

Truthful Fair Division under Stochastic Valuations

DANIEL HALPERN, Google Research

ALEXANDROS PSOMAS, Purdue University and Google Research

SHIRLEY ZHANG, Harvard University

We study no-money mechanisms for allocating indivisible items to strategic agents with additive preferences under a stochastic model. In this model, items' values are drawn from an underlying distribution and mechanisms are evaluated with respect to this draw (e.g., in expectation, or with high probability). Motivated by worst-case impossibilities which show that truthfulness severely restricts fairness and efficiency, we ask whether truthful mechanisms continue to perform poorly on random instances.

We first focus on dominant-strategy incentive compatible (DSIC) mechanisms. For two agents, we obtain a tight picture. Specifically, we show that there exists a distribution under which no DSIC mechanism achieves an expected welfare approximation better than $\frac{2+\sqrt{2}}{4} \approx 0.854$, and we give a DSIC mechanism that matches this bound for all distributions simultaneously. We further show that, for every distribution, there exists a DSIC mechanism that is envy-free with high probability and obtains the same welfare. A key ingredient is a new, tight connection between welfare guarantees of a family of DSIC, no-money mechanisms and i.i.d. prophet inequalities. This connection allows us to generalize to n agents; in particular, we obtain a DSIC mechanism that achieves a ≈ 0.745 approximation to welfare, and another DSIC mechanism achieving a $1/2$ -approximation welfare that is envy-free with high probability.

We then turn to Bayesian incentive compatibility (BIC). Under i.i.d. valuations, we show that BIC comes at essentially no cost: we design a prior-independent BIC mechanism that achieves a $(1 - \epsilon)$ -approximation to the optimal welfare, while being envy-free with high probability. Under independent but non-identical priors, we obtain BIC mechanisms that are $(1 - \epsilon)$ -approximately Pareto efficient and envy-free with high probability.

1 INTRODUCTION

We study the allocation of indivisible items among strategic agents with additive preferences in the absence of monetary transfers. In this classic setting, the gold standard is mechanisms that are simultaneously fair, economically efficient (e.g., Pareto optimal), and strategically robust (e.g., truthful reporting is a dominant strategy for every agent).

While the literature has uncovered a plethora of natural rules that are both fair and efficient, the same cannot be said for combinations with strategic robustness. Impossibility results paint a bleak picture. For instance, deterministic truthful mechanisms cannot guarantee envy-freeness up to one good (EF1) [Amanatidis et al., 2017], a canonical fairness notion for indivisible goods [Lipton et al., 2004]. This can be extended to EF_{-v}^{+u} for any fixed choice of u and v [Bu and Tao, 2024].¹ Moreover, the only mechanism that is both Pareto optimal and truthful is a serial dictatorship [Schummer, 1996], arguably the least fair mechanism imaginable. In short, mechanisms that are simultaneously fair, efficient, and truthful do not exist. As a result, much of the fair division literature has largely set aside issues of incentives and instead focused purely on fairness and efficiency.

However, a core tenet in economics is that agents *are* strategic. If agents can easily manipulate the outcome, how can we trust the promised fairness guarantees of the resulting allocations? This, on the surface, seems to create a fundamental dilemma for the field.

We posit that all hope may not be lost. A closer look at the aforementioned negative results reveals a commonality: they all hold in the *worst case*, relying on carefully constructed pathological instances. To understand the performance beyond the worst case, much of the recent fair division literature has adopted distributional assumptions. For example, in asymptotic fair division it is well understood that fairness guarantees that are unattainable in the worst case (e.g., envy-freeness) are often achievable with high probability on random instances. What is far less understood is whether the same phenomenon holds for incentives. Do truthful mechanisms — that can't meaningfully balance fairness and efficiency under worst-case analysis — perform just as poorly on random instances?

In this paper, we study truthful no-money mechanisms under stochastic valuations. Our main question is whether such mechanisms can provide meaningful fairness and efficiency guarantees with high probability. Put differently, is it possible that by going beyond worst-case analysis we can actually achieve the fairness–efficiency–truthfulness trifecta on most fair division instances?

1.1 Our Contribution

A subtlety arises when considering incentives under stochastic valuations, which is that there is no longer a single canonical notion of truthfulness. The definition depends on whether an agent considers the consequences of their action ex-post or in expectation over the other agents' values. In this paper, we consider both *Dominant Strategy Incentive Compatibility* (DSIC) and *Bayesian Incentive Compatibility* (BIC) to more deeply understand the trade-offs imposed by these constraints.

We begin by focusing on DSIC mechanisms. In our stochastic model, agents' valuations are drawn from a known distribution \mathcal{D} , and the mechanism is evaluated with respect to the random draw of valuations (e.g., in expectation or with high probability).

We start with the two-agent case. In this setting, we can leverage a known characterization of DSIC no-money mechanisms and restrict attention to the class of *picking–exchange* mechanisms [Amanatidis et al., 2017]. We give a complete picture. First, we show that there exists a distribution \mathcal{D}

¹An allocation satisfies EF_{-v}^{+u} if no agent envies another agent after adding u items to their bundle and removing v items from the other agent's bundle.

such that no DSIC mechanism can, in expectation, achieve an α -approximation to the optimal welfare for any $\alpha > \frac{2+\sqrt{2}}{4} \approx 0.854$ (Theorem 3.1). Second, we show that there exists a mechanism that achieves, in expectation, a $\frac{2+\sqrt{2}}{4}$ -approximation to the optimal welfare, for all underlying distributions *simultaneously* (Proposition 3.3). Furthermore, we identify a pair of mechanisms, such that, given a distribution \mathcal{D} , at least one of the two mechanisms will achieve a $\frac{2+\sqrt{2}}{4}$ -approximation to the optimal welfare, as well as be envy-free with high probability (Theorem 3.7).

A notable aspect of our positive results is that they reveal a tight and, to the best of our knowledge, previously unnoticed connection to the rich literature on prophet inequalities. Concretely, in Proposition 3.8 we show that the welfare approximation ratio of a natural class of *picking-sequence* mechanisms (which are DSIC and do not use monetary transfers) coincides with the optimal guarantees in the i.i.d. prophet inequality problem. This correspondence explains the appearance of the constant $\frac{2+\sqrt{2}}{4}$ (which is the best possible ratio in a specific instance of the i.i.d. prophet inequality problem) in the two-agent case, and furthermore provides a principled route for generalizing our positive results to $n > 2$ agents. Specifically, for $n > 2$ agents, we prove that there exists a DSIC mechanism that achieves an ≈ 0.745 approximation to welfare (Corollary 3.9) and that a simple DSIC mechanism gives a $1/2$ approximation to the optimal expected welfare, as well as is envy-free with high probability (Theorem 3.10).

Finally, we turn to BIC mechanisms. Under i.i.d. stochastic valuations, it is known that it is possible to simultaneously achieve optimal welfare and envy-freeness with high probability [Dickerson et al., 2014]. We prove that BIC comes at essentially no cost (!); specifically, we give a BIC mechanism that guarantees a $(1 - \epsilon)$ -approximation to the optimal welfare and is envy-free with high probability (Theorem 4.1). Our mechanism requires no knowledge of the underlying distribution (it just assumes that an underlying distribution exists). This result can also be extended to independent but non-identical distributions per agent. Under the same conditions as [Bai and Gözl, 2022] (see related work), when the underlying distributions are known, we can get a mechanism that is BIC, $(1 - \epsilon)$ -approximately Pareto efficient and envy-free with high probability.

1.2 Related Work

Incentives in Fair Division. Numerous impossibility results in fair division say that “fairness” and “truthfulness” are incompatible for various fairness notions (where “truthfulness” is used as a synonym to DSIC) [Amanatidis et al., 2017, 2016, Bu and Tao, 2024, Lipton et al., 2004, Markakis and Psomas, 2011]. Furthermore, (Pareto) efficiency and truthfulness are known to imply dictatorships [Pápai, 2000, 2001, Schummer, 1996]. In fact, a truthful, non-bossy, and neutral mechanism must be a serial-quota mechanism [Babaioff and Manaker Morag, 2025]. Therefore, much of the work in truthful fair division attempts to achieve fairness, efficiency, and incentive compatibility by restricting the family of instances considered, e.g., dichotomous valuations [Amanatidis et al., 2021, Babaioff et al., 2021, Barman and Verma, 2022, Bu and Tao, 2024, Halpern et al., 2020], Leontief valuations [Friedman et al., 2014, Ghodsi et al., 2011, Parkes et al., 2015], or a small number (e.g., $n = 2$) agents [Bu and Tao, 2024]. Some works use “money burning” i.e., leave some resources unallocated, to recover incentive compatibility in addition to fairness and efficiency [Abebe et al., 2020, Cole et al., 2013, Friedman et al., 2019].

In all aforementioned works, “truthfulness” is a synonym for “dominant strategy incentive compatibility (DSIC).” Some recent works in fair division explore whether the aforementioned impossibility results can be bypassed by relaxing DSIC to truthful in expectation [Mossel and Tamuz, 2010], Bayesian incentive compatibility [Gkatzelis et al., 2024], not obviously manipulable [Ortega and

Segal-Halevi, 2022, Psomas and Verma, 2022], or partial strategyproofness [Mennle and Seuken, 2021]. [Hartman et al., 2025] explores how manipulable different mechanisms are. Finally, another recent, related line of work explores whether the Nash equilibria of non-truthful mechanisms (e.g., the round robin procedure) satisfy fairness guarantees [Amanatidis et al., 2024, 2023].

Asymptotic Fair Division. Dickerson et al. [2014] initiate the study of fair division under stochastic preferences, and proved that the welfare maximizing allocation is envy-free with high probability when the number of goods $m \in \Omega(n \log n)$ (where n is the number of agents) and items' values are drawn i.i.d. from a fixed distribution with bounded PDF. Manurangsi and Suksompong [2021] improve this result, and show that an envy-free allocation exists with high probability when $m \in \Omega(n \log n / \log \log n)$. Bai and Gözl [2022] study the existence of envy-free allocations for the case of independent but non-identically distributed agents' values. They show that under mild restrictions on the distribution (in particular, the distribution PDF values are both lower and upper bounded on the non-zero interval), envy-free and Pareto optimal allocations exist with high probability. Furthermore, such an allocation is achieved by maximizing *weighted* welfare, where the weights depend on the agents' specific distributions. Benadè et al. [2024a] study independent agents with non-additive stochastic valuations, and Benadè et al. [2024b] study non-independent agents in an online setting. We note that all aforementioned works treat the distributions as constants; Halpern et al. [2025] show that envy-free allocations exist with high probability even without this assumption, and even in an online setting.

Closer to our work, and to the best of our knowledge the only other work that studies incentives in the asymptotic regime, Garg et al. [2025] give a truthful in expectation mechanism that outputs envy-free allocations with high probability, even in a model where agents' valuations are correlated. In contrast, in this work, we give DSIC mechanisms (a stronger notion of incentive compatibility) that are envy-free with high probability, as well as approximately maximize welfare; alas, we require agents to be independent.

We will also make use of some results from i.i.d. prophet inequalities. We will introduce the necessary preliminaries and discuss related work when they are used in Section 3.3.

2 MODEL

A set of m indivisible items $\mathcal{M} = \{1, 2, \dots, m\}$ must be allocated among a set of n agents $\mathcal{N} = [n]$ with additive preferences. An *allocation* $A = \{A_1, A_2, \dots, A_n\}$ partitions the m items in \mathcal{M} among the n agents, where the set of items allocated to agent $i \in [n]$ is A_i (we refer to A_i as agent i 's allocation or bundle). Note that each item $j \in \mathcal{M}$ must be assigned to exactly one agent in \mathcal{N} . Each agent $i \in \mathcal{N}$ has a valuation vector $v_i \in \mathbb{R}_{\geq 0}^m$, where $v_{i,j}$ is agent i 's value for item j . The utility of agent i for an allocation A to be $u_i(A) = \sum_{j \in A_i} v_{i,j}$.

We assume that agents' values are drawn i.i.d. from a distribution \mathcal{D} , i.e. $v_{i,j} \sim \mathcal{D}$ for all $i \in \mathcal{N}$ and $j \in \mathcal{M}$. We make some minimal assumptions on \mathcal{D} , namely, that it is nonnegative, bounded, has nonzero variance, and continuous.² We call a distribution satisfying these conditions *valid*. Note that none of our results will depend on the rescaling of \mathcal{D} ; thus, it is without loss of generality to assume that \mathcal{D} is supported on $[0, 1]$.

While values v_i are private and known only to agent i , agent i can strategically report their values as a vector of bids b_i , where $b_{i,j}$ represents agent i 's bid for item j . Note that the agents are strategic

²Strictly speaking, continuous distributions are not necessary for any of our results; we simply assume it for ease of exposition. This ensures quantiles for item values are well-behaved. For discrete distributions, care must be taken to handle tie-breaking.

and therefore may misreport their bids, so b_i is not necessarily equal to v_i . A mechanism is a function $M(\cdot)$ that maps a profile of bids $\mathbf{b} = (b_1, \dots, b_n)$ to an allocation $M(\mathbf{b})$.

A mechanism is *incentive compatible* if it is robust to strategic agents. In this paper, we focus on two formalizations of incentive compatibility. First, a mechanism is *dominant strategy incentive compatible (DSIC)* if truthfully bidding their private values is a dominant strategy for every agent. Formally, a mechanism $M(\cdot)$ is DSIC if, for every agent $i \in \mathcal{N}$, every possible bid b_i , and every possible bids of other agents b_{-i} ,

$$u_i(M(v_i, b_{-i})) \geq u_i(M(b_i, b_{-i})).$$

Second, a mechanism is *Bayesian incentive compatible (BIC)* if every agent truthfully bidding their private values is a Bayesian Nash equilibrium. This means that no agent can achieve higher expected utility by bidding anything other than their true valuations, where the expectation is taken over the random values of all other agents. Formally, a mechanism $M(\cdot)$ is BIC if for every possible bid b_i ,

$$\mathbb{E}_{v_{-i} \sim \mathcal{D}^{n-1}} [u_i(M(v_i, v_{-i}))] \geq \mathbb{E}_{v_{-i} \sim \mathcal{D}^{n-1}} [u_i(M(b_i, v_{-i}))].$$

An allocation A is *envy-free (EF)* if, for every pair of agents $i_1, i_2 \in \mathcal{N}$, agent i_1 has a higher utility for their own bundle than for agent i_2 's bundle. Formally, for all $i, j \in \mathcal{N}$, $u_{i_1}(A_{i_1}) \geq u_{i_1}(A_{i_2})$.

The social welfare of an allocation A is the total utility across all agents, i.e. $\text{sw}^u(A) = \sum_{i \in \mathcal{N}} u_i(A)$. We say a mechanism M achieves an α approximation to welfare if

$$\mathbb{E}_u[\text{sw}^u(M(\mathbf{v}))] \geq \alpha \cdot \mathbb{E}_u[\max_A \text{sw}^u(A)].$$

We will say the mechanism achieves an $\alpha - o(1)$ approximation to welfare if the approximation approaches α as m grows large. Furthermore, we will say a mechanism is envy-free with high probability if the probability it outputs an envy-free allocation approaches 1 as m grows large.

3 DSIC MECHANISMS

3.1 Two Agents

We begin by studying DSIC mechanisms for the case of $n = 2$ agents. In the context of stochastic values, one might hope that the relaxation from worst-case analysis would allow for mechanisms that approximate the optimal welfare arbitrarily well. However, as we show next, this is not the case; there is a constant ceiling on the welfare approximation achievable by any DSIC mechanism.

THEOREM 3.1. *When $n = 2$, for all $\varepsilon > 0$, there exists a valid distribution \mathcal{D} such that no DSIC mechanism achieves a $\frac{2+\sqrt{2}}{4} + \varepsilon$ approximation to welfare.*

The proof of this result builds on a structural characterization of DSIC mechanisms for two agents established by Amanatidis et al. [2017]. At a high level, every DSIC mechanism for two agents can be decomposed into two fundamental types of mechanisms: *picking* mechanisms and *exchange* mechanisms. A picking mechanism is defined by a set of allowable subsets $\mathcal{O} \subseteq \mathcal{P}(\mathcal{M})$ and a designated agent i . Given the reported bids, agent i receives the subset $A \in \mathcal{O}$ that maximizes their perceived value, while the other agent receives the remainder $\mathcal{M} \setminus A$. This is essentially a ‘‘menu’’ approach where one agent chooses their favorite bundle from a restricted list.

An exchange mechanism, by contrast, relies on a fixed partition of items \mathcal{M} into sets S and T . We can think of agent 1 as being ‘‘endowed’’ with S and agent 2 with T . A trade is considered favorable for an agent if they strictly prefer the other agent’s endowment to their own; it is considered unfavorable if they strictly prefer their own endowment to the other agent’s. An exchange mechanism outputs

(T, S) if the trade is favorable for both agents and (S, T) if the trade is unfavorable for either. (In the cases of indifference, the mechanism must still select and output one of (S, T) or (T, S) .)

Amanatidis et al. [2017] prove that all DSIC mechanisms for two agents are *picking-exchange* mechanisms: the items are partitioned into subsets $\mathcal{M} = M_1 \sqcup \dots \sqcup M_\ell$ and each subset M_k is allocated according to a fixed picking or exchange mechanism. Our proof for Theorem 3.1 proceeds by providing a single distribution for which neither of these building blocks can perfectly capture the social welfare. The distribution we consider takes value 0 or 1 most of the time, and with small probability takes a (very) high value H . For this distribution, we then show that no picking or exchange mechanism can achieve higher than a $\frac{2+\sqrt{2}}{4} + \varepsilon$ approximation to welfare.

Intuitively, a picking mechanism fails to achieve a high approximation to welfare because, in a picking mechanism, the allocation of items is determined solely by agent 1. Agent 1 does not know which items agent 2 has a high value for, and therefore, in expectation, agent 1 will pick some items where they have value 1 for the item but agent 2 has value H for the item. This leads to an avoidable loss in expected social welfare for any picking mechanism, which we show is at least a multiplicative factor of $\frac{2+\sqrt{2}}{4} + \varepsilon$.

Bounding the expected welfare of all exchange mechanisms requires a more involved argument. Specifically, we bound the expected welfare contribution of all non-high-value items and all high-value items separately. For the former, we observe that the exchange component (S, T) must output either the allocation (S, T) or (T, S) ; therefore, we can bound the contribution of all non-high-value items by bounding the expected maximum welfare of all non-high-value items in (S, T) or (T, S) . The bound for high-value items is technically more involved. To bound the expected welfare contribution of all high-value items, we first condition on the total number of high-value items and then bound the welfare contribution based on correctly and incorrectly assigned high-value items. Finally, we can combine the expected welfare contribution of the low-value items and the high-value items to show that any exchange mechanism is at best a $\frac{2+\sqrt{2}}{4} + \varepsilon$ approximation.

PROOF OF THEOREM 3.1. Fix $\varepsilon > 0$. We define a family of distributions parameterized by positive constants p and δ . Taking p and δ sufficiently small (in terms of ε) will give the desired result.

To draw a value v , we first sample a “base” value:

$$v_{\text{base}} = \begin{cases} \frac{1-q}{p} & \text{with probability } p \\ 1 & \text{with probability } q \\ 0 & \text{with probability } 1 - q - p \end{cases}$$

where $q = 2 - \sqrt{2} \approx 0.59$. To this base value, we add uniform noise scaled by δ , so we have $v = v_{\text{base}} + \text{Unif}[0, \delta]$. The noise term’s only purpose is to ensure that ties in agents’ valuations for distinct subsets of goods occur with probability 0. We will refer to the above distribution (for sampling a value v) as \mathcal{D} throughout this proof.

The expected value of a single draw from \mathcal{D} is:

$$\mathbb{E}[v] = p \cdot \frac{1-q}{p} + q \cdot 1 + (1 - q - p) \cdot 0 + \frac{\delta}{2} = (1 - q) + q = 1 + \frac{\delta}{2}. \quad (1)$$

The Welfare Benchmark. The expected optimal social welfare per item is the expected maximum of two i.i.d. draws v_1 and v_2 from \mathcal{D} . The probability that at least one draw is a “high value” (i.e.,

$v_{\text{base}} = \frac{1-q}{p}$ is $1 - (1-p)^2 = 2p - p^2$ and the probability that at least one draw is 1 and no draw is a high value is $q^2 + 2q(1-q-p)$.

The expected value of the max of two draws v_1 and v_2 is therefore:

$$\begin{aligned} \mathbb{E}[\max(v_1, v_2)] &= \frac{1-q}{p} (2p - p^2) + 1 \cdot (q^2 + 2q(1-q-p)) + \delta/2 \\ &= (1-q)(2-p) + 2q - q^2 - 2pq + \delta/2 \\ &= 2 - q^2 - qp - p + \delta/2 \\ &= 4\sqrt{2} - 4 - qp - p + \delta/2. \end{aligned}$$

The rest of the proof will focus on showing that no DSIC mechanism can achieve expected welfare greater than $\sqrt{2} + \psi(\delta, p)$ per item, for some function ψ that goes to zero as δ and p go to zero. Since we already showed that the maximum expected welfare is $4\sqrt{2} - 4 - qp - p + \delta/2$ per item, this establishes the desired approximation ratio upper bound of

$$\frac{\sqrt{2} + \psi(\delta, p)}{4\sqrt{2} - 4 - qp - p + \delta/2} \leq \frac{\sqrt{2}}{4\sqrt{2} - 4} + \varepsilon = \frac{2 + \sqrt{2}}{4} + \varepsilon, \quad (2)$$

for sufficiently small δ and p .

Mechanism Decomposition. From the results in [Amanatidis et al., 2017], any DSIC mechanism for two agents partitions the set of items M into components and runs a distinct sub-mechanism on each component: $M = P_1 \sqcup P_2 \sqcup E_1 \sqcup \dots \sqcup E_k$. The first component is *Picking Mechanisms* (for allocating subsets P_1 and P_2), where an agent is allocated their most preferred bundle from a collection of subsets (subsets of P_1 and P_2 , respectively), and the other agent receives the remaining goods. Note that two picking mechanisms suffice; the first one, where agent 1 picks first, allocates the subset P_1 , and the second one, where agent 2 picks first, allocates the subset P_2 . The second component is *Exchange Mechanisms* (for allocating subsets E_1, \dots, E_k), where items are endowed to the agents, and traded if and only if both agents prefer the other's bundle. We will prove that for any component corresponding to a set of items C from this partition (with $|C| = s$), the expected welfare generated by these items is at most

$$\left(\sqrt{2} + \psi(\delta, p)\right) \cdot |C| = \left(\sqrt{2} + \psi(\delta, p)\right) s, \quad (3)$$

for some function $\psi(\delta, p)$ that goes to zero as δ and p go to zero.

Analysis of Picking Mechanisms. Without loss of generality, consider P_1 with $|P_1| = s$. Although the fraction of items that agent 1 picks can be arbitrary, the key property we will use is that the allocation rule for P_1 depends only on agent 1's values. Condition on agent 1's values ($v_{1,j}$ for $j \in P_1$) but assume we have not yet sampled agent 2's values.

Suppose agent 1 is allocated a set $S \subseteq P_1$ and agent 2 gets $P_1 \setminus S$. Let $H = \{j \in P_1 \mid v_{1,j} \geq \frac{1-q}{p}\}$ be the set of items agent 1 has high (base) value for in P_1 , and let \mathcal{E}_H be the event that H is the set of items for which agent 1 has high (base) value for. For any $j \in P_1$, if $j \in H$ and agent 1 is given item j , then this item contributes at least $\frac{1-q}{p}$, and at most $\frac{1-q}{p} + \delta$, to the welfare. For any $j \in P_1$ such that $j \notin H$, if agent 1 is given item j , then this item contributes at most $1 + \delta$ to the welfare. For any item $j \in P_1$, if we give item j to agent 2, then the expected contribution of item j to the welfare,

conditioned on \mathcal{E}_H , is $\mathbb{E}[v_{2,j} \mid \mathcal{E}_H] = 1 + \frac{\delta}{2}$ by Equation (1). Therefore, the conditional expected welfare of the component P_1 is at most

$$\mathbb{E}[\text{welfare of } P_1 \mid \mathcal{E}_H] \leq |H| \cdot \left(\frac{1-q}{p} + \delta \right) + (|P_1 \setminus H|) \cdot (1 + \delta) \leq s + |H| \left(\frac{1-q-p}{p} \right) + \delta s.$$

Taking an expectation over H gives that the expected welfare of the component P_1 is at most

$$\begin{aligned} \mathbb{E}[\text{Welfare of } P_1] &= \mathbb{E}[\mathbb{E}[\text{welfare of } P_1 \mid \mathcal{E}_H]] \\ &= \mathbb{E} \left[s + |H| \left(\frac{1-q-p}{p} \right) + \delta s \right] \\ &= s + \left(\frac{1-q-p}{p} \right) \mathbb{E}[|H|] + \delta s \\ &= s + \left(\frac{1-q-p}{p} \right) ps + \delta s \\ &= s(2-q-p) + \delta s \\ &\leq (2-q)s + \delta s = \sqrt{2}s + \delta s = (\sqrt{2} + \delta)s, \end{aligned}$$

as desired, where we used that $\mathbb{E}[|H|] = ps$. The analysis for P_2 is identical by symmetry. Therefore, we have shown Equation (3) for picking mechanisms.

Analysis of Exchange Mechanisms. Consider an exchange component E where agent 1 is endowed with ℓ items and agent 2 with k items. Let $s = |E| = \ell + k$. The following argument will work for all $s \geq 3$. At the end of this proof, we separately handle the final nondegenerate case of $\ell = k = 1$. (Note that it is without loss of generality that $\ell \geq 1$ and $k \geq 1$, as otherwise these items are always allocated to a single agent and trade never occurs; this can therefore be represented as part of the picking mechanism.)

We bound the welfare by analyzing the contributions from non-high-value and high-value items separately, where a high-value item is an item with base value $\frac{1-q}{p}$.

Contribution from Non-High-Value Items. For any item, if an agent's value for that item is not "high," then the item's value is drawn from a $\text{Bern}\left(\frac{q}{1-p}\right) + \text{Unif}(0, \delta)$ distribution. Let W be the welfare of non-high-value items in the endowed allocation. Let W' be the welfare of the non-high-value items in the "flipped" allocation, where the ℓ items given to agent 1 are given to agent 2 instead, and the k items given to agent 2 are given to agent 1. We can decompose W as $W = W_{base} + Z$, where W_{base} is the sum of the Bernoulli contributions, and Z is the sum of the noise terms from $\text{Unif}(0, \delta)$. Note that W_{base} follows a binomial distribution $W_{base} \sim \text{Bin}\left(s, \frac{q}{1-p}\right)$, and the noise term is bounded by $0 \leq Z \leq s\delta$ with probability 1. Similarly, $W' = W'_{base} + Z'$ where W'_{base} is an independent draw from the same binomial distribution and $0 \leq Z' \leq s\delta$.

The expected welfare of component E from non-high-values can be upper bounded by $\mathbb{E}[\max(W, W')]$. Using the deterministic bound on the noise, we have:

$$\mathbb{E}[\max(W, W')] \leq \mathbb{E}[\max(W_{base} + s\delta, W'_{base} + s\delta)] = \mathbb{E}[\max(W_{base}, W'_{base})] + s\delta.$$

Since W_{base} and W'_{base} are independent and identically distributed from $\text{Bin}\left(s, \frac{q}{1-p}\right)$,

we have:

$$\begin{aligned} \mathbb{E}[\max(W_{base}, W'_{base})] &= \mathbb{E}[W_{base}] + \frac{1}{2}\mathbb{E}[|W_{base} - W'_{base}|] \\ &\leq \mathbb{E}[W_{base}] + \frac{1}{2}\sqrt{\text{Var}(W_{base} - W'_{base})} \\ &= \mathbb{E}[W_{base}] + \sqrt{\frac{\text{Var}(W_{base})}{2}}. \end{aligned}$$

where the second transition follows from Jensen's inequality, and the third follows because the variance of the sum of independent random variables is the sum of their variances. Substituting the mean $\mathbb{E}[W_{base}] = s\frac{q}{1-p}$ and variance $\text{Var}(W_{base}) = s\frac{q}{1-p}(1 - \frac{q}{1-p})$, and combining with the noise error term $s\delta$, we obtain the upper bound:

$$sq + \sqrt{\frac{sq(1-q)}{2}} + s\delta.$$

Contribution from High-Value Items. Let h be the total count of high-value draws in the endowed and non-endowed items by both agents, i.e., $h = |\{(i, j) \mid v_{i,j} \geq \frac{1-q}{p}\}|$. Note that $\mathbb{E}[h] = 2sp$. We will show the following lemma, the proof of which can be found in Appendix C.

LEMMA 3.2. *Conditioned on any value of h , the expected welfare of high-value items in this component is at most*

$$\frac{3}{4} \frac{1-q}{p} h + \delta h.$$

Using that $\mathbb{E}[h] = 2sp$, this gives that the expected welfare contribution of high-value items is at most

$$\mathbb{E}\left[\frac{3}{4} \frac{1-q}{p} h + \delta h\right] = \left(\frac{3}{4} \frac{1-q}{p} + \delta\right) (2sp) = \frac{3}{2}(1-q)s + 2p\delta s.$$

Combining the Bounds. The total expected welfare is bounded by the sum of the contributions from non-high-value and high-value items, which from above is upper bounded by:

$$\left(sq + \sqrt{\frac{sq(1-q)}{2}} + \delta s\right) + \left(\frac{3}{2}(1-q)s + 2p\delta s\right) = \left(\frac{3-q}{2}\right)s + \sqrt{\frac{sq(1-q)}{2}} + s \cdot \psi(\delta, p),$$

where ψ goes to 0 as δ and p go to 0. Our goal is to show that for $s \geq 3$, this is at most $s(\sqrt{2} + \psi(\delta, p))$. In other words, we will show that:

$$\left(\frac{3-q}{2}\right)s + \sqrt{\frac{sq(1-q)}{2}} \leq \sqrt{2}s.$$

Rearranging and plugging in $q = 2 - \sqrt{2}$, it suffices to show that

$$\sqrt{\frac{s(3\sqrt{2}-4)}{2}} \leq \left(\frac{\sqrt{2}-1}{2}\right)s.$$

Isolating s (and using that $s > 0$), we get that the above is equivalent to

$$\frac{2(3\sqrt{2}-4)}{(\sqrt{2}-1)^2} \leq s.$$

The LHS simplifies to $2\sqrt{2} \leq 2.83$. Therefore, for any $s \geq 3 > 2\sqrt{2}$, the desired inequality holds, implying that the total expected welfare of all items from this component is upper bounded by

$$s(\sqrt{2} + \psi(\delta, p)).$$

This completes the proof of Equation (3) for exchange components when $s \geq 3$.

Small Exchanges ($s = 2$). The only non-degenerate case left is a 1-for-1 trade ($\ell = 1, k = 1$).

Consider the expected value of agent 1 for the item they receive. Their expected value for their endowed item is $1 + \delta/2$ (by Equation (1)). Their expected gain from trade, if allowed by agent 2, is $\mathbb{E}[\max(v_{1,2} - v_{1,1}, 0)] = \frac{\mathbb{E}[v_{1,2} - v_{1,1}] + \mathbb{E}[|v_{1,2} - v_{1,1}|]}{2} = \frac{\mathbb{E}[|v_{1,2} - v_{1,1}|]}{2}$. Note that agent 2 is willing to trade with probability $1/2$, independently of agent 1's value. Thus, the expected welfare gain for agent 1 from trading is $\frac{\mathbb{E}[|v_{1,2} - v_{1,1}|]}{4}$. We can compute the expected absolute gap between two draws from the distribution \mathcal{D} as

$$\begin{aligned} \mathbb{E}[|v_{1,2} - v_{1,1}|] &\leq 2pq \left(\frac{1-q}{p} - 1 \right) + 2p(1-p-q) \left(\frac{1-q}{p} \right) + 2q(1-p-q) \cdot 1 + \delta \\ &= 2 - 2q^2 - 2p - 2pq + \delta \\ &\leq 2(1 - q^2) + \delta. \end{aligned}$$

The expected welfare for agent 1's item in this component is therefore upper bounded by

$$\begin{aligned} 1 + \frac{\delta}{2} + \frac{1}{4}(2(1 - q^2) + \delta) &\leq \frac{3 - q^2}{2} + \delta \\ &= \frac{3 - (2 - \sqrt{2})^2}{2} + \delta \\ &= \frac{4\sqrt{2} - 3}{2} + \delta \\ &= \sqrt{2} + \frac{2\sqrt{2} - 3}{2} + \delta \\ &\leq \sqrt{2} + \delta. \end{aligned}$$

This shows that agent 1's expected welfare from the exchange is at most $\sqrt{2} + \delta$. By symmetry, the same holds for agent 2, yielding a $\sqrt{2} + \delta$ per-item bound on the final expected welfare for this component. This shows Equation (3) holds for exchange components when $s = 2$.

Therefore, we have shown that Equation (3) holds for all s and all exchange mechanisms and for all picking mechanisms, and therefore it must hold for all components of any DSIC mechanism. This completes the proof of Equation (2), and therefore the proof of Theorem 3.1. \square

A Matching Bound. Strikingly, the upper bound of Theorem 3.1 is tight. We can match the ≈ 0.854 ratio using a remarkably simple class of rules we call *Pick- r* mechanisms. For $r \in [0, 1]$, a *Pick- r* mechanism for two agents grants agent 1 their $\lfloor rm \rfloor$ most valuable items, leaving the rest for agent 2. This is simply a picking mechanism where the "menu" consists of all subsets of a fixed cardinality; therefore, it is inherently DSIC.

PROPOSITION 3.3. *For two agents and any valid distribution \mathcal{D} , the Pick- r mechanism, for $r = (\frac{2-\sqrt{2}}{2})$, is DSIC and achieves a $\frac{2+\sqrt{2}}{4} - o(1)$ approximation to welfare.*

This establishes that simple picking mechanisms are, in a sense, the “optimal” DSIC mechanisms for two agents under i.i.d. valuations.

The proof of Proposition 3.3 relies on the key observation that for a large number of items, a Pick- r mechanism for two agents is approximately equivalent to the first agent taking any items for which its corresponding quantile (with respect to the underlying distribution) is at least $1 - r$. Therefore, rather than considering each distribution separately, we can instead analyze a Pick- r mechanism via the quantile of the value of the agent who is assigned the item. Concretely, consider any one item. In the optimal allocation, the item is allocated to the agent with the highest value; therefore, the corresponding quantile is the maximum of two Uniform $[0, 1]$ random variables. To bound our approximation ratio, we compare this with the equivalent quantile for a Pick- r mechanism.

The fact that Pick- r mechanisms are approximately equivalent to this quantile version will be useful here, and especially when generalizing to $n > 2$ agents. We first formalize this connection.

Definition 3.4 (Pick- \mathbf{r} mechanisms). Given a vector $\mathbf{r} = (r_1, r_2, \dots, r_n)$ such that (i) $r_i \geq 0$ for all i , and (ii) $\sum_{i=1}^n r_i = 1$, let $t_i = \lfloor m \cdot \sum_{j=1}^i r_j \rfloor - \lfloor m \cdot \sum_{j=1}^{i-1} r_j \rfloor$. The Pick- \mathbf{r} mechanism is a *quota serial dictatorship* [Pápai, 2000] where agent 1 selects their favorite t_1 items, then agent 2 selects their favorite t_2 items from the remainder, and so on. (Note that this choice of t_i ensures that $t_i \in \{\lfloor r_i m \rfloor, \lceil r_i m \rceil\}$ and $\sum_i t_i = m$.)

Definition 3.5 (QT-s mechanism). For a vector $\mathbf{s} = (s_1, \dots, s_n)$ with $s_n = 1$, we define the quantile-threshold-s (QT-s) mechanism as follows: For each item value $v_{i,j}$, we determine its quantile $q_{i,j} \in [0, 1]$.³ The mechanism iterates through agents $1, \dots, n$; item j is allocated to the first agent i such that $q_{i,j} \geq 1 - s_i$.

We now show that these two mechanisms produce nearly identical allocations, and consequently, have similar guarantees.

LEMMA 3.6. Fix a vector $\mathbf{r} = (r_1, \dots, r_n)$ such that (i) $r_i \geq 0$ for all i , and (ii) $\sum_i r_i = 1$, and let $\mathbf{s} = (s_1, \dots, s_n)$ be such that

$$s_i = \frac{r_i}{1 - \sum_{k=1}^{i-1} r_k}.$$

We have that

(1) With probability $1 - O(n/m^2)$, the allocations of Pick- \mathbf{r} and QT-s differ on at most $O(n^2 \sqrt{m \log m})$ items.

Furthermore, fix a valid distribution \mathcal{D} with CDF F , and let X_1, \dots, X_n and X all be i.i.d. draws from \mathcal{D} . Let $\tau_i = 1 - s_i$.

(2) Pick- \mathbf{r} achieves an $\alpha - o(1)$ approximation to the optimal expected welfare, where

$$\alpha = \frac{\sum_{i=1}^n r_i \cdot \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)]}{\mathbb{E}[\max_i X_i]}.$$

(3) Pick- \mathbf{r} is Envy-Free (EF) with high probability if for all agents i and j , if $i < j$,

$$r_i \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] > r_j \mathbb{E}[X \mid X \leq F^{-1}(\tau_i)],$$

and if $i > j$,

$$r_i \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] > r_j \mathbb{E}[X].$$

³Recall that we assumed that distributions are continuous, so this is well-defined.

The proof of Lemma 3.6 can be found in Appendix D. Using Lemma 3.6, we are now ready to prove Proposition 3.3.

PROOF OF PROPOSITION 3.3 . We directly analyze the corresponding QT mechanism from Lemma 3.6. Specifically, Pick- r , for $r = \frac{2-\sqrt{2}}{2}$, behaves near-identically to QT- $(\frac{2-\sqrt{2}}{2}, 1)$ (for a large number of items). Let $\tau_1 = 1 - \frac{2-\sqrt{2}}{2} = \frac{\sqrt{2}}{2}$ and $\tau_2 = 0$. From part (2) of Lemma 3.6 we have that Pick- r achieves an $\alpha - o(1)$ approximation to the optimal expected welfare, where

$$\begin{aligned} \alpha &= \frac{\sum_{i=1}^n r_i \cdot \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)]}{\mathbb{E}[\max_i X_i]} \\ &= \frac{\frac{2-\sqrt{2}}{2} \cdot \mathbb{E}[X \mid X \geq F^{-1}(\frac{\sqrt{2}}{2})] + (1 - \frac{2-\sqrt{2}}{2}) \cdot \mathbb{E}[X \mid X \geq F^{-1}(0)]}{\mathbb{E}[\max\{X_1, X_2\}]} \\ &= \frac{\frac{2-\sqrt{2}}{2} \cdot \mathbb{E}[X \mid X \geq F^{-1}(\frac{\sqrt{2}}{2})] + \frac{\sqrt{2}}{2} \cdot \mathbb{E}[X]}{\mathbb{E}[\max\{X_1, X_2\}]} \\ &= \frac{r \cdot \mathbb{E}[X \mid X \geq F^{-1}(1-r)] + (1-r) \cdot \mathbb{E}[X]}{\mathbb{E}[\max\{X_1, X_2\}]} \end{aligned}$$

For the numerator we have that $r \cdot \mathbb{E}[X \mid X \geq F^{-1}(1-r)] = \mathbb{E}[X \cdot 1\{X \geq F^{-1}(1-r)\}] = \int_{1-r}^1 F^{-1}(x)dx$. We also have that $\mathbb{E}[X] = \int_0^1 F^{-1}(x)dx$. Therefore, the numerator becomes

$$\begin{aligned} \int_{1-r}^1 F^{-1}(x)dx + (1-r) \int_0^1 F^{-1}(x)dx &= \int_{1-r}^1 F^{-1}(x)dx + (1-r) \left(\int_0^{1-r} F^{-1}(x)dx + \int_{1-r}^1 F^{-1}(x)dx \right) \\ &= (1-r) \int_0^{1-r} F^{-1}(x)dx + (2-r) \int_{1-r}^1 F^{-1}(x)dx \\ &= \int_0^1 g(x)F^{-1}(x)dx, \end{aligned}$$

where $g(x) = 1-r$ for $x < 1-r$, and $g(x) = 2-r$ for $x \geq 1-r$. Similarly, the denominator is equal to $\int_0^1 2xF^{-1}(x)dx$. Therefore, we have that

$$\alpha = \frac{\int_0^1 g(x)F^{-1}(x)dx}{\int_0^1 2xF^{-1}(x)dx}.$$

We will prove that $\int_t^1 g(x)dx \geq \frac{2+\sqrt{2}}{4} \cdot \int_t^1 2xdx$, for all $t \in [0, 1]$, which implies that $\alpha \geq \frac{2+\sqrt{2}}{4}$ (since $F^{-1}(x)$ is a non-decreasing function), concluding the proof. We have that $\int_t^1 2xdx = 1 - t^2$. For $t < 1-r$, $\int_t^1 g(x)dx = \int_t^{1-r} 1-r dx + \int_{1-r}^1 2-r dx = (1-r)(1-r-t) + (2-r)r = 1 - (1-r)t = 1 - \frac{\sqrt{2}}{2}t = \frac{2-\sqrt{2}t}{2}$. It is easy to confirm that $\frac{2-\sqrt{2}t}{2} \geq \frac{2+\sqrt{2}}{4}(1-t^2)$. For $t \geq 1-r$, $\int_t^1 g(x)dx = \int_t^1 2-r dx = (2-r)(1-t) = \frac{2+\sqrt{2}}{2}(1-t)$. It is easy to confirm that $\frac{2+\sqrt{2}}{2}(1-t) \geq \frac{2+\sqrt{2}}{4}(1-t^2)$.

This concludes the proof. \square

Incorporating Fairness. While Proposition 3.3 provides a welfare approximation, it says nothing about fairness. Indeed, the parameter choice $r = \frac{2-\sqrt{2}}{2} \approx 0.29$ implies that agent 2 receives over 70% of the items. Under distributions where items have similar values, agent 1 will inevitably be envious. Furthermore, it is straightforward to show that varying r in Pick- r mechanisms to balance the allocation (and restore fairness) leads to suboptimal welfare guarantees.

One might therefore suspect a necessary trade-off between fairness and efficiency among DSIC mechanism. However, strikingly, we can incorporate envy-freeness at essentially no loss to the worst-case welfare guarantee. The key insight is that we don't need to restrict ourselves to a single picking mechanism for all distributions; instead, we show that for any distribution, there exists a choice of r that simultaneously achieves both envy-freeness and optimal welfare guarantees.

THEOREM 3.7. *For two agents and any valid distribution \mathcal{D} , there exists an r for which the Pick- r mechanism is simultaneously DSIC, envy-free with high probability, and achieves a $\frac{2+\sqrt{2}}{4} - o(1)$ approximation to welfare.*

The proof of Theorem 3.7 builds on the result of Proposition 3.3, which showed that the Pick- r mechanism for $r = \frac{2-\sqrt{2}}{2}$ achieves the desired welfare approximation. The key idea for the proof of Theorem 3.7 is separate all possible underlying distribution in to two categories. The first category is the set of distributions for which the Pick- $\frac{2-\sqrt{2}}{2}$ is envy-free with high probability. For these distributions, Proposition 3.3 directly implies the desired result. The second category is the set of distributions for which the Pick- $\frac{2-\sqrt{2}}{2}$ mechanism is not envy-free with high probability. Here, we show that for sufficiently small $\varepsilon > 0$, the Pick- $(1/2 - \varepsilon)$ mechanism achieves at least a $7/8 - o(1) \geq \frac{2+\sqrt{2}}{4} - o(1)$ approximation to welfare and is envy-free with high probability. Therefore, in this case, the Pick- $(1/2 - \varepsilon)$ satisfies the desired property. The proof of Theorem 3.7 can be found in Appendix E.

3.2 Connection to Prophet Inequalities

For a reader that is well-versed in optimal stopping theory, the ratio $\frac{2+\sqrt{2}}{4} \approx 0.854$ from Proposition 3.3 may look familiar. This is no coincidence. It turns out that there is a tight connection between i.i.d. prophet inequalities and the approximation ratios of a class Pick- r (or, more specifically, QT-s mechanisms). To discuss this, we will need to formally introduce the concepts and give a brief background of the area.

Preliminaries on I.I.D. Prophets Inequalities. In the i.i.d. prophet setting, a sequence of n random variables X_1, \dots, X_n are drawn i.i.d. from a fixed, known distribution \mathcal{D} . An online algorithm observes the realized values x_1, \dots, x_n sequentially. Upon observing x_i , the algorithm must make an irrevocable decision to either accept x_i and stop, or reject x_i and continue to step $i + 1$. An algorithm in this problem is characterized by a *stopping time* τ , a random variable taking values in $\{1, \dots, n\}$ such that the event $\{\tau = i\}$ depends only on the history X_1, \dots, X_i . The objective is to maximize the expected reward, $\mathbb{E}[X_\tau]$. The algorithm competes against a ‘‘prophet’’ who knows the realizations of all random variables in advance and selects the maximum. We say an algorithm achieves an α -approximation (or has a competitive ratio of α) if $\mathbb{E}[X_\tau] \geq \alpha \cdot \mathbb{E}[\max_i X_i]$.

This problem was first studied by Hill and Kertz [1982] (while the non-i.i.d. version dates back to Krengel and Sucheston [1978]), who derived a recursive formula for the optimal α for any fixed n , denoted by a_n . For $n = 2$, this bound is $\frac{2+\sqrt{2}}{4} \approx 0.854$.⁴ They further demonstrated that for all n , $a_n \geq 1 - 1/e \approx 0.632$. Subsequently, [Kertz, 1986] proved that as $n \rightarrow \infty$, a_n converges to a constant $\beta \approx 0.745$. This constant is defined such that $1/\beta$ satisfies the integral equation:

$$\int_0^1 \frac{1}{y(1 - \ln y) + \beta - 1} dy = 1.$$

⁴Hill and Kertz originally analyzed the ratio of the prophet to the gambler ($1/\alpha$). We adopt the modern convention of the approximation ratio $\alpha \leq 1$.

While Kertz established β as the asymptotic limit, he did not prove that the sequence of a_n s is monotone, leaving open the possibility that the bound for finite n could be strictly worse than β . This gap was closed by [Correa et al., 2017], who proved that $\beta \approx 0.745$ is a universal lower bound for all $n \geq 1$.

We can now state the equivalence between i.i.d. prophet inequalities and DSIC mechanisms in our setting. In particular, we will define the Quantile-threshold-prophet-s (QTP-s) algorithm for the i.i.d. prophet inequality problem, which, upon observing value x_i , it accepts it if and only if $q_i \geq 1 - s_i$, where q_i is the quantile of x_i (i.e., $q_i = F(x_i)$). We then have the following.

PROPOSITION 3.8. *Fix a valid distribution \mathcal{D} and a number of agents n . Fix $\alpha \in [0, 1]$. The following statements hold.*

- (1) *For any \mathbf{s} , the QT-s algorithm (Definition 3.5) achieves an α -approximation to welfare iff the QTP-s algorithm achieves an α -approximation to the reward of the prophet.*
- (2) *There exists a prophet algorithm achieving an α -approximation iff there exists a vector \mathbf{r} for which Pick- \mathbf{r} achieves an α -approximation.*

PROOF. For (1), fix an arbitrary \mathbf{s} . Consider a single item j with quantiles $q_{i,j} = F(v_{i,j})$. By Definition 3.5, QT-s allocates j to $\tau = \min\{i \in [n] : q_{i,j} \geq 1 - s_i\}$. The QTP-s stopping rule is exactly the same, hence on every realization QT-s and QTP-s select the same index τ and obtain the same realized value. Summing over items, the total welfare of QT-s equals the total reward of running QTP-s independently on each item (with the same \mathbf{s}), and therefore their expectations coincide. Finally, the offline benchmark in our allocation setting is $\sum_j \max_i v_{i,j}$ (each item assigned to the agent with the highest value), which is exactly the prophet benchmark for these per-item instances.

For (2), first note that by Lemma 3.6, every Pick- \mathbf{r} mechanism is equivalent to a QT-s mechanism (with respect to the welfare approximation, up to a $o(1)$ difference), for some threshold vector \mathbf{s} . Second, by part (1) of this proposition, the welfare of a QT-s mechanism is identical to the reward of a prophet algorithm QTP-s. Therefore, it suffices to show that the optimal prophet algorithm must be a QTP-s algorithm.

Consider an arbitrary optimal prophet algorithm, and let z_i be the marginal probability that the algorithm stops at step i . To maximize the expected reward $\sum_i \mathbb{E}[X_i | \text{stop at } i] \cdot z_i$, for these (fixed) marginal stopping probabilities, the optimal strategy at step i (conditioned on reaching step i) must be to accept when X_i takes a large value. This implies a threshold strategy where we accept X_i if and only if it exceeds a certain quantile s_i , and specifically $s_i = z_i / \Pr(\text{reach } i)$. This is a QTP-s strategy, and it achieves the same welfare as the original algorithm. Therefore, applying part (1) of this proposition and Lemma 3.6, we have a correspondence between optimal prophet algorithms and Pick- \mathbf{r} mechanisms. \square

3.3 $n > 2$ Agents

We now turn to the case of $n > 2$ agents. By importing results from the i.i.d. prophet problem [Correa et al., 2017], we can ensure that, for any distribution \mathcal{D} and any number of agents n , there exists a mechanism achieving a $\beta \approx 0.745$ approximation to welfare.

COROLLARY 3.9. *For all n and distributions \mathcal{D} , there exists $\mathbf{r} = (r_1, \dots, r_{n-1})$ such the Pick- \mathbf{r} mechanism achieves a $\beta - o(1) \approx 0.745$ approximation to welfare. Furthermore, there exists a valid distribution \mathcal{D} such that no picking sequence can achieve a higher approximation to welfare.*

PROOF. By Proposition 3.8, the welfare approximation of a Pick- r mechanism is equivalent to the competitive ratio of an i.i.d. prophet inequality algorithm. Since $\beta-o(1) \approx 0.745$ is the optimal bound for the i.i.d. prophet inequality problem [Correa et al., 2017], the existence of the corresponding Pick- r mechanism follows. \square

Incorporating Fairness. The most natural approach to fairness in this framework is to equalize the number of items agents pick, i.e., to use the Pick- $(1/n, \dots, 1/n)$ mechanism. This corresponds to a stopping rule studied in the prophet inequality literature [Arsenis and Kleinberg, 2022], albeit under a different notion of fairness. For this mechanism, we can prove the following guarantee: for all n , the Pick- $(1/n, \dots, 1/n)$ mechanism is DSIC, envy-free in expectation, and achieves a $1/2$ approximation to welfare. By slightly modifying the number of items allocated to each agent, we are able to find a mechanism that can instead achieve envy-freeness with high probability.

THEOREM 3.10. *For all n and all valid distributions \mathcal{D} , for sufficiently small ε , the Pick- r mechanism with $\mathbf{r} = (1/n - \varepsilon, \dots, 1/n - \varepsilon, 1/n + (n-1)\varepsilon)$ is envy-free with high probability and achieves a $1/2 - o(1)$ approximation to welfare.*

PROOF OF THEOREM 3.10. To show the welfare approximation, we will apply part (2) of Lemma 3.6 to show that the welfare approximation is strictly greater than $1/2$ for $\varepsilon = 0$. Since the welfare approximation is continuous in ε (as the thresholds in the expected values are continuous), it will remain greater than $1/2$ for sufficiently small $\varepsilon > 0$.

Using the definition in Lemma 3.6, we have that $s_i = \frac{1/n}{1-(i-1)/n} = \frac{1}{n-i+1}$. Therefore, using part (2) of Lemma 3.6, we have that the approximation ratio for $\varepsilon = 0$ is

$$\frac{\sum_{i=1}^n \frac{1}{n} \cdot \mathbb{E}[X \mid X \geq F^{-1}(1 - \frac{1}{n+1-i})]}{\mathbb{E}[\max_i X_i]} = \frac{\sum_{i=1}^n \frac{1}{n} \cdot \mathbb{E}[X \mid X \geq F^{-1}(1 - \frac{1}{i})]}{\mathbb{E}[\max_i X_i]}.$$

Using the identity $\mathbb{E}[Y] = \int_0^\infty \Pr[Y > x] dx$ for nonnegative random variables, we can rewrite the expectation terms. For a fixed integration variable x , the conditional probability $\Pr[X > x \mid X \geq F^{-1}(1 - 1/i)]$ depends on whether x is below or above the threshold $F^{-1}(1 - 1/i)$.

- If $x \leq F^{-1}(1 - 1/i)$, then the condition $X > x$ is always met given the condition $X \geq F^{-1}(1 - 1/i)$, and thus, the probability is 1.
- If $x > F^{-1}(1 - 1/i)$, then $\Pr[X > x \mid X \geq F^{-1}(1 - 1/i)] = \frac{\Pr[X > x]}{\Pr[X \geq F^{-1}(1 - 1/i)]} = \frac{1 - F(x)}{1/i} = i(1 - F(x))$.

Therefore, $\Pr[X > x \mid X \geq F^{-1}(1 - 1/i)] = \min\{1, i(1 - F(x))\}$. Summing over i we have that $\sum_{i=1}^n \frac{1}{n} \cdot \mathbb{E}[X \mid X \geq F^{-1}(1 - \frac{1}{i})]$ is equal to

$$\sum_{i=1}^n \frac{1}{n} \int_{x=0}^{\infty} \min\{1, i(1 - F(x))\} dx = \int_{x=0}^{\infty} \left(\frac{1}{n} \sum_{i=1}^n \min\{1, i(1 - F(x))\} \right) dx.$$

On the other hand, $\mathbb{E}[\max_i X_i] = \int_0^\infty (1 - F(x)^n) dx$. We will prove that $\frac{1}{n} \sum_{i=1}^n \min\{1, i(1 - F(x))\} > \frac{1}{2}(1 - F(x)^n)$ for all $x \geq 0$.

Case 1: $x > F^{-1}(1 - 1/n)$. In this case, $1 - F(x) < 1/n$, and therefore $i(1 - F(x)) < 1$ for all $i \in [n]$. Therefore,

$$\frac{1}{n} \sum_{i=1}^n \min\{1, i(1 - F(x))\} = \frac{1 - F(x)}{n} \sum_{i=1}^n i = \frac{1 - F(x)}{n} \frac{n(n+1)}{2} = \frac{n+1}{2}(1 - F(x)).$$

Using Bernoulli's inequality, $F(x)^n = (1 - (1 - F(x)))^n \geq 1 - n(1 - F(x))$. This implies that $1 - F(x)^n \leq n(1 - F(x))$. So we overall have

$$\frac{1}{n} \sum_{i=1}^n \min\{1, i(1 - F(x))\} = \frac{n+1}{2}(1 - F(x)) > \frac{n}{2}(1 - F(x)),$$

as desired.

Case 2: $x \leq F^{-1}(1 - 1/n)$. Here we use the trivial bound $1 - F(x)^n \leq 1$, so it suffices to show that $\frac{1}{n} \sum_{i=1}^n \min\{1, i(1 - F(x))\} > 1/2$. Since $x \leq F^{-1}(1 - 1/n)$, $1 - F(x) \geq 1/n$. Therefore,

$$\frac{1}{n} \sum_{i=1}^n \min\{1, i(1 - F(x))\} \geq \frac{1}{n} \sum_{i=1}^n \min\{1, i/n\} = \frac{1}{n} \frac{n(n+1)}{2n} > \frac{1}{2},$$

as desired.

Envy-Freeness: We now verify envy-freeness for the vector $\mathbf{r} = (1/n - \varepsilon, \dots, 1/n - \varepsilon, 1/n + (n-1)\varepsilon)$. Let τ_i be the value threshold for agent i . For $i < n$, $\tau_i \in (0, 1)$ because $s_i < 1$. For agent n , $s_n = 1$ (since they take everything remaining), so $\tau_n = F^{-1}(0) = 0$. There are three cases based on whether or not one of the agents is agent n .

- (1) *Agents $i, j < n$.* Here, $r_i = r_j$. Since X is continuous and $\tau_i > 0$, $\mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] > \mathbb{E}[X] \geq \mathbb{E}[X \mid X < F^{-1}(\tau_i)]$, so the inequality holds regardless of the order of i and j .
- (2) *Agent $j = n$.* Here $r_i = 1/n - \varepsilon$ and $r_j = 1/n + (n-1)\varepsilon$. However, $\mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] > \mathbb{E}[X \mid X < F^{-1}(\tau_i)]$. Thus, for sufficiently small ε , $r_i \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] > r_j \mathbb{E}[X \mid X < F^{-1}(\tau_i)]$.
- (3) *Agent $i = n$.* Here $r_i = 1/n + (n-1)\varepsilon$. Furthermore, $\tau_i = 0$, so $\mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] = \mathbb{E}[X]$. Therefore, for all $\varepsilon > 0$, $r_i \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)] > r_j \mathbb{E}[X]$.

Since all expected value conditions strictly hold, by Lemma 3.6 part (3), the mechanism is envy-free with high probability. \square

4 BIC MECHANISMS

In this section, we study Bayesian Incentive Compatible mechanisms. Consider the *Ranking* algorithm, defined as Algorithm 1. The algorithm, at a high level, works as follows. We interpret the agents' bids as ordinal, i.e., we convert the cardinal (reported) values to a ranking over the items. Then, we give each item to the agent who ranked it highest. We prove that this achieves strong theoretical guarantees.

Specifically, the algorithm will compute the ranking σ_i based on the bids b_i for each agent i , where item j is ranked higher than item k in σ_i if and only if $b_{i,j} \geq b_{i,k}$. The algorithm will then give each item to the agent with the highest rank for that item (with consistent tie-breaking). Let $\text{rank}(\sigma_i, j)$ be the rank of item j in agent i 's list, with 1 being the top ranked item.

THEOREM 4.1. *For all valid distributions \mathcal{D} , the ranking algorithm (Algorithm 1) is BIC, envy-free with high probability, and achieves a $1 - o(1)$ approximation to welfare.*

ALGORITHM 1: Ranking Algorithm

Input: Bids $\{b_{i,j}\}_{i \in \mathcal{N}, j \in \mathcal{M}}$ for each agent, item pair
 $\sigma_i \leftarrow$ Ranking of bids b_i , breaking ties arbitrarily, $\forall i$
 $A_i \leftarrow \emptyset, \forall i$
for $j \leftarrow 1$ **to** m **do**
 $r \leftarrow \min_i \text{rank}(\sigma_i, j)$
 $i^* \leftarrow \min_{i: \text{rank}(\sigma_i, j) = r} i$
 $A_{i^*} \leftarrow A_{i^*} \cup \{j\}$
end

To prove the BIC property, we show that for each position j , there is a fixed probability p_j that an agent receives the item ranked in position j . Furthermore, these probabilities are monotonically decreasing in j . Thus, to maximize their expected value for the items they receive, an agent should report their true ranking, which is induced by truthfully reporting their values. Regarding fairness and efficiency, the i.i.d. assumption implies that Algorithm 1 approximates *welfare maximization*, i.e., the rule that gives each item to the agent with the highest value. Welfare maximization is known to have strong fairness and efficiency guarantees in this setting [Dickerson et al., 2014]. Our proof shows that these guarantees extend to this approximate regime.

Consequently, Algorithm 1 achieves nearly all desired properties. The primary trade-off is the $1 - o(1)$ approximation rather than exact optimality. However, it is straightforward to show that any mechanism achieving a perfect welfare guarantee cannot be incentive compatible: optimal welfare requires allocating every item to the agent with the highest value, creating an unavoidable incentive for agents to exaggerate their bids.

PROOF OF THEOREM 4.1. First, we prove the BIC property. Since item values are i.i.d. across agents, the rankings σ_{-i} induced by other agents are, from i 's perspective, exchangeable across items: for any fixed item, the relative ranks assigned to it by the other agents have the same distribution as for any other item. Consequently, for each position $t \in [m]$ there exists a number $p_t \in [0, 1]$ such that, regardless of which item agent i places in position t , the probability that i wins that item equals p_t . Moreover, these probabilities are monotone: $p_1 \geq p_2 \geq \dots \geq p_m$, since moving an item higher in σ_i can only increase i 's probability of being the highest-ranked agent for that item (tie-breaking is fixed and consistent). Therefore, given σ_i and true valuation v_i , agent i 's expected utility can be written as

$$\mathbb{E}[u_i | \sigma_i, v_i] = \sum_{t=1}^m p_t v_{i, \pi(t)},$$

where $\pi(t)$ is the item agent i placed in position t (according to bids b_i). Since $(p_t)_{t=1}^m$ is a fixed, non-increasing sequence, the above expression is maximized by assigning larger values to smaller positions. Thus, agent i maximizes expected utility by submitting bids whose rankings order the items by decreasing true value (i.e., the ranking induced by truthful bidding). This establishes BIC.

Next, we show that Algorithm 1 matches welfare maximization on up to $O(n\sqrt{m \log m})$ items with probability $1 - O(n/m^2)$.

To this end, let \hat{F}_i be i 's empirical distribution of values. Let \mathcal{E}_i be the event that $\sup_x |\hat{F}_i(x) - F(x)| \leq \sqrt{\log m/m}$. By the DKW inequality [Massart, 1990],

$$\Pr[\mathcal{E}_i] \geq 1 - 2 \exp(-2m(\sqrt{\log m/m})^2) = 1 - 2 \exp(-2 \log m) = 1 - 2/m^2.$$

Conditioned on \mathcal{E}_i for all $i \in \mathcal{N}$, all empirical quantiles are within $\sqrt{\log m/m}$ of their true quantiles. Thus, as long as there is a quantile difference between agents of at least $2\sqrt{\log m/m}$, then the higher-value agent will rank it higher. The difference between the highest quantile and second highest quantile of n i.i.d. draws follows a Beta(1, n) distribution. As the PDF of this distribution is at most n , the probability that it takes a value smaller than $2\sqrt{\log m/m}$ is at most $2n\sqrt{\log m/m}$ for each item, and this is independent across items. The expected number of items that have a small highest-to-second-highest-quantile difference (recall that only those items can be allocated to the “wrong” agent) is therefore $2n\sqrt{m \log m}$. By Hoeffding’s inequality, the number of items that have a small highest-to-second-highest-quantile difference is at most $(2n + 1)\sqrt{m \log m}$ with probability at least $1 - \exp(-2(\sqrt{m \log m})^2/m) = 1 - 1/m^2$. Conditioned on this event along with all of \mathcal{E}_i (which by a union bound occurs with probability $1 - O(n/m^2)$), Algorithm 1 and welfare maximization differ on at most $O(n\sqrt{m \log m})$ items. This implies a $1 - o(1)$ welfare approximation. Furthermore, by Dickerson et al. [2014], for the allocation A^* chosen by the welfare maximization algorithm, $v_i(A_i^*) - v_i(A_j^*) \in \Omega(m/n)$ with high probability (recalling that values are in $[0, 1]$ without loss of generality). Thus, conditioned on this event as well, even moving $O(n\sqrt{m \log m})$ items adversarially, the allocation remains envy-free with high probability. \square

Finally, we show that Algorithm 1 can be extended to a large class of independent but *non*-identical distributions (e.g., different agents have different distributions) using similar tools to [Bai and Gözl, 2022] (who study the existence of envy-free and efficient allocations in the non-i.i.d. stochastic setting). We expand on this in Appendix A.

5 DISCUSSION

In this work, we revisited the problem of truthful fair division through the lens of stochastic valuations. We showed that moving beyond worst-case analysis allows us to bypass various impossibility results. Nonetheless, there are a number of avenues for future work.

Tightness for $n > 2$ Agents. While we provide a complete picture for the DSIC two-agent case, the case for $n > 2$ has some gaps. Our prophet inequality-based mechanisms achieve a ≈ 0.745 -approximation, but we do not know if this is tight. Proving impossibility results for $n > 2$ is significantly more challenging than for the two-agent case. For $n = 2$, we leveraged the structural characterization of Amanatidis et al. [2017] which restricts all DSIC mechanisms to “picking-exchange” rules. No such characterization exists for $n > 2$. In the general case, DSIC mechanisms can be much more complex; for instance, the mechanism used to divide items between agents 1 and 2 could depend on the report of agent 3 in intricate ways. Without a structural foothold, ruling out the existence of high-welfare DSIC mechanisms for more than two agents remains a difficult theoretical hurdle.

Correlated Valuations. Our analysis relied heavily on the independence of item values. A natural and important direction for future work is to extend these results to distributions where values are correlated across agents (although still independent across items), for example, from Benadè et al. [2024b]. One promising angle is to view the distribution of item values as a cake in the classic cake-cutting problem. For two agents, this analogy can be used to show that a Pick- r style mechanism for some ordering of the agents achieves envy-freeness with high probability. We expand on this in Appendix B. However, considering Pareto optimality (a more natural notion of efficiency when agents’ value magnitudes are different) or extending this to more agents seems to require new techniques.

Beyond additive valuations. Finally, it would be very compelling to extend these results to settings beyond additive valuations. In many combinatorial allocation problems, agents exhibit substitutes or complements (e.g., submodular or supermodular valuations). A major challenge here is modeling: defining natural, analytically tractable distributions over complex valuation functions is non-trivial. One recent example in this direction is the work of Benadè et al. [2024a], which studies the existence of envy-free allocations under a beyond-additive stochastic model. However, incorporating incentives in such non-additive stochastic settings remains an open problem.

REFERENCES

- Rediet Abebe, Richard Cole, Vasilis Gkatzelis, and Jason D Hartline. 2020. A truthful cardinal mechanism for one-sided matching. In *Proceedings of the fourteenth annual ACM-SIAM symposium on discrete algorithms*. SIAM, 2096–2113.
- Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. 2017. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 545–562.
- Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris. 2021. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science* 863 (2021), 69–85.
- Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2024. Allocating indivisible goods to strategic agents: Pure nash equilibria and fairness. *Mathematics of operations research* 49, 4 (2024), 2425–2445.
- Georgios Amanatidis, Georgios Birmpas, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2023. Round-robin beyond additive agents: Existence and fairness of approximate equilibria. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 67–87.
- Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. On truthful mechanisms for maximin share allocations. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 31–37.
- Makis Arsenis and Robert Kleinberg. 2022. Individual fairness in prophet inequalities. *arXiv preprint arXiv:2205.10302* (2022).
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5119–5126.
- Moshe Babaioff and Noam Manaker Morag. 2025. On Truthful Mechanisms without Pareto-efficiency: Characterizations and Fairness. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 447–447.
- Yushi Bai and Paul Gözl. 2022. Envy-Free and Pareto-Optimal Allocations for Agents with Asymmetric Random Valuations. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI*.
- Siddharth Barman and Paritosh Verma. 2022. Truthful and fair mechanisms for matroid-rank valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4801–4808.
- Gerdus Benadè, Daniel Halpern, Alexandros Psomas, and Paritosh Verma. 2024a. On the Existence of Envy-Free Allocations Beyond Additive Valuations. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 1287–1287.
- Gerdus Benadè, Aleksandr M Kazachkov, Ariel D Procaccia, Alexandros Psomas, and David Zeng. 2024b. Fair and efficient online allocations. *Operations Research* 72, 4 (2024), 1438–1452.
- Xiaolin Bu and Biaoshuai Tao. 2024. Truthful and Almost Envy-Free Mechanism of Allocating Indivisible Goods: the Power of Randomness. In *Proceedings of the 66th Symposium on Foundations of Computer Science*.
- Richard Cole, Vasilis Gkatzelis, and Gagan Goel. 2013. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 251–268.
- José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 169–186.
- John Dickerson, Jonathan Goldman, Jeremy Karp, Ariel Procaccia, and Tuomas Sandholm. 2014. The computational rise and fall of fairness. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 28.
- Lester E Dubins and Edwin H Spanier. 1961. How to cut a cake fairly. *The American Mathematical Monthly* 68, 1P1 (1961), 1–17.
- Eric Friedman, Ali Ghodsi, and Christos-Alexandros Psomas. 2014. Strategyproof allocation of discrete jobs on multiple machines. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 529–546.
- Eric J Friedman, Vasilis Gkatzelis, Christos-Alexandros Psomas, and Scott Shenker. 2019. Fair and efficient memory sharing: Confronting free riders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1965–1972.
- Jugal Garg, Vishnu V Narayan, and Yuang Eric Shen. 2025. Designing Truthful Mechanisms for Asymptotic Fair Division. *arXiv preprint arXiv:2512.10892* (2025).

- Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *8th USENIX symposium on networked systems design and implementation (NSDI 11)*.
- Vasilis Gkatzelis, Alexandros Psomas, Xizhi Tan, and Paritosh Verma. 2024. Getting more by knowing less: Bayesian incentive compatible mechanisms for fair division. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. 2807–2815.
- Daniel Halpern, Ariel D Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair division with binary valuations: One rule to rule them all. In *International Conference on Web and Internet Economics*. Springer, 370–383.
- Daniel Halpern, Alexandros Psomas, Paritosh Verma, and Daniel Xie. 2025. Online Envy Minimization and Multicolor Discrepancy: Equivalences and Separations. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 188–188.
- Eden Hartman, Erel Segal-Halevi, and Biaoshuai Tao. 2025. It’s Not All Black and White: Degree of Truthfulness for Risk-Avoiding Agents. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 996–1016.
- Theodore P Hill and Robert P Kertz. 1982. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability* (1982), 336–345.
- Robert P Kertz. 1986. Stop rule and supremum expectations of iid random variables: a complete comparison by conjugate duality. *Journal of multivariate analysis* 19, 1 (1986), 88–112.
- Ulrich Krengel and Louis Sucheston. 1978. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces* 4, 197–266 (1978), 1–2.
- Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. 125–131.
- Pasin Manurangsi and Warut Suksompong. 2021. Closing gaps in asymptotic fair division. *SIAM Journal on Discrete Mathematics* 35, 2 (2021), 668–706.
- Evangelos Markakis and Christos-Alexandros Psomas. 2011. On worst-case allocations in the presence of indivisible goods. In *International Workshop on Internet and Network Economics*. Springer, 278–289.
- Pascal Massart. 1990. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability* (1990), 1269–1283.
- Timo Mennle and Sven Seuken. 2021. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory* 191 (2021), 105144.
- Elchanan Mossel and Omer Tamuz. 2010. Truthful fair division. In *International Symposium on Algorithmic Game Theory*. Springer, 288–299.
- Josué Ortega and Erel Segal-Halevi. 2022. Obvious manipulations in cake-cutting. *Social Choice and Welfare* 59, 4 (2022), 969–988.
- Szilvia Pápai. 2000. Strategyproof multiple assignment using quotas. *Review of Economic Design* 5, 1 (2000), 91–105.
- Szilvia Pápai. 2001. Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory* 3, 3 (2001), 257–271.
- David C Parkes, Ariel D Procaccia, and Nisarg Shah. 2015. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)* 3, 1 (2015), 1–22.
- Alexandros Psomas and Paritosh Verma. 2022. Fair and efficient allocations without obvious manipulations. *Advances in Neural Information Processing Systems* 35 (2022), 13342–13354.
- James Schummer. 1996. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare* 14, 1 (1996), 47–56.
- Moshe Shaked and J George Shanthikumar. 2007. *Stochastic orders*. Springer.
- Walter Stromquist. 1980. How to cut a cake fairly. *The American Mathematical Monthly* 87, 8 (1980), 640–644.

A NON-IDENTICAL DISTRIBUTIONS

In this section, we extend our results to the setting where agents' valuations are independent but not necessarily identically distributed. We adopt the model of Bai and Gözl [2022]. Instead of all agents drawing values from the same distribution \mathcal{D} , there are n distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$, and each agent i 's value for item j is drawn independently as $v_{i,j} \sim \mathcal{D}_i$.

We assume (as Bai and Gözl did) that the distributions satisfy regularity conditions: The distributions admit probability density functions (PDFs) f_1, \dots, f_n that satisfy

- (1) Interval Support: For each i , there exists an interval $[a_i, b_i] \subseteq [0, 1]$ such that $f_i(x) > 0$ if and only if $x \in [a_i, b_i]$.
- (2) Bounded PDFs: There exist constants $p, q > 0$ such that for all i and all x where $f_i(x) > 0$, we have $p \leq f_i(x) \leq q$.

In this setting, maximizing social welfare is no longer an appropriate notion of efficiency, as agents' values may be on an arbitrarily different scales. Instead, we focus on *Pareto optimality (PO)*. An allocation (A_1, \dots, A_n) is PO if the values $(u_1(A_1), \dots, u_n(A_n))$ are not *Pareto dominated*, i.e., there is no allocation (A'_1, \dots, A'_n) such that $u_i(A'_i) \geq u_i(A_i)$ for all i , with strict inequality for at least one agent. We say an allocation is α -Pareto-Optimal (α -PO) if $(u_1(A_1)/\alpha, \dots, u_n(A_n)/\alpha)$ is not Pareto dominated.

Bai and Gözl show that under these regularity conditions, there exist weights w_1, \dots, w_n such that maximizing *weighted* welfare, assigning item j to an agent maximizing $w_i \cdot v_{i,j}$, yields an allocation that is PO and EF with high probability as m grows large. These weights are chosen such that the probability that any specific agent i has the highest weighted value for an item is exactly $1/n$. As in Dickerson et al. [2014], they in fact show that not only will the allocation be EF, but $u_i(A_i) - u_i(A_j) \in \Omega(m/n)$ with high probability.

We show that under the same conditions as Bai and Gözl [2022], a weighted version of the ranking algorithm is still BIC with high probability. Note that the algorithm will now depend on the distributions \mathcal{D} .

ALGORITHM 2: Weighted Ranking Algorithm for CDFs F_1, \dots, F_n .

Input: $\{b_i\}$

$w_1, \dots, w_n \leftarrow$ Weights for F_1, \dots, F_n from Bai and Gözl [2022].

$\sigma_i \leftarrow$ Ranking of bids b_i , breaking ties arbitrarily $\forall i$

$A_i \leftarrow \emptyset \quad \forall i$

for $j \leftarrow 1$ **to** m **do**

$\hat{d}_{ij} \leftarrow F^{-1}(1 - \text{rank}(\sigma_i, j)/(m + 1))$

$v \leftarrow \max_i w_i \cdot \hat{d}_{ij}$

$i^* \leftarrow \min_{i: v = w_i \cdot \hat{d}_{ij}} i$

$A_{i^*} \leftarrow A_{i^*} \cup \{j\}$

end

THEOREM A.1. *For all interval-support and PDF-bounded distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$ with CDFs F_1, \dots, F_n , Algorithm 2 is BIC, EF, and $(1 - \epsilon)$ -PO with high probability.*

PROOF. We first establish the BIC property. The logic follows the same as Theorem 4.1. From the perspective of agent i , the values (and thus the rankings) of all other agents are drawn independently across items. Consequently, the distribution of the "competing" weighted virtual

values— $\max_{k \neq i} \{w_k F_k^{-1}(q_{k,j})\}$ —is identical for every item j . This implies that for any rank position $r \in [m]$, there exists a fixed probability p_r that agent i wins the item they place at rank r . Crucially, because the CDF F_i is monotone non-decreasing, placing an item at a higher rank (smaller r) results in a higher implied quantile $q_{i,j}$, a higher virtual value $\tilde{v}_{i,j}$, and strictly higher winning probability p_r . Thus, $p_1 \geq p_2 \geq \dots \geq p_m$. To maximize expected utility $\sum_j v_{i,j} \cdot \mathbb{P}(\text{win } j)$, agent i must assign higher probabilities to higher-valued items. Truthful reporting (ranking by value) is therefore a dominant strategy in the induced game, establishing BIC.

Next, we address fairness and efficiency. We show that the allocation produced by Algorithm 2 differs from the ideal weighted-welfare maximization on only a vanishing fraction of items. The proof is again similar to Theorem 4.1, except we need to handle the different values with slightly more care.

First, we claim it is without loss of generality to assume all weights $w_i = 1$. Indeed, instead of working directly with these distributions, we will consider the *weighted* distributions $w_1 \mathcal{D}_1, \dots, w_n \mathcal{D}_n$. Note that these will all be p' - q' -PDF bounded for $p' = p/\max_i w_i$ and $q' = q/\min_i w_i$. Our proof will hold for these weighted distributions, and both of these lead to indistinguishable algorithms.

Let \hat{F}_i denote the empirical CDF of agent i 's values. By the DKW inequality, with probability $1 - O(1/m^2)$, we have $\sup_x |\hat{F}_i(x) - F_i(x)| \leq \sqrt{(\ln m)/m}$. Because the PDFs are bounded above by q' , the error in the value space is also bounded. Specifically, for any quantile x , the difference between the true value and the value implied by the empirical rank is bounded by $q' \cdot O(\sqrt{(\ln m)/m})$.

Consider an item j . Algorithm 2 assigns j to the agent with the highest virtual value. An error occurs only if the agent with the highest true value, say i^* , loses to some agent k because of the ranking error. This requires the true values $v_{i^*,j}$ and $v_{k,j}$ to be within $O(\sqrt{(\ln m)/m})$ of each other. Since the difference of two independent random variables with bounded PDFs also has a bounded PDF, the probability that $|v_{i^*,j} - v_{k,j}| \leq \epsilon$ is $O(\epsilon)$. Substituting our error bound, the probability that a specific item is allocated "incorrectly" (compared to exact weighted-welfare maximization) is $O(\sqrt{(\ln m)/m})$. Summing over all items and applying Hoeffding's inequality, the total number of mismatched items is $O(\sqrt{m \ln m})$ with high probability. As established by Dickerson et al. [2014] and Bai and Gözl [2022], the ideal weighted allocation is robust to adversarial changes of $o(m)$ items. Thus, Algorithm 2 retains the properties of Envy-Freeness and $(1 - \epsilon)$ -Pareto Optimality with high probability. \square

B A CAKE-CUTTING APPROACH TO CORRELATED VALUES

In this section, we consider the setting where item values are drawn from a correlated distribution \mathcal{D} . That is, the valuation vectors $(v_{i,1}, \dots, v_{i,n}) \sim \mathcal{D}$ are correlated across agents, though we maintain the assumption that values are independent and identically distributed across items. This matches the model of, e.g., Benadè et al. [2024b].

We focus on the case of $n = 2$ agents. We show that under a mild assumption on the correlation structure of \mathcal{D} —specifically, that the agents' quantile rankings are not perfectly dependent—there exists a Pick- r mechanism that yields an allocation that is Envy-Free with high probability.

Correlation Assumption. We assume that for any agent i and any quantile threshold $q \in (0, 1)$, the set of items agent i ranks above the q -th quantile is not identical to the set of items agent j ranks above the q -th quantile with probability 1. Formally, if F_i is the marginal CDF for agent i , we assume that for any q where the upper tail has non-zero probability, $\mathbb{P}(F_j(v_j) \geq q \mid F_i(v_i) \geq q) < 1$.

Two agents. Without loss of generality, normalize the distribution \mathcal{D} such that $\mathbb{E}[v_{i,j}] = 1$ for all agents i . We define a critical quantile threshold $\tau_i \in [0, 1]$ for each agent i . This threshold represents the quantile cutoff required for the agent to achieve exactly half of their total expected value. Formally, τ_i satisfies:

$$\mathbb{E}_{v \sim F_i}[v \cdot \mathbb{I}(F_i(v) \geq \tau_i)] = \frac{1}{2} \mathbb{E}[v] = \frac{1}{2}.$$

Let i be an agent such that $\tau_i \geq \tau_j$ for $j \neq i$. We run the Pick- r mechanism with agent i picking first, setting r to be the fraction of items with mass $(1 - \tau_i) + \varepsilon$ for sufficiently small ε . The key observation is that agent i will value the items they receive at strictly more than half of their total value (and will therefore not envy j). Furthermore, because $\tau_i \geq \tau_j$ and valuations are not perfectly correlated, agent j 's expected value for agent i 's bundle is strictly less than $1/2$. Hence, for sufficiently large m , concentration inequalities ensure that these expected preferences translate to an EF allocation with high probability. Furthermore, for two agents, envy-freeness implies a $1/2$ -approximation to maximum welfare (as all agents are receiving at least $1/2$ of their maximum possible utility).

Extensions to $n > 2$ agents and cake cutting. For $n > 2$ agents, the problem becomes significantly more challenging. Our procedure for $n = 2$ can be viewed as a discrete analogue of the Dubins-Spanier Moving-Knife procedure [Dubins and Spanier, 1961], where a knife moves across “quantile space” rather than across the physical cake. Under mild correlation assumptions, this approach can be generalized to guarantee a proportional allocation using a DSIC mechanism for any number of agents. In standard cake-cutting, *contiguous* envy-free allocations always exist for any n [Stromquist, 1980]. It remains an open question whether this result extends to *quantile-contiguous* allocations, where there is an ordering of the agents $\sigma(1), \dots, \sigma(n)$, where each agent $\sigma(i)$ receives the portion of the cake with highest density among the remaining pieces of $\sigma(i), \dots, \sigma(n)$. The existence of such an allocation would ensure that there is a DSIC algorithm for n agents that is EF with high probability, even with correlated valuations. Furthermore, it is unclear how to combine these properties with efficiency.

C PROOF OF LEMMA 3.2

PROOF OF LEMMA 3.2. We will proceed by casework on h .

For $h = 1$, a single high-value draw exists for one agent-item pair, (i, j) . To realize this value in the final allocation, agent i must receive item j . We will show that with probability at least $1/4$ (over the choice of (i, j) and other item values), agent i does not receive item j .

For intuition, we first show why this is the case for an equal-sized trade, where $\ell = k$. Suppose without loss of generality that $i = 1$. With probability $1/2$, j is not endowed to i . From agent 2's perspective, all items are worth either 0 or 1, and due to the equally sized bundle, they will prefer agent 1's bundle only with probability $1/2$. Thus, with probability $1/4$, the high-value item is not endowed to 1 and agent 2 is not willing to trade, which means that with probability $1/4$ agent i does not receive item j , which proves Lemma 3.2 for the case when $\ell = k$ and $i = 1$.

Next, we extend this to unbalanced trades with $\ell \neq k$ but still with $i = 1$. Let τ be the probability that, if we sample values for bundles of size ℓ and k from $\text{Bern}(q) + \text{Unif}[0, \delta]$, the ℓ -bundle is more valuable than the k -bundle. Since ties occur with probability 0, this means that the k -bundle is more valuable with probability $1 - \tau$. Without loss of generality, assume $\ell \geq k$, which implies that $\tau \geq 1/2$.

Note that if agent 1 has the high value, then agent 2 will be willing to trade with probability $1 - \tau$; conversely, if agent 2 has the high value, then agent 1 will be willing to trade with probability τ .

There are two events of interest: (i) Agent 2 has high value, it is in agent 1's bundle, and agent 1 is not willing to trade and (ii) Agent 1 has the high value, it is in agent 2's bundle, and agent 2 is not willing to trade. The first occurs with probability $\ell/(2s) \cdot \tau$ and the second occurs with probability $k/(2s) \cdot (1 - \tau)$. These are disjoint events, so the probability that at least one occurs is

$$\ell/(2s) \cdot \tau + k/(2s) \cdot (1 - \tau) = \frac{\ell\tau + k(1 - \tau)}{2s} \geq \frac{\ell/2 + k/2}{2s} = 1/4,$$

where the inequality holds because $\tau \geq 1/2$ and $\ell \geq k$.

Thus, the single high-value item i is not received by agent j with probability at least $1/4$. Therefore, the expected welfare contribution of this single item is at most $(1 - 1/4) \left(\frac{1-q}{p} + \delta \right) \leq \frac{3}{4} \cdot h \cdot \frac{1-q}{p} + \delta h$.

We will now show the $h > 1$ case. Fix any $h > 1$. Let r be the number of high-value items in endowed bundles, and $h - r$ be the remaining number in non-endowed bundles. The welfare contribution from high-value items in this case is bounded by

$$\max(r, h - r) \cdot \frac{1 - q}{p} + \delta h = \left(\frac{h}{2} + \left| r - \frac{h}{2} \right| \right) \frac{1 - q}{p} + \delta h. \quad (4)$$

Therefore, to bound the expected welfare contribution from high-value items in this case, we will bound $\mathbb{E}[|r - \frac{h}{2}|] \leq \frac{h}{4}$.

The distribution of r is a hypergeometric with parameters $N = 2s, K = s, n = h$. Indeed, we are sampling h pairs (i, j) without replacement, and exactly half are in endowed bundles. However, for our purposes, it turns out that this distribution is convex dominated by a $r' \sim \text{Bin}(h, 1/2)$ distribution. Since $|r - \frac{h}{2}|$ is convex in r , $\mathbb{E}_{r \sim \text{HypGeo}(2s, s, h)}[|r - \frac{h}{2}|] \leq \mathbb{E}_{r' \sim \text{Bin}(h, 1/2)}[|r' - \frac{h}{2}|]$ (see Theorem 3.A.39 of [Shaked and Shanthikumar, 2007]). We will instead upper bound the binomial version.

For any $h \geq 4$, by Jensen's inequality:

$$\mathbb{E}[|r' - h/2|] \leq \sqrt{\mathbb{E}[(r' - h/2)^2]} = \sqrt{\text{Var}(r')} = \sqrt{h/4} = \sqrt{h}/2 \leq h/4.$$

For $h = 2$: $\mathbb{E}[|r' - 1|] = \frac{1}{4}|0 - 1| + \frac{1}{2}|1 - 1| + \frac{1}{4}|2 - 1| = 2/4$.

For $h = 3$: $\mathbb{E}[|r' - 3/2|] = \frac{1}{8}|0 - 3/2| + \frac{3}{8}|1 - 3/2| + \frac{3}{8}|2 - 3/2| + \frac{1}{8}|3 - 3/2| = 3/4$.

Thus, for all $h > 1$ we have that

$$\mathbb{E}[|r - h/2|] \leq \mathbb{E}[|r' - h/2|] \leq h/4.$$

Therefore, by Equation (4), we can conclude that the expected welfare contributed for high-value items is bounded by

$$\begin{aligned} \mathbb{E} \left[\left(\frac{h}{2} + \left| r - \frac{h}{2} \right| \right) \frac{1 - q}{p} + \delta h \right] &\leq \left(\frac{h}{2} + \mathbb{E}[|r - \frac{h}{2}|] \right) \frac{1 - q}{p} + \delta h \\ &\leq \left(\frac{h}{2} + \mathbb{E}[|r' - \frac{h}{2}|] \right) \frac{1 - q}{p} + \delta h \\ &\leq \left(\frac{h}{2} + \frac{h}{4} \right) \frac{1 - q}{p} + \delta h \\ &= \frac{3h}{4} \frac{1 - q}{p} + \delta h \end{aligned}$$

as desired to complete the proof of Lemma 3.2. \square

D PROOF OF LEMMA 3.6

PROOF OF LEMMA 3.6. First, observe that in QT-s, the probability that agent i receives a specific item is exactly r_i . This follows by strong induction: the probability the item reaches agent i is $1 - \sum_{k=1}^{i-1} r_k$, and conditioned on reaching i , it is taken with probability s_i . The product of these terms is r_i .

We begin with (1). Let (A_1, \dots, A_n) denote the allocation of Pick-r and (A'_1, \dots, A'_n) denote the allocation of QT-s.

By the above argument, we have that $|A'_i|$ follows a binomial distribution with expectation $\mathbb{E}[|A'_i|] = r_i m$. Let $C = \sqrt{m \log(m)} + 1 \in O(\sqrt{m \log(m)})$. Applying Hoeffding's inequality, $\Pr[||A'_i| - r_i \cdot m| > C - 1] \leq 2 \exp(-2(\sqrt{m \log(m)})^2/m) = 2/m^2$. Furthermore, $||A_i| - r_i \cdot m| \leq 1$. Combining these, under the Hoeffding event $(||A'_i| - r_i \cdot m| > C - 1)$, we have that $||A'_i| - |A_i|| \leq C$. By a union bound over n agents, this holds for all agents with probability at least $1 - O(n/m^2)$. We will condition on the event that $||A_i| - |A'_i|| \leq C$, for all i , for the remainder of the proof.

Fix arbitrary quantiles $q_{i,j}$ such that the event holds, and condition on the probability 1 event that they are all distinct.

Fix an agent i . Let U_i and U'_i be the set of available items for agent i in Pick-r and QT-s respectively (where $U_i = \mathcal{M} \setminus \bigcup_{k=1}^{i-1} A_k$ and $U'_i = \mathcal{M} \setminus \bigcup_{k=1}^{i-1} A'_k$). Let $d_i = |U_i \Delta U'_i|$ be the number of items available in one mechanism but not the other. We define $U_i^* = U_i \cap U'_i$ as the set of items available in both mechanisms.

Both mechanisms are greedy with respect to the same values (quantiles). Therefore, when choosing from the common set U_i^* :

- Pick-r selects the top $y_i = |A_i \cap U_i^*|$ items from U_i^* .
- QT-s selects the top $y'_i = |A'_i \cap U_i^*|$ items from U_i^* .

Since both select the best items from the same set, the selected subsets are nested. Specifically, if $y_i \leq y'_i$, then $(A_i \cap U_i^*) \subseteq (A'_i \cap U_i^*)$, and vice versa. The number of disagreements on these common items is exactly the difference in count:

$$|(A_i \cap U_i^*) \Delta (A'_i \cap U_i^*)| = |y_i - y'_i|.$$

We now bound the total allocation disagreement $|A_i \Delta A'_i|$. Let $x_i = |A_i \setminus U_i^*|$ be the items agent i took from $U_i \setminus U'_i$ (items only available in Pick-r). Let $x'_i = |A'_i \setminus U_i^*|$ be the items agent i took from $U'_i \setminus U_i$ (items only available in QT-s). Note that $x_i \leq |U_i \setminus U'_i|$ and $x'_i \leq |U'_i \setminus U_i|$, so $x_i + x'_i \leq d_i$.

The total number of items where allocations A_i and A'_i disagree is the sum of disagreements from the disjoint sets and the common set:

$$|A_i \Delta A'_i| = x_i + x'_i + |y_i - y'_i|.$$

We bound $|y_i - y'_i|$ using the bundle size constraint. We know $|A_i| = x_i + y_i$ and $|A'_i| = x'_i + y'_i$. From our conditioning, we established $||A_i| - |A'_i|| \leq C$. Substituting the sums:

$$|(x_i + y_i) - (x'_i + y'_i)| \leq C \implies |y_i - y'_i| \leq C + |x_i - x'_i|.$$

Substituting this back into the disagreement equation:

$$|A_i \Delta A'_i| = x_i + x'_i + |y_i - y'_i| \leq x_i + x'_i + C + |x_i - x'_i|.$$

Since $x_i, x'_i \geq 0$, we have $|x_i - x'_i| \leq x_i + x'_i$. Thus:

$$|A_i \Delta A'_i| \leq 2(x_i + x'_i) + C \leq 2d_i + C.$$

Furthermore, from these values, we can derive d_{i+1} (the disagreement passed to the next agent), which will allow use to define a recurrence. The set of available items updates as $U_{i+1} = U_i \setminus A_i$ and $U'_{i+1} = U'_i \setminus A'_i$. An item is in the symmetric difference $U_{i+1} \Delta U'_{i+1}$ if it was already a disagreement in d_i and was *not* picked, or if it was in the common set U_i^* and was picked by exactly one mechanism (creating a new disagreement).

$$d_{i+1} = \underbrace{(d_i - (x_i + x'_i))}_{\text{Unpicked input disagreements}} + \underbrace{|y_i - y'_i|}_{\text{New disagreements from } U_i^*}.$$

Substituting $|y_i - y'_i| \leq C + |x_i - x'_i|$:

$$d_{i+1} \leq d_i - (x_i + x'_i) + C + |x_i - x'_i|.$$

Using $|x_i - x'_i| \leq x_i + x'_i$, the terms cancel to satisfy:

$$d_{i+1} \leq d_i + C.$$

Since $d_1 = 0$ (both start with all items), it follows by induction that $d_i \leq (i-1)C$.

Plugging this in above, we have that $|A_i \Delta A'_i| \leq (2i-1)C$. Summing over all n agents:

$$\sum_{i=1}^n |A_i \Delta A'_i| \leq \sum_{i=1}^n (2i-1)C = n^2 C.$$

Substituting $C = O(\sqrt{m \log m})$, the total number of disagreement items is $O(n^2 \sqrt{m \log m})$.

We now turn to (2). Consider the expected welfare of QT-s. Recall that we showed that each agent i receives each item with probability r_i . These allocation decisions are made independently. Conditioned on item j being allocated to agent i , the value v_{ij} is guaranteed to be at least $F^{-1}(1-s_i)$. Thus, the expected welfare contribution of a single item is:

$$w_{\text{avg}} = \sum_{i=1}^n r_i \cdot \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)].$$

By linearity of expectation, the total expected welfare of QT-s is $m \cdot w_{\text{avg}} = \alpha \cdot m \cdot \mathbb{E}[\max_i X_i]$ (by using the equality $\alpha = (\sum_{i=1}^n r_i \cdot \mathbb{E}[X \mid X \geq F^{-1}(\tau_i)]) / \mathbb{E}[\max_i X_i]$).

Now, consider Pick-r. Conditioned on the event that the allocations differ by at most $O(n^2 \sqrt{m \log m})$ items, the difference in welfare can be at most $O(n^2 \sqrt{m \log m})$. Furthermore, in the remaining case (that occurs with $O(n/m^2)$ probability), the welfare difference is at most m . Hence, the contribution to the expected welfare is at most $m \cdot O(n/m^2) = O(n/m)$.

Dividing by the expected maximum welfare of $m \cdot \mathbb{E}[\max_i X_i]$, we get an approximation of

$$\frac{\alpha \cdot m \cdot \mathbb{E}[\max_i X_i] - O(n^2 \sqrt{m \log m}) - O(n/m)}{m \cdot \mathbb{E}[\max_i X_i]} = \alpha - o(1).$$

Finally, we prove (3). We must show that for any pair of agents i, j , the probability that agent i envies agent j approaches 0 as $m \rightarrow \infty$.

We first analyze the valuations in QT-s. Since valuations are independent across items, we calculate the expected value agent i assigns to a single random item allocated to i versus one allocated to j .

- *Value for Own Bundle (A_i):* In QT-s, agent i receives an item with probability r_i . Conditioned on receipt, the item's value is at least τ_i . Thus, the expected total value is:

$$\mathbb{E}[v_i(A_i)] = m \cdot r_i \cdot \mathbb{E}[X \mid X \geq \tau_i].$$

- *Value for Other Bundle (A_j) when $i < j$:* Since i precedes j in the priority sequence, j can only receive an item if i has already rejected it (implying $v_{i,k} < \tau_i$) and j subsequently accepts it. Since valuations are independent across agents, the event that j accepts the item provides no information about i 's value other than the fact that i rejected it. Therefore, i 's expected value for an item in j 's bundle is conditioned only on $X < \tau_i$:

$$\mathbb{E}[v_i(A_j)] = m \cdot r_j \cdot \mathbb{E}[X \mid X < \tau_i].$$

- *Value for Other Bundle (A_j) when $i > j$:* Here, j considers items before i . If j selects an item, it is removed before i sees it. Since i 's valuation is independent of j 's, the event that j picked the item is independent of i 's specific value for it. Thus, i 's expected value for items in A_j follows the unconditional distribution:

$$\mathbb{E}[v_i(A_j)] = m \cdot r_j \cdot \mathbb{E}[X].$$

The conditions given in the lemma statement correspond exactly to the requirement that agent i strictly prefers their expected bundle to agent j 's. Let $\Delta_{ij} = \mathbb{E}[v_i(A_i)] - \mathbb{E}[v_i(A_j)]$. In both cases, the strict inequality implies that $\Delta_{ij} = c \cdot m$ for some constant $c > 0$.

In QT-s, both $v_i(A_i)$ and $v_i(A_j)$ are sums of independent random variables bounded in $[0, 1]$. By Hoeffding's inequality, the realized difference $v_i(A_i) - v_i(A_j)$ concentrates around its expectation Δ_{ij} with deviations bounded by $O(\sqrt{m \log m})$ with high probability.

Finally, we transfer this result to the Pick- r mechanism. From Part (1), we know that with high probability, the allocations of Pick- r and QT-s differ on at most $K = O(n^2 \sqrt{m \log m})$ items. Since values are bounded by 1, this difference affects the total valuation by at most K . Thus, for the Pick- r mechanism:

$$v_i(A_i) - v_i(A_j) \geq \underbrace{c \cdot m}_{\text{Expected Gap}} - \underbrace{O(\sqrt{m \log m})}_{\text{Concentration Noise}} - \underbrace{O(n^2 \sqrt{m \log m})}_{\text{Coupling Noise}}.$$

For sufficiently large m , the linear term $c \cdot m$ dominates the sub-linear noise terms. Therefore, $v_i(A_i) > v_i(A_j)$ holds with probability approaching 1. \square

E PROOF OF THEOREM 3.7

PROOF OF THEOREM 3.7. For the rest of this proof, let $r = \frac{2-\sqrt{2}}{2}$. Fix any distribution \mathcal{D} .

By Part 3 of Lemma 3.6, the Pick- r mechanism is envy-free with high probability if

$$r \cdot \mathbb{E}_{X \sim \mathcal{D}}[X \mid F_{\mathcal{D}}(X) \geq 1 - r] > (1 - r) \cdot \mathbb{E}_{X \sim \mathcal{D}}[X \mid F_{\mathcal{D}}(X) < 1 - r],$$

Note that because $r < 0.5$, the other condition from Part 3 of Lemma 3.6 holds immediately.

By Proposition 3.3, if the above equation holds then we have that Pick- $(\frac{2-\sqrt{2}}{2})$ is both envy-free with high probability and a $\frac{2+\sqrt{2}}{4} - o(1)$ approximation to welfare, proving the desired result in this case.

Now suppose that

$$r \mathbb{E}_{X \sim \mathcal{D}}[X \mid F_{\mathcal{D}}(X) \geq 1 - r] \leq (1 - r) \mathbb{E}_{X \sim \mathcal{D}}[X \mid F_{\mathcal{D}}(X) < 1 - r].$$

We will show (later in the proof) that the above equation implies that Pick- $(\frac{1}{2} - \frac{1}{m^{0.25}})$ must be a $7/8 - o(1) \geq \frac{2+\sqrt{2}}{4} - o(1)$ approximation to optimal social welfare. Furthermore, for sufficiently large m and any distribution \mathcal{D} with non-0 variance, the Pick- $(\frac{1}{2} - \frac{1}{m^{0.25}})$ mechanism is envy-free with high probability by Hoeffding's Inequality. This is because agent 2 is allocated an extra $2m^{0.75} \pm o(1)$ items and therefore prefers their bundle with high probability by Hoeffding's inequality. Agent 1 is getting $2m^{0.75} \pm o(1)$ fewer items than agent 2, but we can still show that they are envy-free with high probability as follows. The non-zero variance of \mathcal{D} means that there must be some constant $\varepsilon > 0$ such that in expectation, agent 1 values their $(\frac{1}{2} - \frac{1}{m^{0.25}})m$ chosen items εm more in expectation than agent 1 values the items they leave for agent 2. Again by Hoeffding's Inequality, with high probability, this means that agent 1 strictly prefers their bundle even if agent 2 receives $2m^{0.75}$ more items (because $\varepsilon m > O(m^{0.75})$ for sufficiently large m). This means that with high probability, both agents do not envy each other, so Pick- $(\frac{1}{2} - \frac{1}{m^{0.25}})$ is EF with high probability, thus completing the proof of the main theorem in this case as well.

The rest of the proof will be spent showing that

$$r \mathbb{E}_{X \sim \mathcal{D}}[X \mid F_{\mathcal{D}}(X) \geq 1 - r] \leq (1 - r) \mathbb{E}_{X \sim \mathcal{D}}[X \mid F_{\mathcal{D}}(X) < 1 - r].$$

implies that Pick- $(\frac{1}{2} - \frac{1}{m^{0.25}})$ must be a $7/8 - o(1)$ approximation to optimal social welfare

Rewriting the expectations as integrals gives

$$r \int_0^\infty \Pr(X > s \mid F_{\mathcal{D}}(X) \geq 1 - r) ds \leq (1 - r) \int_0^\infty \Pr(X > s \mid F_{\mathcal{D}}(X) < 1 - r) ds,$$

which simplifies to

$$\int_0^\infty \min(r, 1 - F_{\mathcal{D}}(s)) ds \leq \int_0^\infty \max(0, 1 - r - F_{\mathcal{D}}(s)) ds.$$

Finally, subtracting the RHS gives

$$\int_0^\infty (r - |F_{\mathcal{D}}(s) - (1 - r)|) ds \leq 0. \quad (5)$$

Define A as the LHS of this equation, i.e.

$$A := \int_0^\infty (r - |F_{\mathcal{D}}(s) - (1 - r)|) ds.$$

Let Q_{ALG} be the Pick- $(\frac{1}{2} - \frac{1}{m^{0.25}})$ mechanism and Q_{OPT} be the welfare-maximizing mechanism. To prove that Pick- $(\frac{1}{2} - \frac{1}{m^{0.25}})$ is a $7/8 - o(1)$ approximation to optimal social welfare, it suffices to show that

$$\int_0^\infty \Pr(Q_{ALG} > F_{\mathcal{D}}(s)) ds \geq \left(\frac{7}{8} - o(1)\right) \int_0^\infty \Pr(Q_{OPT} > F_{\mathcal{D}}(s)) ds.$$

Because $A \leq 0$, to show the above equation it is sufficient to show that

$$\int_0^\infty \Pr(Q_{ALG} > F_{\mathcal{D}}(s)) + \frac{A}{4} \geq \left(\frac{7}{8} - o(1)\right) \int_0^\infty \Pr(Q_{OPT} > F_{\mathcal{D}}(s)) ds.$$

Also, as in Proposition 3.3, to show the above equation, it is sufficient to show that for all $t \in [0, 1]$,

$$\Pr(Q_{ALG} > t) + \frac{r - |t - (1-r)|}{4} \geq \left(\frac{7}{8} - o(1)\right)(1-t^2). \quad (6)$$

By the same logic as in Proposition 3.3,

$$\Pr(Q_{ALG} > t) = \begin{cases} 1/2 + (1-t)/2 \pm o(1) & \text{if } t < 1/2, \\ 1-t + (1-t)/2 \pm o(1) & \text{if } t \geq 1/2. \end{cases}$$

Therefore,

$$\Pr(Q_{ALG} > t) + \frac{r - |t - (1-r)|}{4} = \begin{cases} \frac{3+2r-t}{4} \pm o(1), & 0 \leq t < \frac{1}{2}, \\ \frac{5+2r-5t}{4} \pm o(1), & \frac{1}{2} \leq t < 1-r, \\ \frac{7-7t}{4} \pm o(1), & 1-r \leq t \leq 1. \end{cases} \quad (7)$$

Define

$$S(t) = \begin{cases} \frac{3+2r-t}{4}, & 0 \leq t < \frac{1}{2}, \\ \frac{5+2r-5t}{4}, & \frac{1}{2} \leq t < 1-r, \\ \frac{7-7t}{4}, & 1-r \leq t \leq 1. \end{cases}$$

We will spend the rest of the proof showing that $S(t) \geq \frac{7}{8}(1-t^2)$ for all $t \in [0, 1]$, which together with Equation (7) will imply Equation (6), thereby completing the proof. Define $D(T) = S(T) - \frac{7}{8}(1-t^2)$.

On each piece, $S(t)$ is affine, so $D(t)$ is a convex quadratic. This means to show that $D(T)$ is non-negative, it suffices to check its minimum on each piece.

1) $0 \leq t < \frac{1}{2}$.

$$D(t) = \frac{7}{8}t^2 - \frac{1}{4}t + \left(\frac{5-\sqrt{2}}{4} - \frac{7}{8}\right), \quad D'(t) = \frac{7}{4}t - \frac{1}{4}.$$

The critical point is at $t = \frac{1}{7} \in [0, \frac{1}{2})$, hence

$$D_{\min} = D\left(\frac{1}{7}\right) = \frac{10-7\sqrt{2}}{28} > 0.$$

2) $\frac{1}{2} \leq t < 1-r$.

$$D(t) = \frac{7}{8}t^2 - \frac{5}{4}t + \left(\frac{5+2r}{4} - \frac{7}{8}\right), \quad D'(t) = \frac{7}{4}t - \frac{5}{4}.$$

The root of $D'(t)$ is $t = \frac{5}{7} > 1-r = \frac{\sqrt{2}}{2}$, so $D'(t) < 0$ on this interval and the minimum occurs at the right endpoint:

$$D_{\min} = D(1-r) = \frac{7}{8}(r)^2 = \frac{7}{8}\left(\frac{3}{2} - \sqrt{2}\right) > 0.$$

3) $1-r \leq t \leq 1$.

$$D(t) = \frac{7}{4}(1-t) - \frac{7}{8}(1-t^2) = \frac{7}{8}(t-1)^2 \geq 0,$$

with equality only at $t = 1$.

Combining the three cases above, we have shown that $D(t) \geq 0$ for all $t \in [0, 1]$, completing the proof. \square