
On the Robustness of Mechanism Design under Total Variation Distance

Anuran Makur
Purdue University
amakur@purdue.edu

Marios Mertzaniadis
Purdue University
mmertzan@purdue.edu

Alexandros Psomas
Purdue University
apsomas@cs.purdue.edu

Athina Terzoglou
Purdue University
aterzogl@purdue.edu

Abstract

We study the problem of designing mechanisms when agents’ valuation functions are drawn from unknown and correlated prior distributions. In particular, we are given a prior distribution \mathcal{D} , and we are interested in designing a (truthful) mechanism that has good performance for all “true distributions” that are close to \mathcal{D} in Total Variation (TV) distance. We show that DSIC and BIC mechanisms in this setting are strongly robust with respect to TV distance, for any bounded objective function \mathcal{O} , extending a recent result of Brustle et al. ([BCD20], EC 2020). At the heart of our result is a fundamental duality property of total variation distance. As direct applications of our result, we (i) demonstrate how to find approximately revenue-optimal and approximately BIC mechanisms for weakly dependent prior distributions; (ii) show how to find correlation-robust mechanisms when only “noisy” versions of marginals are accessible, extending recent results of Bei et al. ([BGLT19], SODA 2019); (iii) prove that prophet-inequality type guarantees are preserved for correlated priors, recovering a result of Dütting and Kesselheim ([DK19], EC 2019) as a special case; (iv) give a new necessary condition for a correlated distribution to witness an infinite separation in revenue between simple and optimal mechanisms, complementing recent results of Psomas et al. ([PSCW22], NeurIPS 2022); (v) give a new condition for simple mechanisms to approximate revenue-optimal mechanisms for the case of a single agent whose type is drawn from a correlated distribution that can be captured by a Markov Random Field, complementing recent results of Cai and Oikonomou ([CO21], EC 2021).

1 Introduction

Mechanism design studies optimization in strategic settings. The designer’s task is to create a mechanism that interacts with strategic agents, each with their own, private preferences over the mechanism’s output. The inability to provide meaningful guarantees for important objectives — such as revenue — when studying mechanism design problems through the lens of worst-case analysis, has motivated the study of Bayesian mechanisms. In the Bayesian setting, there is a probability distribution, typically known to the designer, from which agents’ private information — their *types* — is drawn, and the designer seeks to maximize an objective function in expectation over the randomness of the types, and, at the same time, incentivize the agents to report their type truthfully.

While we have greatly deepened our understanding of mechanism design under the Bayesian setting, when taking this approach to practice, it is natural to ask what happens if the designer has only partial information about the agents’ type distributions. In recent years, a growing literature studies

the *robustness* of mechanisms with respect to inaccurate priors. The term “robustness” can mean many things in this context, such as mechanism design using only sampling access from the underlying prior [CR14, HMR15, DHP16, MR16, CD17, GHZ19, GW21], or mechanism design using only parametric knowledge of the underlying prior [AMDW13, Car17, GL18, BGLT19, GPTD23]. Arguably, the most ambitious line of work in this thread strives to, given a prior distribution \mathcal{D} , design a mechanism that provides good guarantees simultaneously for all “true distributions” $\widehat{\mathcal{D}}$ that are close to \mathcal{D} under some notion of statistical distance. Recent, compelling positive results [BS11, CD17, DK19, LLY19, BCD20] show that this endeavor is, in fact, possible.

In this work, we demonstrate that a large class of such robustness results can be obtained from a fundamental duality property of total variation distance. We distill this key property into various forms that are directly applicable to mechanism design, and then illustrate how they imply a variety of robustness results for both Dominant Strategy Incentive Compatible (DSIC) and Bayesian Incentive Compatible (BIC) mechanisms. In particular, we recover a litany of known results, such as the revenue robustness of BIC mechanisms of [BCD20] and robustness of prophet inequalities of [DK19], as well as prove several new results, such as a new necessary condition for a prior distribution to exhibit a large gap between simple and optimal mechanisms (complementing recent results of [PSCW22]).

Our Contributions. We consider the problem of a designer that wants to design a mechanism that maximizes a bounded objective function \mathcal{O} , when allocating m items to n strategic agents. Agent i has a private type $t_i \in \mathcal{T}_i$, drawn from a probability distribution \mathcal{D}_i , which specifies her value for every subset of items $S \subseteq [m]$. We write \mathcal{D} for the joint distribution over agents’ types.

We start in Section 3 by proving a simple, but crucial lemma, Lemma 2, about non-truthful mechanisms in this setting. Lemma 2 states that the expected performance, with respect to the objective \mathcal{O} , of a (possibly non-truthful) mechanism \mathcal{M} is stable with respect to small changes in *total variation (TV) distance* to the underlying prior distribution. This lemma follows from the argument for Kantorovich-Rubinstein duality of TV distance [LPW09, Vi09].

Appropriate applications of Lemma 2 allow us to prove our first robustness result: *DSIC mechanisms are strongly robust*. Specifically, given a mechanism \mathcal{M} that is ex-post individually rational (ex-post IR), DSIC, and an α approximation to the optimal mechanism for a distribution \mathcal{D} (w.r.t. the objective \mathcal{O}), we construct an ex-post IR and DSIC mechanism $\widehat{\mathcal{M}}$ that is agnostic to $\widehat{\mathcal{D}}$ and α approximates, under distribution $\widehat{\mathcal{D}}$, the performance of the *optimal* mechanism for $\widehat{\mathcal{D}}$ (w.r.t. \mathcal{O}) minus a small error which depends on the TV distance, $d_{TV}(\mathcal{D}, \widehat{\mathcal{D}})$, of \mathcal{D} and $\widehat{\mathcal{D}}$ (**Theorem 1**). Notably, \mathcal{D} and $\widehat{\mathcal{D}}$ need not be product distributions, and if they have the same support, $\widehat{\mathcal{M}}$ is the same mechanism as \mathcal{M} .

We proceed to study the robustness of BIC mechanisms. As opposed to the DSIC case, the BIC property is not maintained under small perturbations of a prior, or even by small changes in the support of a distribution. Rather surprisingly, however, in Lemma 4 we show that a mechanism \mathcal{M} that is BIC w.r.t. a distribution \mathcal{D} , is also approximately BIC w.r.t. a distribution $\widehat{\mathcal{D}}$, where the approximation depends on the distance between \mathcal{D} and $\widehat{\mathcal{D}}$. Combined with Lemma 2, we show our main robustness result for BIC mechanisms (**Theorem 2**): *a mechanism \mathcal{M} that is BIC w.r.t. \mathcal{D} is approximately BIC w.r.t. $\widehat{\mathcal{D}}$* , and its performance for an objective function \mathcal{O} is similar under \mathcal{D} and $\widehat{\mathcal{D}}$, as long as \mathcal{D} and $\widehat{\mathcal{D}}$ have the same support. Compared to the corresponding robustness result of [BCD20], Theorem 2 only holds if \mathcal{D} and $\widehat{\mathcal{D}}$ have the same support. Furthermore, \mathcal{M} is not ε -BIC (as in [BCD20]), but (ε, q) -BIC (see Section 2). On the other hand, our robustness results hold for *arbitrary objectives*, not just revenue, and *arbitrary distributions*, not just product distributions.

In Section 4 we show numerous applications of our robustness framework. Specifically, we give five applications. First, in Section 4.1, we extend our strong DSIC robustness to BIC mechanisms, albeit only for the revenue objective and product distributions. That is, we show that given a distribution \mathcal{D} that is close to a product distribution \mathcal{D}^p , approximately optimal (w.r.t. revenue) and BIC mechanisms for \mathcal{D}^p can be transformed into approximately optimal and approximately-BIC mechanisms for \mathcal{D} (**Theorem 3**). En route, we prove a self-reduction, that constructs an ε' -BIC mechanism w.r.t. \mathcal{D} , given an (ε, q) -BIC mechanism w.r.t. \mathcal{D} , that might be of independent interest.

In Section 4.2 we study the correlation-robust framework of [BGLT19] (first introduced in [Car17]). At a high level, we are given the marginals of n dependent agents for a single item. The goal is to design a mechanism that maximizes (among all feasible mechanisms) the minimum, over all possible

joint distributions consistent with the given marginals, expected revenue. We extend our robustness results to this setting (**Theorem 4**), and get implications for the max-min revenue performance of sequential posted prices and the Lookahead auction when the given marginals are inaccurate.

In Section 4.3 we consider prophet inequalities. In the simplest version of this problem, we are shown n non-negative numbers, t_1, \dots, t_n , one at a time, and upon arrival we need to decide (immediately and irrevocably) whether we should accept t_i , or keep going. t_i s are drawn independently from known distributions \mathcal{D}_i . A celebrated result, known as the prophet inequality, states that a simple threshold policy τ (“pick the first x_i bigger than τ ”) has expected reward at least half of a prophet, who knows the values of all x_i s in advance (and can therefore get reward $\max_i t_i$). Prophet inequalities have numerous applications in mechanism design, and are the main tool in the analysis of posted price mechanisms. Observing that prophet inequalities correspond to sequential posted prices, which are DSIC mechanisms, and applying Lemma 2 for the welfare objective, we get robustness of prophet inequalities (**Corollary 1**); as a special case of our application we recover a result of Dütting and Kesselheim [DK19] on the TV robustness of prophet inequalities for product distributions.

Finally, in Section 4.4 we apply our robustness framework to study revenue gaps between simple and optimal mechanisms. Here, [BCKW15, HN19] construct correlated distributions whose revenue is infinite, but such that simple mechanisms (e.g., those with finite menus) cannot extract a lot of revenue. Recently, Psomas et al. [PSCW22] provide (arguably, complex) conditions that a distribution \mathcal{D} should satisfy in order to witness such large gaps; our framework readily provides new, simple, and necessary conditions (**Corollary 2**), complementing the results of [PSCW22]. On a similar note, Cai and Oikonomou [CO21] escape the negative results of [BCKW15, HN19] by considering distributions described by a Markov Random Field (MRF). Their revenue guarantees for simple mechanisms are controlled by Δ , a parameter of the MRF, that is determined by how much the value of an item can be influenced by values of other items. We show that a bound on the TV distance to a product distribution also suffices to provide guarantees on the revenue of simple mechanisms (**Proposition 2**). Furthermore, under conditions on the MRF, bounds on Δ imply bounds on this distance (**Proposition 4**), but the other direction is not necessarily true (**Proposition 3**); thus, getting bounds on the TV distance of an MRF to a product distribution is a meaningful endeavor.

Related work. A number of recent works study robustness in mechanism design; see [Car17] for a survey. The paper (thematically and technically) closest to ours is [BCD20], which proves the robustness of the revenue objective for BIC mechanisms under various notions of statistical distance, including TV and Prokhorov, as well as Lévy and Kolmogorov for single parameter settings. Specifically, the goal is to show that, given a distribution \mathcal{D} , one can construct a mechanism that is (approximately) BIC, and performs well, when executed on any distribution $\widehat{\mathcal{D}}$ that is close to \mathcal{D} (under any of the above notions of distance). Here, we focus on TV distance, and recover their result for this case, slightly extending it for the case of revenue, and significantly extending it under some minor assumptions. Dütting and Kesselheim [DK19] study prophet inequalities with inaccurate priors. Specifically, given a product distribution $\mathcal{D} = \times_{i \in [n]} \mathcal{D}_i$, [DK19] study policies that perform well when executed on a product distribution $\widehat{\mathcal{D}} = \times_{i \in [n]} \widehat{\mathcal{D}}_i$, such that \mathcal{D}_i is close to $\widehat{\mathcal{D}}_i$ for all $i \in [n]$, under various statistical distances. We recover their result for TV as a special case of an application of robustness for DSIC mechanisms, extending it to non-product distributions. Further related work is discussed in Appendix A.

2 Preliminaries

We examine the problem of a central designer who seeks to create a mechanism that maximizes an objective function given some prior knowledge of the universe. Consider a set $[n] \triangleq \{1, \dots, n\}$ of selfish agents and a finite set $[m] \triangleq \{1, \dots, m\}$ of items. Each agent $i \in [n]$ has a type t_i belonging to a set \mathcal{T}_i of possible types. We assume that, for all $i \in [n]$, there exists a special type $\perp \in \mathcal{T}_i$, interpreted as the option of not participating in the designer’s mechanism. Let $\mathcal{T} = \mathcal{T}_1 \times \dots \times \mathcal{T}_n$. We use \mathcal{X} to denote the set of all possible allocations of the items. (In particular, the sets \mathcal{X} and \mathcal{T}_i are typically standard Borel spaces, e.g., finite-dimensional Euclidean spaces. Hence, our analysis with couplings and total variation distance in the sequel do not require measure theoretic arguments.)

The goal of the designer is to construct a mechanism $\mathcal{M} = (x, p)$ which consists of (1) an allocation rule $x : \mathcal{T} \mapsto \Delta(\mathcal{X})$, which maps reported types $(t_1, \dots, t_n) \in \mathcal{T}$ to a distribution over allocations,

and (2) a payment function $p : \mathcal{T} \mapsto [-H, H]^n$ which maps reported types to (bounded) payments for each agent, for some fixed constant $H > 0$. We say that a mechanism $\mathcal{M} = (x, p)$ is *defined* on types \mathcal{T} if the domain of x and p is \mathcal{T} . We write $\mathcal{M}(t) = (x(t), p(t))$ for the outcome, i.e., the allocation and payments, of mechanism \mathcal{M} on input $t \in \mathcal{T}$.

Each agent i has a valuation function $v_i : \mathcal{T}_i \times \mathcal{X} \mapsto [0, H]$, which specifies their value for an allocation, given their type. We assume that agents are quasi-linear, i.e., the *utility* of agent i , with type t_i , for an allocation A and a payment p_i is equal to $v_i(t_i, A) - p_i$. We overload notation and write $t_i(\mathcal{M}(t_i, t_{-i}))$ for the utility of agent i with type $t_i \in \mathcal{T}_i$ for the outcome of the mechanism \mathcal{M} on input (t_i, t_{-i}) . We use \mathcal{F} for the set of possible outcomes of a mechanism \mathcal{M} . We assume that there exists a probability distribution \mathcal{D} supported on \mathcal{T} from which agents' types are drawn.

Mechanism Design considerations. When faced with a mechanism \mathcal{M} , each agent i reports a type $b_i \in \mathcal{T}_i$ to the mechanism. We aim to design mechanisms that (approximately) incentivize agents to report their true types, so we typically have that $b_i = t_i$, for all $i \in [n]$. Given a mechanism \mathcal{M} , we write $u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i})$ for the difference in utility of agent i when she reports $t'_i \in \mathcal{T}_i$ instead of her true type $t_i \in \mathcal{T}_i$, and all other agents report according to $t_{-i} \in \mathcal{T}_{-i}$. That is, $u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i}) = t_i(\mathcal{M}(t_i, t_{-i})) - t_i(\mathcal{M}(t'_i, t_{-i}))$.

We consider four, increasingly weaker, notions of incentive compatibility. First, we say that a mechanism \mathcal{M} is *Dominant Strategy Incentive Compatible* (henceforth, DSIC) if an agent is better off reporting her true type, no matter what other agents report, i.e., for all $i \in [n]$, every type $t_i \in \mathcal{T}_i$, possible misreport $t'_i \in \mathcal{T}_i$, and types $t_{-i} \in \mathcal{T}_{-i}$ for the remaining agents, it holds that $u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i}) \geq 0$. Second, a mechanism \mathcal{M} is *Bayesian Incentive Compatible* (henceforth BIC) with respect to a distribution \mathcal{D} , if an agent is better off reporting her true type in expectation over the other agents' reports, i.e. if for all $i \in [n]$, every type $t_i \in \mathcal{T}_i$ and possible misreport $t'_i \in \mathcal{T}_i$, it holds that $\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i} | t_i} [u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i})] \geq 0$. Third, a mechanism \mathcal{M} is ε -BIC w.r.t. \mathcal{D} if for all $i \in [n]$, and all $t_i, t'_i \in \mathcal{T}_i$, it holds that $\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i} | t_i} [u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i})] \geq -\varepsilon$. Fourth, a mechanism \mathcal{M} is (ε, q) -BIC w.r.t. \mathcal{D} if it is ε -BIC with probability at least $1 - q$, i.e., if for all $i \in [n]$, and all $t_i, t'_i \in \mathcal{T}_i$, $\Pr_{t_i \sim \mathcal{D}_i} [\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i} | t_i} [u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i})] \geq -\varepsilon] \geq 1 - q$. An $(\varepsilon, 0)$ -BIC mechanism is simply ε -BIC, and a 0-BIC is simply BIC.

Finally, a mechanism \mathcal{M} is *ex-post Individually Rational* (henceforth ex-post IR) if for every agent $i \in [n]$, every type $t_i \in \mathcal{T}_i$ and types $t_{-i} \in \mathcal{T}_{-i}$ for the remaining agents, it holds that $t_i(\mathcal{M}(t_i, t_{-i})) \geq 0$.

The designer's objective. The designer has an objective function $\mathcal{O}(\cdot, \cdot) : \mathcal{T} \times \mathcal{F} \mapsto [a, b]$ that takes as input agents' (reported) types $t \in \mathcal{T}$ and mechanism outcomes $\mathcal{M}(t) \in \mathcal{F}$ (noting that the mechanism's outcome might be a randomized allocation) and outputs a real number in the interval $[a, b]$. Let $V = b - a$. The task of the designer is to find an ex-post IR and truthful mechanism (under one of the aforementioned notions of truthfulness) that maximizes this objective function in expectation over the randomness of \mathcal{D} . We denote the optimal value of the objective \mathcal{O} under distribution \mathcal{D} by $OPT_{\mathcal{O}}(\mathcal{D})$. We specify in context whether this is with respect to DSIC, or BIC, or ε -BIC mechanisms. A mechanism $\mathcal{M}_{\mathcal{D}}^{\alpha}$ is an α approximation to the optimal mechanism under \mathcal{D} , with respect to \mathcal{O} , if $\mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}^{\alpha}(t))] \geq \alpha OPT_{\mathcal{O}}(\mathcal{D})$.

Some of our results hold for arbitrary, bounded objectives \mathcal{O} . Two objectives of specific interest to us will be welfare and revenue. The welfare objective, denoted by $Val(\cdot)$, is simply the sum of agents' valuations of an outcome, i.e. $Val(t, \mathcal{M}(t)) = \sum_{i \in [n]} v_i(t_i, x_i(t))$, where $x_i(t)$ is the allocation of agent i in the outcome $\mathcal{M}(t)$. The revenue objective, denoted by $Rev(\cdot)$, is the sum of agents' payments, i.e. $Rev(t, \mathcal{M}(t)) = \sum_{i \in [n]} p_i(t)$, where $p_i(t)$ is the payment of agent i in the outcome $\mathcal{M}(t)$. We often overload notation and write $Rev(\mathcal{M}, \mathcal{D})$ for the expected revenue of mechanism \mathcal{M} under distribution \mathcal{D} , i.e. $Rev(\mathcal{M}, \mathcal{D}) = \mathbb{E}_{t \sim \mathcal{D}} [Rev(t, \mathcal{M}(t))]$. To maintain consistency with the (vast) literature on mechanism design we further denote the optimal revenue as $Rev(\mathcal{D})$, and the optimal welfare as $Val(\mathcal{D})$, under distribution \mathcal{D} .

Statistical distance. Throughout the paper, we are interested in how mechanisms behave under different distributions that are not "too far" from each other. Our notion of distance in this paper is total variation distance.

Definition 1 (Total Variation Distance). *The total variation (TV) distance between any two probability distributions P and Q on a sample space Ω is defined as*

$$d_{\text{TV}}(P, Q) \triangleq \sup_{E \subseteq \Omega} |P(E) - Q(E)|,$$

where the supremum is over all Borel measurable subsets $E \subseteq \Omega$, and $P(E)$ (resp. $Q(E)$) denotes the probability of the event E with respect to the distribution P (resp. Q).

We note that Ω is either discrete or a subset of a finite-dimensional Euclidean space in our analysis, and hence, it is always a standard Borel space. For any probability distributions P and Q on Ω , let $\Pi(P, Q)$ be the (non-empty) set of all *couplings* of P and Q , i.e., all joint probability distributions $\gamma = P_{X,Y}$ of two random variables $X, Y \in \Omega$ such that the marginal distributions are $P_X = P$ and $P_Y = Q$, respectively. The following definition will be of utility in the sequel [Dob70, LPW09].

Definition 2 (Optimal Coupling). *For any two probability distributions P and Q on Ω , we define the optimal coupling of P and Q as the joint distribution $\gamma^* = \arg \min_{\gamma \in \Pi(P, Q)} \mathbb{E}_{(X, Y) \sim \gamma} [\mathbb{1}\{X \neq Y\}]$ of two random variables (X, Y) that has marginal distributions P and Q of X and Y , respectively, and minimizes the probability that X is different to Y .*

We note that such an optimal coupling always exists [Dob70, Vil09]. The next lemma presents several useful characterizations of TV distance including an optimal coupling representation (which demonstrates how TV distance is a Wasserstein distance with respect to the discrete metric [Vil09]); see [Mak19, Section 2.2.1] and [PW22, Theorem 7.7] for a compilation of other characterizations.

Lemma 1 (Equivalent Characterizations of TV Distance [Str65, Dob70, LPW09, Vil09]). *For any two probability distributions P and Q on Ω , we have*

$$\begin{aligned} d_{\text{TV}}(P, Q) &= \frac{1}{2} \|P - Q\|_1 \\ &= \min_{\gamma \in \Pi(P, Q)} \Pr_{(X, Y) \sim \gamma} [X \neq Y] \\ &= \sup_{f: \Omega \rightarrow [-\frac{1}{2}, \frac{1}{2}]} \mathbb{E}_{X \sim P} [f(X)] - \mathbb{E}_{Y \sim Q} [f(Y)], \end{aligned}$$

where $\|P - Q\|_1$ is the \mathcal{L}^1 -distance between P and Q , the second equality is the optimal coupling characterization which minimizes $\Pr_{(X, Y) \sim \gamma} [X \neq Y] = \mathbb{E}_{(X, Y) \sim \gamma} [\mathbb{1}\{X \neq Y\}]$ over all couplings γ of P and Q , and the third equality is the Kantorovich-Rubinstein dual characterization which takes the supremum over all (measurable) functions $f: \Omega \rightarrow \mathbb{R}$ bounded by $\sup_{t \in \Omega} |f(t)| \leq \frac{1}{2}$.

3 Robustness of Mechanisms Under Total Variation Distance

In this section, we prove our robustness results for DSIC and BIC mechanisms. Missing proofs throughout the section are deferred to Appendix B.

At the heart of our approach lies the following lemma, which shows that, even for correlated prior distributions, assuming truthful bidding, a mechanism's performance with respect to an arbitrary objective function \mathcal{O} is stable under small perturbations to the prior.

Lemma 2. *Let P and Q be two arbitrary probability distributions supported on \mathcal{T} and let \mathcal{M} be any mechanism. Assuming truthful bidding, for all objective functions $\mathcal{O}(\cdot, \cdot) \in [a, b]$, letting $V = b - a$, it holds that $\mathbb{E}_{t \sim P} [\mathcal{O}(t, \mathcal{M}(t))] - \mathbb{E}_{t' \sim Q} [\mathcal{O}(t', \mathcal{M}(t'))] \leq V d_{\text{TV}}(P, Q)$.*

We note that \mathcal{M} in the statement of Lemma 2 may not be DSIC or BIC, for neither P nor Q .

Proof of Lemma 2. The objective function is lower bounded by a and upper bounded by b ; thus, we have that for any $t, t' \in \mathcal{T}$, $\mathcal{O}(t, \mathcal{M}(t)) - \mathcal{O}(t', \mathcal{M}(t')) \leq V \mathbb{1}\{t \neq t'\}$. Since this inequality holds for all $t, t' \in \mathcal{T}$, it also holds after taking an expectation with respect to any coupling γ of P and Q , and specifically for the optimal coupling γ^* between P and Q (see Definition 2):

$$\mathbb{E}_{(t, t') \sim \gamma^*} [\mathcal{O}(t, \mathcal{M}(t)) - \mathcal{O}(t', \mathcal{M}(t'))] \leq V \mathbb{E}_{(t, t') \sim \gamma^*} [\mathbb{1}\{t \neq t'\}]. \quad (1)$$

Using the second form of Lemma 1, the RHS of (1) is equal to $V d_{\text{TV}}(P, Q)$. For the LHS of (1), using linearity of expectation, the fact that $\mathcal{O}(t, \mathcal{M}(t))$ does not depend on t' , and the fact that $\mathcal{O}(t', \mathcal{M}(t'))$ does not depend on t , we have that $\mathbb{E}_{(t, t') \sim \gamma^*} [\mathcal{O}(t, \mathcal{M}(t)) - \mathcal{O}(t', \mathcal{M}(t'))] = \mathbb{E}_{t \sim P} [\mathcal{O}(t, \mathcal{M}(t))] - \mathbb{E}_{t' \sim Q} [\mathcal{O}(t', \mathcal{M}(t'))]$. Putting everything together, we get the desired inequality. \square

In the remainder of this section, we show that this lemma can be used to prove strong robustness results for DSIC and BIC mechanisms.

3.1 DSIC mechanisms

Our main robustness result for DSIC mechanisms is stated as follows.

Theorem 1 (Robustness for DSIC). *Let \mathcal{D} and $\widehat{\mathcal{D}}$ be two arbitrary distributions supported on \mathcal{T} , respectively, such that $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq \delta$, and let $\mathcal{O}(\cdot, \cdot) \in [a, b]$ be an objective function. Let $\mathcal{M}_{\mathcal{D}}^{\alpha}$ be an ex-post IR, DSIC, and α -approximate mechanism (under \mathcal{D}), with respect to the optimal (under \mathcal{D}) ex-post IR and DSIC mechanism for \mathcal{O} . Then, letting $V = b - a$, it holds that*

$$\mathbb{E}_{t \sim \widehat{\mathcal{D}}} [\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}^{\alpha}(t))] \geq \alpha \text{OPT}_{\mathcal{O}}(\widehat{\mathcal{D}}) - (1 + \alpha)V\delta,$$

where $\text{OPT}_{\mathcal{O}}(\widehat{\mathcal{D}})$ is the performance of the optimal (for \mathcal{O}) DSIC mechanism.

Proof. We apply Lemma 2 twice. First, choosing $P = \mathcal{D}$, $Q = \widehat{\mathcal{D}}$, and $\mathcal{M} = \mathcal{M}_{\mathcal{D}}^{\alpha}$, Lemma 2 implies

$$\mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}^{\alpha}(t))] - \mathbb{E}_{t' \sim \widehat{\mathcal{D}}} [\mathcal{O}(t', \mathcal{M}_{\mathcal{D}}^{\alpha}(t'))] \leq V d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq V \delta.$$

Now, let $\mathcal{M}_{\widehat{\mathcal{D}}}^*$ be the optimal (under $\widehat{\mathcal{D}}$) ex-post IR and DSIC mechanism for \mathcal{O} . Using the definition of $\mathcal{M}_{\widehat{\mathcal{D}}}^*$, and re-arranging we have

$$\begin{aligned} \mathbb{E}_{t' \sim \widehat{\mathcal{D}}} [\mathcal{O}(t', \mathcal{M}_{\widehat{\mathcal{D}}}^*(t'))] &\geq \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}^{\alpha}(t))] - V \delta \\ &\geq \alpha \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}^*(t))] - V \delta \\ &\geq \alpha \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}_{\widehat{\mathcal{D}}}^*(t))] - V \delta, \end{aligned}$$

where the last inequality because $\mathcal{M}_{\widehat{\mathcal{D}}}^*$ is feasible (i.e. ex-post IR and DSIC) for \mathcal{D} . Second, choosing $P = \widehat{\mathcal{D}}$, $Q = \mathcal{D}$, and $\mathcal{M} = \mathcal{M}_{\widehat{\mathcal{D}}}^*$, by Lemma 2, plus re-arranging, we have that $\mathbb{E}_{t' \sim \mathcal{D}} [\mathcal{O}(t', \mathcal{M}_{\widehat{\mathcal{D}}}^*(t'))] \geq \mathbb{E}_{t \sim \widehat{\mathcal{D}}} [\mathcal{O}(t, \mathcal{M}_{\widehat{\mathcal{D}}}^*(t))] - V \delta$. Combining these inequalities we have:

$$\mathbb{E}_{t \sim \widehat{\mathcal{D}}} [\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}^{\alpha}(t))] \geq \alpha \mathbb{E}_{t \sim \widehat{\mathcal{D}}} [\mathcal{O}(t, \mathcal{M}_{\widehat{\mathcal{D}}}^*(t))] - (1 + \alpha)V\delta. \quad \square$$

Intuitively, Theorem 1 states that, if a mechanism is approximately optimal for \mathcal{D} , then it is also approximately optimal for all $\widehat{\mathcal{D}}$ that are close in total variation distance, paying a small additive error. Note that neither \mathcal{D} nor $\widehat{\mathcal{D}}$ needs to be a product distribution. That is, if \mathcal{D} and $\widehat{\mathcal{D}}$ have the same support, a DSIC mechanism $\mathcal{M}_{\mathcal{D}}^{\alpha}$ performs approximately-optimally under $\widehat{\mathcal{D}}$. Finally, notice that in the above theorem, we assume that the \mathcal{D} and $\widehat{\mathcal{D}}$ share the same support; in Appendix B.1 we relax this assumption, and show how to modify the mechanism $\mathcal{M}_{\mathcal{D}}^{\alpha}$, in a way that is agnostic to $\widehat{\mathcal{D}}$, and provide exactly the same guarantee.

3.2 Bayesian Incentive Robustness

In this section, we study BIC mechanisms. As opposed to DSIC mechanisms, arguing about whether a BIC mechanism remains BIC after perturbing the prior distribution is a lot more involved. Our goal in this section is to prove that the BIC property degrades gracefully as a function of the TV distance, even for arbitrary objectives, albeit, with two small technical caveats (compared to the DSIC robustness): (1) \mathcal{D} and $\widehat{\mathcal{D}}$ must share the same support, and (2) the incentive guarantees of our mechanisms also degrade. The second requirement is necessary; in Section 4.1 we show how to bypass the first requirement for the revenue objective, recovering a slightly stronger version of a TV robustness result of Brustle et al. [BCD20].

First, we need the following technical lemma that states that when two conditional $P_{X,Y}$, $Q_{X,Y}$ are close in TV distance, then, with high probability, the conditional distributions $P_{Y|X=x}$, $Q_{Y|X=x}$ are also close in TV.

Lemma 3. Let $P_{X,Y}, Q_{X,Y}$ be two probability distributions for the (possibly multivariate) random variables X and Y . Let $P_{Y|X=x}$ (resp. $Q_{Y|X=x}$) be the probability distribution of $P_{X,Y}$ (resp. $Q_{X,Y}$) conditioned on $X = x$, and let Q_X be the marginal probability distribution of X as dictated by $Q_{X,Y}$. Then, for all $q \in [0, 1]$, $\Pr_{x \sim Q_X} \left[d_{\text{TV}}(P_{Y|X=x}, Q_{Y|X=x}) > \frac{2 d_{\text{TV}}(P_{X,Y}, Q_{X,Y})}{q} \right] \leq q$.

Next, we prove that a BIC mechanism for \mathcal{D} is (ε, q) -BIC for $\widehat{\mathcal{D}}$, assuming that \mathcal{D} and $\widehat{\mathcal{D}}$ have the same support and small TV distance.

Lemma 4. Let \mathcal{D} and $\widehat{\mathcal{D}}$ be two probability distributions supported on \mathcal{T} , with $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq \delta$. If \mathcal{M} is an ex-post IR and BIC mechanism w.r.t. \mathcal{D} then it is also an ex-post IR and $(\frac{8H\delta}{q}, q)$ -BIC mechanism w.r.t. $\widehat{\mathcal{D}}$, for all $q \in [0, 1]$.

To prove the above lemma we leverage Lemma 3. We know that with high probability the perception of each agent over the distributions of the rest of the agents is very close under \mathcal{D} and $\widehat{\mathcal{D}}$. That is why, if agent i cannot gain by misreporting under \mathcal{D} , then with high probability, she cannot significantly increase her utility by misreporting under $\widehat{\mathcal{D}}$.

As a direct implication of Lemma 2 and Lemma 4 we get our main robustness result for BIC.

Theorem 2 (Robustness for BIC). Let \mathcal{D} and $\widehat{\mathcal{D}}$ be two arbitrary distributions supported on \mathcal{T} , such that $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq \delta$, and let $\mathcal{O}(\cdot, \cdot) \in [a, b]$ be an objective function. Let $\mathcal{M}_{\mathcal{D}}$ be a mechanism that is ex-post IR and BIC w.r.t. \mathcal{D} . Then $\mathcal{M}_{\mathcal{D}}$ is also ex-post IR and $(\frac{8H\delta}{q}, q)$ -BIC w.r.t. $\widehat{\mathcal{D}}$, for all $q \in [0, 1]$. Also, letting $V = b - a$, it holds that $\mathbb{E}_{t \sim \widehat{\mathcal{D}}}[\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}(t))] \geq \mathbb{E}_{t \sim \mathcal{D}}[\mathcal{O}(t, \mathcal{M}_{\mathcal{D}}(t))] - V\delta$.

Comparing to the corresponding result of Brustle et al. [BCD20], Theorem 2 only holds if \mathcal{D} and $\widehat{\mathcal{D}}$ have the same support. Also, the guarantee on incentives is weaker: $\mathcal{M}_{\mathcal{D}}$ is not ε -BIC (as in [BCD20]), but (ε, q) -BIC. However, our robustness results holds for *arbitrary objectives*, not just revenue, and *arbitrary distributions*, not just product distributions.

4 Applications

In this section, we show a number of applications of Lemma 2, and Theorems 1 and 2.

4.1 BIC mechanisms

We start by showing applications of our robustness results for the revenue objective of BIC mechanisms. Our goal is to extend our DSIC robustness result (Theorem 1) to BIC mechanisms. That is, we'd like, given an approximately optimal BIC mechanism for a distribution \mathcal{D} , to get an approximately optimal and approximately BIC mechanism for distribution $\widehat{\mathcal{D}}$. We will achieve this goal for the revenue objective and product distributions; missing proofs can be found in Appendix C.

Towards our main result for this section, we prove the following lemma, which might be of independent interest. Intuitively, the lemma shows that one can turn an (ε, q) -BIC mechanism into a $O(\varepsilon + nqH)$ -BIC mechanism, paying a small loss in revenue.

Lemma 5. For any product distribution $\mathcal{D} = \times_{i \in [n]} \mathcal{D}_i$, given a mechanism \mathcal{M} that is ex-post IR and (ε, q) -BIC w.r.t. \mathcal{D} , we can design a mechanism $\widehat{\mathcal{M}}$ that is ex-post IR and $O(\varepsilon + nqH)$ -BIC w.r.t. \mathcal{D} , such that $\text{Rev}(\widehat{\mathcal{M}}, \mathcal{D}) \geq \text{Rev}(\mathcal{M}, \mathcal{D}) - nqV$.

Our main theorem for this stage is stated as follows.

Theorem 3. Let \mathcal{D} be a probability distribution supported on \mathcal{T} , and let \mathcal{D}^p be a product distribution such that $d_{\text{TV}}(\mathcal{D}, \mathcal{D}^p) \leq \delta$. Let $\mathcal{M}_{\mathcal{D}^p}^\alpha$ be an ex-post IR, BIC, and α -approximate mechanism (under \mathcal{D}^p), with respect to the revenue optimal (under \mathcal{D}^p) ex-post IR and BIC mechanism. Then, $\mathcal{M}_{\mathcal{D}^p}^\alpha$ is ex-post IR and $(\frac{8H\delta}{q}, q)$ -BIC with respect to \mathcal{D} , for all $q \in (0, 1]$. Furthermore, $\text{Rev}(\mathcal{M}_{\mathcal{D}^p}^\alpha, \mathcal{D}) \geq \alpha \text{OPT}(\mathcal{D}) - O\left((1 + \alpha)V \sqrt{n\sqrt{\delta}}\right)$.

Of course, in order to get the full benefits of Theorem 3, one needs the product distribution \mathcal{D}^p that is the closest (in TV distance) to the original distribution \mathcal{D} . The following proposition shows that, for every distribution \mathcal{D} , the distance of \mathcal{D} to the product of its marginals is bounded by $(n + 1)\varepsilon$, and therefore, Theorem 3 holds for every distribution \mathcal{D} for $\delta = (n + 1)\varepsilon$.

Proposition 1. *Let $\mathcal{D} = \times_{i \in [n]} \mathcal{D}_i$ be a product distribution supported on $\times_{i \in [n]} \mathcal{T}_i$, and let $\widehat{\mathcal{D}}$ be a joint distribution supported on $\times_{i \in [n]} \mathcal{T}_i$, whose marginal over \mathcal{T}_i is exactly \mathcal{D}_i . Then, if there exists a product distribution \mathcal{D}^p such that $d_{\text{TV}}(\widehat{\mathcal{D}}, \mathcal{D}^p) \leq \varepsilon$, it holds that $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq (n + 1)\varepsilon$.*

4.2 Marginal Robustness

As a second application, we consider the setting of Bei et. al. [BGLT19]. In this problem, we want to sell one item to n dependent agents, and we only know the marginal distribution of each agent i . Overloading notation, let $t_i \in \mathcal{T}_i \subseteq \mathbb{R}_+$ be the valuation of agent i for the item. There is a distribution \mathcal{D} , supported on $\times_{i \in [n]} \mathcal{T}_i$, from which agents' valuations are sampled from. For the sake of simplicity, we only argue about discrete distributions; however, our results easily extend to the continuous case. Let \mathcal{D}_i be the marginal distribution of agent i , i.e., $\Pr_{t_i \sim \mathcal{D}_i}[t_i = v_i] = \sum_{v_{-i} \in \mathcal{T}_{-i}} \Pr_{t \sim \mathcal{D}}[t = (v_i, v_{-i})]$. The designer knows marginal distributions \mathcal{D}_i for each i but not \mathcal{D} . Given a set of marginal distributions $(\mathcal{D}_1, \dots, \mathcal{D}_n)$ let $\Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)$ be the set of all distributions consistent with such marginals, i.e. $\Pi(\mathcal{D}_1, \dots, \mathcal{D}_n) = \{\mathcal{D}' | \Pr_{t_i \sim \mathcal{D}_i}[t_i = v_i] = \sum_{v_{-i} \in \mathcal{T}_{-i}} \Pr_{t \sim \mathcal{D}'}[t = (v_i, v_{-i})], \forall i \in [n], \forall t_i \in \mathcal{T}_i\}$. Our goal is to find a mechanism \mathcal{M} such that $\mathcal{M} = \arg \max_{\mathcal{M}'} \min_{\mathcal{D}' \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}'}[\mathcal{O}(t, \mathcal{M}'(t))]$ where \mathcal{M}' is taken over all possible ex-post IR and DSIC mechanisms. We can prove the following theorem.

Theorem 4. *Given a set of marginals $(\mathcal{D}_1, \dots, \mathcal{D}_n)$ and a DSIC and ex-post IR mechanism \mathcal{M}^α such that $\min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}}[\mathcal{O}(t, \mathcal{M}^\alpha(t))] \geq \alpha \max_{\mathcal{M}'} \min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}}[\mathcal{O}(t, \mathcal{M}'(t))]$, then for any set of marginals $(\mathcal{D}'_1, \dots, \mathcal{D}'_n)$ such that for all $i \in [n]$, $d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}'_i) \leq \varepsilon$, it holds that $\min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'}[\mathcal{O}(t, \mathcal{M}^\alpha(t))] \geq \alpha \max_{\mathcal{M}'} \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'}[\mathcal{O}(t, \mathcal{M}'(t))] - (1 + \alpha)n\varepsilon V$.*

The input to our problem is a set of marginals; however, we do not sample from these marginals to compute our objective. Instead, these marginals are used in order to derive a new distribution from which we will sample. Therefore, we cannot immediately “black box” the results we have shown until now. To prove Theorem 4 we first relate, and bound the “distance” between $\Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)$ and $\Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)$. Then, we need to re-prove arguments equivalent to the ones used for our robustness results so far, while taking into consideration the min max nature of the problem. We postpone the formal proof of Theorem 4 to Appendix D.

[BGLT19] show that Sequential Posted Prices Mechanisms are a 4.78-approximation and there exists a Lookahead Auction that is a 2-approximation with respect to revenue for the above problem. Using Theorem 4, we can readily get implications for the robustness of those two auctions.

4.3 Prophet inequalities

Here, we study the prophet inequality problem. Recall that in this problem, agents arrive over time; in the i -th step, we need to (immediately and irrevocably) decide on the allocation of agent i , whose type t_i is drawn from a marginal distribution \mathcal{D}_i . Our goal is to design a prophet inequality: a policy that competes with the optimal in hindsight welfare maximizing allocation. It is known that, in fairly general domains, one can achieve this goal using posted prices, i.e., set a price p_j for each item j , and let users pick their utility-maximizing subset of items [SC84, FGL14, RS17, KW19, ANSS19, DKL20, DFKL20]; see [Luc17] for a survey. Using Lemma 2 for the welfare objective we prove robustness for prophet inequalities, noting that posted price mechanisms are ex-post IR and DSIC.

Corollary 1. *For a posted price mechanism \mathcal{M} , and distributions $\mathcal{D}, \widehat{\mathcal{D}}$ such that $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq \delta$, it holds that $\text{Val}(\mathcal{M}, \mathcal{D}) \geq \text{Val}(\mathcal{M}, \widehat{\mathcal{D}}) - V\delta$.*

Using Corollary 1 for $\mathcal{D} = \times_{i \in [n]} \mathcal{D}_i$ and $\widehat{\mathcal{D}} = \times_{i \in [n]} \widehat{\mathcal{D}}_i$ that are product distributions, we can get the TV robustness result of [DK19] as a special case, noting that (1) if $d_{\text{TV}}(\mathcal{D}_i, \widehat{\mathcal{D}}_i) \leq \varepsilon$, then $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq n\varepsilon$, and (2) if valuations are normalized to $[0, 1]$ (as in their setting), $V = n$.

4.4 Gaps between simple and optimal mechanisms

Here, we show applications of our robustness results to the study of simple and approximately optimal mechanisms. Missing proofs are deferred to Appendix E.

Motivated by numerous negative results for revenue optimal auctions (see [Das15] for a survey), a major research thread in mechanism design studies the performance of simple mechanisms [CHK07, CHMS10, CMS15, Yao15, RW15, CM16, CDW16, CZ17, KW19, BILW20]. Two canonical mechanisms that are considered simple in this literature are: (1) the mechanism that sells each item separately (and optimally), and (2) the mechanism that optimally sells all items as a grand bundle; let $SRev(\mathcal{D})$ and $BRev(\mathcal{D})$ be the expected revenue of these mechanisms, respectively, under prior \mathcal{D} . On the flip side, product distributions are known to witness simple mechanisms that are approximately optimal. For instance, Babaioff et. al. [BILW20] prove that for a single agent and product distribution \mathcal{D}^p , $\max\{SRev(\mathcal{D}^p), BRev(\mathcal{D}^p)\} \geq \frac{1}{6} Rev(\mathcal{D}^p)$, even though the individual approximation factors for $SRev(\mathcal{D}^p)$ and $BRev(\mathcal{D}^p)$ are $\tilde{O}(1/\log(m))$ and $1/m$, respectively [HN13, LY13]. The following observation is an immediate implication of Theorem 1.

Observation 1. *Let \mathcal{SM} be a family of mechanisms, such that $\mathbb{E}_{t \sim \hat{\mathcal{D}}^p} [\mathcal{O}(t, \mathcal{SM}_{\hat{\mathcal{D}}^p}(t))] \geq \alpha OPT_{\mathcal{O}}(\hat{\mathcal{D}}^p)$ for some objective \mathcal{O} , where $\mathcal{SM}_{\hat{\mathcal{D}}^p}$ is a mechanism in \mathcal{SM} parameterized by the product distribution $\hat{\mathcal{D}}^p$. Then, for any (possibly non-product) distribution \mathcal{D} that is close to some product distributions \mathcal{D}^p , and specifically, $d_{TV}(\mathcal{D}, \mathcal{D}^p) \leq \frac{\alpha OPT_{\mathcal{O}}(\mathcal{D})}{2(1+\alpha)V}$, we have that $\mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{SM}_{\mathcal{D}^p}(t))] \geq \frac{\alpha}{2} OPT_{\mathcal{O}}(\mathcal{D})$.*

A construction of Hart and Nisan [HN13] shows that there exist non-product distribution \mathcal{D} such that $SRev(\mathcal{D}) \leq 2m$, $BRev(\mathcal{D}) \leq 2m$ but $Rev(\mathcal{D}) \rightarrow \infty$. Recently, [PSCW22] show necessary and sufficient conditions for a distribution \mathcal{D} to exhibit such infinite gaps between the revenue of simple and optimal mechanisms. Unfortunately, these conditions are rather complex (namely, they are conditions on the existence of infinite sequences of points with certain properties). Here, leveraging Observation 1 for the revenue objective, we complement results of [PSCW22] by giving a new, simple necessary condition that a “pathological” construction must satisfy:

Corollary 2. *For a single agent and any distribution \mathcal{D} such that $\frac{Rev(\mathcal{D})}{BRev(\mathcal{D})} \geq 2m$ it must be the case that for any product distribution \mathcal{D}^p , $d_{TV}(\mathcal{D}, \mathcal{D}^p) \geq \frac{Rev(\mathcal{D})}{4V}$. Furthermore, if $\frac{Rev(\mathcal{D})}{SRev(\mathcal{D})} \in \Omega(\log(m))$ it must be the case that for any product distribution \mathcal{D}^p , $d_{TV}(\mathcal{D}, \mathcal{D}^p) \geq \frac{Rev(\mathcal{D})}{4V}$.*

As a final application of our framework, we give new positive results on simple mechanisms.

Proposition 2. *Let \mathcal{D} be a distribution supported on \mathcal{T} , and let \mathcal{D}^p be a product distribution such that $d_{TV}(\mathcal{D}, \mathcal{D}^p) < \delta$. Then, we have $\max\{SRev(\mathcal{D}), BRev(\mathcal{D})\} \geq \frac{1}{6} Rev(\mathcal{D}) - \frac{7}{6} H\delta$.*

Proposition 2 implies that the better of bundling and selling each item separately for a distribution \mathcal{D} is a good approximation to $Rev(\mathcal{D})$, as long as \mathcal{D} is close to a product distribution. Cai and Oikonomou [CO21] prove that the same mechanism is a good approximation for dependent distributions that can be captured by a Markov Random Field (MRF); see Appendix E for basic definitions regarding MRFs. The approximation ratio of [CO21] is controlled by Δ , a parameter of the MRF that is determined by how much the value of an item can be influenced by the values of the other items. Specifically, [CO21] prove that $\max\{SRev(\mathcal{D}), BRev(\mathcal{D})\} \geq \frac{1}{12e^{4\Delta}} Rev(\mathcal{D})$. In Proposition 3, we prove that there exist distributions such that the object of interest for Proposition 2, i.e., the distance to a product distribution, is bounded, while the objective of interest for [CO21], the parameter Δ , is arbitrarily large. At the same time, in Proposition 4 we show that when the MRF parameter Δ is bounded, we can bound the distance of a distribution \mathcal{D} to a product distribution, for distributions that can be represented by MRFs that only have pairwise edges.

Proposition 3. *For every $k > 0$, there exists a distribution \mathcal{D} produced by an MRF with parameter Δ and a product distribution \mathcal{D}^p such that $d_{TV}(\mathcal{D}, \mathcal{D}^p) \leq 1/2$ and $\Delta \geq k$.*

Proposition 4. *Let \mathcal{D} be a distribution produced by an MRF with only pairwise edges. Then, there exists a product distribution \mathcal{D}^p such that $d_{TV}(\mathcal{D}, \mathcal{D}^p) \leq \min\{\sqrt{m\Delta}/4, \sqrt{1 - e^{-m\Delta/2}}\}$, where m is the number of items.*

References

- [AMDW13] Pablo Azar, Silvio Micali, Constantinos Daskalakis, and S Matthew Weinberg. Optimal and efficient parametric auctions. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 596–604. SIAM, 2013.
- [ANSS19] Nima Anari, Rad Niazadeh, Amin Saberi, and Ali Shameli. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 91–92, 2019.
- [BCD20] Johannes Brustle, Yang Cai, and Constantinos Daskalakis. Multi-item mechanisms without item-independence: Learnability via robustness. In *Proceedings of the 21st ACM Conference on Economics and Computation, EC '20*, page 715–761, New York, NY, USA, 2020. Association for Computing Machinery.
- [BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015.
- [BGLT19] Xiaohui Bei, Nick Gravin, Pinyan Lu, and Zhihao Gavin Tang. Correlation-robust analysis of single item auction. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 193–208. SIAM, 2019.
- [BH78] J. L. Bretagnolle and Catherine Huber. Estimation des densités: risque minimax. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 47:119–137, 1978.
- [BILW20] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)*, 67(4):1–40, 2020.
- [BS11] Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.
- [Car17] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, 2017.
- [CD17] Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 516–527. IEEE, 2017.
- [CDW16] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, STOC '16*, page 926–939, New York, NY, USA, 2016. Association for Computing Machinery.
- [CHK07] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07*, pages 243–251, New York, NY, USA, 2007. ACM.
- [CHMS10] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.
- [CM16] Shuchi Chawla and J Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 579–596. ACM, 2016.
- [CMS15] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- [CO21] Yang Cai and Argyris Oikonomou. On simple mechanisms for dependent items. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 242–262, 2021.

- [COVZ21] Yang Cai, Argyris Oikonomou, Grigoris Veleghkas, and Mingfei Zhao. An efficient ε -bic to bic transformation and its application to black-box reduction in revenue maximization. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1337–1356. SIAM, 2021.
- [CR14] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 243–252, 2014.
- [CZ17] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 170–183. ACM, 2017.
- [Das15] Constantinos Daskalakis. Multi-item auctions defying intuition? *ACM SIGecom Exchanges*, 14(1):41–75, 2015.
- [DFK11] Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. Optimal auctions with correlated bidders are easy. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 129–138, New York, NY, USA, 2011. Association for Computing Machinery.
- [DFKL20] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582, 2020.
- [DG85] Luc Devroye and László Györfi. *Nonparametric Density Estimation: The L_1 View*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, NY, USA, 1985.
- [DHP16] Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 426–439, 2016.
- [DK19] Paul Dütting and Thomas Kesselheim. Posted pricing and prophet inequalities with inaccurate priors. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 111–129, New York, NY, USA, 2019. Association for Computing Machinery.
- [DKL20] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. An $o(\log \log m)$ prophet inequality for subadditive combinatorial auctions. *ACM SIGecom Exchanges*, 18(2):32–37, 2020.
- [Dob70] Roland L. Dobrushin. Prescribing a system of random variables by conditional distributions. *Theory of Probability and Its Applications*, 15(3):469–497, 1970.
- [Dob71] Roland L. Dobrushin. Markov processes with a large number of locally interacting components: Existence of a limit process and its ergodicity. *Problemy Peredachi Informatsii*, 7(2):70–87, 1971.
- [Doe38] Wolfgang Doeblin. Exposé de la théorie des chaînes simples constantes de Markov à un nombre fini d'états. *Revue Mathématique de l'Union Interbalkanique*, 2:77–105, 1938.
- [Dud68] Richard M. Dudley. Distances of probability measures and random variables. *The Annals of Mathematical Statistics*, 39(5):1563–1572, October 1968.
- [DW12] Constantinos Daskalakis and Seth Matthew Weinberg. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, page 370–387, New York, NY, USA, 2012. Association for Computing Machinery.
- [FGL14] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 123–135. SIAM, 2014.

- [GHZ19] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673, 2019.
- [GL18] Nick Gravin and Pinyan Lu. Separation in correlation-robust monopolist problem with budget. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2069–2080. SIAM, 2018.
- [GPTD23] Yiannis Giannakopoulos, Diogo Poças, and Alexandros Tsigonias-Dimitriadis. Robust revenue maximization under minimal statistical information. *ACM Transactions on Economics and Computation*, 10(3):1–34, 2023.
- [Gri75] David Griffeath. A maximal coupling for markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 31:95–106, June 1975.
- [GW21] Yannai A Gonczarowski and S Matthew Weinberg. The sample complexity of up-to- ϵ multi-dimensional revenue maximization. *Journal of the ACM (JACM)*, 68(3):1–28, 2021.
- [HJW15] Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Minimax estimation of discrete distributions under ℓ_1 loss. *IEEE Transactions on Information Theory*, 61(11):6343–6354, November 2015.
- [HMR15] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 45–60, 2015.
- [HN13] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. *ACM Conference on Electronic Commerce*, 04 2013.
- [HN19] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019.
- [Kal21] Olav Kallenberg. *Foundations of Modern Probability*, volume 99 of *Probability Theory and Stochastic Modelling*. Springer, New York, NY, USA, third edition, 2021.
- [Kan60] Leonid V. Kantorovich. Mathematical methods of organizing and planning production. *Management Science*, 6(4):366–422, July 1960.
- [KS80] Ross Kindermann and Laurie Snell. *Markov random fields and their applications*, volume 1. American Mathematical Society, 1980.
- [KW19] Robert Kleinberg and S Matthew Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113:97–115, 2019.
- [LLY19] Yingkai Li, Pinyan Lu, and Haoran Ye. Revenue maximization with imprecise distribution. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1582–1590, 2019.
- [LPW09] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI, USA, first edition, 2009.
- [Luc17] Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.
- [LY13] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.
- [Mak19] Anuran Makur. *Information Contraction and Decomposition*. Sc.D. thesis in Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA, May 2019.

- [MR16] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pages 1298–1318. PMLR, 2016.
- [PSCW22] Alexandros Psomas, Ariel Schwartzman Cohenca, and S Weinberg. On infinite separations between simple and optimal mechanisms. *Advances in Neural Information Processing Systems*, 35:4818–4829, 2022.
- [PSW19] Alexandros Psomas, Ariel Schwartzman, and S Matthew Weinberg. Smoothed analysis of multi-item auctions with correlated values. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 417–418. ACM, 2019.
- [PW22] Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning*. Cambridge University Press Preprint, New York, NY, USA, 2022.
- [RS17] Aviad Rubinfeld and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687. SIAM, 2017.
- [RW15] Aviad Rubinfeld and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394. ACM, 2015.
- [RW18] Aviad Rubinfeld and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Trans. Econ. Comput.*, 6(3–4), oct 2018.
- [SC84] Ester Samuel-Cahn. Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *The Annals of Probability*, 12(4):1213 – 1216, 1984.
- [SK75] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35:1792–1796, Dec 1975.
- [Sko56] Anatoliy V. Skorokhod. Limit theorems for stochastic processes. *Theory of Probability and Its Applications*, 1(3):261–290, 1956.
- [ST04] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2004.
- [Str65] Volker Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, April 1965.
- [Tsy08] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008.
- [Vil09] Cédric Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, Germany, 2009.
- [Wic70] Michael J. Wichura. On the construction of almost uniformly convergent random variables with given weakly convergent image laws. *The Annals of Mathematical Statistics*, 4141(1):284–291, February 1970.
- [Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109. Society for Industrial and Applied Mathematics, 2015.

A Further Related Work

A related, but different, approach is to assume sample access to the underlying distribution [CR14, HMR15, DHP16, MR16, CD17, GHZ19, GW21]. In this line of work, the goal is bound the number of samples necessary to design a near-optimal mechanism, or, given a fixed number of samples, design the best mechanism possible. Robustness results of the former style, i.e., the current paper or [BCD20], sometimes imply sample complexity results, e.g., by arguing that using the samples, one can learn a distribution that is close to the real distribution, and then applying robustness results [BCD20]. We conjecture that similar results can be shown using our robustness framework, combined with estimation in TV distance [HJW15, DG85]. Another line of work [AMDW13, Car17, GL18, BGLT19, GPTD23] assumes partial knowledge of the true prior distribution, e.g., its mean or the marginals of a correlated distribution. The goal is to find the mechanism that (approximately) maximizes the worst-case performance with respect to the missing details (e.g., the CDF consistent with the mean or the joint distribution that respects the given marginals). Our results have implications for such settings as well. Finally, certain works consider the robustness of pathological examples in mechanism design, e.g., constructions of distributions that have infinite revenue gaps between simple and optimal mechanisms [BCKW15, HN19]. For example, [PSW19] uses the lens of smoothed analysis ([ST04]) to reason about the robustness of the [HN19] constructions. Our framework has implications about these constructions; specifically, we give a new necessary condition for a distribution to be “pathological,” complementing a recent result of [PSCW22].

From a probability theory perspective, we heavily use coupling techniques. Couplings are a general proof technique in probability theory with several historically notable uses [Doe38, Sko56, Kan60, Str65, Dud68, Wic70, Dob70] (also see [Vil09, Kal21]). Given two marginal distributions, the basic idea of coupling is to construct a consistent joint distribution on a common probability space in order to deduce certain relationships between the marginals. The key coupling used in this work — Dobrushin’s optimal coupling — minimizes the probability that two random variables with given marginal distributions are different, and has been historically utilized to develop sharper results on Markov chain ergodicity, cf. [Dob70, Dob71, Gri75]. It turns out that this optimal coupling also defines total variation distance, cf. [LPW09, Vil09]. Such distances exhibit Kantorovich-Rubinstein duality and are characterized by the maximal difference of expected values with respect to the given marginal distributions [Kan60, LPW09, Vil09]. In this work, we distill how optimal couplings and duality for total variation can be used in yet another setting: mechanism design.

B Proofs missing from Section 3

The following simple proposition will also be useful in multiple proofs throughout this appendix.

Proposition 5. *Let \mathcal{M} be an ex-post IR mechanism. Then, $-H \leq u_i^{\mathcal{M}}(t_i \leftarrow t'_i, t_{-i}) \leq 3H$, for all $i \in [n], t_i, t'_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$.*

Proof of Proposition 5. Since \mathcal{M} is ex-post IR, we have that $t_i(\mathcal{M}(t_i, t_{-i})) \geq 0$, for all $i \in [n], t_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$. Furthermore, since payments are lower bounded by $-H$, and since the valuations are bounded and quasi-linear, we have that $t_i(\mathcal{M}(t'_i, t_{-i})) \leq 2H$, for all $i \in [n], t_i, t'_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$. Since payments are also upper bounded by H (due to the ex-post IR constraint), and valuations are non-negative, we also have $t_i(\mathcal{M}(t'_i, t_{-i})) \geq -H$, for all $i \in [n], t_i, t'_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$. Combining these inequalities we have $-H \leq u_i(t_i \leftarrow t'_i, t_{-i}) \leq 3H$, for all $i \in [n], t_i, t'_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$. \square

B.1 Relaxing the assumptions in Theorem 1

We start by showing that, in sharp contrast to BIC, the DSIC property is much easier to “propagate” from a small set of types to a larger set, using the following construction.

Definition 3 (DSIC extension of a mechanism). *Let $\mathcal{T}_i^+ \subseteq \mathcal{T}_i$ be a subset of possible types for agent $i \in [n]$, such that $\perp \in \mathcal{T}_i^+$, and let $\mathcal{M} = (x, p)$ be a mechanism defined on types $\times_{i \in [n]} \mathcal{T}_i^+$. The extension of \mathcal{M} to \mathcal{T} is the mechanism $\widehat{\mathcal{M}} = (\widehat{x}, \widehat{p})$, where for reported types $t = (t_1, \dots, t_n)$:*

1. *If $\times_{i \in [n]} \mathcal{T}_i^+$, then $\widehat{x}(t) = x(t)$ and $\widehat{p}(t) = p(t)$.*

2. If there exists i , such that $t_i \notin \mathcal{T}_i^+$ and $\forall j \in [n]/\{i\} : t_j \in \mathcal{T}_j^+$ then $\hat{x}_i(t) = x_i(t'_i, t_{-i})$ and $\hat{p}_i(t) = \hat{p}_i(t'_i, t_{-i})$, where $t'_i = \arg \max_{z_i \in \mathcal{T}_i^+} t_i(\mathcal{M}(z_i, t_{-i}))$. For each $j \in [n]/\{i\}$ we have that $\hat{x}_j(t) = 0$ and $\hat{p}_j(t) = 0$ (They receive nothing, and pay nothing).
3. If there exist i, i' such that $i \neq i'$ and $t_i \notin \mathcal{T}_i^+$ and $t_{i'} \notin \mathcal{T}_{i'}^+$, then nobody receives and pays nothing (i.e. $x(t) = 0, \hat{p}(t) = 0$).

A similar construction appears in [DFK11], in the context of implementing the solution of a linear program as a DSIC auction.

Lemma 6. Let $\mathcal{T}_i^+ \subseteq \mathcal{T}_i$ be a subset of possible types for agent $i \in [n]$, such that $\perp \in \mathcal{T}_i^+$, and let $\mathcal{M} = (x, p)$ be a DSIC and ex-post IR mechanism defined on types $\mathcal{T}^+ = \times_{i \in [n]} \mathcal{T}_i^+$. Then, the extension of \mathcal{M} to \mathcal{T} , $\widehat{\mathcal{M}} = (\widehat{x}, \widehat{p})$, is DSIC and ex-post IR.

Proof of Lemma 6. The fact that $\widehat{\mathcal{M}}$ is ex-post IR is trivial for cases 1 and 3 of Definition 3. For case 2, it is trivial that it is ex-post IR for all $j \in [n]/\{i\}$. Also since $\perp \in \mathcal{T}_i^+$ we have that $\max_{z_i \in \mathcal{T}_i^+} t_i(\mathcal{M}(z_i, t_{-i})) \geq t_i(\mathcal{M}(\perp, t_{-i})) \geq 0$, which implies that the mechanism is ex-post IR for agent i .

Next, we argue that $\widehat{\mathcal{M}}$ is DSIC. If $t \in \mathcal{T}^+$, then any misreport t'_i of agent i will also get mapped to a type in \mathcal{T}_i^+ ; since \mathcal{M} is DSIC, agent i cannot increase her utility by deviating. If t falls into the second case, an agent $j \in [n]/\{i\}$ receives nothing and pays nothing, no matter what she reports. If agent i misreports a type t'_i , she either receives utility $t_i(\mathcal{M}(t'_i, t_{-i}))$, if $t'_i \in \mathcal{T}_i^+$, or $t_i(\mathcal{M}((t^*)', t_{-i}))$, where $(t^*)' = \arg \max_{z_i \in \mathcal{T}_i^+} t'_i(\mathcal{M}(z_i, t_{-i}))$, if $t'_i \notin \mathcal{T}_i^+$, both of which are (weakly) worse than $\max_{z_i \in \mathcal{T}_i^+} t_i(\mathcal{M}(z_i, t_{-i}))$, her utility when reporting t_i . Finally, in case 3, every agent i always receives nothing and pays nothing, even after unilaterally changing her report. \square

Thus without loss of generality, we can always assume that DSIC mechanism defined on a subset of the type space $\mathcal{T}^+ \subseteq \mathcal{T}$ is DSIC on all bids in \mathcal{T} .

B.2 Proofs missing from Section 3.2

Proof of Lemma 3.

$$\begin{aligned}
2 d_{\text{TV}}(P_{X,Y}, Q_{X,Y}) &= \sum_x \sum_y |P_{X,Y}(x, y) - Q_{X,Y}(x, y)| \\
&\geq \sum_{x: Q_X(x) > 0} \sum_y |P_{X,Y}(x, y) - Q_{X,Y}(x, y)| \\
&= \sum_{x: Q_X(x) > 0} Q_X(x) \sum_y \left| P_{Y|X=x}(y) \frac{P_X(x)}{Q_X(x)} - Q_{Y|X=x}(y) - P_{Y|X=x}(y) + P_{Y|X=x}(y) \right| \\
&\geq \sum_{x: Q_X(x) > 0} Q_X(x) \sum_y \left(|P_{Y|X=x}(y) - Q_{Y|X=x}(y)| - P_{Y|X=x}(y) \left| 1 - \frac{P_X(x)}{Q_X(x)} \right| \right) \\
&= \sum_{x: Q_X(x) > 0} Q_X(x) \left(2 d_{\text{TV}}(P_{Y|X=x}, Q_{Y|X=x}) - \frac{|Q_X(x) - P_X(x)|}{Q_X(x)} \right) \\
&\geq \left(2 \sum_x Q_X(x) d_{\text{TV}}(P_{Y|X=x}, Q_{Y|X=x}) \right) - 2 d_{\text{TV}}(Q_X, P_X).
\end{aligned}$$

Re-arranging, we have that

$$\mathbb{E}_{x \sim Q_X} [d_{\text{TV}}(P_{Y|X=x}, Q_{Y|X=x})] \leq d_{\text{TV}}(P_{X,Y}, Q_{X,Y}) + d_{\text{TV}}(Q_X, P_X).$$

The data processing inequality gives us that $d_{\text{TV}}(Q_X, P_X) \leq d_{\text{TV}}(P_{X,Y}, Q_{X,Y})$ [PW22, Theorem 7.4], and thus we have $\mathbb{E}_{x \sim Q_X} [d_{\text{TV}}(P_{Y|X=x}, Q_{Y|X=x})] \leq 2 d_{\text{TV}}(P_{X,Y}, Q_{X,Y})$, as desired. For distributions supported over continuous sets, the proof follows with similar arguments.

So far, we have established that $\mathbb{E}_{x \sim Q_X} [d_{\text{TV}}(P_{Y|X=x}, Q_{Y|X=x})] \leq d_{\text{TV}}(P_{X,Y}, Q_{X,Y}) + d_{\text{TV}}(Q_X, P_X)$. Using Markov's inequality completes the proof of Lemma 3. \square

Proof of Lemma 4. \mathcal{M} is ex-post IR for \mathcal{D}' , by definition. Let $\mathcal{D}_{-i|t_i}$ be the probability distribution for the valuations of every agent except i , conditioned on the event that the type of agent i is $t_i \in \mathcal{T}_i$. Proposition 5 implies that $u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i}) \in [-H, 3H]$, for all $i \in [n], t_i, w_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$, and therefore $u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i}) - u_i^{\mathcal{M}}(t_i \leftarrow w_i, t'_{-i}) \leq 4H \mathbb{1}\{t_{-i} \neq t'_{-i}\}$. Thus, for any coupling γ of $\mathcal{D}_{-i|t_i}$ and $\mathcal{D}'_{-i|t_i}$, and specifically for the optimal coupling γ^* between $\mathcal{D}_{-i|t_i}$ and $\mathcal{D}'_{-i|t_i}$ (see Definition 2), we have:

$$\begin{aligned} \mathbb{E}_{(t_{-i}, t'_{-i}) \sim \gamma^*} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i}) - u_i^{\mathcal{M}}(t_i \leftarrow w_i, t'_{-i})] &\leq 4H \mathbb{E}_{(t_{-i}, t'_{-i}) \sim \gamma^*} [\mathbb{1}\{t_{-i} \neq t'_{-i}\}] \\ &\leq 4H d_{\text{TV}}(\mathcal{D}_{-i|t_i}, \mathcal{D}'_{-i|t_i}). \end{aligned}$$

Using linearity of expectation and re-arranging we have:

$$-\mathbb{E}_{t'_{-i} \sim \mathcal{D}'_{-i|t_i}} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t'_{-i})] \leq 4H d_{\text{TV}}(\mathcal{D}_{-i|t_i}, \mathcal{D}'_{-i|t_i}) - \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i|t_i}} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i})].$$

By setting $Q_X = \mathcal{D}'_i$, $P_{Y|X=x} = \mathcal{D}_{-i|t_i}$, and $Q_{Y|X=x} = \mathcal{D}'_{-i|t_i}$ in Lemma 3 we have that, with probability at least $1 - q$, $d_{\text{TV}}(\mathcal{D}_{-i|t_i}, \mathcal{D}'_{-i|t_i}) \leq \frac{2}{q} d_{\text{TV}}(\mathcal{D}, \mathcal{D}') \leq 2\frac{\delta}{q}$. Therefore, with probability at least $1 - q$:

$$\begin{aligned} -\mathbb{E}_{t'_{-i} \sim \mathcal{D}'_{-i|t_i}} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t'_{-i})] &\leq 4H d_{\text{TV}}(\mathcal{D}_{-i|t_i}, \mathcal{D}'_{-i|t_i}) - \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i|t_i}} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i})] \\ &\leq 8H \frac{\delta}{q} - \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i|t_i}} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i})] \\ &\leq \frac{8H\delta}{q}, \end{aligned}$$

where the last inequality uses the fact that \mathcal{M} is BIC. Replacing with the definition of $u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i})$ we get $-\mathbb{E}_{t_{-i} \sim \mathcal{D}'_{-i|t_i}} [t_i(\mathcal{M}(t_i, t_{-i}))] + \mathbb{E}_{t_{-i} \sim \mathcal{D}'_{-i|t_i}} [t_i(\mathcal{M}(w_i, t_{-i}))] \leq \frac{8H\delta}{q}$, with probability at least $1 - q$. Re-arranging we get the desired (ε, q) BIC constraint. \square

C Proofs missing from Section 4.1

In order to prove Lemma 5, it will be convenient to define the following notion of an extension of a BIC mechanism.

Definition 4 (BIC extension of a mechanism). *Let $\mathcal{T}_i^+ \subseteq \mathcal{T}_i$ be a subset of types for agent $i \in [n]$ such that $\perp \in \mathcal{T}_i^+$, and let $\mathcal{M} = (x, p)$ be a mechanism defined on types in $\times_{i \in [n]} \mathcal{T}_i^+$. Let $\mathcal{T}_i^- = \mathcal{T}_i - \mathcal{T}_i^+$, and consider the mapping*

$$\tau_i(t_i) = \begin{cases} t_i, & \text{if } t_i \in \mathcal{T}_i^+ \\ \arg \max_{z \in \mathcal{T}_i^+} \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [t_i(\mathcal{M}(z, t_{-i}))], & \text{if } t_i \in \mathcal{T}_i^- \end{cases}$$

The extension of \mathcal{M} to \mathcal{T} is the mechanism $\widehat{\mathcal{M}} = (\widehat{x}, \widehat{p})$, where $\widehat{x}(t) = x(\tau(t))$, and for all $i \in [n]$,

$$\widehat{p}_i(t_i, t_{-i}) = \begin{cases} p_i(t_i, t_{-i}), & \text{if } t_i \in \mathcal{T}_i^+ \\ v_i(\widehat{x}(t_i, t_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau_i(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x(\tau_i(t_i), t_{-i}))]}, & \text{if } t_i \in \mathcal{T}_i^- \end{cases}$$

We prove the following technical lemma.

Lemma 7. *Let $\mathcal{T}_i^+ \subseteq \mathcal{T}_i$ be a subset of types for agent $i \in [n]$ such that $\perp \in \mathcal{T}_i^+$, and let $\mathcal{D} = \times_{i \in [n]} \mathcal{D}_i$ be a product distribution, where each \mathcal{D}_i is supported on \mathcal{T}_i . Let $\mathcal{M} = (x, p)$ be an ex-post IR mechanism which satisfies $\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [u_i^{\mathcal{M}}(t_i \leftarrow w_i, t_{-i})] \geq -\varepsilon$, for all $t_i \in \mathcal{T}_i^+, w_i \in \mathcal{T}_i$.*

Then, for any product distribution $\widehat{\mathcal{D}} = \times_{i \in [n]} \widehat{\mathcal{D}}_i$ such that $d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \leq \delta$, the extension of \mathcal{M} to \mathcal{T} (as defined in Definition 4) is ex-post IR and $O(\varepsilon + (\beta n + \delta)H)$ -BIC with respect to $\widehat{\mathcal{D}}$, where $\beta = 1 - \Pr_{t_i \sim \widehat{\mathcal{D}}_i} [t_i \in \mathcal{T}_i^+]$. Furthermore, $\text{Rev}(\widehat{\mathcal{M}}, \widehat{\mathcal{D}}) \geq \text{Rev}(\mathcal{M}, \mathcal{D}) - V(\beta n + \delta)$.

Proof of Lemma 7. Let $\widehat{\mathcal{M}} = (\widehat{x}, \widehat{p})$ be the extension of \mathcal{M} to \mathcal{T} . First, we argue that $\widehat{\mathcal{M}}$ is ex-post IR. Since \mathcal{M} is ex-post IR, the ex-post IR condition for $\widehat{\mathcal{M}}$ is satisfied for all $t_i \in \mathcal{T}_i^+$, by construction. For a type $t_i \in \mathcal{T}_i^-$, since $\perp \in \mathcal{T}_i^+$ and $\tau_i(t_i) \in \mathcal{T}_i^+$, we have that $\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [t_i(\mathcal{M}(\tau_i(t_i), t_{-i}))] \geq \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [t_i(\mathcal{M}(\perp, t_{-i}))] = 0$. Therefore, $\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau_i(t_i), t_{-i})] \leq \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x(\tau_i(t_i), t_{-i}))]$, which implies that $v_i(\widehat{x}(t)) - \widehat{p}_i(t) = v_i(\widehat{x}(t)) - v_i(\widehat{x}(t)) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau_i(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x(\tau_i(t_i), t_{-i}))]} \geq 0$.

Next, we prove the BIC guarantee of $\widehat{\mathcal{M}}$. Towards this, first define $\tau(\widehat{\mathcal{D}})$ as the distribution induced by first sampling from $\widehat{\mathcal{D}}$, and then apply mapping $\tau(\cdot)$, as defined in Definition 4. The tensorization property of TV distance [LPW09, Chapter 4] implies that $d_{\text{TV}}(\widehat{\mathcal{D}}, \tau(\widehat{\mathcal{D}})) \leq \beta n$, and thus from the triangle inequality, $d_{\text{TV}}(\mathcal{D}, \tau(\widehat{\mathcal{D}})) \leq \delta + \beta n$. Our goal is to prove the following lower bound:

$$\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [u_i^{\widehat{\mathcal{M}}}(t_i \leftarrow w_i, t_{-i})] \geq - \left(4 \left(\frac{3}{2} \delta + \beta n \right) H + 4\delta H + \varepsilon \right).$$

We first prove the following intermediate bound:

$$\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [u_i^{\widehat{\mathcal{M}}}(t_i \leftarrow w_i, t_{-i})] \geq \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_{-i})] - 4 \left(\frac{3}{2} \delta + \beta n \right) H$$

Generally, our bounds will be trivial when $t_i \in \mathcal{T}_i^+$ due to the nature of $\widehat{\mathcal{M}}$. So the main focus of the analysis is to prove those bounds for $t_i \in \mathcal{T}_i^-$.

First, we prove two inequalities that will be useful in our analysis.

$$\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))] \leq \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [\widehat{x}_i(t_i, t_{-i})] + H \beta n. \quad (2)$$

$$\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})] \geq \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [\widehat{p}_i(t_i, t_{-i})] - H \beta n. \quad (3)$$

For inequality (2), using Lemma 2 we can get:

$$\begin{aligned} \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))] &\leq \mathbb{E}_{t_{-i} \sim \tau(\mathcal{D}_{-i})} [v_i(x_i(\tau(t_i), t_{-i}))] + H d_{\text{TV}}(\mathcal{D}_{-i}, \tau(\mathcal{D}_{-i})) \\ &\leq \mathbb{E}_{t_{-i} \sim \tau(\mathcal{D}_{-i})} [v_i(x_i(\tau(t_i), t_{-i}))] + H d_{\text{TV}}(\mathcal{D}, \tau(\mathcal{D})) \\ &\leq \mathbb{E}_{t_{-i} \sim \tau(\mathcal{D}_{-i})} [v_i(x_i(\tau(t_i), t_{-i}))] + H \beta n \\ &\leq \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [x_i(\tau(t_i), \tau(t_{-i}))] + H \beta n \\ &\leq \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [\widehat{x}_i(t_i, t_{-i})] + H \beta n. \end{aligned}$$

Similarly, for inequality (3):

$$\begin{aligned} \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})] &= \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})] \frac{\mathbb{E}_{t'_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t'_{-i}))]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \\ &= \mathbb{E}_{t'_{-i} \sim \mathcal{D}_{-i}} \left[v_i(x_i(\tau(t_i), t'_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \right]. \end{aligned}$$

We've already shown, when arguing the ex-post IR property, that $\frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \leq 1$ and

thus $v_i(x_i(\tau(t_i), t'_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \in [0, H]$. Therefore, we can use Lemma 2 for \mathcal{D}_{-i}

and $\tau(\mathcal{D}_{-i})$ on this function (as the objective) to get:

$$\begin{aligned}
\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})] &= \mathbb{E}_{t'_{-i} \sim \mathcal{D}_{-i}} \left[v_i(x_i(\tau(t_i), t'_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \right] \\
&\geq \mathbb{E}_{t'_{-i} \sim \tau(\mathcal{D}_{-i})} \left[v_i(x_i(\tau(t_i), t'_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \right] - H d_{\text{TV}}(\mathcal{D}_{-i}, \tau(\mathcal{D}_{-i})) \\
&\geq \mathbb{E}_{t'_{-i} \sim \tau(\mathcal{D}_{-i})} \left[v_i(x_i(\tau(t_i), t'_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \right] - H d_{\text{TV}}(\mathcal{D}, \tau(\mathcal{D})) \\
&\geq \mathbb{E}_{t'_{-i} \sim \tau(\mathcal{D}_{-i})} \left[v_i(x_i(\tau(t_i), t'_{-i})) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \right] - H \beta n \\
&= \mathbb{E}_{t'_{-i} \sim \mathcal{D}_{-i}} \left[v_i(x_i(\tau(t_i), \tau(t'_{-i}))) \frac{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [p_i(\tau(t_i), t_{-i})]}{\mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [v_i(x_i(\tau(t_i), t_{-i}))]} \right] - H \beta n \\
&= \mathbb{E}_{t'_{-i} \sim \mathcal{D}_{-i}} [\widehat{p}_i(t_i, t'_{-i})] - H \beta n.
\end{aligned}$$

With inequalities (2) and (3) at hand, we are ready to show the following, for all $t_i \in \mathcal{T}_i^-$:

$$\begin{aligned}
\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(t_i), t_{-i}))] &\leq^{(\text{Lemma 2})} \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [t_i(\mathcal{M}(\tau(t_i), t_{-i}))] + 2\delta H \\
&= \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [(v_i(x_i(\tau(t_i), t_{-i})) - p_i(\tau(t_i), t_{-i}))] + 2\delta H \\
&\leq^{(\text{Ineq. (2) and (3)})} \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [\widehat{x}_i(t_i, t_{-i})] - \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} [\widehat{p}_i(t_i, t_{-i})] + 2(\delta + \beta n) H \\
&= \mathbb{E}_{t_{-i} \sim \mathcal{D}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(t_i, t_{-i}) \right) \right] + 2(\delta + \beta n) H \\
&\leq^{(\text{Lemma 2})} \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(t_i, t_{-i}) \right) \right] + 2 \left(\frac{3}{2} \delta + \beta n \right) H.
\end{aligned}$$

Whenever $t_i \in \mathcal{T}_i^+$ we can directly argue that:

$$\begin{aligned}
\mathbb{E}_{y_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(t_i), t_{-i}))] &\leq \mathbb{E}_{t_{-i} \sim \tau(\widehat{\mathcal{D}}_{-i})} [t_i(\mathcal{M}(\tau(t_i), t_{-i}))] + \beta n H \\
&= \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(t_i), \tau(t_{-i})))] + \beta n H \\
&= \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(t_i, t_{-i}) \right) \right] + \beta n H.
\end{aligned}$$

Similarly, we get that $\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(w_i), t_{-i}))] \geq \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(w_i, t_{-i}) \right) \right] - 2(\frac{3}{2}\delta + \beta n) H$ for all $w_i \in \mathcal{T}_i$. Combining we get that for $t_i \in \mathcal{T}_i^-, w_i \in \mathcal{T}_i$:

$$\begin{aligned}
\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(t_i, t_{-i}) \right) \right] - \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(w_i, t_{-i}) \right) \right] &\geq \\
\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(t_i), t_{-i}))] - \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(w_i), t_{-i}))] - 4 \left(\frac{3}{2} \delta + \beta n \right) H,
\end{aligned}$$

and for $t_i \in \mathcal{T}_i^+, w_i \in \mathcal{T}_i$ we can get that $\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [t_i(\mathcal{M}(\tau(t_i), t_{-i}))] \geq \mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} \left[t_i \left(\widehat{\mathcal{M}}(t_i, t_{-i}) \right) \right] - \beta n H$.

This concludes the proof of the intermediate bound. To conclude the proof for the BIC guarantee we need to show that:

$$\mathbb{E}_{t_{-i} \sim \widehat{\mathcal{D}}_{-i}} [u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_{-i})] \geq -4H\delta - \varepsilon.$$

By Proposition 5, $u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_{-i}) \in [-H, 3H]$, for all $i \in [n], t_i, w_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$, and hence $u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_{-i}) - u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t'_{-i}) \leq 4H \mathbb{1}\{t_{-i} \neq t'_{-i}\}$. Thus, for any coupling γ of \mathcal{D}_{-i} and $\widehat{\mathcal{D}}_{-i}$, and thus for the optimal coupling γ^* between \mathcal{D}_{-i} and $\widehat{\mathcal{D}}_{-i}$, we get

$$\begin{aligned}
\mathbb{E}_{(t_{-i}, t'_{-i}) \sim \gamma^*} [u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_{-i}) - u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t'_{-i})] &\leq 4H d_{\text{TV}}(\mathcal{D}_{-i}, \widehat{\mathcal{D}}_{-i}) \\
&\leq 4H d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) \\
&\leq 3H \delta.
\end{aligned}$$

Using linearity of expectation and the fact that the chosen coupling maintains the marginals, by re-arranging we have:

$$\begin{aligned} -\mathbb{E}_{t'_i \sim \widehat{\mathcal{D}}_i} [u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t'_i)] &\leq 4H\delta - \mathbb{E}_{t_i \sim \mathcal{D}_i} [u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_i)] \\ &\leq 4H\delta + \varepsilon, \end{aligned}$$

where in the last inequality we used the fact that, since $\tau(t_i) \in \mathcal{T}_i^+$, from the definition of \mathcal{M} , for all $w_i, t_i \in \mathcal{T}_i$, we have $\mathbb{E}_{t_i \sim \mathcal{D}_i} [u_i^{\mathcal{M}}(\tau(t_i) \leftarrow \tau(w_i), t_i)] \geq -\varepsilon$.

We will now prove the revenue guarantee of the lemma. The tensorization property of TV distance [LPW09, Chapter 4] implies that $d_{\text{TV}}(\widehat{\mathcal{D}}, \tau(\widehat{\mathcal{D}})) \leq \beta n$, and thus from the triangle inequality, $d_{\text{TV}}(\mathcal{D}, \tau(\widehat{\mathcal{D}})) \leq \delta + \beta n$. Now notice from triangle inequality that $d_{\text{TV}}(\mathcal{D}, \tau(\widehat{\mathcal{D}})) \leq d_{\text{TV}}(\mathcal{D}, \widehat{\mathcal{D}}) + d_{\text{TV}}(\widehat{\mathcal{D}}, \tau(\widehat{\mathcal{D}}))$. Let $t \sim \mathcal{D}$ and $\widehat{t} \sim \tau(\widehat{\mathcal{D}})$. Since $d_{\text{TV}}(\mathcal{D}, \tau(\widehat{\mathcal{D}})) \leq \beta n + \delta$ there exists a coupling where $t \neq \widehat{t}$ with probability less than $\beta n + \delta$. Whenever $t = \widehat{t}$ the two mechanisms make exactly the same revenue. Whenever they are not, their difference is bounded by V . The desired inequality follows. \square

Lemma 5 is then a simple corollary of Lemma 7.

Proof of Lemma 5. For an (ε, q) -BIC mechanism \mathcal{M} , one can split the type space \mathcal{T}_i of each agent i into two disjoint sets, \mathcal{T}_i^G and \mathcal{T}_i^B , such that when $t_i \in \mathcal{T}_i^G$ agent i ε -maximizes her utility by reporting t_i , and $\Pr_{t_i \sim \mathcal{D}} [t_i \in \mathcal{T}_i^B] \leq q$. Noting that $\perp \in \mathcal{T}_i^G$, the corollary is an immediate implication of Lemma 7. \square

Proof of Theorem 3. The (ε, q) -BIC property is an immediate consequence of Lemma 4.

Applying Lemma 2, with \mathcal{O} as the revenue objective (which is lower bounded by $-V/2$ and upper bounded by $V/2$), and setting $P = \mathcal{D}^p$, $Q = \mathcal{D}$, and $\mathcal{M} = \mathcal{M}_{\mathcal{D}^p}^a$, we have that $\text{Rev}(\mathcal{M}_{\mathcal{D}^p}^a, \mathcal{D}) \geq \text{Rev}(\mathcal{M}_{\mathcal{D}^p}^a, \mathcal{D}^p) - 2V\delta \geq \alpha \text{OPT}(\mathcal{D}^p) - 2V\delta$. Our main goal will be to lower bound $\text{OPT}(\mathcal{D}^p)$.

Let $\mathcal{M}_{\mathcal{D}}^*$ be the revenue optimal mechanism for \mathcal{D} . By Lemma 4, $\mathcal{M}_{\mathcal{D}}^*$ is an ex-post IR and $(\frac{8H\delta}{q}, q)$ -BIC mechanism for \mathcal{D}^p (for all $q \in [0, 1]$). Therefore, Lemma 5 implies that there exists a mechanism $\widehat{\mathcal{M}}$ that is ex-post IR and $O(\frac{H\delta}{q} + nqH)$ -BIC with respect to \mathcal{D}^p , such that $\text{Rev}(\widehat{\mathcal{M}}, \mathcal{D}^p) \geq \text{Rev}(\mathcal{M}_{\mathcal{D}}^*, \mathcal{D}^p) - nqV$.

Next, we apply the ε -BIC to BIC reduction of [COVZ21], on the mechanism $\widehat{\mathcal{M}}$. Specifically, we use the following lemma.

Lemma 8 ([DW12], [RW18], [COVZ21]). *In any n agent setting where the valuations of agents are bounded by H , for any mechanism \mathcal{M} with payments in $[-H, H]$, that is ex-post IR and ε -BIC with respect to some product distribution \mathcal{D} , there exists a mechanism \mathcal{M}' with payments in $[-H, H]$,¹ that is ex-post IR and BIC with respect to \mathcal{D} , such that, assuming truthful bidding $\text{Rev}(\mathcal{M}', \mathcal{D}) \geq \text{Rev}(\mathcal{M}, \mathcal{D}) - O(n\sqrt{H\varepsilon})$.*

So, Lemma 8 implies that there exists a mechanism \mathcal{M}' that is ex-post IR and BIC with respect to \mathcal{D}^p such that $\text{Rev}(\mathcal{M}', \mathcal{D}^p) \geq \text{Rev}(\widehat{\mathcal{M}}, \mathcal{D}^p) - O(n\sqrt{H(\frac{H\delta}{q} + nqH)})$. Combining all the ingredients

¹In the reduction payments are only scaled by a value less than 1. Thus if \mathcal{M} had payments in $[-H, H]$, then \mathcal{M}' also has payments in that range.

so far, we have

$$\begin{aligned}
\text{Rev}(\mathcal{M}_{\mathcal{D}^p}^a, \mathcal{D}) &\geq \text{Rev}(\mathcal{M}_{\mathcal{D}^p}^a, \mathcal{D}^p) - V\delta \\
&\geq \alpha \text{OPT}(\mathcal{D}^p) - V\delta \\
&\geq \alpha \text{Rev}(\mathcal{M}', \mathcal{D}^p) - V\delta \\
&\geq \alpha \text{Rev}(\widehat{\mathcal{M}}, \mathcal{D}^p) - O\left(\alpha n \sqrt{H\left(\frac{H\delta}{q} + nqH\right)} + V\delta\right) \\
&\geq \alpha \text{Rev}(\mathcal{M}_{\mathcal{D}}^*, \mathcal{D}^p) - O\left(\alpha n \sqrt{H\left(\frac{H\delta}{q} + nqH\right)} + V(\delta + \alpha nq)\right) \\
&= \alpha \text{Rev}(\mathcal{M}_{\mathcal{D}}^*, \mathcal{D}^p) - O\left(\alpha n H \sqrt{\frac{\delta}{q} + nq} + V(\delta + \alpha nq)\right)
\end{aligned}$$

Applying Lemma 2 again, with $P = \mathcal{D}$, $Q = \mathcal{D}^p$, and $\mathcal{M} = \mathcal{M}_{\mathcal{D}}^*$ we have $\text{Rev}(\mathcal{M}_{\mathcal{D}}^*, \mathcal{D}^p) \geq \text{OPT}(\mathcal{D}) - V\delta$. Combining with the previous inequality, we have $\text{Rev}(\mathcal{M}_{\mathcal{D}^p}^a, \mathcal{D}) \geq \alpha \text{OPT}(\mathcal{D}) - O\left(\alpha n H \sqrt{\frac{\delta}{q} + nq} + \alpha nqV + (1 + \alpha)V\delta\right)$. Picking $q = \sqrt{\delta/n}$, and noting that $V \leq 2nH$, we have: $\text{Rev}(\mathcal{M}_{\mathcal{D}^p}^a, \mathcal{D}) \geq \alpha \text{OPT}(\mathcal{D}) - O\left(\alpha V(n\delta)^{1/4} + \alpha V(n\delta)^{1/2} + (1 + \alpha)V\delta\right) \geq \alpha \text{OPT}(\mathcal{D}) - O\left((1 + \alpha)V \sqrt{n\sqrt{\delta}}\right)$. \square

Proof of Proposition 1. The marginal distributions for \mathcal{D}^p and \mathcal{D} are close in total variation distance, and specifically, $d_{\text{TV}}(\widehat{\mathcal{D}}_i, \mathcal{D}_i^p) \leq d_{\text{TV}}(\widehat{\mathcal{D}}, \mathcal{D}^p) \leq \varepsilon$. Therefore, $d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}_i^p) \leq \varepsilon$, which implies that $d_{\text{TV}}(\mathcal{D}, \mathcal{D}^p) \leq n\varepsilon$. Applying the triangle inequality completes the proof. \square

D Proofs missing from Section 4.2

Proof of Theorem 4. In order to prove this theorem we will first need to prove two intermediate lemmas. Recall that $\Pi(\mathcal{D}_1, \dots, \mathcal{D}_n) = \{\mathcal{D}' | \Pr_{t_i \sim \mathcal{D}_i}[t_i = v_i] = \sum_{v_{-i} \in \mathcal{T}_{-i}} \Pr_{t \sim \mathcal{D}'}[t = (v_i, v_{-i})], \forall i \in [n], \forall t_i \in \mathcal{T}_i\}$.

Lemma 9. *For any distribution $\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)$ there exists a distribution $\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)$ such that $d_{\text{TV}}(\mathcal{D}, \mathcal{D}') \leq n\varepsilon$, where for all i , $d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}'_i) \leq \varepsilon$.*

Proof. We will prove an intermediate step that will then immediately yield the desired outcomes. More precisely we will first show that for any distribution $\mathcal{D}^{(i-1)} \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_{i-1}, \mathcal{D}_i, \dots, \mathcal{D}_n)$ there exists a distribution $\mathcal{D}^{(i)} \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_{i-1}, \mathcal{D}'_i, \dots, \mathcal{D}_n)$ such that $d_{\text{TV}}(\mathcal{D}^{(i-1)}, \mathcal{D}^{(i)}) \leq \varepsilon$, where $d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}'_i) \leq \varepsilon$. To prove this we will leverage the \mathcal{L}^1 -distance characterization of TV distance.

Our proof will be constructive through a simple ‘‘moving mass’’ argument. For simplicity let’s assume that there exist $v_i, v'_i \in \mathcal{T}_i$ such that $\Pr_{t_i \sim \mathcal{D}_i}[t_i = v_i] = \Pr_{t'_i \sim \mathcal{D}'_i}[t'_i = v_i] + \varepsilon$ and $\Pr_{t_i \sim \mathcal{D}_i}[t_i = v'_i] = \Pr_{t'_i \sim \mathcal{D}'_i}[t'_i = v'_i] - \varepsilon$. Extending the following procedure for arbitrary $\mathcal{D}_i, \mathcal{D}'_i$ such that $d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}'_i) \leq \varepsilon$ will be immediate. Given $\mathcal{D}^{(i-1)}$, construct $\mathcal{D}^{(i)}$ as follows:

1. Set $\varepsilon_{\text{cur}} = \varepsilon$ and $\mathcal{D}^{(i-1)} = \mathcal{D}^{(i)}$.
2. As long as $\varepsilon_{\text{cur}} > 0$ do the following process:
 - (a) Find $v_{-i} \in \mathcal{T}_{-i}$ such that $\Pr_{t' \sim \mathcal{D}^{(i)}}[t' = (v_i, v_{-i})] > 0$ and let γ be the minimum of $\Pr_{t' \sim \mathcal{D}^{(i)}}[t' = (v_i, v_{-i})]$ and ε_{cur} .
 - (b) Change $\mathcal{D}^{(i)}$ such that $\Pr_{t' \sim \mathcal{D}^{(i)}}[t' = (v_i, v_{-i})] - \gamma$ and $\Pr_{t' \sim \mathcal{D}^{(i)}}[t' = (v'_i, v_{-i})] + \gamma$.
 - (c) Set $\varepsilon_{\text{cur}} = \varepsilon_{\text{cur}} - \gamma$
3. Output $\mathcal{D}^{(i)}$

From our construction of $\mathcal{D}^{(i)}$ it is immediate that $\mathcal{D}^{(i)} \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_{i-1}, \mathcal{D}'_i, \dots, \mathcal{D}'_n)$ and $d_{\text{TV}}(\mathcal{D}^{(i-1)}, \mathcal{D}^{(i)}) \leq \varepsilon$. Chaining up the resulting inequalities and using triangle inequality concludes the proof. \square

Leveraging the above we can prove the following:

Lemma 10. *For any mechanism \mathcal{M} and sets of marginals $(\mathcal{D}_1, \dots, \mathcal{D}_n)$ and $(\mathcal{D}'_1, \dots, \mathcal{D}'_n)$ such that for all $i \in [n]$, $d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}'_i) \leq \varepsilon$ we have that:*

$$\min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}(t))] \geq \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t' \sim \mathcal{D}'} [\mathcal{O}(t', \mathcal{M}(t'))] - n\varepsilon V$$

Proof. We will prove this using a contradiction. Assume that

$$\min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}(t))] < \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t' \sim \mathcal{D}'} [\mathcal{O}(t', \mathcal{M}(t'))] - n\varepsilon V.$$

Lets call $\mathcal{D}^* = \arg \min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}(t))]$. Now using Lemma 9 we have that there exists $\widehat{\mathcal{D}}^* \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)$ such that $d_{\text{TV}}(\mathcal{D}^*, \widehat{\mathcal{D}}^*) \leq n\varepsilon$. Using Lemma 2 we have that $\mathbb{E}_{t \sim \mathcal{D}^*} [\mathcal{O}(t, \mathcal{M}(t))] \geq \mathbb{E}_{t \sim \widehat{\mathcal{D}}^*} [\mathcal{O}(t, \mathcal{M}(t))] - n\varepsilon V$. Chaining the above inequalities we get that:

$$\mathbb{E}_{t \sim \widehat{\mathcal{D}}^*} [\mathcal{O}(t, \mathcal{M}(t))] - n\varepsilon V \leq \mathbb{E}_{t \sim \mathcal{D}^*} [\mathcal{O}(t, \mathcal{M}(t))] < \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t' \sim \mathcal{D}'} [\mathcal{O}(t', \mathcal{M}(t'))] - n\varepsilon V$$

However, $\min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t' \sim \mathcal{D}'} [\mathcal{O}(t', \mathcal{M}(t'))] - n\varepsilon V \leq \mathbb{E}_{t \sim \widehat{\mathcal{D}}^*} [\mathcal{O}(t, \mathcal{M}(t))] - n\varepsilon V$ which concludes the contradiction. \square

Now we have all the components to prove the main theorem.

First by using Lemma 10 on \mathcal{M}^α we have that $\min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'} [\mathcal{O}(t, \mathcal{M}^\alpha(t))] \geq \min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}^\alpha(t))] - n\varepsilon V$.

Now lets call $\mathcal{M}^* = \arg \max_{\mathcal{M}'} \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'} [\mathcal{O}(t, \mathcal{M}'(t))]$. By applying Lemma 10 on \mathcal{M}^* we have that $\min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}^*(t))] \geq \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'} [\mathcal{O}(t, \mathcal{M}^*(t))]$. Chaining all of the above we have that:

$$\begin{aligned} \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'} [\mathcal{O}(t, \mathcal{M}^\alpha(t))] &\geq \min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}^\alpha(t))] - n\varepsilon V \\ &\geq \alpha \max_{\mathcal{M}'} \min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}'(t))] - n\varepsilon V \\ &\geq \alpha \min_{\mathcal{D} \in \Pi(\mathcal{D}_1, \dots, \mathcal{D}_n)} \mathbb{E}_{t \sim \mathcal{D}} [\mathcal{O}(t, \mathcal{M}^*(t))] - n\varepsilon V \\ &\geq \alpha \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'} [\mathcal{O}(t, \mathcal{M}^*(t))] - (1 + \alpha)n\varepsilon V \\ &= \alpha \max_{\mathcal{M}'} \min_{\mathcal{D}' \in \Pi(\mathcal{D}'_1, \dots, \mathcal{D}'_n)} \mathbb{E}_{t \sim \mathcal{D}'} [\mathcal{O}(t, \mathcal{M}'(t))] - (1 + \alpha)n\varepsilon V. \end{aligned}$$

\square

E Proofs missing from Section 4.4

Proof of Proposition 2. Let $S_{\mathcal{D}}$ be the mechanism that implements the better of bundling and selling separately, as computed on a prior \mathcal{D} . $S_{\mathcal{D}^p}$ is a DISC and ex-post IR mechanism, and $\text{Rev}(S_{\mathcal{D}^p}, \mathcal{D}^p) \geq \frac{1}{6} \text{Rev}(\mathcal{D}^p)$. Thus, applying Theorem 1 we have that $\text{Rev}(S_{\mathcal{D}^p}, \mathcal{D}) \geq \frac{1}{6} \text{Rev}(\mathcal{D}) - \frac{7}{6} H\delta$. The mechanism $S_{\mathcal{D}^p}$ is either selling each item separately, or it is setting a posted price for the grand bundle. If the former case occurs, then running $S_{\mathcal{D}^p}$ on \mathcal{D} makes (weakly) less revenue than $S_{\text{Rev}}(\mathcal{D})$; if the latter case occurs, running $S_{\mathcal{D}^p}$ on \mathcal{D} makes (weakly) less revenue than $B\text{Rev}(\mathcal{D})$. Therefore, we overall have that $\text{Rev}(S_{\mathcal{D}}, \mathcal{D}) \geq \text{Rev}(S_{\mathcal{D}^p}, \mathcal{D})$. Combining with the previous inequality we get $\text{Rev}(S_{\mathcal{D}}, \mathcal{D}) \geq \frac{1}{6} \text{Rev}(\mathcal{D}) - \frac{7}{6} H\delta$. \square

MRFs. We state some basic definitions for Markov Random Fields.

Definition 5 (Markov Random Field [SK75],[KS80],[CO21]). A Markov Random Field (MRF) is defined by a hypergraph $G = (V, E)$. Associated with every vertex $v \in V$ is a random variable X_v taking values in some alphabet Σ_v , as well as a potential function $\psi_v : \Sigma_v \rightarrow \mathbb{R}$. Associated with every hyperedge $e \subseteq E$ is a potential function $\psi_e : \Sigma_e \rightarrow \mathbb{R}$. In terms of these potentials, we define a probability distribution \mathcal{D} associating to each vector $\mathbf{c} \in \times_{v \in V} \Sigma_v$ probability $\mathcal{D}(\mathbf{c})$ satisfying: $\mathcal{D}(\mathbf{c}) \propto \prod_{v \in V} e^{\psi_v(c_v)} \prod_{e \in E} e^{\psi_e(c_e)}$, where Σ_e denotes $\times_{v \in e} \Sigma_v$ and \mathbf{c}_e denotes $\{c_v\}_{v \in e}$.

Definition 6 ([CO21]). Given a random variable/type \mathbf{t} generated by an MRF over a hypergraph $G = ([m], E)$, we define **weighted degree** of item i as: $d_i := \max_{x \in \mathcal{T}} |\sum_{e \in E: i \in e} \psi_e(x_e)|$ and the **maximum weighted degree** as $\Delta := \max_{i \in [m]} d_i$.

Lemma 11 (Lemma 2[CO21]). Let random variable t be generated by an MRF. For any i and any set $\mathcal{E} \subseteq \mathcal{T}_i$ and set $\mathcal{E}' \subseteq \mathcal{T}_{-i}$:

$$\exp(-4\Delta) \leq \frac{\Pr_{t \sim \mathcal{D}} [t_i \in \mathcal{E} \wedge t_{-i} \in \mathcal{E}']}{\Pr_{t_i \sim \mathcal{D}_i} [t_i \in \mathcal{E}] \Pr_{t_{-i} \sim \mathcal{D}_{-i}} [t_{-i} \in \mathcal{E}']} \leq \exp(4\Delta)$$

Proof of Proposition 3. Let us consider the case where $m = 2$. For the distribution \mathcal{D} that is produced by an MRF, let \mathcal{D}_1 and \mathcal{D}_2 be the marginal distributions. Using Lemma 11 we get that for any profile $t = (t_1, t_2)$ we have that $-4\Delta \leq \log(\frac{\mathcal{D}(t)}{\mathcal{D}_1(t_1)\mathcal{D}_2(t_2)}) \leq 4\Delta$. Therefore, the Kullback-Leibler (KL) divergence is $D_{KL}(\mathcal{D}_1\mathcal{D}_2||\mathcal{D}) = \sum_{t=(t_1, t_2) \in \mathcal{T}} \mathcal{D}_1(t)\mathcal{D}_2(t) \log(\frac{\mathcal{D}_1(t_1)\mathcal{D}_2(t_2)}{\mathcal{D}(t)}) \leq 4\Delta$.

Now consider the case where $\mathcal{T} = \{0, 1\}^2$. Let $\mathcal{D} = [p, 0.5 - p; 0.5 - p, p]$ which is a valid MRF and has marginals $\mathcal{D}_1 = \mathcal{D}_2 = [0.5, 0.5]$. Then, we have $d_{TV}(\mathcal{D}, \mathcal{D}_1\mathcal{D}_2) = 0.5 - 2p \leq 0.5$. However, when p approaches 0 we have $\lim_{p \rightarrow 0} D_{KL}(\mathcal{D}_1\mathcal{D}_2||\mathcal{D}) = \lim_{p \rightarrow 0} [0.5(\log(1/4p) + \log(0.25/(0.5 - p)))] = \infty$, and hence, $\Delta \rightarrow +\infty$.

Therefore, there exists a sequence of MRF distributions D_k with associated product-of-marginal distributions D_k^p and max-weighted degrees Δ_k such that $\sup_k d_{TV}(D_k||D_k^p) \leq 0.5$ and $\lim_{k \rightarrow \infty} \Delta_k = +\infty$. □

Proof of Proposition 4. As a first step, we are going to bound the Kullback-Leibler (KL) divergence between the distribution \mathcal{D} and a product distribution \mathcal{D}^p . Then we are going to use Pinsker's inequality [Tsy08] and the Bretagnolle-Huber inequality [Tsy08, BH78] to bound the TV distance using KL divergence.

Let $\mathcal{D}(t) = \frac{1}{Z_1} \prod_{v \in V} e^{\psi_v(t_v)} \prod_{e \in E} e^{\psi_e(t_e)}$, where Z_1 is the partition function. Let \mathcal{D}^p be product distribution such that $\mathcal{D}^p(t) = \frac{1}{Z_2} \prod_{v \in V} e^{\psi_v(t_v)}$, where Z_2 is the partition function.

The KL divergence is between \mathcal{D} and \mathcal{D}^p is:

$$\begin{aligned} D_{KL}(\mathcal{D}||\mathcal{D}^p) &= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \log \frac{\mathcal{D}(t)}{\mathcal{D}^p(t)} \\ &= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \log \frac{Z_2 \prod_{v \in V} e^{\psi_v(t_v)} \prod_{e \in E} e^{\psi_e(t_e)}}{Z_1 \prod_{v \in V} e^{\psi_v(t_v)}} \\ &= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \log \frac{Z_2}{Z_1} \prod_{e \in E} e^{\psi_e(t_e)} \\ &= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \left(\log \frac{Z_2}{Z_1} + \sum_{e \in E} \psi_e(t_e) \right) \\ &\leq \sum_{t \in \mathcal{T}} \mathcal{D}(t) \left(\log \frac{Z_2}{Z_1} + \frac{m}{2} \Delta \right) \\ &= \frac{m}{2} \Delta + \log \frac{Z_2}{Z_1} \end{aligned}$$

Since KL divergence is not symmetric, we can also compute: $D_{KL}(\mathcal{D}^p||\mathcal{D})$:

$$\begin{aligned}
D_{KL}(\mathcal{D}^p||\mathcal{D}) &= \sum_{t \in \mathcal{T}} \mathcal{D}^p(t) \log \frac{\mathcal{D}^p(t)}{\mathcal{D}(t)} \\
&= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \log \frac{Z_1 \prod_{v \in V} e^{\psi_v(t_v)}}{Z_2 \prod_{v \in V} e^{\psi_v(t_v)} \prod_{e \in E} e^{\psi_e(t_e)}} \\
&= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \log \frac{Z_1}{Z_2} \prod_{e \in E} e^{-\psi_e(t_e)} \\
&= \sum_{t \in \mathcal{T}} \mathcal{D}(t) \left(\log \frac{Z_1}{Z_2} - \sum_{e \in E} \psi_e(t_e) \right) \\
&\leq \sum_{t \in \mathcal{T}} \mathcal{D}(t) \left(\log \frac{Z_1}{Z_2} + \frac{m}{2} \Delta \right) \\
&= \frac{m}{2} \Delta - \log \frac{Z_2}{Z_1}
\end{aligned}$$

We can get that $\sum_{e \in E} \psi_e(t_e) \in [-\frac{m}{2} \Delta, \frac{m}{2} \Delta]$ as follows. $\sum_e \psi_e(t_e) = \frac{1}{2} \sum_{i \in [m]} \sum_{e \in E: i \in e} \psi_e(t_e) \leq \frac{1}{2} \sum_{i \in [m]} d_i \leq \frac{m \Delta}{2}$. Similarly, we can lower bound $\sum_{e \in E} \psi_e(t_e) \geq -\frac{m \Delta}{2}$ since the definition of d_i is $d_i := \max_{x \in \mathcal{T}} |\sum_{e \in E: i \in e} \psi_e(x_e)|$.

From the above inequalities we have that $\min\{D_{KL}(\mathcal{D}^p||\mathcal{D}), D_{KL}(\mathcal{D}||\mathcal{D}^p)\} \leq \frac{m}{2} \Delta$. From Pinsker's inequality we get $d_{TV}(\mathcal{D}, \mathcal{D}^p) \leq \sqrt{\frac{m \Delta}{4}}$, and from the Bretagnolle-Huber inequality we get $d_{TV}(\mathcal{D}, \mathcal{D}^p) \leq \sqrt{1 - \exp(-m \Delta / 2)}$. Combining these inequalities we have the desired bound on the TV distance. \square