LEARNING SECOND-BEST MECHANISM IN MATCHING MARKETS

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the sample complexity of learning the Bayesian second-best gains from trade in matching markets with unknown value and cost distributions. For independent distributions, we give a sample-based learning algorithm that always outputs a universally dominant-strategy incentive-compatible and ex-post individually rational mechanism and, with high probability, outputs an ex-ante weakly budget-balanced mechanism whose expected gains from trade are within ϵ of the second-best benchmark. For a market with d unknown marginal distributions supported on $[0, h]$ and maximum feasible matching size r, we establish a per-marginal sample complexity of $\widetilde{\Theta}\left(drh^2/\epsilon^2\right)$, which is tight up to logarithmic factors.

The upper bound uses a critical-contact decomposition to localize the effect of marginal perturbations, with the aggregate exposure controlled by the maximum feasible matching size r. For the lower bound, we construct hard instances on star forests and show that the dependence on both d and r is unavoidable.

1 INTRODUCTION

Two-sided matching markets involve multiple buyers and sellers with private information. Let B and S be the sets of buyers and sellers, and let $E \subseteq B \times S$ be the public bipartite graph of possible trades. Feasible matchings belong to a downward-closed family $\mathcal{F} \subseteq 2^E$. We write $d := |B \cup S|$ for the number of agents and $r := \max_{N \in \mathcal{F}} |N|$ for the maximum feasible matching size. Each buyer i has a private value $v_i \in [0, h]$, each seller j has a private cost $c_j \in [0, h]$, and their types are drawn independently from a product distribution D. The platform elicits these types, selects a feasible matching, and determines transfers. With public types, the natural objective is the *first-best gains from trade* (GFT), the maximum total surplus attainable by a feasible matching. Private information introduces incentive, participation, and budget constraints. Even in bilateral trade, incentive compatibility (IC), individual rationality (IR), and weak budget balance (WBB) generally preclude the first best (Myerson & Satterthwaite, 1983). The relevant Bayesian benchmark is therefore the *second-best* (SB): the maximum expected GFT achievable by a Bayesian incentive-compatible (BIC), interim individually rational (interim-IR), and ex-ante weakly budget-balanced (ex-ante-WBB) mechanism when D is known.

This framework arises naturally in online trading platforms that bring together many buyers and sellers (Einav et al., 2016). An agent may have several potential trading partners, so assigning one trade can eliminate other feasible opportunities and requires the platform to coordinate transactions across the market. Known-distribution results provide strong GFT guarantees in increasingly general two-sided market settings and, more recently, under general matching constraints (Babaioff et al., 2018; 2026; Bei et al., 2026). In practice, however, the platform may have access only to historical samples rather than the underlying distributions, making it natural to ask whether a near-optimal mechanism can be learned from data.

For bilateral trade, Shi et al. (2026) characterize the sample complexity of learning the second-best mechanism from samples. Moving to matching markets introduces a new statistical difficulty: bilateral trade contains a single trading relationship, whereas many potential trades interact through a common feasibility constraint in a matching market. Estimation errors are therefore local to individual type distributions but can have global consequences: a small perturbation in one agent's distribution may change the selected matching and propagate across the market. It is therefore unclear whether learning must pay for the full combinatorial size of the market, or only for a much smaller set of

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Market setting	Benchmark	Guarantee	Sample complexity
General bipartite matching market (Babaioff et al., 2026)	First-best GFT	$\frac{1}{3.15}\,\mathrm{FB}$	—
General bipartite matching market (Bei et al., 2026)	Second-best GFT under SBB	$\frac{1}{2}\,\mathrm{FB}$	—
Small market with buyers and sellers (Castiglioni et al., 2026)	Fixed-price mechanism	$\mathrm{OPT}_{\mathrm{FP}} - \varepsilon$	$\frac{1}{\varepsilon^3}\log\left(\frac{nL}{\delta\varepsilon}\right)^{O(n)}$
Regular bilateral trade (Shi et al., 2026)	Second-best GFT	$\mathrm{SB} - \varepsilon$	$\widetilde{\Theta}\left(\frac{h^2}{\varepsilon^2}\right)$
General bipartite matching market This work, Theorem 1.1	Second-best GFT	$\mathrm{SB} - \varepsilon$	$\widetilde{\Theta}\left(\frac{drh^2}{\varepsilon^2}\right)$

Table 1: Comparison with closely related results.

structural parameters. This leads to our central question: *What is the sample complexity of learning the second-best GFT in a matching market?*

1.1 Our Contributions

We answer this question with matching upper and lower bounds, tight up to logarithmic factors.

Theorem 1.1. *There is a universal constant $c_0 > 0$ such that, for every $0 < \varepsilon \le c_0 rh$, the worst-case number of samples required from each individual type distribution to learn the second-best GFT is $\widetilde{\Theta}\left(\frac{drh^2}{\varepsilon^2}\right)$.*

The upper bound holds for every such market and every product distribution with Borel individual type distributions supported on $[0, h]$, with success probability at least $1 - \gamma$ for arbitrary $\gamma \in (0, 1)$ and logarithmic dependence on $1/\gamma$. Table 1 compares our result with the closest recent results on gains from trade and its statistical learning in two-sided markets.

We next explain the main ideas behind our results. The straightforward estimate-then-design approach is brittle: even a small estimation error for one agent can change the globally optimal matching. Instead of tracking this unstable optimizer, we allow the optimal matching to change with the distribution and compare the resulting optimal Lagrangian values. To do so, we introduce a one-dimensional Lagrangian frontier for the ex-ante WBB constraint:

$$\mathrm{SB}(D) = \min_{\lambda \ge 0} \sup_{M} \{\mathrm{GFT}_D(M) + \lambda\,\mathrm{Rev}_D(M)\}.$$

Global-to-local reduction. A small perturbation of one type distribution may change the globally optimal matching, so controlling the learned mechanism directly would appear to require uniform control over a large combinatorial family. Our analysis instead compares optimal frontier values along a sequence of product distributions that differ in one coordinate at a time, reoptimizing after every replacement. For each coordinate, conditioning on the other reports reduces the comparison to a one-dimensional threshold event for a single agent. We then converts the resulting loss into a local distributional error, weighted by the agent's exposure in the selected matching. Since at most r trades can be active simultaneously, these local contributions aggregate according to the matching size rather than the number of feasible edges or matchings.

Calibrating the learned frontier. Ex-ante WBB constrains only the platform's expected revenue, so some realizations may still run at a deficit. This creates an additional learning difficulty: controlling the Lagrangian frontier does not by itself guarantee that a learned frontier mechanism has nonnegative revenue under the true distribution. We first show that whenever the second-best GFT is nontrivial, the learned frontier contains a mechanism with sufficiently positive expected revenue. An independent validation sample is used only to determine the revenue signs of candidate frontier mechanisms under the true distribution. We locate two nearby mechanisms whose validated revenues have opposite

signs and mix between them before reports are observed. The mixture is ex-ante WBB while losing only $O(\varepsilon)$ additional GFT.

The lower bound. For the lower bound, we first construct a hard instance on a single star-like market. One seller faces several competing buyers, and small perturbations of the buyers' distributions encode hidden local decisions about which trade should be selected. The two signs of each perturbation are statistically difficult to distinguish, but choosing the wrong side leads to a large enough loss in GFT. We then combine r such stars into a star forest, distributing many hidden decisions across the market. A direct-sum argument is not immediate under ex-ante WBB, since revenue generated in one star may subsidize another. Our construction avoids this interaction by aligning the local instances so that their losses add under the single global budget constraint. Given this market construction, we use standard information theoretical arguments to derive the matching lower bound.

1.2 Related Work

Gains from trade and matching markets. The impossibility of achieving first-best GFT under incentive compatibility, individual rationality, and budget balance goes back to Myerson & Satterthwaite (1983), motivating a large literature on approximately efficient trade. Early work develops truthful and budget-balanced mechanisms for bilateral and more general two-sided markets (McAfee, 1992; Colini-Baldeschi et al., 2016; Dütting et al., 2017). A closely related line studies simple and approximately efficient mechanisms for gains from trade, including virtual-surplus, delegated-pricing, and fixed-price approaches (Deng et al., 2022; Fei, 2022; Kang et al., 2022; Cai & Wu, 2023; Liu et al., 2023; Hartline & Wang, 2025). These questions have also been extended beyond bilateral trade to richer feasibility constraints and matching environments (Brustle et al., 2017; Babaioff et al., 2018; Cai et al., 2021; Babaioff et al., 2026). More recently, Liu et al. (2026) show that, under strong budget balance, the second-best GFT in bilateral trade is always at least one half of the first-best benchmark, and that this factor is tight. Bei et al. (2026) extend the corresponding one-half guarantee to general matching markets under strong budget balance.

Learning and sample complexity in mechanism design. Sample-based mechanism design has a well-developed theory for revenue maximization, including nearly tight sample-complexity bounds in single- and multidimensional settings and related structural formulations (Cole & Roughgarden, 2014; Roughgarden & Schrijvers, 2016; Gonczarowski & Nisan, 2017; Guo et al., 2019; Chen et al., 2023; Gonczarowski & Weinberg, 2021; Cai et al., 2012; 2016). For GFT, prior work studies limited-information approximation (Dütting et al., 2021), sample-based learning in small two-sided markets and for fixed-price mechanisms (Babaioff et al., 2024; Castiglioni et al., 2026), and sample-based approximation and private learning in bilateral trade (Deng et al., 2025; Di Gregorio et al., 2026). A complementary line studies online bilateral and two-sided trade under regret objectives, with particular emphasis on fixed-price mechanisms, feedback models, and budget-balance constraints (Cesa-Bianchi et al., 2024a;b; Azar et al., 2024; Bernasconi et al., 2024; Chen et al., 2026; Jin, 2026). Related extensions incorporate contextual information, multiple buyers, relaxed budget constraints, and more general mechanism classes (Gaucher et al., 2025; Lunghi et al., 2025; 2026; Di Gregorio et al., 2025). These works study sequential regret under online feedback, whereas sample complexity asks a different statistical question: how much information about the unknown distributions is needed to learn a mechanism approaching the Bayesian benchmark. Most closely related, Shi et al. (2026) study the sample complexity of learning the second-best GFT in bilateral trade under regularity assumptions. We allow arbitrary Borel type distributions and extend the problem to matching markets with globally coupled allocations.

2 Preliminaries

2.1 Model and Benchmark

Let B be the buyers, S the sellers, $A := B \cup S$, and $E \subseteq B \times S$ the public bipartite graph of possible trades. Feasible matchings form a nonempty downward-closed family $\mathcal{F} \subseteq 2^E$. Write $d := |A|$ and $r := \max_{N \in \mathcal{F}} |N|$, and assume $r \geq 1$. Buyer i has value $v_i \in [0, h]$ and seller j has cost $c_j \in [0, h]$. Types are mutually independent, with product distribution $D := \bigotimes_{i \in B} F_i^B \otimes \bigotimes_{j \in S} F_j^S$, where F_i^B is the value distribution of buyer i and F_j^S is the cost distribution of seller j.

A direct randomized mechanism M maps every reported profile to a distribution over outcomes (N, p^B, p^S), where $N \subseteq E$, p_i^B is the payment made by buyer i, and p_j^S is the transfer received by seller j. We say that M is *feasible* if it selects a matching $N \in \mathcal{F}$ with probability one. For a realized outcome, the utilities of buyer i and seller j are defined as $u_i = v_i \mathbf{1}\{i \in N\} - p_i^B$ and $u_j = p_j^S - c_j \mathbf{1}\{j \in N\}$.

A mechanism M is called *Bayesian incentive compatible* (BIC) under D if, for every agent and every type, truthful reporting maximizes the agent's expected utility when all other agents report truthfully, where the expectation is over the other agents' types and the internal randomization of M. The mechanism is *interim individually rational* (interim IR) if, for every agent and every type, this expected utility under truthful reporting is nonnegative. We will also use the stronger notions of *dominant-strategy incentive compatibility* (DSIC), which requires truthful reporting to maximize each agent's expected utility for every fixed report profile of the other agents, and *ex-post individual rationality* (ex-post IR), which requires truthful expected utility to be nonnegative for every realized type profile, where the expectation in both cases is over the internal randomization of M. A randomized mechanism is *universally DSIC* if its internal randomization is drawn before reports and every deterministic realization is DSIC. For a fixed type profile, let N_M be the selected matching of M and define, with expectation over M's internal randomization,

$$\mathrm{gft}_M := \mathbb{E}_M\left[\sum_{(i,j)\in N_M} (v_i - c_j)\right], \qquad \mathrm{rev}_M := \mathbb{E}_M\left[\sum_{i\in B} p_i^B - \sum_{j\in S} p_j^S\right].$$

Write $\mathrm{GFT}_D(M) := \mathbb{E}_D[\mathrm{gft}_M]$ and $\mathrm{Rev}_D(M) := \mathbb{E}_D[\mathrm{rev}_M]$. The mechanism M is called *ex-ante weakly budget balanced* (ex-ante WBB) under D if $\mathrm{Rev}_D(M) \geq 0$.

We define the second-best GFT by

$$\mathrm{SB}(D) := \sup\{\mathrm{GFT}_D(M) : M \text{ is feasible, BIC, interim IR, and ex-ante WBB under } D\}.$$

The learner receives m independent samples from each individual type distribution. Given $\varepsilon > 0$ and $\gamma \in (0,1)$, it must always output a universally DSIC and ex-post-IR mechanism $\widehat{M}$ and, with probability at least $1-\gamma$ over the samples, satisfy $\mathrm{Rev}_D(\widehat{M}) \geq 0$ and $\mathrm{GFT}_D(\widehat{M}) \geq \mathrm{SB}(D) - \varepsilon$. The sample complexity is the minimum m sufficient to guarantee this uniformly over the specified class of markets and product distributions.

2.2 Known-Distribution Frontier

We first characterize the second-best GFT when the type distribution is known. For now, we suppose that D has finite support; the extension to arbitrary Borel distributions is given in Section 4.

Unify the notation of buyers and sellers. We place buyers and sellers in a common increasing type order by setting $y_i := v_i$ for $i \in B$ and $y_j := h - c_j$ for $j \in S$, and let F_a denote the distribution of y_a. For a report profile $y = (y_a)_{a\in A}$, write $y_{-a} := (y_b)_{b\neq a}$ for the reports of all agents other than a. Let $x_e(y)$ be the probability of selecting edge $e \in E$ and define $x_a(y) := \sum_{e\ni a} x_e(y)$, which is the probability that the agent a is selected in a trade. With expected transfers $p_i^B(y)$ and $p_j^S(y)$, define $\pi_i(y) := p_i^B(y)$ and $\pi_j(y) := h x_j(y) - p_j^S(y)$. Then every agent has utility $u_a(y) = y_a x_a(y) - \pi_a(y)$, and

$$\mathrm{gft}_M(y) = \sum_{a\in A} y_a x_a(y) - h\sum_{e\in E} x_e(y), \qquad \mathrm{rev}_M(y) = \sum_{a\in A} \pi_a(y) - h\sum_{e\in E} x_e(y). \tag{2.1}$$

Since ex-ante WBB is a single scalar constraint, define for $\lambda \geq 0$

$$\Phi_D(\lambda) := \sup_M \left\{\mathrm{GFT}_D(M) + \lambda\,\mathrm{Rev}_D(M)\right\}, \tag{2.2}$$

where the supremum ranges over feasible BIC, interim-IR mechanisms. Here λ is the shadow price of budget balance: $\lambda = 0$ optimizes GFT, while larger λ places greater weight on revenue.

Lemma 2.1 (cf. Liu et al. 2026, Lemma 4.1)**.** *For every finite product distribution D,*

$$\mathrm{SB}(D) = \min_{\lambda\geq 0} \Phi_D(\lambda). \tag{2.3}$$

Score of agents and edges. For a fixed multiplier $\lambda \geq 0$, let $\phi_{a;D}$ be the virtual value function associated with F_a. We define the generalized score of agent a as $s^{\lambda}_{a;D}(y) := \frac{1}{1+\lambda} \operatorname{Iron}_{F_a}(y + \lambda\phi_{a;D}(y))$, where $\operatorname{Iron}_{F_a}$ denotes the standard probability-weighted ironing operator, its finite-support construction is recalled in Appendix A.2. Accordingly, for a trade edge $e = (i, j) \in E$ with $i \in B$ and $j \in S$, define its generalized weight by $w^{\lambda}_{e;D}(y) := s^{\lambda}_{i;D}(y_i) + s^{\lambda}_{j;D}(y_j) - h$.

Lemma 2.2 (cf. Giannakopoulos & Hahn 2024, Theorem 4.1)**.** *For every finite product distribution D and $\lambda \geq 0$, let $M_{D,\lambda}$ choose a feasible matching of maximum total weight under $w^{\lambda}_{e;D}$. Then $M_{D,\lambda}$ admits a dominant-strategy incentive-compatible, ex-post-IR implementation and satisfies*

$$\mathrm{GFT}_D(M_{D,\lambda}) + \lambda \,\mathrm{Rev}_D(M_{D,\lambda}) = \Phi_D(\lambda). \tag{2.4}$$

Thus the known-distribution problem reduces to the one-parameter family $\{M_{D,\lambda}\}_{\lambda\geq 0}$, which forms the basis of our learning algorithm.

3 Learning the Lagrangian Frontier

Equation (2.2) reduces the known-distribution problem to a one-parameter family of Lagrangian optimization problems. When D is known, Lemma 2.2 shows that $\Phi_D(\lambda)$ is attained by the canonical score mechanism $M_{D,\lambda}$. We now ask whether this family can be learned from samples. Throughout this section, we assume that the marginal distributions have finite support. From the samples, we will construct a product distribution $\widetilde{D}$ and use $M_\lambda := M_{\widetilde{D},\lambda}$ as the learned mechanism for multiplier λ. Our goal is to control $\Phi_D(\lambda) - (\mathrm{GFT}_D(M_\lambda) + \lambda \,\mathrm{Rev}_D(M_\lambda))$ simultaneously for all $\lambda \geq 0$. The main difficulty is that the optimal matching is a global object. A small perturbation of one marginal distribution may change which matching is optimal and, in particular, may change an agent's trading partner. We therefore avoid comparing optimal matchings directly. Instead, we first show that after fixing all reports except one agent's, the dependence on that agent's report has a simple one-dimensional form.

For each agent $a \in A$, let $Y^1_a, \ldots, Y^{m_{\mathrm{tr}}}_a \sim F_a$ be independent training samples, where m_{tr} is the number of training samples per agent. Let $T_a(t) := \Pr_{Y_a\sim F_a}[Y_a \geq t]$ and $\widehat{T}_a(t) := \frac{1}{m_{\mathrm{tr}}} \sum_{\ell=1}^{m_{\mathrm{tr}}} \mathbf{1}\{Y^\ell_a \geq t\}$ denote the true and empirical upper-tail probabilities, respectively. Set $\Gamma_{\mathrm{tr}} := \log(32dm_{\mathrm{tr}}/\gamma)$, where $\gamma \in (0, 1)$ is the target failure probability. Rather than using the empirical distribution directly, we slightly decrease its upper tail. Define $\widetilde{T}_a(0) := 1$ and for $t \in (0, h]$,

$$\widetilde{T}_a(t) := \left[\widehat{T}_a(t) - C_{\mathrm{sh}} \left(\sqrt{\frac{\widehat{T}_a(t)\Gamma_{\mathrm{tr}}}{m_{\mathrm{tr}}}} + \frac{\Gamma_{\mathrm{tr}}}{m_{\mathrm{tr}}}\right)\right]_+, \tag{3.1}$$

where $C_{\mathrm{sh}} > 0$ is a sufficiently large universal constant and $[x]_+ := \max\{x, 0\}$. Let $\widetilde{F}_a$ be the distribution induced by $\widetilde{T}_a$, with the removed mass placed at zero, and define

$$\widetilde{D} := \bigotimes_{a\in A} \widetilde{F}_a. \tag{3.2}$$

We use the canonical score mechanisms $\{M_{\widetilde{D},\lambda}\}_{\lambda\geq 0}$ of Section 2.2 as the learned Lagrangian frontier.

Critical contacts. We next isolate the effect of changing one marginal distribution. Fix a multiplier λ, a finite product distribution D', an agent a, and the reports y_{-a} of all other agents. As y_a increases, every generalized edge weight incident to a changes through the same score $s^{\lambda}_{a;D'}(y_a)$. The matching structure from Babaioff et al. (2026) implies the following simple behavior. If agent a can ever be matched in this fixed external context, then there is a lowest winning report $\tau_a \in [0, h]$, and whenever $y_a \geq \tau_a$, agent a is matched to the same partner a'. Thus, once y_{-a} is fixed, the interaction between a and the globally optimal matching is summarized by the threshold event $\{y_a \geq \tau_a\}$. We refer to this event as the *critical contact* of agent a in the given external context. Once the winning partner a' is fixed, all contributions of the selected edge other than a's own score are constant as y_a varies. We collect these terms into $b_a := s^{\lambda}_{a';D'}(y_{a'}) - h$. The generalized score $s^{\lambda}_{a;D'}$ depends on the distribution D', so directly comparing its expectation under two different marginals is inconvenient. The next identity removes this distribution-dependent score from the relevant expectation.

Lemma 3.1. *For every fixed profile y_{-a} and λ for which a can be matched, define*

$$g_a(y) := \left(\frac{y + \lambda\tau_a}{1+\lambda} + b_a\right)\mathbf{1}\{y \geq \tau_a\}. \tag{3.3}$$

Then g_a is nondecreasing, takes values in $[0, h]$, and

$$\mathbb{E}_{Y_a \sim F'_a}\left[(s^{\lambda}_{a;D'}(Y_a) + b_a)\mathbf{1}\{Y_a \geq \tau_a\}\right] = \mathbb{E}_{Y_a \sim F'_a}[g_a(Y_a)]\,. \tag{3.4}$$

Lemma 3.1 expresses each conditional contribution as the expectation of a bounded nondecreasing function, allowing us to apply tail estimates. To aggregate these contributions, fix a report profile y, let N be the matching selected by $M = M_{D',\lambda}$, and write W for its total generalized weight. Counting both endpoints of each selected edge gives

$$\sum_{a \in V(N)} \left(s^{\lambda}_{a;D'}(y_a) + b_a\right) = 2W, \qquad \sum_{a \in V(N)} g_a(y_a) \quad = W + \frac{\mathrm{gft}_M(y) + \lambda\,\mathrm{rev}_M(y)}{1+\lambda}. \tag{3.5}$$

We next bound how these expectations change under shading.

Lemma 3.2. *With probability at least $1 - \gamma/2$ over the training samples, simultaneously for every agent $a \in A$ and every nondecreasing function $g : [0, h] \to [0, h]$, letting $\mu := \mathbb{E}_{Y_a \sim F_a}[g(Y_a)]$, we have*

$$0 \leq \mathbb{E}_{Y_a \sim F_a}[g(Y_a)] - \mathbb{E}_{Y_a \sim \widetilde{F}_a}[g(Y_a)] \leq O\left(\sqrt{\frac{h\mu\Gamma_{\mathrm{tr}}}{m_{\mathrm{tr}}}} + \frac{h\Gamma_{\mathrm{tr}}}{m_{\mathrm{tr}}}\right). \tag{3.6}$$

For the learned mechanism $M_\lambda = M_{\widetilde{D},\lambda}$, we decompose the loss as

$$\Phi_D(\lambda) - \mathrm{GFT}_D(M_\lambda) - \lambda\,\mathrm{Rev}_D(M_\lambda) = \underbrace{\Phi_D(\lambda) - \Phi_{\widetilde{D}}(\lambda)}_{\text{contraction}} + \underbrace{\Phi_{\widetilde{D}}(\lambda) - \mathrm{GFT}_D(M_\lambda) - \lambda\,\mathrm{Rev}_D(M_\lambda)}_{\text{deployment}}. \tag{3.7}$$

The first term measures the decrease in optimal value caused by shading. We shade buyer marginals first and seller marginals second, retaining trades that still clear their original critical thresholds at each stage. Lemmas 3.1 and 3.2 control the local losses, while (3.5) bounds the total expected contact contribution. For the second term, we keep M_λ fixed and compare its value under $\widetilde{D}$ and D. The same lemmas control the difference between its score and contact contributions under D. The second identity in (3.5) translates this difference into a bound on its actual GFT and revenue. Combining the two bounds gives the following stability guarantee.

Lemma 3.3. *With probability at least $1 - \gamma/2$, simultaneously for every $\lambda \geq 0$,*

$$0 \leq \Phi_D(\lambda) - \mathrm{GFT}_D(M_\lambda) - \lambda\,\mathrm{Rev}_D(M_\lambda) \;\leq C\left(\sqrt{\frac{(1+\lambda)dh\,\Phi_D(\lambda)\Gamma_{\mathrm{tr}}}{m_{\mathrm{tr}}}} + \frac{(1+\lambda)dh\Gamma_{\mathrm{tr}}}{m_{\mathrm{tr}}}\right), \tag{3.8}$$

where $C > 0$ is a universal constant.

4 From the Learned Frontier to the Second Best

Lemma 3.3 shows that, for every multiplier λ, the learned mechanism $M_\lambda = M_{\widetilde{D},\lambda}$ nearly attains the frontier value $\Phi_D(\lambda)$ under D. This guarantee alone does not ensure the existence of a learned mechanism with nonnegative revenue. The goal of this section is to turn these near-optimal frontier mechanisms into an ex-ante-WBB mechanism while incurring only $O(\varepsilon)$ additional GFT loss. We first give the construction for finite-support type distributions and remove this assumption at the end of the section.

Let $R^\star_D$ denote the maximum expected revenue achievable under D by a feasible DSIC, ex-post-IR mechanism. The next lemma shows that $R^\star_D$ is nearly at the same scale as the second-best GFT.

Lemma 4.1. *There is a universal constant $c_{\mathrm{rev}} > 0$ such that, whenever $\mathrm{SB}(D) > \varepsilon$,*

$$R^\star_D \geq c_{\mathrm{rev}}\frac{\mathrm{SB}(D)}{\Lambda_\varepsilon}, \qquad \Lambda_\varepsilon := 1 + \log\left(2 + \frac{d^2 h}{\varepsilon}\right).$$

Algorithm 1 SECOND-BEST LEARNER

Initialize $\lambda_{\max} := \frac{8\Lambda_\varepsilon}{c_{\mathrm{rev}}}$, $K := \lceil \log_2\big(1 + \frac{8\lambda_{\max} rh}{\varepsilon}\big) \rceil$
Construct $\widetilde{D}$ according to (3.2)
Construct $M_0 := M_{\widetilde{D},0}$, and $M^{\mathrm{rev}} := M^{\mathrm{rev}}_{\widetilde{D}}$, and compute R_0 and R_{rev}
if $R_0 \geq 0$ **then return** M_0
if $R_{\mathrm{rev}} \leq 0$ **then return** the no-trade mechanism
Set $(\lambda_-, M_-, R_-) = (0, M_0, R_0)$ and $(\lambda_+, M_+, R_+) = (\lambda_{\max}, M^{\mathrm{rev}}, R_{\mathrm{rev}})$
for $t = 1, \ldots, K$ **do** // Use binary search to narrow the revenue bracket
  Set $\lambda = (\lambda_- + \lambda_+)/2$
  Construct $M_\lambda := M_{\widetilde{D},\lambda}$ and compute R_λ
  if $R_\lambda < 0$ **then** Set $(\lambda_-, M_-, R_-) = (\lambda, M_\lambda, R_\lambda)$
  else Set $(\lambda_+, M_+, R_+) = (\lambda, M_\lambda, R_\lambda)$
end
Set $\alpha := R_+/(R_+ - R_-)$
return the pre-report mixture $\alpha M_- + (1-\alpha) M_+$

Let m_{val} denote the number of independent validation samples from each individual type distribution and write $m := \min\{m_{\mathrm{tr}}, m_{\mathrm{val}}\}$. Using the training sample, construct $M_0 := M_{\widetilde{D},0}$ by Lemma 2.2. For any finite product law D', let $M^{\mathrm{rev}}_{D'}$ denote the canonical maximum-weight score mechanism for the pure-revenue objective, and construct $M^{\mathrm{rev}} := M^{\mathrm{rev}}_{\widetilde{D}}$ under the same shaded product law $\widetilde{D}$. Thus M_0 is the $\lambda = 0$ endpoint of the learned frontier, while M^{rev} corresponds to its pure-revenue limit as $\lambda \to \infty$. Set $\lambda_{\max} := \frac{8\Lambda_\varepsilon}{c_{\mathrm{rev}}}$ and $K := \lceil \log_2\big(1 + \frac{8\lambda_{\max} rh}{\varepsilon}\big) \rceil$, and define

$$\mathcal{L}_K := \left\{ \frac{q\lambda_{\max}}{2^K} : q = 0, \ldots, 2^K \right\}, \quad \mathcal{C}_K := \{M_\lambda : \lambda \in \mathcal{L}_K\} \cup \{M^{\mathrm{rev}}\}, \quad \Gamma_{\mathrm{val}} := \log\left(\frac{32|\mathcal{C}_K|}{\gamma}\right).$$

For every $M \in \mathcal{C}_K$, let $\widehat{R}_M$ and $\widehat{V}_{R,M}$ denote the empirical mean and sample variance of its realized revenue on the independent validation sample. For a sufficiently large universal constant $C_{\mathrm{val}} > 0$, define

$$\mathrm{rad}_{R,M} := C_{\mathrm{val}} \left(\sqrt{\frac{\widehat{V}_{R,M}\Gamma_{\mathrm{val}}}{m_{\mathrm{val}}}} + \frac{rh\Gamma_{\mathrm{val}}}{m_{\mathrm{val}}} \right), \qquad R_M := \widehat{R}_M - \mathrm{rad}_{R,M}\,.$$

We write $R_\lambda := R_{M_\lambda}$ for $\lambda \in \mathcal{L}_K$, and in particular $R_0 := R_{M_0}$ and $R_{\mathrm{rev}} := R_{M^{\mathrm{rev}}}$. We present our learning algorithm in Algorithm 1. The algorithm maintains two candidates with opposite validated revenue signs and uses binary search to halve the multiplier interval between them. It then returns a pre-report mixture of the two endpoint mechanisms whose validated revenues cancel.

Computational complexity. Assuming a maximum-weight optimization oracle over $\mathcal{F}$, its running time is polynomial in the sample size and the market size, and it uses only $O(\log(1 + \lambda_{\max} rh/\varepsilon))$ multiplier-search rounds.

Algorithm analysis. For the validation step, observe that every multiplier queried by Algorithm 1 lies in the fixed dyadic grid $\mathcal{L}_K$. Thus, although the sequence of queries is adaptive, it suffices to establish the required guarantees simultaneously for the fixed family $\mathcal{C}_K$. For each $M \in \mathcal{C}_K$, let $\widehat{G}_M$ denote its empirical GFT on the validation sample and define the lower-confidence estimate

$$G_M := \widehat{G}_M - rh\sqrt{\frac{\Gamma_{\mathrm{val}}}{2m_{\mathrm{val}}}}.$$

For $\lambda \in \mathcal{L}_K$, write $G_\lambda := G_{M_\lambda}$, and let $G_{\mathrm{rev}} := G_{M^{\mathrm{rev}}}$. The next lemma gives these guarantees uniformly over $\mathcal{C}_K$.

Lemma 4.2. *For $m = \widetilde{\Omega}(drh^2/\varepsilon^2)$, on the training event used in Lemma 3.3, with probability at least $1 - \gamma/2$ over the validation sample, simultaneously $G_M \leq \mathrm{GFT}_D(M)$ and $R_M \leq \mathrm{Rev}_D(M)$ for every $M \in \mathcal{C}_K$ and, for every $\lambda \in \mathcal{L}_K$,*

$$G_\lambda + \lambda R_\lambda \geq \mathrm{SB}(D) - \eta_m, \quad \textit{where} \quad \eta_m = \widetilde{O}\left(h\sqrt{\frac{dr(1+\lambda_{\max})}{m}} + \frac{dh(1+\lambda_{\max})}{m} \right).$$

Moreover, if $\mathrm{SB}(D) > \varepsilon$*, then* $R_{\mathrm{rev}} > 0$ *and* $G_{\mathrm{rev}} + \lambda_{\max} R_{\mathrm{rev}} \geq \mathrm{SB}(D) - \eta_m$.

Because Lemma 4.2 holds simultaneously over the fixed family $\mathcal{C}_K$, its guarantees remain valid along the adaptive execution path of Algorithm 1. Conditioned on this event, the rest of the analysis is deterministic.

We next analyze the search and mixing step. If neither early return applies, the search maintains $R_- < 0 \leq R_+$ while halving the multiplier interval in every round, and therefore ends with $\lambda_+ - \lambda_- \leq \varepsilon/(8rh)$. The final mixture is chosen so that the validated revenues cancel. Since both endpoint mechanisms satisfy the directional guarantee of Lemma 4.2, this cancellation converts their Lagrangian guarantees into

$$\mathrm{GFT}_D(\widehat{M}) \geq \mathrm{SB}(D) - \eta_m - O(rh(\lambda_+ - \lambda_-)) \geq \mathrm{SB}(D) - \eta_m - O(\varepsilon).$$

The following lemma formalizes this calibration argument, including the two early-return cases.

Lemma 4.3. *Suppose the conclusions of Lemma 4.2 hold. Then Algorithm 1 outputs a universally DSIC, ex-post-IR, and ex-ante-WBB mechanism* $\widehat{M}$*. If the algorithm returns the no-trade mechanism, then* $\mathrm{SB}(D) \leq \varepsilon$*; otherwise,*

$$\mathrm{GFT}_D(\widehat{M}) \geq \mathrm{SB}(D) - \eta_m - \frac{\varepsilon}{8}.$$

From finite support to Borel distributions. Fix a deterministic grid of mesh Δ on $[0, h]$ and round every common-coordinate type y_a down to the nearest grid point; equivalently, buyer values are rounded down and seller costs are rounded up. Let D^{grid} denote the resulting finite-support product distribution. Any feasible BIC, interim-IR mechanism under D can be contracted to D^{grid} with at most $O(r\Delta)$ loss in GFT and $O(d\Delta)$ loss in revenue. When $\mathrm{SB}(D) > \varepsilon$, the resulting revenue deficit can be repaired by a positive-revenue posted-spread mechanism on the grid, at an additional $O(d\Delta\Lambda_\varepsilon)$ GFT cost. Consequently,

$$\mathrm{SB}(D^{\mathrm{grid}}) \geq \mathrm{SB}(D) - O(d\Delta\Lambda_\varepsilon).$$

We then apply the finite-support learner to D^{grid} and lift its output cellwise to the original type space. This lift preserves expected revenue exactly and weakly increases GFT, while preserving universal DSIC and ex-post IR. Taking $\Delta = \Theta(\varepsilon/(d\Lambda_\varepsilon))$ therefore extends the finite-support guarantee to arbitrary Borel individual type distributions with only $O(\varepsilon)$ additional loss. If $\mathrm{SB}(D) \leq \varepsilon$, the guarantee is immediate because the mechanisms returned by the learner have nonnegative GFT, a property preserved by the cellwise lift. The full argument is given in Appendix C.4.

Combining the finite-support analysis with the discretization argument above gives the claimed sample complexity for arbitrary Borel individual type distributions.

Theorem 4.4 (Upper bound)**.** *There is a universal constant* $c_0 > 0$ *such that, for every product distribution* D *with Borel individual type distributions supported on* $[0, h]$*, every* $0 < \varepsilon \leq c_0 rh$*, and every* $\gamma \in (0, 1)$*, there is a learner using* $\widetilde{O}(drh^2/\varepsilon^2)$ *samples from each individual type distribution that always outputs a universally DSIC and ex-post-IR mechanism* $\widehat{M}$ *and, with probability at least* $1 - \gamma$*, satisfies*$\mathrm{Rev}_D(\widehat{M}) \geq 0$ *and* $\mathrm{GFT}_D(\widehat{M}) \geq \mathrm{SB}(D) - \varepsilon$.

5 The Lower Bound

We first construct a family of one-seller star markets and show that achieving near-second-best GFT requires a number of samples proportional to the number of buyers. We then combine r such stars to obtain the lower bound for a market with d agents and maximum matching size r.

Star market construction. Consider a family of markets $\{D_\zeta\}$ indexed by sign vectors $\zeta \in \{\pm 1\}^k$. Two instances whose sign vectors differ in a single coordinate i differ only in buyer i's distribution, and only by slightly redistributing probability between two adjacent values; such a shift is hard to detect from samples. At the same time, the sign ζ_i determines whether one particular trade raises or lowers the Lagrangian objective, so a mechanism that cannot tell the two signs apart behaves suboptimally on one of them. A learner that achieves near-second-best GFT on every D_ζ must

therefore correctly identify a large fraction of the coordinates of ζ, and the per-coordinate statistical cost aggregates into the claimed bound. Every D_ζ is a star-like market: one seller faces k buyers, and at most one trade is feasible.

Buyer 0, the *anchor*, has a positive value with constant probability. Each buyer $i \geq 1$ is *rare*: its value is 0 with probability $1-p$ and positive with probability p, where we let $(k-1)p \leq 1/10$. The coordinate ζ_0 perturbs the anchor's distribution and ζ_i perturbs that of rare buyer i.

Fix $\zeta \in \{\pm 1\}^k$. Under D_ζ, the seller's cost is $c_L := h/10$ or $c_H := 4h/5$ with equal probability, independently of ζ. The anchor's value is

$$\Pr[V_0 = w] = \tfrac{3}{5}, \qquad \Pr[V_0 = u] = \tfrac{2}{5}\big(\tfrac{1}{2} - \zeta_0\delta\big), \qquad \Pr[V_0 = h] = \tfrac{2}{5}\big(\tfrac{1}{2} + \zeta_0\delta\big),$$

with $w := 101h/700$ and $u := 7h/8$. Each rare buyer i is inactive (value 0) with probability $1-p$ and otherwise spreads over two positive values, that is,

$$\Pr[V_i = 0] = 1-p, \qquad \Pr[V_i = v_L] = p\big(\tfrac{1}{2} - \zeta_i\delta\big), \qquad \Pr[V_i = v_H] = p\big(\tfrac{1}{2} + \zeta_i\delta\big),$$

with $v_L := 11h/100$ and $v_H := 11h/50$. Those constants place certain trading decisions near indifference, so the two signs require opposite choices. The anchor also allows the optimal mechanism's expected revenue to be adjusted to zero, a property we use when combining stars.

Lemma 5.1. *For $k \geq 1$ and small enough η, set $p = 1/(10(k+1))$ and $\delta = C\eta/h$. Learning the second-best GFT to additive accuracy η on this star family requires $m = \Omega(kh^2/\eta^2)$ independent samples per marginal.*

The dependence on k comes from the rare buyers. For $k \geq 2$, a sample from any one of them reveals information about its sign only with probability $p = \Theta(1/k)$. Flipping that sign has one-sample KL divergence $O(p\delta^2)$, while choosing the wrong trade contributes a loss of order $p\delta h$. When $mp\delta^2$ is small, testing bounds force errors across a constant fraction of the $k-1$ rare coordinates, giving a total loss at scale δh. Taking $\delta = \Theta(\eta/h)$ yields the bound in Lemma 5.1. When $k = 1$, the anchor alone gives the same bound.

From one star to a star forest. We now use the single-star construction to obtain a lower bound for markets with d agents and maximum matching size r. For $d \geq 2r$, take r disjoint stars, use one seller per star, and distribute the remaining $d-r$ buyers as evenly as possible. If star u has k_u buyers, then $\sum_u (k_u + 1) = d$ and $k_u \geq 1$. Use the construction above independently in every star, with a common perturbation size δ and $p := r/(10d)$; since $k_u - 1 \leq d/r$, we have $(k_u - 1)p \leq 1/10$ for every star u. The resulting market has exactly d agents and maximum matching size r.

Combining the single-star lower bounds requires care because global budget balance permits revenue from one star to subsidize another. Our construction ensures that, for the same multiplier $\lambda_0 = 1/10$, every local law satisfies $\mathrm{SB}(D_{\zeta_u}) = \Phi_{D_{\zeta_u}}(\lambda_0)$ and has a zero-revenue optimizer. These optimizers can run simultaneously, giving

$$\mathrm{SB}(D_\zeta) = \sum_{u=1}^{r} \mathrm{SB}(D_{\zeta_u}) = \sum_{u=1}^{r} \Phi_{D_{\zeta_u}}(\lambda_0).$$

The proof also bounds the local losses while retaining each star's revenue term. After summation, these terms depend only on total revenue, so global WBB suffices to make the losses add. This gives the following extension of Lemma 5.1; the budget argument is proved in Section D.2.

Theorem 5.2 (Lower bound)**.** *For small enough ε, set $\delta = C\varepsilon/(rh)$ in the star forest above. Learning the second-best GFT to additive accuracy ε requires $m = \Omega(drh^2/\varepsilon^2)$ independent samples per marginal.*

To see the scaling, consider r stars with $k \geq 2$ buyers each, so $d = r(k+1)$. Their combined loss is at scale $r\delta h$, while distinguishing the rare-buyer perturbations requires $\Omega(k/\delta^2)$ samples from each individual type distribution. Choosing $\delta = \Theta(\varepsilon/(rh))$ gives $\Omega(kr^2h^2/\varepsilon^2) = \Omega(drh^2/\varepsilon^2)$.

AI USE STATEMENT

We used generative AI tools as an interactive aid in developing and checking parts of the proofs, as well as in improving the presentation of the paper. All mathematical arguments and claims were independently reviewed and verified by the authors, who take full responsibility for the content of the paper.

REFERENCES

Yossi Azar, Amos Fiat, and Federico Fusco. An α-regret analysis of adversarial bilateral trade. *Artificial Intelligence*, 337:104231, 2024. doi: 10.1016/j.artint.2024.104231.

Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pp. 373, 2018. doi: 10.1145/3219166.3219203.

Moshe Babaioff, Amitai Frey, and Noam Nisan. Learning to maximize gains from trade in small markets. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pp. 195, 2024. doi: 10.1145/3670865.3673463.

Moshe Babaioff, Aviad Rubinstein, Xizhi Tan, and Kangning Wang. Approximating gains-from-trade in matching markets. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pp. 710–721, 2026. doi: 10.1145/3798129.3800786.

Xiaohui Bei, Bo Li, Wenhao Wu, and Shengwei Zhou. Second-best gains from trade in matching markets. *arXiv preprint arXiv:2609.18724*, 2026.

Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Federico Fusco. No-regret learning in bilateral trade via global budget balance. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pp. 247–258, 2024. doi: 10.1145/3618260.3649653.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pp. 589–590, 2017. doi: 10.1145/3033274.3085148.

Yang Cai and Jinzhao Wu. On the optimal fixed-price mechanism in bilateral trade. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pp. 737–750, 2023. doi: 10.1145/3564246.3585171.

Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pp. 130–139, 2012. doi: 10.1109/FOCS.2012.88.

Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality-based unified approach to bayesian mechanism design. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pp. 926–939, 2016. doi: 10.1145/2897518.2897645.

Yang Cai, Kira Goldner, Steven Ma, and Mingfei Zhao. On multi-dimensional gains from trade maximization. In *Proceedings of the 2021 Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1079–1098, 2021. doi: 10.1137/1.9781611976465.67.

Matteo Castiglioni, Anna Lunghi, and Alberto Marchesi. The sample complexity of uniform approximation for multi-dimensional CDFs and fixed-price mechanisms. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pp. 1353–1364, 2026. doi: 10.1145/3798129.3800845.

Nicolò Cesa-Bianchi, Tommaso Cesari, Roberto Colomboni, Federico Fusco, and Stefano Leonardi. Bilateral trade: A regret minimization perspective. *Mathematics of Operations Research*, 49(1): 171–203, 2024a. doi: 10.1287/moor.2023.1351.

Nicolò Cesa-Bianchi, Tommaso Cesari, Roberto Colomboni, Federico Fusco, and Stefano Leonardi. Regret analysis of bilateral trade with a smoothed adversary. *Journal of Machine Learning Research*, 25(234):1–36, 2024b.

Houshuang Chen, Yaonan Jin, Pinyan Lu, and Chihao Zhang. Tight regret bounds for fixed-price bilateral trade. In *53rd International Colloquium on Automata, Languages, and Programming*, volume 374 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 57:1–57:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026. doi: 10.4230/LIPIcs.ICALP.2026.57.

Ziyun Chen, Zhiyi Huang, Dorsa Majdi, and Zipeng Yan. Strong revenue (non-)monotonicity of single-parameter auctions. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 452–471, 2023. doi: 10.1145/3580507.3597745.

Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pp. 243–252, 2014. doi: 10.1145/2591796.2591867.

Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the 2016 Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1424–1443, 2016. doi: 10.1137/1.9781611974331.ch98.

Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing*, pp. 718–721, 2022. doi: 10.1145/3519935.3520054.

Yuan Deng, Jieming Mao, Balasubramanian Sivan, Kangning Wang, and Jinzhao Wu. Approximately efficient bilateral trade with samples. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pp. 206–223, 2025. doi: 10.1145/3736252.3742519.

Simone Di Gregorio, Paul Dütting, Federico Fusco, and Chris Schwiegelshohn. Nearly tight regret bounds for profit maximization in bilateral trade. In *Proceedings of the 66th IEEE Annual Symposium on Foundations of Computer Science*, pp. 1570–1594, 2025. doi: 10.1109/FOCS63196.2025.00083.

Simone Di Gregorio, Federico Fusco, Stefano Leonardi, and Chris Schwiegelshohn. Private learning in bilateral trade. *arXiv preprint arXiv:2606.02050*, 2026.

Paul Dütting, Inbal Talgam-Cohen, and Tim Roughgarden. Modularity and greed in double auctions. *Games and Economic Behavior*, 105:59–83, 2017. doi: 10.1016/j.geb.2017.06.008.

Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing*, pp. 1452–1465, 2021. doi: 10.1145/3406325.3451076.

Liran Einav, Chiara Farronato, and Jonathan Levin. Peer-to-peer markets. *Annual Review of Economics*, 8:615–635, 2016. doi: 10.1146/annurev-economics-080315-015334.

Yumou Fei. Improved approximation to first-best gains-from-trade. In *Web and Internet Economics*, volume 13778 of *Lecture Notes in Computer Science*, pp. 204–218. Springer, 2022. doi: 10.1007/978-3-031-22832-2_12.

Solenne Gaucher, Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Vianney Perchet. Feature-based online bilateral trade. In *The Thirteenth International Conference on Learning Representations*, 2025.

Yiannis Giannakopoulos and Johannes Hahn. Discrete single-parameter optimal auction design. In *Algorithmic Game Theory*, volume 15156 of *Lecture Notes in Computer Science*, pp. 165–183. Springer, 2024. doi: 10.1007/978-3-031-71033-9_10.

Yannai A. Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, pp. 856–868, 2017. doi: 10.1145/3055399.3055427.

Yannai A. Gonczarowski and S. Matthew Weinberg. The sample complexity of up-to-ϵ multidimensional revenue maximization. *Journal of the ACM*, 68(3):15:1–15:28, 2021. doi: 10.1145/3439722.

Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, pp. 662–673, 2019. doi: 10.1145/3313276.3316325.

Jason Hartline and Kangning Wang. A geometric analysis of gains from trade. *arXiv preprint arXiv:2508.06469*, 2025.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.

Yaonan Jin. Tight regret bounds for bilateral trade under semi feedback. *arXiv preprint arXiv:2601.16412*, 2026.

Zi Yang Kang, Francisco Pernice, and Jan Vondrák. Fixed-price approximations in bilateral trade. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2964–2985, 2022. doi: 10.1137/1.9781611977073.115.

Zhengyang Liu, Zeyu Ren, and Zihe Wang. Improved approximation ratios of fixed-price mechanisms in bilateral trades. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pp. 751–760, 2023. doi: 10.1145/3564246.3585160.

Zhengyang Liu, Ying Qin, Zeyu Ren, and Zihe Wang. Second-best bilateral trade is $1/2$ efficient. *arXiv preprint arXiv:2606.03849*, 2026.

Anna Lunghi, Matteo Castiglioni, and Alberto Marchesi. Online two-sided markets: Many buyers enhance learning. *arXiv preprint arXiv:2503.01529*, 2025.

Anna Lunghi, Matteo Castiglioni, and Alberto Marchesi. Better regret rates in bilateral trade via sublinear budget violation. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 6494–6536, 2026. doi: 10.1137/1.9781611978971.233.

Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009.

R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2): 434–450, 1992. doi: 10.1016/0022-0531(92)90091-U.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. doi: 10.1287/moor.6.1.58.

Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983. doi: 10.1016/0022-0531(83)90048-0.

Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pp. 1–18, 2016. doi: 10.1145/2940716.2940723.

Qiaoyun Shi, Shengxin Liu, and Zongqi Wan. Sample complexity of the second-best bilateral trade. In *Proceedings of the 22nd Conference on Web and Internet Economics*, 2026.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York, 2009.