij'} = T$ for some $j' \neq j$, conditioned on $v_{ij} < T$ is $\Omega(m/T^2)$.

Finally, the probability that at least one bidder i has both $v_{ij} \geq mn/T$ and $v_{ij'} = T$ for some $j' \neq j$ is $\Omega(nm/T^2) = \Omega(1)$.[12] This means that the additional revenue per item gained by the Naive Auction over selling separately is $\Omega((1-1/T)^n \cdot (mn/T)) = \Omega((1-1/T)^n \cdot \sqrt{mn})$, and multiplying by m items establishes the result. □

Corollary 3.5. *There exists an absolute constant c' such that $\text{Rev}^{NA}_n(\mathcal{ER}^m_{\leq T}) \geq \text{SRev}_{n+c'\sqrt{mn}}(\mathcal{ER}^m_{\leq T})$.*

[10] Recall that the Naive Auction is not BIC – this analysis is just to supply intuition for our later (more involved) computations.
[11] This follows as $1-(1-1/x)^y = \Omega(y/x)$ when $y < x$, and that $m < T$.
[12] This again follows from the fact that $1-(1-1/x)^y = \Omega(y/x)$ when $y < x$, and that $n < T^2/m$.

Corollary 3.5 establishes that our (not at all BIC) Naive Auction would witness the desired lower bound on $\text{SSComp}_{\mathcal{A}_m}$, if only it were BIC. One obvious problem with the Naive Auction is that it extracts full welfare from buyers with value T for all items, and yet awards items with non-zero probability to buyers with lower values. Our next step is to adjust the payments to address this specific issue (this will still not result in a BIC auction, but it is the first of two steps).

3.2 Step Two: A Closer-to-BIC Auction

Our next step is to address the obvious issue with the Naive Auction by keeping the same allocation rule with less problematic payments. This will still not yet result in a BIC auction, but will get close to the correct format.

Definition 3.6 (The Less-Naive Auction)**.** The Less-Naive Auction allocates each item separately as follows.

(1) Use the same *allocation rule* as the Naive Auction. Let a_0 denote the interim allocation probability of winning item j conditioned on reporting $v_{ij} = T$, and b_0 denote the interim allocation probability of winning item j conditioned on reporting $v_{ij} \in [mn/T, T)$ and $v_{ij'} = T$ for some $j' \neq j$.
(2) If the bidder i receiving item j has $v_{ij} < T$, charge mn/T (as in the Naive Auction).
(3) If the bidder i receiving item j has $v_{ij} = T$, and also has $v_{ij'} < T$ for all $j' < j$, charge T (as in the Naive Auction).
(4) If the bidder i receiving item j has $v_{ij} = T$, and also has $v_{ij'} = T$ for some $j' < j$,[13] then charge a price of $T - \frac{b_0}{a_0}\left(T - \frac{mn}{T}\right)$. Think of $\frac{b_0}{a_0}\left(T - \frac{mn}{T}\right)$ as a subsidy.
(5) Otherwise, do not allocate or elicit payments for item j.
(6) This results in the following possible interim allocations/probabilities for each bidder:
   - Receive any single item with interim probability a_0, paying interim price $a_0 T$.
   - Receive any non-empty set H of items with interim probability a_0 and any (possibly empty) set L of items with interim probability b_0, paying interim price $b_0 \cdot |L| \cdot \frac{mn}{T} + a_0 \cdot |H| \cdot T - b_0\left(T - \frac{mn}{T}\right) \cdot (|H| - 1)$.
   - Receive nothing and pay nothing.

Let us first observe that a_0 is small, but not terribly small ($\Theta(T/n) = \Theta(\sqrt{m/n})$ – roughly the inverse of the expected number of bidders with value T for a single item). b_0, on the other hand, is exponentially small – at most $(1 - 1/T)^{n-1} = e^{-\Omega(n/T)} = e^{-\Omega(\sqrt{n/m})}$. This initially seems like good news – even if we return a subsidy on *every* item, the subsidies are exponentially small, and therefore the revenue of the Less-Naive Auction falls short of the Naive Auction by at most an exponentially small amount.

However, recall that the Naive Auction's gains over selling separately are also exponentially small, so this exceptionally simple argument doesn't quite suffice, and these subsidies roughly cancel the gains over selling separately *if we pay them out every time.*[14] However, there is one key case where we don't need to pay a subsidy: if bidder i values exactly one item at T. Indeed, Lemma 3.7 formalizes this intuition and shows that the Less-Naive Auction's gain in revenue comes precisely from selling items $\neq j$ for cheap to bidders who value a single item j at T.

[13] Note that we use $j' < j$ to ensure that *exactly* one such item valued at T is *not* subsidized – the lowest-indexed such item.
[14] Essentially, paying the subsidies every time amounts to selling each item j using a randomized-but-still-single-dimensional auction, which again cannot outperform selling separately.

Lemma 3.7. *The Less-Naive Auction satisfies:*[15]

$$\textsc{Rev}_n^{LNA}(\mathcal{ER}^m_{\leq T}) = \textsc{SRev}_n(\mathcal{ER}^m_{\leq T}) + b_0 mn^2 \left(\tfrac{T}{mn} - \tfrac{1}{T}\right)\left(1 - \left(1-\tfrac{1}{T}\right)^m - \tfrac{m}{T}\left(1-\tfrac{1}{T}\right)^{m-1}\right).$$

Proof. The revenue extracted by the Less-Naive Auction from selling item j to bidder i is

$$a_0 T \cdot \mathbb{1}\left(v_{ij} = T, \max_{j'<j} v_{ij'} < T\right) + a_0\left(T - \frac{b_0}{a_0}\left(T - \frac{mn}{T}\right)\right) \cdot \mathbb{1}\left(v_{ij} = T, \max_{j'<j} v_{ij'} = T\right) + \frac{b_0 mn}{T} \cdot \mathbb{1}\left(v_{ij} \in [\tfrac{mn}{T}, T), \max_{j'\neq j} v_{ij'} = T\right)$$

Meanwhile, SRev extracts $a_0 T \cdot \mathbb{1}(v_{ij} = T)$. Thus, the Less-Naive Auction obtains

$$\frac{b_0 mn}{T} \cdot \mathbb{1}\left(v_{ij} \in [\tfrac{mn}{T}, T), \max_{j'\neq j} v_{ij'} = T\right) - b_0\left(T - \frac{mn}{T}\right) \cdot \mathbb{1}\left(v_{ij} = T, \max_{j'<j} v_{ij'} = T\right)$$

more revenue from selling item j to bidder i than $\textsc{SRev}_n$ does. Across all items, the Less-Naive Auction obtains in expectation

$$\begin{aligned}
&\frac{b_0 mn}{T}\sum_{j\in[m]} \mathbb{P}\left[v_{ij} \in [\tfrac{mn}{T}, T), \max_{j'\neq j} v_{ij'} = T\right] - b_0\left(T - \frac{mn}{T}\right)\sum_{j\in[m]} \mathbb{P}\left[v_{ij} = T, \max_{j'<j} v_{ij'} = T\right] \\
&= \frac{b_0 mn}{T}\sum_{j\in m}\left(\frac{T}{mn} - \frac{1}{T}\right)\left(1 - \left(1 - \frac{1}{T}\right)^{m-1}\right) - b_0\left(T - \frac{mn}{T}\right)\sum_{j\in[m]}\frac{1}{T}\left(1 - \left(1-\frac{1}{T}\right)^{j-1}\right) \\
&= \frac{b_0 m^2 n}{T}\left(\frac{T}{mn} - \frac{1}{T}\right)\left(1 - \left(1 - \frac{1}{T}\right)^{m-1}\right) - b_0\left(T - \frac{mn}{T}\right)\underbrace{\left(\frac{m}{T} - 1 + \left(1 - \frac{1}{T}\right)^m\right)}_{\text{expected number of subsidies}} \quad \text{(geometric sum)} \\
&= b_0 mn\left(\frac{T}{mn} - \frac{1}{T}\right)\left(1 - \left(1 - \frac{1}{T}\right)^m - \frac{m}{T}\left(1 - \frac{1}{T}\right)^{m-1}\right)
\end{aligned}$$

more revenue from bidder i than $\textsc{SRev}_n$ does. Summing over all bidders yields the lemma. □

We highlight that since $m \leq n$, the expected number of subsidies per bidder is strictly less than 1. In fact, when $m \ll n$, we expect to pay almost no subsidies. Thus, we expect our extra revenue to come from selling items for which no bidders have value T to bidders who value exactly one other item at T.

Finally, let us revisit incentives of the Less-Naive Auction. The Less-Naive Auction is *almost* BIC. Indeed, any bidder who values at least one item at T is incentivized to report truthfully (we will prove this formally in the subsequent Section 3.3). Moreover, any bidder who values all items far from T will also prefer to take no items and pay nothing. However, a bidder with v_{ij} *extremely close* to T for item j and $v_{ij'} > mn/T$ for another (and $v_{ij''} < T$ for all j'') *may* prefer to misreport that $v_{ij} = T$ (taking a small negative utility on item j) in order to receive item j' with non-zero probability. Therefore, the Less-Naive Auction is not BIC.

However, because b_0 is exponentially small, the maximum possible gain of such a misreport is also exponentially-small. Therefore, only types with v_{ij} *inverse exponentially-close* to T will even consider this misreport, and there is hope that a slight modification to the Less-Naive Auction might work. Indeed, we now show an auction whose interim menu takes the same format as the

[15] Again recall that the Less-Naive Auction is not BIC – this analysis is just to supply intuition for our later (even more involved) computations.

Less-Naive Auction, but with parameters $a \approx a_0$ and $b \approx b_0$ that is BIC and again has essentially the same expected revenue.

3.3 Step Three: A BIC Auction

Now, we introduce a final set of modifications to make the Less-Naive Auction *fully* BIC. Recall that the Less-Naive Auction is not BIC because a bidder with no values equal to T but sufficiently high $v_{ij} \approx T, v_{ij'} > mn/T$ may choose to misreport and take a small loss on item j in order to receive item j' (indeed, Corollary 3.10 formalizes this intuition). It follows that the interim allocation probabilities a_0, b_0 are not actually feasible, because there are more bidders that want to purchase items than we can keep our promises to.

To fix this, we maintain the same menu format, but lower the interim probabilities (and prices accordingly) to a, b. This fixes the incentive issues of the Less-Naive Auction by making misreporting less attractive, but since $a \approx a_0$ and $b \approx b_0$ we still attain approximately the same revenue. Observe that even a small change in a_0 and b_0 works, because of two simultaneous effects at play: (1) lowering the allocation probabilities inherently increases the number of bidders to which it is feasible to allocate an item; (2) lowering b_0 in particular reduces the set of types that may prefer to misreport. Balancing these two effects so that the feasibility constraint is tight (we would like to allocate the item as much as possible, so that we can extract as much revenue as possible) results in a system with a fixed point (a, b) that is not too far from the original (a_0, b_0).

Definition 3.8 (The Not-So-Naive Auction)**.** The Not-So-Naive Auction allocates the items according to the following menu of interim allocations/probabilities for each bidder:

- Receive any single item with interim probability a, paying interim price aT.
- Receive any non-empty set H of items with interim probability a and any (possibly empty) set L of items with interim probability b, paying interim price $b \cdot |L| \cdot \frac{mn}{T} + a \cdot |H| \cdot T - b\left(T - \frac{mn}{T}\right) \cdot (|H| - 1)$.
- Receive nothing and pay nothing.

Again, we highlight that if we had $a = a_0, b = b_0$, this menu just describes the Less-Naive Auction. In Lemma 3.9 and Corollary 3.10, we characterize the incentive properties of *any* such menu of the above form parametrized by $a \geq b$; following this, we proceed to set a and b specifically so that the Not-So-Naive Auction is feasible.

Lemma 3.9. *Suppose $T \geq \sqrt{mn}$ and $a \geq b$. Let $v \in [1, T]^m$ and let $j^* \in \arg\max_j v_j$. Define $H := \{j : v_j = T\} \cup \{j^*\}$ and $L := \{j : v_j \geq mn/T\} \setminus H$. For all $H' \in 2^{[m]} \setminus \{\varnothing\}$ and $L' \subseteq [m] \setminus H'$, a bidder with type v prefers the menu option (H, L) over the menu option (H', L').*

Lemma 3.9 says that a bidder with type v prefers (H, L) over any other option that allocates an item with some positive probability. In particular, it does *not* say whether a bidder with type v would prefer (H, L) over not getting any items at all.

Proof. The utility of a bidder with type v for the menu option (H', L') is

$$a \sum_{j \in H'} v_j + b \sum_{j \in L'} v_j - \left(|H'|\, aT + |L'|\, b\frac{mn}{T} - (|H'| - 1)b\left(T - \frac{mn}{T}\right)\right)$$
$$= \sum_{j \in H'} \left(av_j - \left((a-b)T + b\frac{mn}{T}\right)\right) + \sum_{j \in L'} b\left(v_j - \frac{mn}{T}\right) - b\left(T - \frac{mn}{T}\right).$$

We show that (H, L) maximizes this utility over all $H' \in 2^{[m]} \setminus \{\varnothing\}$ and $L' \subseteq [m] \setminus H'$.

Consider the difference in utility of getting item j with probability a and getting the same item with probability b: $av_j - \left((a-b)T + b\frac{mn}{T}\right) - b\left(v_j - \frac{mn}{T}\right) = (a-b)(v_j - T)$. It is clear that if $a \geq b$, then there are only two cases in which a bidder who values item j at v_j (and who is forced to get at least one item) would prefer to get item j with probability a: either (1) $v_j = T$, or (2) she does not value any item at T, but $j = j^*$ (if she must pay T for *some* item, her utility is least negative when j is the item she values the most rather than some other item). Thus, the only items that a bidder with type v prefers to get with probability a rather than b are the items in H. Getting any item outside of H with probability a strictly decreases utility.

Of the items not in H, a bidder with type v would only choose to get those with values at least mn/T with some positive probability, since all options cost at least mn/T, so the bidder would be overpaying for any item valued less than mn/T (which strictly decreases utility). Note that the items with values at least mn/T that are not in H are precisely those in L.

As discussed before, the only items that a bidder with type v prefers to get with probability a rather than b are the items in H, so such a bidder prefers to get the items in L with probability b instead of a. Thus, the bidder's most preferred menu item is exactly (H, L). □

Corollary 3.10. *Suppose $T \geq \sqrt{mn}$ and $a \geq b$. Let $v \in [1, T]^m$ and let $j^* \in \arg\max_j v_j$. Define $H := \left\{j : v_j = T\right\} \cup \{j^*\}$ and $L := \left\{j : v_j \geq mn/T\right\} \setminus H$.*

- *A bidder who values some item at T prefers the menu option (H, L) over any other option in the menu (including not receiving any items).*
- *A bidder who does not value any item at T prefers the menu option $(\{j^*\}, L)$ over any other option in the menu (including not receiving any items) if and only if $av_{j^*} + b\sum_{j \in L} v_j \geq aT + |L|\, b\frac{mn}{T}$.*

Now, we define a and b such that the resulting menu is feasible. We do so implicitly. Recall that setting $a = a_0$ and $b = b_0$ is infeasible because there are bidders who do not value any items at T yet value a subset of the items enough to be willing to pay T for an item in order to be eligible to get additional items at lower prices. The probability that there exists a bidder who values each item less than T yet is willing to purchase the menu option $(\{j^*\}, L)$ is

$$q_\ell := \mathop{\mathbb{P}}_v\left[v_{j^*} = \max_j v_j < T, \min_{j \in \{j^*\} \cup L} v_j \geq \tfrac{mn}{T}, \max_{j \notin \{j^*\} \cup L} v_j < \tfrac{mn}{T}, av_{j^*} + b\textstyle\sum_{j \in L} v_j \geq aT + |L|\, b\tfrac{mn}{T}\right].$$

Note that for a given $\ell \geq 1$, the above probability is the same for any choice of $j^* \in [m]$ and $L \subseteq [m] \setminus \{j^*\}$ such that $|L| = \ell$, so we may denote it by q_ℓ.

Since there are more bidders who want to purchase items than just those with T values, the interim allocation probabilities a and b must be smaller than a_0 and b_0 to accommodate these bidders. More specifically, if we term bidders who are willing to receive item j with probability a as "high" and those who are only willing to receive item j with probability b as "low," and we want to allocate each item uniformly at random to the high bidders before allocating uniformly at random to the low bidders, then a and b must satisfy the following implicit definitions.

$$a = \mathop{\mathbb{E}}_{v_{-i}}\left[\tfrac{1}{1+\sum_{k \neq i} \mathbb{1}(i \text{ high})}\right] = \textstyle\sum_{k=0}^{n-1} \frac{\binom{n-1}{k}}{k+1} \mathbb{P}[\text{high}]^k\, \mathbb{P}[\text{not high}]^{n-k-1} = \frac{1-(1-\mathbb{P}[\text{high}])^n}{n\mathbb{P}[\text{high}]},$$

$$b = \mathop{\mathbb{E}}_{v_{-i}}\left[\tfrac{\mathbb{1}(\nexists \text{ high})}{1+\sum_{k \neq i} \mathbb{1}(i \text{ low})}\right] = \frac{\mathbb{P}_{v_{-i}}\left[\nexists \text{ high}\right](1-(1-\mathbb{P}[\text{low} \mid \text{not high}])^n)}{n\mathbb{P}[\text{low} \mid \text{not high}]},$$

(by bidder independence, $\mathbb{P}\left[\text{low} \,\middle|\, \nexists \text{ high}\right] = \mathbb{P}[\text{low} \mid \text{not high}]$)

where

$$\mathbb{P}[\text{high}] = \mathop{\mathbb{P}}_v\left[v_j = T\right] + \sum_{L \subseteq [m] \setminus \{j\}} q_{|L|} = \frac{1}{T} + \sum_{\ell=1}^{m-1} \binom{m-1}{\ell} q_\ell,$$

$$\mathbb{P}[\text{low}] = \mathop{\mathbb{P}}_v\begin{bmatrix}\max_{j' \neq j} v_{j'} = T, \\ \frac{mn}{T} \leq v_j < T\end{bmatrix} + \sum_{j^* \neq j} \sum_{\substack{L \subseteq [m] \setminus \{j^*\}: \\ j \in L}} q_{|L|} = \left(1 - \left(1 - \frac{1}{T}\right)^{m-1}\right)\left(\frac{T}{mn} - \frac{1}{T}\right) + (m-1)\sum_{\ell=1}^{m-1} \binom{m-2}{\ell-1} q_\ell.$$

The expression for $\mathbb{P}$ [high] follows from the fact that by Corollary 3.10, a bidder is high for an item if and only if (1) she has value T for it or (2) it is her favorite item and she has sufficiently high values for the other items. Similarly, a bidder is low for an item if and only if her value for it is in $[mn/T, T)$ and either (1) she has value T for some other item or (2) she has sufficiently high values for the other items. We point out that our definitions of a and b are indeed implicit since $q_1, \dots, q_{m-1}$ depend on a and b. We now show that $a \geq b$, so that Lemma 3.9 holds with our definitions of a and b.

Lemma 3.11. *If $T \geq \sqrt{mn}$, then $\frac{b}{a} \leq \frac{n}{T} e^{-\frac{n}{T}} \left(1 - e^{-\frac{n}{T}}\right)^{-1}$.*

Corollary 3.12. *If $T \geq \sqrt{mn}$, then $a \geq b$.*

Proof. A direct consequence of Lemma 3.11 and the fact that $xe^{-x} \leq 1 - e^{-x}$ for $x \geq 0$. □

3.4 Step Four: Comparing the Revenue of the Not-So-Naive Auction to the Revenue of the Less-Naive Auction

The ultimate goal is to compare the revenue of the Not-So-Naive Auction against the revenue of selling separately. We proceed via an intermediate comparison between the Not-So-Naive Auction and the Less-Naive Auction (which we have already compared to $\text{SRev}_n(\mathcal{ER}^m_{\leq T})$ in Section 3.2).

Lemma 3.13. *If bidders report their values truthfully in the Less-Naive Auction, then $\text{Rev}^{NSN}_n(\mathcal{ER}^m_{\leq T})$ exceeds $\text{Rev}^{LNA}_n(\mathcal{ER}^m_{\leq T})$ by at least*

$$(b - b_0)mn^2 \left(\frac{T}{mn} - \frac{1}{T}\right) \left(1 - \left(1 - \frac{1}{T}\right)^m - \frac{m}{T}\left(1 - \frac{1}{T}\right)^{m-1}\right) + \frac{bm^2(m-1)n^2}{T} \sum_{\ell=1}^{m-1} \binom{m-2}{\ell-1} q_\ell.$$

Corollary 3.14. *$\text{Rev}^{NSN}_n(\mathcal{ER}^m_{\leq T})$ exceeds $\text{SRev}_n(\mathcal{ER}^m_{\leq T})$ by at least*

$$bmn^2 \left(\left(\frac{T}{mn} - \frac{1}{T}\right) \left(1 - \left(1 - \frac{1}{T}\right)^m - \frac{m}{T}\left(1 - \frac{1}{T}\right)^{m-1}\right) + \frac{m(m-1)}{T} \sum_{\ell=1}^{m-1} \binom{m-2}{\ell-1} q_\ell\right).$$

Proof. A direct consequence of Lemmas 3.7 and 3.13. □

3.5 Step Five: Bounding $\text{Comp}_{\mathcal{ER}^m_{\leq T}}(n)$

Notice that $\text{Comp}_{\mathcal{ER}^m_{\leq T}}(n)$ is *at least* the smallest c such that $\text{SRev}_{n+c}(\mathcal{ER}^m_{\leq T})$ exceeds $\text{Rev}^{NSN}_n(\mathcal{ER}^m_{\leq T})$. Combining Lemma 3.3 and Corollary 3.14 along with the fact that all $q_\ell \geq 0$, this occurs only if[16]

$$c \geq \left(1 - \frac{1}{T}\right)^{-n} bn^2 \left(\frac{T}{mn} - \frac{1}{T}\right) \left(1 - \left(1 - \frac{1}{T}\right)^m - \frac{m}{T}\left(1 - \frac{1}{T}\right)^{m-1}\right).$$

By the union bound, we expect the RHS to behave like $bmn/(T(1-\frac{1}{T})^n)$. If we set $T \sim \sqrt{mn}$, then to show that the competition complexity is $c = \Omega(\sqrt{mn})$, it suffices to show that $b = \Omega\left(\left(1 - \frac{1}{T}\right)^n\right)$. Intuitively, if the probability of the types of bidders for which we had to modify the Less-Naive Auction to the Not-So-Naive Auction is sufficiently small, then $b \approx b_0 = \Omega((1 - \frac{1}{T})^n)$ since the probability of valuing item j in $[mn/T, T)$ and some other item at T is at most $O(1/n)$ if $T \sim \sqrt{mn}$. We show that all of this is indeed the case in Lemma 3.15 and prove that the competition complexity is $\Omega(\sqrt{mn})$ in Theorem 3.16.

Lemma 3.15. *If $T = \lambda\sqrt{mn}$ for some constant $\lambda > 1$, then $b = \Omega\left(\left(1 - \frac{1}{T}\right)^n\right)$.*

Theorem 3.16. *If $T = \lambda\sqrt{mn}$ for some constant $\lambda > 1$, then $\text{Comp}_{\mathcal{ER}^m_{\leq T}}(n) = \Omega(\sqrt{mn})$.*

[16] We ignore the expected revenue gained from selling items to "low"-type bidders with no T values since this term is irrelevant to the remainder of our analysis.

This concludes the proof of our main result. We have explicitly defined a BIC auction (the Less-Naive Auction) for $(\mathcal{ER}^m_{\leq T})^n$ whose revenue exceeds $\text{SRev}_{\mathcal{ER}^m_{\leq T}}(n + c\sqrt{nm})$ for some absolute constant $c > 0$, and all $m \geq 2$, and all n.

4 Further Preliminaries: Dual Flow Benchmarks

In the following sections, we provide revenue upper bounds using the [Cai et al., 2016] framework. We briefly state the minimal preliminaries necessary to get started, and will fully flesh out terminology as needed during proofs. As needed, we will clarify what is a 'useful dual flow' so that the statement of Theorem 4.1 is fully self-contained. Theorem 4.2 is already self-contained. Their framework establishes the following revenue benchmark in terms of the induced *virtual values* of bidder i for item j, $\Phi^\lambda_{ij}(\vec{v}_i)$ as a function of a 'useful dual flow' λ:

Theorem 4.1 ([Cai et al., 2016], Theorem 6). *Let λ be any useful dual flow, and $M = (\pi, p)$ be a BIC mechanism. The revenue of M is less than or equal to the virtual welfare of π with respect to the virtual value function Φ^λ; that is: $\text{Rev}^M_n(D) \leq \sum_{i=1}^n \sum_{j=1}^m \mathbb{E}_{\vec{v}\leftarrow D^n}[\pi_{ij}(\vec{v}_i) \cdot \Phi^\lambda_{ij}(\vec{v}_i)]$.*

Taking the supremum over all feasible M then provides an upper bound on $\text{Rev}_n(D)$. In particular, when λ is the *canonical flow* that divides the type space into regions R_j based on the favorite item j, then uses a Myersonian-like flow within each region (refer to [Cai et al., 2016] for a precise definition), the following *relaxation* of the benchmark is useful:

Theorem 4.2 ([Cai et al., 2016], Corollary 28). *Let $R_j := \{\vec{v} \mid \arg\max_{\ell\in[m]}\{v_\ell\} = j\}$ (with ties broken lexicographically). Then $\text{Rev}_n(D) \leq \sum_{j=1}^m \mathbb{E}_{\vec{v}\leftarrow D^n}[\max_{i\in[n]}\{\bar{\varphi}_j(v_{ij}) \cdot \mathbb{1}\left(\vec{v}_i \in R_j\right) + v_{ij} \cdot \mathbb{1}\left(\vec{v}_i \notin R_j\right)\}]$.*

5 A Reduction from $\mathcal{A}^R_m$ to $\mathcal{ERC}_m$ via Stochastic Dominance

We now consider the competition complexity of an arbitrary distribution $D \in \mathcal{A}^R_m$. The key idea is that the Theorem 4.2 benchmark can be written entirely in terms of (ironed) virtual values, which then allows a direct comparison to the virtual value obtained by SRev. If SRev (with c additional bidders) always obtains a virtual value of a higher quantile than the optimal mechanism (without additional bidders), then it also achieves higher revenue. To tighten the previous analysis from [Beyhaghi and Weinberg, 2019, Cai et al., 2016], we first use the style of Theorem 4.2 but *applied to a specific allocation*, then take the supremum over all feasible and BIC mechanisms.

Fix a deterministic allocation rule x, and let $x_j(\vec{v})$ denote the winner of item j under x on input $\vec{v}$ (if no one wins the item, let $x_j(\vec{v}) = \perp$). Sample $\vec{v} \leftarrow D^n$, and choose a bidder i for item j according to x. Define the following quantities:

- $CDW^x_j(\vec{v}) := \sum_i x_{ij}(\vec{v}) \cdot \Phi^\lambda_{ij}(\vec{v}_i) = \sum_i x_{ij}(\vec{v}) \cdot (\bar{\varphi}_j(v_{ij}) \cdot \mathbb{1}\left(\vec{v}_i \in R_j\right) + v_{ij} \cdot \mathbb{1}\left(\vec{v}_i \notin R_j\right))$, the virtual value (using the canonical flow) of the recipient of item j on valuation profile $\vec{v}$,
- $Q^x_j(\vec{v})$, a random variable sampled as follows:
  - If $\vec{v}_i \in R_j$, output $Q^x_j = q_j(\vec{v}_i) = F_j(v_{ij})$.
  - Else, output $Q^x_j \leftarrow U[q_j(\vec{v}_i), 1]$.
- $S_{n+c} := \max_{\vec{q}\leftarrow U[0,1]^{n+c}} q_i$, the maximum of $n + c$ independently drawn quantiles.

Following the rest of the argument from [Beyhaghi and Weinberg, 2019, Cai et al., 2016] gives the following *refined* benchmark (note that this essentially interchanges the expectation and the maximum from Theorem 4.2, so it indeed furnishes a tighter bound):

$$\text{Rev}_n(D) \leq \sup_{\text{feasible BIC } x} \sum_{j=1}^m \mathbb{E}_{\vec{v}\leftarrow D^n}\left[CDW^x_j(\vec{v})\right].$$

Using this improved benchmark and the language of quantile space, we now compare $\text{Rev}_n(D)$ to $\text{SRev}_{n+c}(D)$. Along the way, we appeal to the specific form of the virtual values of distributions in $\mathcal{ERC}_m$, which establishes the non-trivial direction of the equality between $\text{Comp}_{\mathcal{A}_m^R}(n)$ and $\text{Comp}_{\mathcal{ERC}_m}(n)$. Full proofs of all results are provided in the Appendix.

Proposition 5.1. *For all allocation rules x and all items j,* $\mathbb{E}_{\vec{v}\leftarrow D_r^n}[\bar{\varphi}_j(F_j^{-1}(Q_j^x(\vec{v})))] \geq \mathbb{E}_{\vec{v}\leftarrow D^n}[CDW_j^x(\vec{v})]$.

Proposition 5.2. *If* $S_{n+c} \gtrsim Q_j^x(\vec{v})$*, then* $SRev_{n+c}(D) \geq \mathbb{E}_{\vec{v}\leftarrow D^n}[\sum_j CDW_j^x(\vec{v})]$.

Observation 5.3. *For all* $q \in [0,1)$*, the distribution satisfying (up to scaling by a constant)* $\bar{\varphi}(x) = \mathbb{1}(F(x) \geq q)$ *is* $\mathcal{ER}_{\leq \frac{1}{1-q}}$.

Proposition 5.4. $S_{n+c} \gtrsim Q_j^x(\vec{v})$ *if and only if* $SRev_{n+c}(\mathcal{ER}_{\leq T}) \geq \mathbb{E}_{\vec{v}\leftarrow(\mathcal{ER}_{\leq T})^n}[\sum_j CDW_j^x(\vec{v})]$ *for all truncations* $T \in [1,\infty)$.

Corollary 5.5. *If* $SRev_{n+c}(\mathcal{ER}_{\leq T}) \geq \mathbb{E}_{\vec{v}\leftarrow(\mathcal{ER}_{\leq T})^n}[\sum_j CDW_j^x(\vec{v})]$ *for all* $T \in [1,\infty)$*, then the competition complexity for* any *distribution* $D \in \mathcal{A}_m^R$ *is* $O(c)$.

This reduction implies that to establish a bound on $\text{Comp}_{\mathcal{A}_m^R}(n)$, it suffices to just study $\text{Comp}_{\mathcal{ERC}_m}(n)$. Although this claim is not directly used for our earlier main result, we still present it as a technique of general interest, potentially useful for future work, and illuminating as to the context and further implications of our main result focusing on $\mathcal{ER}_{\leq T}$.

6 Upper Bound on $\text{Rev}_n(\mathcal{ER}^m) = nm + O(m^2 \ln n)$ when $n > m$

In this section, we show that we cannot obtain more than $nm+O(m^2 \ln n)$ in revenue from n bidders with additive valuations for m items drawn iid from $\mathcal{ER}$. This upper bound is interesting for two reasons. First, it shows that the untruncated Equal Revenue curve does *not* witness the worst-case competition complexity when $n \geq m$. However, in Section 6.4, we show that if bidders cannot lie about their favorite item, then the untruncated Equal Revenue curve *does* witness the worst-case competition complexity when $n \geq m$. Thus, a "region-separated" flow provably cannot give a tight upper bound on the revenue obtainable for this setting. In Sections 6.1 and 6.2, we demonstrate how to circumvent this impossibility by taking advantage of certain cross-region constraints. We show that our upper bound is nearly tight in Section 6.3.

6.1 Tight Bound for $m = 2$: $\text{Rev}_n(\mathcal{ER}^2) = 2n + \Theta(\ln n)$

To establish a tight bound on $\text{Rev}_n(\mathcal{ER}^2)$, we start from Theorem 4.1. Rather than relaxing all the way to Theorem 4.2, we obtain an upper bound on the optimal revenue by first providing a further characterization of feasible and BIC mechanisms M over which we take the supremum of the virtual welfare.

As established in [Cai et al., 2016], the expected virtual welfare from bidders who are awarded their favorite item is $2n$; we seek to understand the expected virtual welfare from bidders who win their non-favorite item. We begin with some motivating observations (proofs of which are provided in the Appendix):

Observation 6.1. *Fix* $v_1 > v_2$*. It is feasible to have each type* (v_1, v_2) *with* $v_2 \geq \ln^2 n$ *receive their non-favorite item (item 2) with probability* $\min\left\{\frac{1}{2}, \frac{v_1}{10n}\right\}$.

Observation 6.2. *If every type with* $v_N \geq \ln^2 n$ *receives their non-favorite item with probability at least* $\min\left\{\frac{1}{2}, \frac{v_F}{10n}\right\}$*, the benchmark gets at least* $2n + \Omega(\ln^2 n)$.

Observation 6.3. *It is* not *feasible to allocate both items with probability* $\min\left\{\frac{1}{2}, \frac{v_1}{10n}\right\}$ *to* all *types.*

Combining Observations 6.1 and 6.3 suggests that we are only in trouble if it is somehow possible to give items only to players with big v_N without also giving items to players with small v_N. But, this is difficult if items are mostly awarded based on having large v_F (because v_N is generally much smaller than v_F when v_F is large). That is, to get expected virtual welfare $2n + \Omega(\ln^2 n)$, we need to have things like "$(n/100, \sqrt{n})$ gets item 2 with probability $1/100$, but $(n/10, 2)$ gets item 1 with probability 0." So, we seek to show that this is *not* possible by appealing to BIC and IR constraints in Lemma 6.4 in addition to feasibility in Corollary 6.5; see the Appendix for full details.

Lemma 6.4. *Let* (v_1, v_2) *get item 2 with probability* $q := \pi_2(v_1, v_2)$. *Let also* $v_2' \leq v_2$. *Then* $(3v_1, v_2')$ *gets item 1 with probability* $\pi_1(3v_1, v_2') \geq q/4$.

Corollary 6.5. *Let* $y \geq 2$. *Then* $\mathbb{E}_{v_1}\left[\pi_2(v_1, v_2) \mid v_2 = y\right] = \int_y^\infty \pi_2(v_1, y) \cdot \frac{y}{v_1^2} dv_1 \leq \frac{24y}{n}$. *That is, the probability of getting item 2 conditioned on having* $v_2 = y$ *is at most* $\frac{24y}{n} = O(\frac{y}{n})$.

Lemma 6.6. *The expected contribution to the virtual welfare from the non-favorite item is* $O(\ln n)$.

Proof. First, note that a bidder only contributes to the virtual welfare if they are awarded the item (which occurs with probability $q_2 = \max\{O(\frac{y}{n}), 1\}$, in which case they contribute their virtual value (which is at most v_2). Then, we can apply the law of total expectation to compute:

$$\begin{aligned}
\mathbb{E}_{v_1,v_2}\left[v_2 q_2\right] = \mathbb{E}_{v_2}\left[\mathbb{E}_{v_1}\left[v_2 q_2 \mid v_2\right]\right] &= O\left(\mathbb{E}_{v_2}\left[v_2 \cdot \max\left\{\tfrac{v_2}{n}, 1\right\}\right]\right) \\
&= O\left(\int_{x=1}^{n} \tfrac{x^2}{n} \mathbb{P}(v_2 = x) dx + \int_{x=n}^{\infty} x \cdot \mathbb{P}(v_2 = x) dx\right) \\
&= O\left(\int_{x=1}^{n} \tfrac{x^2}{n} \cdot \tfrac{2}{x^3} dx + \int_{x=n}^{\infty} \tfrac{2}{x^2} dx\right) \\
&= O\left(\tfrac{2\ln n}{n} + \tfrac{2}{n}\right) = O\left(\tfrac{\ln n}{n}\right).
\end{aligned}$$

Summing over all n bidders gives a total of $O(\ln n)$. □

Combining this with the lower bound of $\text{Rev}_n(\mathcal{ER}^2) = 2n + \Omega(\ln n)$ due to [Beyhaghi and Weinberg, 2019] establishes that this bound is tight.

Theorem 6.7. $\text{Rev}_n(\mathcal{ER}^2) = 2n + \Theta(\ln n)$.

6.2 Generalizing to $m > 2$: $\text{Rev}_n(\mathcal{ER}^m) = nm + O(m^2 \ln n)$

We generalize the analysis from Section 6.1 to general m, thereby improving the upper bound on the competition complexity of n bidders with additive values drawn i.i.d. from $\mathcal{ER}^m$ to $O(m \ln n)$. We introduce the following additional notation when considering a particular bidder:

- E_j denotes the event that item j is the favorite item.
- E_{-j} denotes the event that item j is a non-favorite item.

Our approach exactly mirrors that of the $m = 2$ analysis but requires much more involved calculations; as such, we defer all proofs to the Appendix.

Lemma 6.8. *Let* (π, p) *be a BIC mechanism for n bidders with m additive valuations drawn i.i.d. from* $\mathcal{ER}$. *The expected contribution of each non-favorite item to the virtual welfare is at most* $\mathbb{E}\left[v_j \pi_j(v) \,\middle|\, E_{-j}\right] \leq O\left(\frac{m \ln n}{n}\right)$.

Summing over all $m - 1$ non-favorite items and all n bidders gives an upper bound on the virtual welfare of $O(m^2 \ln n)$, in addition to nm from the favorite item ([Cai et al., 2016]), and thus:

Theorem 6.9. *$\text{Rev}_n(\mathcal{ER}^m) = nm + O(m^2 \ln n)$, and therefore, for some absolute constant c,*

$$VCG_{n+cm\ln(n)}(\mathcal{ER}^m) \geq \text{Rev}_n(\mathcal{ER}^m).$$

6.3 Selling to $\mathcal{ER}^m$ via the Grand Bundle: $\text{Rev}_n(\mathcal{ER}^m) \geq nm + \Omega(m \ln n)$

In this section, we show that the upper bound on $\text{Rev}_n(\mathcal{ER}^m)$ proved in Section 6.2 is nearly tight. More specifically, we show that selling the grand bundle via a second-price auction obtains $nm + \Theta(m \ln(mn))$ in revenue. Note that we improve upon the previous best lower bound of $nm + \Omega(\ln n)$ from [Beyhaghi and Weinberg, 2018].

Theorem 6.10. *$\text{Rev}^{SPA\text{-}GB}(\mathcal{ER}^m) = nm + \Theta(m \ln(mn))$.*

Since the revenue of a second-price auction is given by the second highest value for the item being sold, we give upper and lower bounds on the second highest value for the grand bundle. Unfortunately, sums of random variables are difficult to work with, so we seek a good proxy for the second highest bundle value that is more straightforward to analyze. We claim that the bundle value of the bidder with the second highest value for her favorite item is a good proxy.

To see why, first note that we expect each bidder's value for the grand bundle to be dominated by her value for her favorite item: conditioned on the value for the favorite item, we expect the value for each non-favorite item to be exponentially smaller than her value for her favorite item.

Lemma 6.11. $\mathbb{E}_{x\sim\mathcal{ER}}\left[x \mid x \leq v\right] = \frac{\ln v}{1-1/v}$.

Thus, we expect the bidder with the highest value for any item to also have the highest value for the grand bundle, the bidder with the second highest value for her favorite item to have the second highest value for the grand bundle, and so on. In particular, we expect the bidder with the second highest value for her favorite item to set the price for the grand bundle.

Let $v_{(i),(j)}$ denote the j-th highest value possessed by the bidder with the i-th highest value for her favorite item. Expressed in this notation, our intuition is that $\sum_j v_{(2),(j)}$ traces the second highest value for the grand bundle. Our proof of Theorem 6.10 shows that this is precisely the case. We will also see that the expectation of $v_{(2),(1)}$ is approximately nm.

Lemma 6.12. *$nm - O(m) \leq \mathbb{E}[v_{(2),(1)}] \leq nm$.*

By Lemmas 6.11 and 6.12, we expect the second highest value for the grand bundle to be around $\mathbb{E}\left[\sum_j v_{(2),(j)}\right] = \mathbb{E}\left[v_{(2),(1)} + m\ln(v_{(2),(1)})\right] = nm + \Theta(m\ln(mn))$.

6.4 "Region-separated" Flows

Finally, we discuss the class of "region-separated" flows, which do not cross any axes between different favorite-item regions R_j. These correspond to auctions that respect all BIC constraints between bidders with the same favorite item, but not necessarily between bidders with different favorite items; we term such auctions Knows-Favorite BIC (KF-BIC). We design a KF-BIC auction that achieves revenue $nm + \Omega(m\sqrt{nm})$ from n bidders with values drawn i.i.d. from $\mathcal{ER}^m$.

In addition to potentially being of independent interest, this result further highlights our upper bound on $\text{Rev}_n(\mathcal{ER}^m)$ from Section 3 as interesting because it provably cannot follow from an "region-separated" flow, and establishes that the cross-diagonal BIC constraints are *necessary* to achieve the optimal bound.

Definition 6.13 (KF-BIC)**.** We say that an auction is Knows-Favorite Bayesian Incentive Compatible if for all types $\vec{v}$ with distinct values for every item, $\vec{v}$ does not wish to misreport any other $\vec{w}$ with the same favorite item. That is, if S_j denotes the subset of valuations in the support of D such that $v_j > v_{j'}$ for all $j' \neq j$, and also $v_{j'} \neq v_{j''}$ for all j', j'', a mechanism with interim allocation rule $\vec{\pi}(\cdot), p(\cdot)$ is KF-BIC if: $\forall j,\ \forall \vec{v}, \vec{w} \in S_j,\ \vec{v} \cdot \vec{\pi}(\vec{v}) - p(\vec{v}) \geq \vec{v} \cdot \vec{\pi}(\vec{w}) - p(\vec{w})$.

Importantly, note that a KF-BIC auction is not necessarily BIC. This is both because there are no constraints that involve types with the same value for multiple items, and also because the KF-BIC constraints only guarantee that bidders do not wish to misreport *while keeping their favorite item the same* (they may wish to misreport their favorite item).

6.4.1 The Knows-Favorite Auction (KFA).

Definition 6.14 (Knows-Favorite Auction)**.** The Knows-Favorite Auction (KFA) proceeds as follows:
Let S denote the set of bidders with distinct values for all m items, and S_j denote the subset of S consisting of bidders with favorite item j. Each item j is auctioned as follows:

- If any bidder $i \in S_j$ has $v_{ij} \geq H = e^{nm}$, the item is awarded to a uniformly random such bidder, and they are charged H.
- If no bidder in S_j has value at least H, then the item is offered to bidders in $S \setminus S_j$ at price $L = \sqrt{nm}$ (that is, as long as any bidder in $S \setminus S_j$ is willing to pay L, a uniformly random such bidder is given the item and charged L).

Observation 6.15. *KFA is KF-BIC: for all j, no bidder in S_j wishes to misreport any other type in S_j.*

Lemma 6.16. *$\textsc{Rev}^{KFA}(\mathcal{ER}^{nm}) = nm + \Omega(m\sqrt{nm})$. That is, the expected revenue (assuming bidders tell the truth) of KFA is $nm + \Omega(m\sqrt{nm})$.*

7 Conclusion

We settle the competition complexity of n bidders with additive valuations over m independent items at $\Theta(\sqrt{nm})$ in the "Big n" regime. As the "Little n" regime is previously settled by [Beyhaghi and Weinberg, 2019, Feldman et al., 2018], this settles the competition complexity for additive bidders over independent items (up to constant factors). On the technical front, we design an explicit BIC-but-not-DSIC mechanism outperforming selling separately (even with additional bidders) in a regime where selling separately is already a $(1 - o(1))$-approximation.

We also provide results of independent interest accumulated from our journey: the competition complexity of additive bidders is exactly equal to the competition complexity when restricted to iid truncated equal revenue curves, and despite this the untruncated Equal Revenue curve witnesses an exponentially-suboptimal lower bound.

As our work now settles the key remaining open problem for competition complexity of exceeding the optimal BIC mechanism by VCG, there are two important directions for future work:

- What about the competition complexity of exceeding the optimal DSIC auction? Our BIC auctions cannot be made DSIC, and it initially seems as though BIC auctions may strictly outperform DSIC auctions for the instances that yield our main result. We suspect that our Independent Result I will be useful for upper bounds on this front (if indeed improved upper bounds are possible).
- What about the competition complexity of other simple auctions? There is limited work in this direction so far, which so far still loses some (small) fraction of revenue rather truly exceeding the optimum [Cai and Saxena, 2021, Feldman et al., 2018].

Acknowledgments

The authors would like to thank Mel Shu and Becky Barber for their contributions to preliminary investigations related to this work, including the development of code for optimal auction design, which led the authors to conjecture that $\textsc{Rev}_n(\mathcal{ER}^2) = 2n + \Theta(\ln n)$. This work was supported in part by NSF grant CCF-1942497 (CAREER) and the NSF Graduate Research Fellowship Program under Grant No. DGE-2139899.

References

Mohammad Akbarpour, Yeganeh Alimohammadi, Shengwu Li, and Amin Saberi. 2022. The Value of Excess Supply in Spatial Matching Markets. In *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, David M. Pennock, Ilya Segal, and Sven Seuken (Eds.). ACM, 62. https://doi.org/10.1145/3490486.3538375

Mohammad Akbarpour, Suraj Malladi, and Amin Saberi. 2018. Diffusion, Seeding, and the Value of Network Information. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, Éva Tardos, Edith Elkind, and Rakesh Vohra (Eds.). ACM, 641. https://doi.org/10.1145/3219166.3219225

Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2020. A Simple and Approximately Optimal Mechanism for an Additive Buyer. *J. ACM* 67, 4 (2020), 24:1–24:40. https://dl.acm.org/doi/10.1145/3398745

Siddharth Barman, Shuchi Chawla, and Seeun Umboh. 2012. A Bicriteria Approximation for the Reordering Buffer Problem. In *Algorithms - ESA 2012 - 20th Annual European Symposium, Ljubljana, Slovenia, September 10-12, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7501)*, Leah Epstein and Paolo Ferragina (Eds.). Springer, 157–168. https://doi.org/10.1007/978-3-642-33090-2_15

Xiaohui Bei and Zhiyi Huang. 2011. Bayesian Incentive Compatibility via Fractional Assignments. In *the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Hedyeh Beyhaghi and S. Matthew Weinberg. 2018. Optimal (and Benchmark-Optimal) Competition Complexity for Additive Buyers over Independent Items. https://doi.org/10.48550/ARXIV.1812.01794

Hedyeh Beyhaghi and S. Matthew Weinberg. 2019. Optimal (and Benchmark-Optimal) Competition Complexity for Additive Buyers over Independent Items. In *Proceedings of the 51st ACM Symposium on Theory of Computing Conference (STOC)*.

Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2015. Pricing lotteries. *J. Economic Theory* 156 (2015), 144–174. https://doi.org/10.1016/j.jet.2014.04.011

Johannes Brustle, José R. Correa, Paul Dütting, and Victor Verdugo. 2022. The Competition Complexity of Dynamic Pricing. In *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, David M. Pennock, Ilya Segal, and Sven Seuken (Eds.). ACM, 303–320. https://doi.org/10.1145/3490486.3538366

Jeremy Bulow and Paul Klemperer. 1996. Auctions Versus Negotiations. *The American Economic Review* (1996), 180–194.

Linda Cai and Raghuvansh R. Saxena. 2021. 99% Revenue with Constant Enhanced Competition. In *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, Péter Biró, Shuchi Chawla, and Federico Echenique (Eds.). ACM, 224–241. https://doi.org/10.1145/3465456.3467647

Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. 2016. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the 48th ACM Conference on Theory of Computation(STOC)*.

Yang Cai, Argyris Oikonomou, Grigoris Velegkas, and Mingfei Zhao. 2021. An Efficient ϵ-BIC to BIC Transformation and Its Application to Black-Box Reduction in Revenue Maximization. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, Dániel Marx (Ed.). SIAM, 1337–1356. https://doi.org/10.1137/1.9781611976465.81

Yang Cai, Argyris Oikonomou, and Mingfei Zhao. 2022. Computing simple mechanisms: Lift-and-round over marginal reduced forms. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, Stefano Leonardi and Anupam Gupta (Eds.). ACM, 704–717. https://doi.org/10.1145/3519935.3520029

Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 170–183. https://doi.org/10.1145/3055399.3055465

Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. 2013. Prior-Independent Mechanisms for Scheduling. In *Proceedings of 45th ACM Symposium on Theory of Computing (STOC)*.

Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.

Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.

Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. 2015. The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior* 91 (2015), 297–317. https://doi.org/10.1016/j.geb.2012.08.010

Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*. 579–596. https://doi.org/10.1145/2940716.2940756

Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 2022. On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms for a Unit-Demand Buyer. *SIAM J. Comput.* 51, 3 (2022), 492–548. https://doi.org/10.1137/17M1136481

Edward H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11, 1 (1971), 17–33.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2014. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2017. Strong Duality for a Multiple-Good Monopolist. *Econometrica* 85, 3 (2017), 735–767.

Constantinos Daskalakis and S. Matthew Weinberg. 2012. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, June 4-8, 2012*. 370–387. https://doi.org/10.1145/2229012.2229042

Shaddin Dughmi, Jason D. Hartline, Robert Kleinberg, and Rad Niazadeh. 2017. Bernoulli factories and black-box reductions in mechanism design. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 158–169. https://doi.org/10.1145/3055399.3055492

Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 343. https://doi.org/10.1145/3033274.3085115

Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2021. A Simple and Approximately Optimal Mechanism for a Buyer with Complements. *Oper. Res.* 69, 1 (2021), 188–206. https://doi.org/10.1287/opre.2020.2039

Michal Feldman, Ophir Friedler, and Aviad Rubinstein. 2018. 99% Revenue via Enhanced Competition. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*. 443–460. https://doi.org/10.1145/3219166.3219202

Hu Fu, Christopher Liaw, and Sikander Randhawa. 2019. The Vickrey Auction with a Single Duplicate Bidder Approximates the Optimal Revenue. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, Anna Karlin, Nicole Immorlica, and Ramesh Johari (Eds.). ACM, 419–420. https://doi.org/10.1145/3328526.3329597

Yiannis Giannakopoulos and Elias Koutsoupias. 2014. Duality and optimality of auctions for uniform distributions. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*. 259–276. https://doi.org/10.1145/2600057.2602883

Yiannis Giannakopoulos and Elias Koutsoupias. 2015. Selling Two Goods Optimally. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*. 650–662. https://doi.org/10.1007/978-3-662-47666-6_52

Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631.

Nima Haghpanah and Jason Hartline. 2015. Reverse Mechanism Design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*.

Sergiu Hart and Noam Nisan. 2013. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*.

Sergiu Hart and Noam Nisan. 2017. Approximate revenue maximization with multiple items. *J. Economic Theory* 172 (2017), 313–347. https://doi.org/10.1016/j.jet.2017.09.001

Sergiu Hart and Philip J. Reny. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics* 10, 3 (2015), 893–922.

Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. 2011. Bayesian Incentive Compatibility via Matchings. In *the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Jason D. Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*. 225–234.

Xinye Li and Andrew Chi-Chih Yao. 2013. On Revenue Maximization for Selling Multiple Independently Distributed Items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.

Siqi Liu and Christos-Alexandros Psomas. 2018. On the Competition Complexity of Dynamic Mechanism Design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*. 2008–2025. https://doi.org/10.1137/1.9781611975031.131

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

Gregory Pavlov. 2011. Optimal Mechanism for Selling Two Goods. *The B.E. Journal of Theoretical Economics* 11, 3 (2011).

Alexandros Psomas, Ariel Schvartzman, and S. Matthew Weinberg. 2019. Smoothed Analysis of Multi-Item Auctions with Correlated Values. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*. 417–418. https://doi.org/10.1145/3328526.3329563

Alexandros Psomas, Ariel Schvartzman, and S. Matthew Weinberg. 2022. On Infinite Separations Between Simple and Optimal Mechanisms. In *NeurIPS*. http://papers.nips.cc/paper_files/paper/2022/hash/1eeacdf8770e6dd5164cdeec8bcfa8cc-Abstract-Conference.html

Jean-Charles Rochet and Philippe Chone. 1998. Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66, 4 (1998), 783–826.

Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. 2012. Supply-limiting mechanisms. In *13th ACM Conference on Electronic Commerce (EC)*.

Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (2002), 236–259. https://doi.org/10.1145/506147.506153

Aviad Rubinstein and S. Matthew Weinberg. 2015. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 377–394. https://doi.org/10.1145/2764468.2764510

Daniel Dominic Sleator and Robert Endre Tarjan. 1985. Amortized Efficiency of List Update and Paging Rules. *Commun. ACM* 28, 2 (1985), 202–208. https://doi.org/10.1145/2786.2793

John Thanassoulis. 2004. Haggling over substitutes. *Journal of Economic Theory* 117 (2004), 217–245.

William Vickrey. 1961. Counterspeculations, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 1 (1961), 8–37.

S. Matthew Weinberg and Zixin Zhou. 2022. Optimal Multi-Dimensional Mechanisms are not Locally-Implementable. In *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, David M. Pennock, Ilya Segal, and Sven Seuken (Eds.). ACM, 875–896. https://doi.org/10.1145/3490486.3538334

Andrew Chi-Chih Yao. 2017. Dominant-Strategy versus Bayesian Multi-item Auctions: Maximum Revenue Determination and Comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin (Eds.). ACM, 3–20. https://doi.org/10.1145/3033274.3085120

Andrew Chi-Chih Yao. 2018. An Incentive Analysis of some Bitcoin Fee Designs. *CoRR* abs/1811.02351 (2018). arXiv:1811.02351 http://arxiv.org/abs/1811.02351

Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

A Deferred proofs

Here, we provide complete proofs from Section 3, our main result. The remaining deferred proofs can be found in the full appendix at https://arxiv.org/abs/2403.03937.

A.1 Proofs from Section 3

Lemma 3.1. *$\textsc{SRev}_{n'}(\mathcal{ER}^m_{\leq T}) = m \cdot T \cdot \left(1 - (1 - 1/T)^{n'}\right)$. One mechanism achieving this sells each item separately with a second-price auction at reserve T.*[17]

Proof. Observe that $\mathcal{ER}_{<T}$ is a regular distribution, with

$$\bar{\varphi}_j(v_{ij}) = v_{ij} - \frac{1 - (1 - 1/v_{ij})}{1/v_{ij}^2} = 0 \qquad \forall v_{ij} < T,$$
$$\bar{\varphi}_j(T) = T - 0 = T,$$

so the optimal auction for each item allocates the item to a bidder with value T for price T. □

Corollary A.1. *Suppose $T \geq \sqrt{mn}$ and $a \geq b$. Let $v \in [1, T]^m$ and suppose there exists some item with value T. Define $H := \{j : v_j = T\}$ and $L := \{j : v_j \geq mn/T\} \setminus H$. A bidder with type v prefers the menu option (H, L) over any other option in the menu (including not receiving any items).*

Proof. By Lemma 3.9, a bidder with such a type prefers (H, L) over any other option that allocates an item with some positive probability. It remains to show that the utility of such a bidder

[17] In fact, a second-price auction with any reserve $\leq T$ achieves this.

for (H, L) is non-negative:

$$\begin{aligned}
&a\sum_{j\in H} v_j + b\sum_{j\in L} v_j - \left(|H|\,aT + |L|\,b\frac{mn}{T} - (|H|-1)b\left(T - \frac{mn}{T}\right)\right) \\
&\qquad = b\sum_{j\in L}\left(v_j - \frac{mn}{T}\right) + (|H|-1)b\left(T - \frac{mn}{T}\right) && (v_j = T \text{ for all } j\in H) \\
&\qquad \geq (|H|-1)b\left(T - \frac{mn}{T}\right) && (v_j \geq mn/T \text{ for all } j \in L) \\
&\qquad \geq 0. && (T^2 \geq mn)
\end{aligned}$$

□

Corollary A.2. *Suppose $T \geq \sqrt{mn}$ and $a \geq b$. Let $v \in [1, T]^m$. Let $j^* \in \arg\max_j v_j$ and suppose $v_{j^*} < T$. Define $L := \{j : v_j \geq mn/T\} \setminus \{j^*\}$. A bidder with type v prefers the menu option $(\{j^*\}, L)$ over any other option in the menu (including not receiving any items) if and only if $av_{j^*} + b\sum_{j\in L} v_j \geq aT + |L|\,b\frac{mn}{T}$.*

Proof. By Lemma 3.9, a bidder with such a type prefers $(\{j^*\}, L)$ over any other option that allocates an item with some positive probability. To conclude, note that the utility for $(\{j^*\}, L)$ is non-negative if and only if the inequality in the lemma statement holds. □

Lemma 3.11. *If $T \geq \sqrt{mn}$, then $\frac{b}{a} \leq \frac{n}{T}e^{-\frac{n}{T}}\left(1 - e^{-\frac{n}{T}}\right)^{-1}$.*

Proof. We have

$$\begin{aligned}
\frac{b}{a} &= \frac{\mathbb{P}_{v_{-i}}\left[\nexists \text{ high}\right](1 - (1 - \mathbb{P}[\text{low} \mid \text{not high}])^n)}{n\mathbb{P}[\text{low} \mid \text{not high}]} \cdot \frac{n\mathbb{P}[\text{high}]}{1 - (1 - \mathbb{P}[\text{high}])^n} \\
&\leq \frac{\mathbb{P}_{v_{-i}}\left[\nexists \text{ high}\right] \cdot n\mathbb{P}[\text{low} \mid \text{not high}]}{n\mathbb{P}[\text{low} \mid \text{not high}]} \cdot \frac{n\mathbb{P}[\text{high}]}{1 - (1 - \mathbb{P}[\text{high}])^n} && \text{(union bound)} \\
&= \frac{n\mathbb{P}_{v_{-i}}\left[\nexists \text{ high}\right]\mathbb{P}[\text{high}]}{1 - (1 - \mathbb{P}[\text{high}])^n} \\
&= \frac{n\left(1 - \frac{1}{T} - \sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)^{n-1}\left(\frac{1}{T} + \sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)}{1 - \left(1 - \frac{1}{T} - \sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)^n} \\
&\leq \frac{n\left(1 - \frac{1}{T}\right)^n}{T\left(1 - \left(1 - \frac{1}{T}\right)^n\right)} && \text{(replace summation with 0)} \\
&\leq \frac{\frac{n}{T}e^{-\frac{n}{T}}}{1 - e^{-\frac{n}{T}}}.
\end{aligned}$$

The second inequality follows from the fact that $\left(1 - \frac{1}{T} - x\right)^{n-1}\left(\frac{1}{T} + x\right) / \left(1 - \left(1 - \frac{1}{T} - x\right)^n\right)$ is decreasing in x for $x \in [0, 1 - 1/T]$ (note that $1 - \frac{1}{T} - \sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell \leq 1$ because the left-hand expression is a probability, so indeed $\sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell \leq 1 - \frac{1}{T}$):

$$\frac{\partial}{\partial x}\frac{\left(1 - \frac{1}{T} - x\right)^{n-1}\left(\frac{1}{T} + x\right)}{1 - \left(1 - \frac{1}{T} - x\right)^n} = \frac{\left(1 - \frac{1}{T} - x\right)^{n-2}\left(\left(1 - \left(1 - \frac{1}{T} - x\right)^n\right) - n\left(\frac{1}{T} + x\right)\right)}{\left(1 - \left(1 - \frac{1}{T} - x\right)^n\right)^2} \leq 0.$$

(union bound; $x \in [0, 1 - 1/T]$)

□

Lemma 3.13. *If bidders report their values truthfully in the Less-Naive Auction, then* $\textsc{Rev}_n^{NSN}(\mathcal{ER}^m_{\leq T})$ *exceeds* $\textsc{Rev}_n^{LNA}(\mathcal{ER}^m_{\leq T})$ *by at least*

$$(b-b_0)mn^2\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)\left(1-\left(1-\tfrac{1}{T}\right)^m-\tfrac{m}{T}\left(1-\tfrac{1}{T}\right)^{m-1}\right)+\tfrac{bm^2(m-1)n^2}{T}\sum_{\ell=1}^{m-1}\tbinom{m-2}{\ell-1}q_\ell.$$

Proof. By Corollaries A.1 and A.2, the revenue *per bidder* of the Not-So-Naive Auction is

$$\begin{aligned}
&\sum_{j^*\in[m]}\sum_{L\subseteq[m]\setminus\{j^*\}}\left(aT+|L|\,b\tfrac{mn}{T}\right)\left(\tfrac{1}{T}\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)^{|L|}\left(1-\tfrac{T}{mn}\right)^{m-1-|L|}+q_{|L|}\right)\\
&+\sum_{\substack{H\subseteq[m]:\\|H|\geq 2}}\sum_{L\subseteq[m]\setminus H}\left(|H|\,aT+|L|\,b\tfrac{mn}{T}-(|H|-1)b\left(T-\tfrac{mn}{T}\right)\right)\tfrac{1}{T^{|H|}}\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)^{|L|}\left(1-\tfrac{T}{mn}\right)^{m-|H|-|L|}\\
&=\sum_{k=1}^{m}\sum_{\ell=0}^{m-k}\tbinom{m}{k}\tbinom{m-k}{\ell}\left(kaT+\ell b\tfrac{mn}{T}-(k-1)b\left(T-\tfrac{mn}{T}\right)\right)\tfrac{1}{T^k}\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)^{\ell}\left(1-\tfrac{T}{mn}\right)^{m-k-\ell}\\
&\quad+m\sum_{\ell=1}^{m-1}\tbinom{m-1}{\ell}\left(aT+\ell b\tfrac{mn}{T}\right)q_\ell\\
&=am+bmn\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)\left(1-\left(1-\tfrac{1}{T}\right)^m-\tfrac{m}{T}\left(1-\tfrac{1}{T}\right)^{m-1}\right)+m\sum_{\ell=1}^{m-1}\tbinom{m-1}{\ell}\left(aT+\ell b\tfrac{mn}{T}\right)q_\ell.
\end{aligned}$$

Recall that the revenue extracted by the Less-Naive Auction from selling item j to bidder i is

$$\begin{aligned}
a_0T\cdot\mathbb{1}\left(v_{ij}=T,\max_{j'<j}v_{ij'}<T\right)+a_0\left(T-\frac{b_0}{a_0}\left(T-\frac{mn}{T}\right)\right)\cdot\mathbb{1}\left(v_{ij}=T,\max_{j'<j}v_{ij'}=T\right)\\
+\frac{b_0mn}{T}\cdot\mathbb{1}\left(v_{ij}\in[\tfrac{mn}{T},T),\max_{j'\neq j}v_{ij'}=T\right).
\end{aligned}$$

Taking the expectation over the randomness of bidder i's type and summing over all items yields that the revenue extracted by the Less-Naive Auction *per bidder* is

$$a_0m+b_0mn\left(\frac{T}{mn}-\frac{1}{T}\right)\left(1-\left(1-\frac{1}{T}\right)^m-\frac{m}{T}\left(1-\frac{1}{T}\right)^{m-1}\right).$$

Now, recall that a_0 is the interim probability of winning an item when bidding T under the allocation rule that allocates items uniformly at random to bidders with value T, so

$$a_0=\mathop{\mathbb{E}}_{v_{-i}}\left[\frac{1}{1+\sum_{k\neq i}\mathbb{1}(v_{kj}=T)}\right]=\sum_{k=0}^{n-1}\frac{\binom{n-1}{k}}{k+1}\frac{1}{T^k}\left(1-\frac{1}{T}\right)^{m-k-1}=\frac{T}{n}\left(1-\left(1-\frac{1}{T}\right)^n\right).$$

Note how a_0 compares to a.

$$a\left(1+T\sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right) = \frac{\left(1-\left(1-\frac{1}{T}-\sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)^n\right)}{n\left(\frac{1}{T}+\sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)}\left(1+T\sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right) \quad \text{(definition of } a\text{)}$$

$$= \frac{T}{n}\left(1-\left(1-\frac{1}{T}-\sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)^n\right)$$

$$\geq \frac{T}{n}\left(1-\left(1-\frac{1}{T}\right)^n\right)$$

$$= a_0 \quad \text{(definition of } a_0\text{)}$$

Thus, the revenue *per bidder* of the Not-So-Naive Auction exceeds that of the Less-Naive Auction by

$$(a-a_0)m + (b-b_0)mn\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)\left(1-\left(1-\tfrac{1}{T}\right)^m-\tfrac{m}{T}\left(1-\tfrac{1}{T}\right)^{m-1}\right) + m\sum_{\ell=1}^{m-1}\tbinom{m-1}{\ell}\left(aT+\ell b\tfrac{mn}{T}\right)q_\ell$$

$$= \left(a\left(1+T\sum_{\ell=1}^{m-1}\tbinom{m-1}{\ell}q_\ell\right)-a_0\right)m + (b-b_0)mn\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)\left(1-\left(1-\tfrac{1}{T}\right)^m-\tfrac{m}{T}\left(1-\tfrac{1}{T}\right)^{m-1}\right)$$

$$+ \tfrac{bm^2(m-1)n}{T}\sum_{\ell=1}^{m-1}\tbinom{m-2}{\ell-1}q_\ell$$

$$\geq (b-b_0)mn\left(\tfrac{T}{mn}-\tfrac{1}{T}\right)\left(1-\left(1-\tfrac{1}{T}\right)^m-\tfrac{m}{T}\left(1-\tfrac{1}{T}\right)^{m-1}\right) + \tfrac{bm^2(m-1)n}{T}\sum_{\ell=1}^{m-1}\tbinom{m-2}{\ell-1}q_\ell.$$

Summing over all bidders yields the lemma. □

Lemma A.3. *If $T \geq \sqrt{mn}$, then for all $\ell \in [m-1]$,*

$$q_\ell \leq \left(\frac{T}{mn}-\frac{1}{T}\right)^\ell\left(1-\frac{T}{mn}\right)^{m-\ell-1}\frac{\ell b}{Ta}.$$

Proof. Observe that

$$q_\ell = \mathbb{P}_v\left[v_1 = \max_j v_j < T, \min_{j\in[\ell+1]} v_j \geq \tfrac{mn}{T}, \max_{j\notin[\ell+1]} v_j < \tfrac{mn}{T}, av_1 + b\sum_{j=2}^{\ell+1} v_j \geq aT + \ell b\tfrac{mn}{T}\right]$$

$$= \tfrac{\left(\frac{T}{mn}-\frac{1}{T}\right)^{\ell+1}\left(1-\frac{T}{mn}\right)^{m-\ell-1}}{\ell+1}\mathbb{P}\left[av_1 + b\sum_{j=2}^{\ell+1} v_j \geq aT + \ell b\tfrac{mn}{T}\,\middle|\, v_1 = \max_j v_j < T, \min_{j\in[\ell+1]} v_j \geq \tfrac{mn}{T} > \max_{j\notin[\ell+1]} v_j\right]$$

$$\leq \tfrac{\left(\frac{T}{mn}-\frac{1}{T}\right)^{\ell+1}\left(1-\frac{T}{mn}\right)^{m-\ell-1}}{\ell+1}\mathbb{P}\left[v_1 \geq \tfrac{aT+\ell b\frac{mn}{T}}{a+\ell b}\,\middle|\, v_1 = \max_j v_j < T, \min_{j\in[\ell+1]} v_j \geq \tfrac{mn}{T} > \max_{j\notin[\ell+1]} v_j\right]$$

$$= \tfrac{\left(\frac{T}{mn}-\frac{1}{T}\right)^{\ell+1}\left(1-\frac{T}{mn}\right)^{m-\ell-1}}{\ell+1}\left(1-\left(1-\frac{\frac{1}{T-\frac{\ell b}{a+\ell b}\left(T-\frac{mn}{T}\right)}-\frac{1}{T}}{\frac{T}{mn}-\frac{1}{T}}\right)^{\ell+1}\right)$$

$$\leq \left(\tfrac{T}{mn}-\tfrac{1}{T}\right)^\ell\left(1-\tfrac{T}{mn}\right)^{m-\ell-1}\left(\frac{1}{T-\frac{\ell b}{a+\ell b}\left(T-\frac{mn}{T}\right)}-\frac{1}{T}\right) \quad \text{(union bound)}$$

$$\leq \left(\tfrac{T}{mn}-\tfrac{1}{T}\right)^\ell\left(1-\tfrac{T}{mn}\right)^{m-\ell-1}\tfrac{\ell b}{Ta}.$$

□

Lemma A.4. *If $T = \lambda\sqrt{mn}$ for some constant $\lambda > 1$, then the probability of a "high" bidder with all values below T can be upper bounded as*

$$\sum_{\ell=1}^{m-1} \binom{m-1}{\ell} q_\ell \leq \left(1 - \frac{1}{\lambda^2}\right) \frac{b(m-1)}{amn} \left(1 - \frac{1}{T}\right)^{m-2}.$$

Proof.

$$\begin{aligned}
\sum_{\ell=1}^{m-1} \binom{m-1}{\ell} q_\ell &\leq \frac{b(m-1)}{Ta} \sum_{\ell=1}^{m-1} \binom{m-2}{\ell-1} \left(\frac{T}{mn} - \frac{1}{T}\right)^{\ell} \left(1 - \frac{T}{mn}\right)^{m-\ell-1} && \text{(Lemma A.3)} \\
&= \frac{b(m-1)}{Ta} \left(\frac{T}{mn} - \frac{1}{T}\right) \sum_{\ell=0}^{m-2} \binom{m-2}{\ell} \left(\frac{T}{mn} - \frac{1}{T}\right)^{\ell} \left(1 - \frac{T}{mn}\right)^{m-\ell-2} \\
&= \frac{b(m-1)}{Ta} \left(\frac{T}{mn} - \frac{1}{T}\right) \left(1 - \frac{1}{T}\right)^{m-2} \\
&= \left(1 - \frac{1}{\lambda^2}\right) \frac{b(m-1)}{amn} \left(1 - \frac{1}{T}\right)^{m-2}.
\end{aligned}$$

□

Lemma A.5. *If $T = \lambda\sqrt{mn}$ for some constant $\lambda > 1$, then the probability of a "low" bidder with all values below T can be upper bounded as*

$$(m-1) \sum_{\ell=1}^{m-1} \binom{m-2}{\ell-1} q_\ell \leq \left(1 - \frac{1}{\lambda^2}\right) \frac{b(m-1)}{amn} \left(1 - \frac{1}{T}\right)^{m-3} \left(1 - \frac{1}{\lambda\sqrt{mn}} + \left(\lambda - \frac{1}{\lambda}\right) \frac{m-2}{\sqrt{mn}}\right).$$

Proof.

$$\begin{aligned}
(m-1) \sum_{\ell=1}^{m-1} \binom{m-2}{\ell-1} q_\ell &\leq \frac{b(m-1)}{Ta} \sum_{\ell=1}^{m-1} \ell \binom{m-2}{\ell-1} \left(\frac{T}{mn} - \frac{1}{T}\right)^{\ell} \left(1 - \frac{T}{mn}\right)^{m-\ell-1} \\
&= \frac{b(m-1)}{Ta} \left(\frac{T}{mn} - \frac{1}{T}\right) \left(1 - \frac{1}{T}\right)^{m-2} \\
&\quad + \frac{b(m-1)(m-2)}{Ta} \sum_{\ell=2}^{m-1} \binom{m-3}{\ell-2} \left(\frac{T}{mn} - \frac{1}{T}\right)^{\ell} \left(1 - \frac{T}{mn}\right)^{m-\ell-1} \\
&= \frac{b(m-1)}{Ta} \left(\frac{T}{mn} - \frac{1}{T}\right) \left(1 - \frac{1}{T}\right)^{m-3} \left(1 - \frac{1}{T} + (m-2)\left(\frac{T}{mn} - \frac{1}{T}\right)\right) \\
&= \left(1 - \frac{1}{\lambda^2}\right) \frac{b(m-1)}{amn} \left(1 - \frac{1}{T}\right)^{m-3} \left(1 - \frac{1}{\lambda\sqrt{mn}} + \left(\lambda - \frac{1}{\lambda}\right) \frac{m-2}{\sqrt{mn}}\right).
\end{aligned}$$

□

Lemma 3.15. *If $T = \lambda\sqrt{mn}$ for some constant $\lambda > 1$, then $b = \Omega\left(\left(1 - \frac{1}{T}\right)^n\right)$.*

Proof. Note that

$$b := \frac{\mathbb{P}_{v_{-i}}\left[\nexists \text{ high}\right] \left(1 - (1 - \mathbb{P}\left[\text{low} \mid \text{not high}\right])^n\right)}{n\mathbb{P}\left[\text{low} \mid \text{not high}\right]} = \frac{\mathbb{P}_{v}\left[\nexists \text{ high}\right] \left(1 - (1 - \mathbb{P}\left[\text{low} \mid \text{not high}\right])^n\right)}{n\mathbb{P}\left[\text{low}\right]}$$

We bound each part of b:

$$\begin{aligned}
\mathbb{P}_v\left[\nexists \text{ high}\right] &= \left(1 - \frac{1}{T} - \sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell\right)^n \\
&\geq \left(1 - \frac{1}{T} - \frac{O(1)}{n}\right)^n && \text{(Lemmas 3.11 and A.4)} \\
&= \Omega\left(\left(1 - \frac{1}{T}\right)^n\right), && (T = \lambda\sqrt{mn},\ m \leq n)
\end{aligned}$$

$$\begin{aligned}
\mathbb{P}\left[\text{low}\right] &= \left(1 - \left(1 - \frac{1}{T}\right)^{m-1}\right)\left(\frac{T}{mn} - \frac{1}{T}\right) + (m-1)\sum_{\ell=1}^{m-1}\binom{m-2}{\ell-1}q_\ell \\
&\leq \frac{m-1}{T}\left(\frac{T}{mn} - \frac{1}{T}\right) + \frac{O(1)}{n} && \text{(union bound; Lemmas 3.11 and A.5)} \\
&= \left(1 - \frac{1}{\lambda^2}\right)\frac{m-1}{mn} + \frac{O(1)}{n} && (T = \lambda\sqrt{mn}) \\
&= O\left(\frac{1}{n}\right),
\end{aligned}$$

$$\begin{aligned}
\mathbb{P}\left[\text{low} \mid \text{not high}\right] &= \frac{\left(1 - \left(1 - \frac{1}{T}\right)^{m-1}\right)\left(\frac{T}{mn} - \frac{1}{T}\right) + (m-1)\sum_{\ell=0}^{m-2}\binom{m-2}{\ell}q_{\ell+1}}{1 - \frac{1}{T} - \sum_{\ell=1}^{m-1}\binom{m-1}{\ell}q_\ell} \\
&\geq \frac{\left(1 - \left(1 - \frac{1}{T}\right)^{m-1}\right)\left(\frac{T}{mn} - \frac{1}{T}\right)}{1 - \frac{1}{T}} \\
&\geq \frac{\left(\frac{m-1}{T} - \frac{\binom{m-1}{2}}{T^2}\right)\left(\frac{T}{mn} - \frac{1}{T}\right)}{1 - \frac{1}{T}} && \text{(inclusion-exclusion)} \\
&= \frac{\left(1 - \frac{1}{\lambda^2}\right)\frac{m-1}{\sqrt{mn}}\left(\frac{1}{\sqrt{mn}} - \frac{1}{2\lambda n}\right)}{1 - \frac{1}{\lambda\sqrt{mn}}} && (T = \lambda\sqrt{mn}) \\
&= \Omega\left(\frac{1}{n}\right).
\end{aligned}$$

Thus,

$$\begin{aligned}
b &= \frac{\mathbb{P}_v\left[\nexists \text{ high}\right]\left(1 - \left(1 - \mathbb{P}\left[\text{low} \mid \text{not high}\right]\right)^n\right)}{n\mathbb{P}\left[\text{low}\right]} \\
&= \frac{\Omega\left(\left(1 - \frac{1}{T}\right)^n\right)\left(1 - \left(1 - \Omega\left(\frac{1}{n}\right)\right)^n\right)}{nO\left(\frac{1}{n}\right)} \\
&= \Omega\left(\left(1 - \frac{1}{T}\right)^n\right).
\end{aligned}$$

□

Theorem 3.16. *If $T = \lambda\sqrt{mn}$ for some constant $\lambda > 1$, then $\textsc{Comp}_{\mathcal{ER}^m_{\leq T}}(n) = \Omega(\sqrt{mn})$.*

Proof. Recall that by Corollary 3.14 and Lemma 3.3, SRev_{n+c} exceeds the revenue of the Not-So-Naive Auction only if

$$c \geq \frac{bn^2\left(\frac{T}{mn}-\frac{1}{T}\right)\left(1-\left(1-\frac{1}{T}\right)^m-\frac{m}{T}\left(1-\frac{1}{T}\right)^{m-1}\right)}{\left(1-\frac{1}{T}\right)^n}.$$

Note that if $m \in [n/4, n]$, then

$$\begin{aligned}
c &\geq \frac{bn^2\left(\frac{T}{mn}-\frac{1}{T}\right)\left(1-e^{-\frac{m}{T}}-\frac{m}{T}e^{-\frac{m-1}{T}}\right)}{\left(1-\frac{1}{T}\right)^n} \\
&\geq \frac{b\frac{n^{3/2}}{\sqrt{m}}\left(\lambda-\frac{1}{\lambda}\right)\left(1-e^{-\frac{1}{\lambda}\sqrt{\frac{m}{n}}}-\frac{1}{\lambda}\sqrt{\frac{m}{n}}e^{-\frac{1}{2\lambda}\sqrt{\frac{m}{n}}}\right)}{\left(1-\frac{1}{T}\right)^n} && (T=\lambda\sqrt{mn}) \\
&\geq \frac{bn\left(\lambda-\frac{1}{\lambda}\right)\left(1-e^{-\frac{1}{2\lambda}}-\frac{1}{2\lambda}e^{-\frac{1}{4\lambda}}\right)}{\left(1-\frac{1}{T}\right)^n} && (1-e^{-x}-xe^{-x/2} \text{ is increasing in } x,\ m\in[n/4,n]) \\
&= \frac{\Omega\left(\left(1-\frac{1}{T}\right)^n\right)n}{\left(1-\frac{1}{T}\right)^n} && \text{(Lemma 3.15)} \\
&= \Omega(\sqrt{mn}). && (m\in[n/4,n])
\end{aligned}$$

If $m \leq n/4$, then

$$\begin{aligned}
c &\geq \frac{bn^2\left(\frac{T}{mn}-\frac{1}{T}\right)\frac{m(m-1)}{2T^2}\left(1-\frac{m-2}{T}\right)}{\left(1-\frac{1}{T}\right)^n} && \text{(inclusion-exclusion)} \\
&\geq \frac{b\frac{(m-1)\sqrt{mn}}{2\lambda m}\left(1-\frac{1}{\lambda^2}\right)\left(1-\frac{1}{\lambda}\sqrt{\frac{m}{n}}\right)}{\left(1-\frac{1}{T}\right)^n} && (T=\lambda\sqrt{mn}) \\
&= \frac{b\frac{(m-1)\sqrt{mn}}{2\lambda m}\left(1-\frac{1}{\lambda^2}\right)\left(1-\frac{1}{2\lambda}\right)}{\left(1-\frac{1}{T}\right)^n} && (m\leq n/4) \\
&= \frac{\Omega\left(\left(1-\frac{1}{T}\right)^n\right)\sqrt{mn}}{\left(1-\frac{1}{T}\right)^n} && \text{(Lemma 3.15)} \\
&= \Omega(\sqrt{mn}).
\end{aligned}$$

□