

Bernoulli Factories and Black-box Reductions in Mechanism Design

SHADDIN DUGHMI, Department of Computer Science, University of Southern California, USA
 JASON HARTLINE, Computer Science Department, Northwestern University, USA
 ROBERT D. KLEINBERG, Department of Computer Science, Cornell University, USA
 RAD NIAZADEH, Chicago Booth School of Business, University of Chicago, USA

We provide a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design for welfare maximization problems. Unlike prior results, our reduction achieves exact incentive compatibility for problems with multi-dimensional and continuous type spaces. The key technical barrier preventing exact incentive compatibility in prior black-box reductions is that repairing violations of incentive constraints requires understanding the distribution of the mechanism's output, which is typically #P-hard to compute. Reductions that instead estimate the output distribution by sampling inevitably suffer from sampling error, which typically precludes exact incentive compatibility. We overcome this barrier by employing and generalizing the computational model in the literature on *Bernoulli Factories*. In a Bernoulli factory problem, one is given a function mapping the bias of an "input coin" to that of an "output coin," and the challenge is to efficiently simulate the output coin given only sample access to the input coin. This is the key ingredient in designing an incentive compatible mechanism for bipartite matching, which can be used to make the approximately incentive compatible reduction of Hartline et al. [18] exactly incentive compatible.

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

Additional Key Words and Phrases: Bayesian mechanism design, black-box reductions, bayesian incentive compatible (BIC) mechanisms, bernoulli factories

ACM Reference format:

Shaddin Dughmi, Jason Hartline, Robert D. Kleinberg, and Rad Niazadeh. 2020. Bernoulli Factories and Black-box Reductions in Mechanism Design. *J. ACM* 68, 2, Article 10 (January 2021), 30 pages.

<https://doi.org/10.1145/3440988>

The first author was supported by NSF CAREER Award CCF-1350900. The second author was supported in part by NSF grant CCF 1618502. The third author was partially supported by NSF grant CCF-1512964.

Authors' addresses: S. Dughmi, Department of Computer Science, University of Southern California, SAL 234, 941 Bloom Walk, Los Angeles, CA 90089; email: shaddin@usc.edu; J. Hartline, Computer Science Department, Northwestern University, Seeley Mudd 3015, 2233 Tech Drive, IL 60208; email: hartline@eecs.northwestern.edu; R. D. Kleinberg, Department of Computer Science, Cornell University, Ithaca, 317 Gates Hall, NY 14853; email: rdk@cs.cornell.edu; R. Niazadeh, Chicago Booth School of Business, University of Chicago, HC-303, 5807 South Woodlawn Ave., Chicago, IL 60637; email: rad.niazadeh@chicagobooth.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0004-5411/2020/01-ART10 \$15.00

<https://doi.org/10.1145/3440988>

1 INTRODUCTION

We resolve an open question from Hartline et al. [17, 18], which is considered as one of the fundamental algorithmic questions in the Bayesian mechanism design: *There is a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design for welfare maximization problems.* A Bayesian algorithm is one that performs well in expectation when the input is drawn from a known distribution. By polynomial time, we mean polynomial in the number of agents and the combined “size” of their type spaces. The key distinction between our result and those of Hartline et al. [17, 18] is that both (a) the agents’ preferences can be multi-dimensional and from a continuous space (rather than single-dimensional or from a discrete space), and (b) the resulting mechanism is exactly Bayesian incentive compatible (rather than approximately Bayesian incentive compatible).

A mechanism solicits preferences from agents, i.e., how much each agent prefers each outcome, and then chooses an outcome. *Incentive compatibility* of a mechanism requires that, though agents could misreport their preferences, it is not in any agent’s best interest to do so. A quintessential research problem at the intersection of mechanism design and approximation algorithms is to identify black-box reductions from approximation mechanism design to approximation algorithm design. The key algorithmic property that makes a mechanism incentive compatible is that, from any individual agent’s perspective, it must be *maximal-in-range*, specifically, the outcome selected maximizes the agent’s utility less some cost that is a function of the outcome (e.g., this cost function can depend on other agents’ reported preferences).¹

The black-box reductions from Bayesian mechanism design to Bayesian algorithm design in the literature are based on obtaining an understanding of the distribution of outcomes produced by the algorithm through simulating the algorithm on samples from agents’ preferences. Notice that, even for structurally simple problems, calculating the exact probability that a given outcome is selected by an algorithm can be #P-hard. For example, Hartline et al. [18] show such a result for calculating the probability that a matching in a bipartite graph is optimal for a simple explicitly given distribution of edge weights. A black-box reduction for mechanism design must therefore produce exactly maximal-in-range outcomes merely from samples. This challenge motivates new questions for algorithm design from samples.

The Expectations from Samples Model. In traditional algorithm design, the inputs are specified to the algorithm exactly. In this article, we formulate the *expectations from samples* model. This model calls for drawing an outcome from a distribution that is a precise function of the expectations of some random sources that are given only by sample access. Formally, a problem for this model is described by a function $f : [0, 1]^n \rightarrow \Delta(X)$ where X is an abstract set of feasible outcomes and $\Delta(X)$ is the family of probability distributions over X . For any n input distributions on support $[0, 1]$ with unknown expectations $\mu = (\mu_1, \dots, \mu_n)$, an algorithm for such a problem, with only sample access to each of the n input distributions, must produce sample outcome from X that is distributed exactly according to $f(\mu_1, \dots, \mu_n)$.

Producing an outcome that is approximately drawn according to the desired distribution can typically be done from estimates of the expectations formed from sample averages (a.k.a., Monte Carlo sampling). However, exact implementation of many natural functions f is either impossible for information theoretic reasons or requires sophisticated techniques. Impossibility generally follows, for example, when f is discontinuous. The literature on *Bernoulli Factories* (e.g., Keane and O’Brien [23]), which inspires our generalization to the expectations from samples model and

¹In general, one can think of maximal-in-range for all agents, meaning that the outcome selected maximizes agents’ social welfare less some cost that is a function of the outcome among a particular range of feasible outcomes.

provides some of the basic building blocks for our results, considers the special case where the input distribution and output distribution are both Bernoullis (i.e., supported on $\{0, 1\}$).

We propose and solve two fundamental problems for the expectations from samples model. The first problem considers the biases $\mathbf{p} = (p_1, \dots, p_m)$ of m Bernoulli random variables as the marginal probabilities of a distribution on $\{1, \dots, m\}$ (i.e., \mathbf{p} satisfies $\sum_i p_i = 1$) and asks to sample from this distribution. We develop an algorithm that we call the Bernoulli Race to solve this problem.

The second problem corresponds to the “soft maximum” problem given by a regularizer that is a multiple $1/\lambda$ of the Shannon entropy function $H(\mathbf{p}) = -\sum_i p_i \log p_i$. The marginal probabilities on outcomes that maximize the expected value of the distribution over outcomes less the cost of the negative entropy regularizer are given by exponential weights, i.e., the function outputs i with probability proportional to $e^{\lambda p_i}$ (this is a standard relationship that has, for example, been employed in previous work in mechanism design, e.g., Huang and Kannan [20]). A straightforward exponentiation and then reduction to the Bernoulli Race above does not have polynomial sample complexity. We develop an algorithm that we call the Fast Exponential Bernoulli Race to solve this problem.

Black-box Reductions in Mechanism Design. A special case of the problem that we must solve to apply the standard approach to black-box reductions is the *single-agent multiple-urns problem*. In this setting, a single agent faces a set X of urns, and each urn contains a random object whose distribution is unknown, but can be sampled. The agent’s type determines his utility for each object; fixing this type, urn i is associated with a random real-valued reward with unknown expectation μ_i . Our goal is to allocate the agent his favorite urn or close to it.

As described above, incentive compatibility requires an algorithm for selecting a high-value urn that is maximal-in-range. If we could exactly calculate the expected values μ_1, \dots, μ_n from the agent’s type, this problem is trivial both algorithmically and from a mechanism design perspective: Simply solicit the agent’s type t then allocate him the urn with the maximum $\mu_i = \mu_i(t)$. As described above, with only sample access to the expected values of each urn, we cannot implement the exact maximum. Our solution is to apply the Fast Exponential Bernoulli Race as a solution to the regularized maximization problem in the expectations from samples model. This algorithm—with only sample access to the agent’s values for each urn—will assign the agent to a random urn with a high expected value and is maximal-in-range.

The multi-agent reduction from Bayesian mechanism design to Bayesian algorithm design of Hartline et al. [17, 18] is based on solving a matching problem between multiple agents and outcomes, where an agent’s value for an outcome is the expectation of a random variable that can be accessed only through sampling. We should also assert that Bei and Huang [7] independently discovered a similar reduction based on solving a fractional assignment problem. Their reduction applies to finite, discrete type spaces and is approximately Bayesian incentive compatible. Specifically, this problem generalizes the above-described single-agent multiple-urns problem to the problem of matching agents to urns with the goal of approximately maximizing the total weight of the matching (the social welfare). Again, for incentive compatibility, we require this expectations from samples algorithm to be maximal-in-range from each agent’s perspective. Using methods from Agrawal and Devanur’s [1] work on stochastic online convex optimization, we reduce this matching problem to the single-agent multiple-urns problem.

As stated in the opening paragraph, our main result—obtained through the approach outlined above—is a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design. The analysis assumes that agents’ values are normalized to the $[0, 1]$ interval and gives additive loss in the welfare. The reduction is an approximation scheme and the dependence of the runtime on the additive loss is inverse polynomial. The reduction depends

polynomially on a suitable notion of the size of the space of agent preferences. For example, applied to environments where agents have preferences that lie in high-dimensional spaces, the runtime of the reduction depends polynomially on the number of points necessary to approximately cover each agent's space of preferences. More generally, the bounds we obtain are polynomial in the bounds of Hartline et al. [17, 18] but the resulting mechanism, unlike in the proceeding work, is exactly Bayesian incentive compatible.

Organization. The organization of the article separates the development of the expectations from samples model and its application to black-box reductions in Bayesian mechanism design. Section 2 introduces Bernoulli factories and reviews basic results from the literature. Section 3 defines two central problems in the expectations from samples model, sampling from outcomes with linear weights and sampling from outcomes with exponential weights, and gives algorithms for solving them. We return to mechanism design problems in Section 4 and solve the single-agent multiple urns problem. In Section 5, we give our main result, the reduction from Bayesian mechanism design to Bayesian algorithm design.

2 BASICS OF BERNOULLI FACTORIES

We use the terms *Bernoulli* and *coin* to refer to distributions over $\{0, 1\}$ and $\{\text{heads}, \text{tails}\}$ interchangeably. The Bernoulli factory problem is about generating new coins from old ones.

Definition 2.1 (23). Given function $f : (0, 1) \rightarrow (0, 1)$, the *Bernoulli factory* problem is to output a sample of a Bernoulli variable with bias $f(p)$ (i.e., an $f(p)$ -coin), given black-box access to independent samples of a Bernoulli distribution with bias $p \in (0, 1)$ (i.e. a p -coin).

To illustrate the Bernoulli factory model, consider the examples of $f(p) = p^2$ and $f(p) = e^{p-1}$. For the former one, it is enough to flip the p -coin twice and output 1 if both flips are 1, and 0 otherwise. For the latter one, the Bernoulli factory is still simple but more interesting: Draw K from the Poisson distribution with parameter $\lambda = 1$ (remind that the Poisson distribution with parameter λ has probability of $K = k$ as $\lambda^k e^{-\lambda}/k!$), flip the p -coin K times, and output 1 if all coin flips were 1, and 0 otherwise (see below).

The question of characterizing functions f for which there is an algorithm from sampling $f(p)$ -coins from p -coins has been the main subject of interest in this literature (e.g., Keane and O'Brien [23], Nacu and Peres [28]). In particular, Keane and O'Brien [23] provide necessary and sufficient conditions for f , under which an algorithm for the Bernoulli factory exists. Moreover, Nacu and Peres [28] suggest an algorithm for simulating an $f(p)$ -coin based on polynomial envelopes of f . The canonical challenging problem of Bernoulli factories—and a primitive in the construction of more general Bernoulli factories—is the *Bernoulli Doubling* problem: $f(p) = 2p$ for $p \in (0, 1/2)$. See Łatuszyński [25] for a survey on this topic.

Questions in Bernoulli factories can be generalized to multiple input coins. Given $f : (0, 1)^m \rightarrow (0, 1)$, the goal is sample from a Bernoulli with bias $f(p_1, \dots, p_m)$ given sample access to m independent Bernoulli variables with unknown biases $\mathbf{p} = (p_1, \dots, p_m)$. Linear functions f were studied and solved by Huber [21]. For example, the special case $m = 2$ and $f(p_1, p_2) = p_1 + p_2$, a.k.a., *Bernoulli Addition*, can be solved by reduction to the Bernoulli Doubling problem (formalized below).

Questions in Bernoulli factories can be generalized to allow input distributions over real numbers on the unit interval $[0, 1]$ (rather than Bernoullis over $\{0, 1\}$). In this generalization the question is to produce a Bernoulli with bias $f(\mu)$ with sample access to draws from a distribution supported on $[0, 1]$ with expectation μ . These problems can be easily solved by reduction to the Bernoulli factory problem:

0. *Continuous to Bernoulli*: Can implement Bernoulli with bias μ with one sample from distribution \mathcal{D} with expectation μ . Algorithm:
 - Draw $Z \sim \mathcal{D}$ and $P \sim \text{Bern}[Z]$.
 - Output P .

Below are enumerated the important building blocks for Bernoulli factories.

- (1) *Bernoulli Down Scaling*: Can implement $f(p) = \lambda \cdot p$ for $\lambda \in [0, 1]$ with one sample from $\text{Bern}[p]$. Algorithm:
 - Draw $\Lambda \sim \text{Bern}[\lambda]$ and $P \sim \text{Bern}[p]$.
 - Output $\Lambda \cdot P$ (i.e., 1 if both coins are 1, otherwise 0).
- (2) *Bernoulli Doubling*: Can implement $f(p) = 2p$ for $p \in (0, 1/2 - \delta]$ with $O(1/\delta)$ samples from $\text{Bern}[p]$ in expectation. The algorithm is complicated; see Nacu and Peres [28].
- (3) *Bernoulli Probability Generating Function*: Can implement $f(p) = \mathbf{E}_{k \sim \mathcal{D}}[p^k]$ for distribution \mathcal{D} over non-negative integers with $\mathbf{E}_{K \sim \mathcal{D}}[K]$ samples from $\text{Bern}[p]$ in expectation. Algorithm:
 - Draw $K \sim \mathcal{D}$ and $P_1, \dots, P_K \sim \text{Bern}[p]$ (i.e., K samples).
 - Output $\prod_i P_i$ (i.e., 1 if all K coins are 1, otherwise 0).
- (4) *Bernoulli Exponentiation*: Can implement $f(p) = \exp(\lambda(p - 1))$ for $p \in [0, 1]$ and non-negative constant λ with λ samples from $\text{Bern}[p]$ in expectation. Algorithm: Apply the Bernoulli Probability Generating Function algorithm for the Poisson distribution with parameter λ .
- (5) *Bernoulli Averaging*: Can implement $f(p_1, p_2) = (p_1 + p_2)/2$ with one sample from $\text{Bern}[p_1]$ or $\text{Bern}[p_2]$. Algorithm:
 - Draw $Z \sim \text{Bern}[1/2]$, $P_1 \sim \text{Bern}[p_1]$, and $P_2 \sim \text{Bern}[p_2]$.
 - Output P_{Z+1} .
- (6) *Bernoulli Addition*: Can implement $f(p_1, p_2) = p_1 + p_2$ for $p_1 + p_2 \in [0, 1 - \delta]$ with $O(1/\delta)$ samples from $\text{Bern}[p_1]$ and $\text{Bern}[p_2]$ in expectation. Algorithm: Apply Bernoulli Doubling to Bernoulli Averaging.

It may seem counterintuitive that Bernoulli Doubling is much more challenging than Bernoulli Down Scaling. Notice, however, that for a coin with bias $p = 1/2$, Bernoulli Doubling with a finite number of coin flips is impossible. The doubled coin must be deterministically heads, while any finite sequence of coin flips of $\text{Bern}[1/2]$ has non-zero probability of occurring. However, a coin with probability $p = 1/2 - \delta$ for some small δ has a similar probability of each sequence, but Bernoulli Doubling must sometimes output tails. Thus, Bernoulli Doubling must require a number of coin flips that goes to infinity as δ goes to zero.

3 THE EXPECTATIONS FROM SAMPLES MODEL

The expectations from samples model is a combinatorial generalization of the Bernoulli factory problem. The goal is to select an outcome from a distribution that is a function of the expectations of a set of input distributions. These input distributions can be accessed only by sampling.

Definition 3.1. Given function $f : (0, 1)^n \rightarrow \Delta(X)$ for domain X , the *expectations from samples* problem is to output a sample from $f(\mu)$ given black-box access to independent samples from n distributions supported on $[0, 1]$ with expectations $\mu = (\mu_1, \dots, \mu_n) \in (0, 1)^n$.

Without loss of generality, by the Continuous to Bernoulli construction of Section 2, the input random variables can be assumed to be Bernoullis and, thus, this expectations of samples model can be viewed as a generalization of the Bernoulli factory question to output spaces X beyond

$\{0, 1\}$. In this section, we propose and solve two fundamental problems for the expectations of samples model. In these problems the outcomes are a finite set of m outcomes $X = \{1, \dots, m\}$ and the input distributions are m Bernoulli distributions with biases $\mathbf{p} = (p_1, \dots, p_m)$.

In the first problem, biases correspond to the marginal probabilities with which each of the outcomes should be selected. The goal is to produce random i from X so the probability of i is exactly its marginal probability p_i . More generally, if the biases do not sum to one, this problem is equivalently the problem of *random selection with linear weights*.

The second problem we solve corresponds to a regularized maximization problem, or specifically *random selection from exponential weights*. For this problem the biases of the m Bernoulli input distributions correspond to the weights of the outcomes. The goal is to produce a random i from X according to the distribution given by exponential weights, i.e., the probability of selecting i from X is $e^{\lambda p_i} / \sum_j e^{\lambda p_j}$.

3.1 Random Selection with Linear Weights

Definition 3.2 (Random Selection with Linear Weights). The *random selection with linear weights* problem is to sample from the probability distribution $f(\mathbf{v})$ defined by $\Pr_{I \sim f(\mathbf{v})}[I = i] = v_i / \sum_j v_j$ for each i in $\{1, \dots, m\}$ with only sample access to distributions with expectations $\mathbf{v} = (v_1, \dots, v_m)$.

We solve the random selection with linear weights problem by an algorithm that we call the *Bernoulli race* (Algorithm 1). The algorithm repeatedly picks a coin uniformly at random and flips it. The winning coin is the first one to come up heads in this process.

ALGORITHM 1: Bernoulli Race

- 1: **input** sample access to m coins with biases v_1, \dots, v_m .
 - 2: **loop**
 - 3: Draw I uniformly from $\{1, \dots, m\}$ and draw P from input distribution I .
 - 4: If P is heads then output I and halt.
 - 5: **end loop**
-

THEOREM 3.3. *The Bernoulli Race (Algorithm 1) samples with linear weights (Definition 3.2) with an expected $m / \sum_i v_i$ samples from input distributions with biases v_1, \dots, v_n .*

PROOF. At each iteration, the algorithm terminates if the flipped coin outputs 1 and iterates otherwise. Since the coin is chosen uniformly at random, the probability of termination at each iteration is $\frac{1}{m} \sum_i v_i$. The total number of iterations (and number of samples) is therefore a geometric random variable with expectation $m / \sum_i v_i$.

The selected outcome also follows the desired distribution, as shown below.

$$\begin{aligned} \Pr[i \text{ is selected}] &= \sum_{k=1}^{\infty} \Pr[i \text{ is selected at time } k] \Pr[\text{algorithm reaches time } k] \\ &= \frac{v_i}{m} \sum_{k=1}^{\infty} \left(1 - \frac{1}{m} \sum_j v_j\right)^{k-1} = \frac{\frac{v_i}{m}}{\frac{1}{m} \sum_j v_j} = \frac{v_i}{\sum_j v_j}. \quad \square \end{aligned}$$

3.2 Random Selection with Exponential Weights

Definition 3.4 (Random Selection with Exponential Weights). For parameter $\lambda > 0$, the *random selection with exponential weights* problem is to sample from the probability distribution $f(\mathbf{v})$ defined by $\Pr_{I \sim f(\mathbf{v})}[I = i] = \exp(\lambda v_i) / \sum_j \exp(\lambda v_j)$ for each i in $\{1, \dots, m\}$ with only sample access to distributions with expectations $\mathbf{v} = (v_1, \dots, v_m)$.

ALGORITHM 2: The Basic Exponential Bernoulli Race (with parameter $\lambda > 0$)

- 1: **input** Sample access to m coins with biases v_1, \dots, v_m .
- 2: For each i , apply Bernoulli Exponentiation to coin i to produce coin with bias $\tilde{v}_i = \exp(\lambda(v_i - 1))$.
- 3: Run the Bernoulli Race on the coins with biases $\tilde{\mathbf{v}} = (\tilde{v}_1, \dots, \tilde{v}_m)$.

The *Basic Exponential Bernoulli Race*, below, samples from the exponential weights distribution. The algorithm follows the paradigm of picking one of the input distributions, exponentiating it, sampling from the exponentiated distribution, and repeating until one comes up heads. While this algorithm does not generally run in polynomial time, it is a building block for one that does.

THEOREM 3.5. *The Basic Exponential Bernoulli Race (Algorithm 2) samples with exponential weights (Definition 3.4) with an expected $\lambda m e^{\lambda(1-v_{\max})}$ samples from input distributions with biases v_1, \dots, v_n and $v_{\max} = \max_i v_i$.*

PROOF. The correctness and runtime follows from the correctness and runtimes of Bernoulli Exponentiation and the Bernoulli Race. \square

3.3 The Fast Exponential Bernoulli Race

Sampling from exponential weights is typically used as a “soft maximum” where the parameter λ controls how close the selected outcome is to the true maximum. For such an application, exponential dependence on λ in the runtime would be prohibitive. Unfortunately, when v_{\max} is bounded away from one, the runtime of the Basic Logistic Bernoulli Race (Algorithm 2; Theorem 3.5) is exponential in λ . A simple observation allows the resolution of this issue: The exponential weights distribution is invariant to any uniform additive shift of all weights. This section applies this idea to develop the *Fast Logistic Bernoulli Race*.

Observe that for any given parameter ϵ , we can easily implement a Bernoulli random variable Z whose bias z is within an additive ϵ of v_{\max} . Note that, unlike the other algorithms in this section, a precise relationship between z and v_1, \dots, v_m is not required.

LEMMA 3.6. *For parameter $\epsilon \in (0, 1]$, there is an algorithm for sampling from a Bernoulli random variable with bias $z \in [v_{\max} - \epsilon, v_{\max} + \epsilon]$, where $v_{\max} = \max_i v_i$, with $O(\frac{m}{\epsilon^2} \cdot \log(\frac{m}{\epsilon}))$ samples from input distributions with biases v_1, \dots, v_m .*

PROOF. The algorithm is as follows: Sample $\frac{4}{\epsilon^2} \log(\frac{4m}{\epsilon})$ times from each of the m coins, let \hat{v}_i be the empirical estimate of coin i 's bias obtained by averaging, then apply the Continuous to Bernoulli algorithm (Section 2) to map $\hat{v}_{\max} = \max_i \hat{v}_i$ to a Bernoulli random variable.

Standard tail bounds (e.g., Chernoff-Hoeffding bound) imply that $|\hat{v}_{\max} - v_{\max}| < \epsilon/2$ with probability at least $1 - \epsilon/2$, and therefore $z = \mathbf{E}[\hat{v}_{\max}] \in [v_{\max} - \epsilon, v_{\max} + \epsilon]$. \square

Since we are interested in a fast logistic Bernoulli race as λ grows large, we restrict attention to $\lambda > 4$. We set $\epsilon = 1/\lambda$ in the estimation of v_{\max} (by Lemma 3.6). This estimate will be used to boost the bias of each distribution in the input so the maximum bias is at least $1 - 3\epsilon$. The boosting of the bias is implemented with Bernoulli Addition, which, to be fast, requires the cumulative bias be bounded away from one. Thus, the probabilities are scaled down by a factor of $1 - 2\epsilon > 1/2$ (due to the fact that $\lambda > 4$); this scaling is subsequently counterbalanced by adjusting the parameter λ . The formal details are given below.

THEOREM 3.7. *The Fast Exponential Bernoulli Race (Algorithm 3) samples with exponential weights (Definition 3.4) with an expected $O(\lambda^4 m^2 \log(\lambda m))$ samples from the input distributions.*

PROOF. The correctness and runtime follow from the correctness and runtimes of the Basic Exponential Bernoulli Race, Bernoulli Doubling, Lemma 3.6 (for estimate of v_{\max}), and the fact

ALGORITHM 3: Fast Exponential Bernoulli Race (with parameter $\lambda > 4$)

- 1: **input** Sample access to m coins with biases v_1, \dots, v_m .
- 2: Let $\epsilon = 1/\lambda$.
- 3: Construct a coin with bias $z \in [v_{\max} - \epsilon, v_{\max} + \epsilon]$ (from Lemma 3.6).
- 4: Apply Bernoulli Down Scaling to a coin with bias $1 - z$ to implement a coin with bias $(1 - 2\epsilon)(1 - z)$.
- 5: For all i , apply Bernoulli Down Scaling to implement a coin with bias $(1 - 2\epsilon)v_i$.
- 6: For all i , apply Bernoulli Addition to implement coin with bias $v'_i = (1 - 2\epsilon)v_i + (1 - 2\epsilon)(1 - z)$.
- 7: Run the Basic Exponential Bernoulli Race with parameter $\lambda' = \frac{\lambda}{1-2\epsilon}$ on the coins with bias v'_1, \dots, v'_m .

$\lambda'v'_i = \lambda(v_i + 1 - z)$ and the distribution given by exponential weights is invariant to additive shifts of all weights.

A detailed analysis of the runtime follows. Since the algorithm builds a number of sampling subroutines in a hierarchy, we analyze the runtime of the algorithm and the various subroutines in a bottom-up fashion. Steps 3 and 4 implement a coin with bias $(1 - 2\epsilon)(1 - z)$ with runtime $O(\lambda^2 m \cdot \log(\lambda m))$ per sample, as per the bound of Lemma 3.6. The coin implemented in Step 5 is sampled in constant time. Observe that $v'_i \leq (1 - 2\epsilon)(1 + v_i - v_{\max} + \epsilon) \leq 1 - \epsilon$, and the runtime of Bernoulli Doubling implies that $O(\lambda)$ samples from the coins of Steps 4 and 5 suffice for sampling $\text{Bern}[v'_i]$; we conclude that a v'_i -coin can be sampled in time $O(\lambda^3 m \cdot \log(\lambda m))$. Finally, note that for $v'_{\max} = \max_i v'_i$, we have $v'_{\max} \geq 1 - 3\epsilon$; Theorem 3.5 then implies that the Basic Exponential Bernoulli Race samples at most $\lambda' m e^{\lambda' 3\epsilon} \leq 2e^6 \lambda m = O(\lambda m)$ times from the v' -coins; we conclude the claimed runtime. \square

4 THE SINGLE-AGENT MULTIPLE-URNS PROBLEM

We investigate incentive compatible mechanism design for the *single-agent multiple-urns* problem. Informally, the mechanism needs to assign an agent to one of many urns. Each urn contains objects and the agent's value for being assigned to an urn is taken in expectation over objects from the urn. The problem asks for an incentive compatible mechanism with good welfare (i.e., the value of the agent for the assigned urn).

4.1 Problem Definition and Notations

A single agent with type t from type space \mathcal{T} desires an object o from outcome space \mathcal{O} . The agent's value for an outcome o is a function of her type t and denoted by $v(t, o) \in [0, 1]$. The agent is a risk-neutral quasi-linear utility maximizer with utility $\mathbf{E}_o[v(t, o)] - p$ for randomized outcome o and expected payment p . There are m urns. Each urn j is given by a distribution \mathcal{D}_j over outcomes in \mathcal{O} . If the agent is assigned to urn j she obtains an object from the urn's distribution \mathcal{D}_j .

A mechanism can solicit the type of the agent (who may misreport if she desires). We further assume (1) the mechanism has black-box access to evaluate $v(t, o)$ for any type t and outcome o , (2) the mechanism has sample access to the distribution \mathcal{D}_j of each urn j . The mechanism may draw objects from urns and evaluate the agent's reported value for these objects, but then must ultimately assign the agent to a single urn and charge the agent a payment. The urn and payment that the agent is assigned are random variables in the mechanism's internal randomization and randomness from the mechanisms potential samples from the urns' distributions.

The distribution of the urn the mechanism assigns to an agent, as a function of her type t , is denoted by $\mathbf{x}(t) = (x_1(t), \dots, x_m(t))$ where $x_j(t)$ is the marginal probability that the agent is assigned to urn j . Denote the expected value of the agent for urn j by $v_j(t) = \mathbf{E}_{o \sim \mathcal{D}_j}[v(t, o)]$. The expected welfare of the mechanism is $\sum_j v_j(t) x_j(t)$. The expected payment of this agent is denoted by $p(t)$. The agent's utility for the outcome and payment of the mechanism is given by $\sum_j v_j(t) x_j(t) - p(t)$.

Incentive compatibility is defined by the agent with type t preferring her outcome and payment to that assigned to another type t' .

Definition 4.1. A single-agent mechanism (\mathbf{x}, p) is *incentive compatible* if, for all $t, t' \in \mathcal{T}$:

$$\sum_j v_j(t) x_j(t) - p(t) \geq \sum_j v_j(t) x_j(t') - p(t'). \quad (1)$$

A multi-agent mechanism is Bayesian Incentive Compatible (BIC) if Equation (1) holds for the outcome of the mechanism in expectation over the truthful reports of the other agents.

4.2 Incentive Compatible Approximate Scheme

If the agent's expected value for each urn is known, or equivalently mechanism designer knows the distributions \mathcal{D}_j for all urns j rather than only sample access, this problem would be easy and admits a trivial optimal mechanism: Simply select the urn maximizing the agent's expected value $v_j(t)$ according to her reported type t , and charge her a payment of zero. What makes this problem interesting is that the designer is restricted to only *sample* the agent's value for an urn. In this case, the following Monte Carlo adaptation of the trivial mechanism is tempting: Sample from each urn sufficiently many times to obtain a close estimate $\tilde{v}_j(t)$ of $v_j(t)$ with high probability (up to any desired precision $\delta > 0$), then choose the urn j maximizing $\tilde{v}_j(t)$ and charge a payment of zero. This mechanism is not incentive compatible, as illustrated by a simple example.

Example 4.2. Consider two urns. Urn A contains only outcome o_2 , whereas B two contains a mixture of outcomes o_1 and o_3 , with o_1 slightly more likely than o_3 . Now consider an agent who has (true) values 0, 1, and 2 for outcomes o_1 , o_2 , and o_3 , respectively. If this agent reports her true type, the trivial Monte Carlo mechanism—instantiated with any desired finite degree of precision—assigns her urn A most of the time, but assigns her urn B with some nonzero probability. The agent gains by misreporting her value of outcome o_3 as 0, since this guarantees her preferred urn A .

The above example might seem counter-intuitive, since the trivial Monte Carlo mechanism appears to be doing its best to maximize the agent's utility, up to the limits of (unavoidable) sampling error. One intuitive rationalization is the following: An agent can slightly gain by procuring (by whatever means) more precise information about the distributions \mathcal{D}_j than that available to the mechanism, and using this information to guide her strategic misreporting of her type. This raises the following question:

Question. Is there an incentive-compatible mechanism for the single-agent multiple-urns problem that achieves welfare within ϵ of the optimal and samples only $\text{poly}(m, \frac{1}{\epsilon})$ times (in expectation) from the urns?

We resolve the above question in the affirmative. We present approximation scheme for this problem that is based on our solution to the problem of random selection with exponential weights (Section 3.2). The solution to the single-agent multiple-urns problem is a main ingredient in the Bayesian mechanism that we propose in Section 5 as our black-box reduction mechanism.

To explain the approximate scheme, we start by recalling the following standard theorem in mechanism design (e.g., see Groves [14] and Nisan and Ronen [29]):

THEOREM 4.3. *For outcome rule \mathbf{x} , there exists payment rule p so single-agent mechanism (\mathbf{x}, p) is incentive compatible if and only if \mathbf{x} is maximal in range, i.e., $\mathbf{x}(t) \in \text{argmax}_{\mathbf{x}'} \sum_j v_j(t) x'_j - c(\mathbf{x}')$, for some cost function $c(\cdot)$.*

Remark. The “only if” direction of this theorem requires that the type space \mathcal{T} be rich enough so the induced space of values across the urns is $\{(v_1(t), \dots, v_m(t)) : t \in \mathcal{T}\} = [0, 1]^m$.

The payments that satisfy Theorem 4.3 can be easily calculated with black-box access to outcome rule $\mathbf{x}(\cdot)$. For a single-agent problem, this payment can be calculated in two calls to the function $\mathbf{x}(\cdot)$, one on the agent's reported type t and the other on a type randomly drawn from the path between the origin and t . Further discussion and details are given later in Section 6. It suffices, therefore, to identify a mechanism that samples from urns and assigns the agent to an urn that induces an outcome rule $\mathbf{x}(\cdot)$ that is good for welfare, i.e., $\sum_i v_j(t) x_j(t)$, and is maximal in range. The following theorem solves the problem:

THEOREM 4.4. *There is an incentive-compatible mechanism for the single-agent multiple-urns problem that achieves an additive ϵ -approximation to the optimal welfare in expectation and runs in time $O(m^2(\frac{\log m}{\epsilon})^5)$ in expectation.*

PROOF. Consider the problem of selecting a distribution over urns to optimize welfare plus (a scaling of) the Shannon entropy function, i.e., $\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x}'} \sum_j v_j(t) x'_j - (1/\lambda) \sum_j x'_j \log x'_j$. The additive entropy term can be interpreted as a negative cost vis-à-vis Theorem 4.3. It is well known that the optimizer $\mathbf{x}(t)$ is given by exponential weights, i.e., the marginal probability of assigning the j th urn is given by $x_j(t) = \exp(\lambda v_j(t)) / \sum_{j'} \exp(\lambda v_{j'}(t))$, a fact that can also be verified easily from the first-order conditions. In Section 3.3, we gave a polynomial time algorithm for sampling from exponential weights, specifically, the Fast Exponential Bernoulli Race (Algorithm 3). Proper choice of the parameter λ controls the trade off between a faster run time and a decreased loss due to the entropy term. The entropy is maximized at the uniform distribution $\mathbf{x}' = (1/m, \dots, 1/m)$ with entropy $\log m$. Thus, choosing $\lambda = \log m/\epsilon$ guarantees that the welfare is within an additive ϵ of the optimal welfare $\max_j v_j(t)$. The bound of the theorem then follows from the analysis of the Fast Exponential Bernoulli Race (Theorem 3.7) with this choice of λ . \square

5 A BAYESIAN INCENTIVE COMPATIBLE BLACK-BOX REDUCTION

A central question at the interface between algorithms and economics is on the existence of black-box reductions for mechanism design. Given black-box access to any algorithm that maps inputs to outcomes, can a mechanism be constructed that induces agents to truthfully report the inputs and produces an outcome that is as good as the one produced by the algorithm? The mechanism must be computationally tractable, specifically, making no more than a polynomial number of elementary operations and black-box calls to the algorithm.

5.1 Basics of Bayesian Mechanism Design

Before formalizing this problem, we provide further details on Bayesian mechanism design and our set of notations in this article, which are mostly based on those in Hartline et al. [18].

5.1.1 Multi-parameter Bayesian Setting. Suppose there are n agents, where agent k has private type t^k from type space \mathcal{T}^k . The type profile of all agents is denoted by $\mathbf{t} = (t^1, \dots, t^n) \in \mathcal{T}^1 \times \dots \times \mathcal{T}^n$. Moreover, we assume types are drawn independently from known prior distributions. For agent k , let F^k be the distribution of $t^k \in \mathcal{T}^k$ and $\mathbf{F} = F^1 \times \dots \times F^n$ be the joint distribution of types. Suppose there is an outcome space denoted by \mathcal{O} . Agent k with type t^k has valuation $v(t^k, o)$ for outcome $o \in \mathcal{O}$, where $v : (\mathcal{T}^1 \cup \dots \cup \mathcal{T}^n) \times \mathcal{O} \rightarrow [0, 1]$ is a fixed function. Note that we assume agent values are non-negative and bounded, and without loss of generality in $[0, 1]$. Finally, we allow charging agents with non-negative money payments and we assume agents are quasi-linear, i.e., an agent with private type t has utility $u = v(t, o) - p$ for the outcome-payment pair (o, p) .

5.1.2 Algorithms, Mechanisms, and Interim Rules. An allocation algorithm \mathcal{A} is a mapping from type profiles \mathbf{t} to outcome space \mathcal{O} . A (direct revelation) mechanism \mathcal{M} is a pair of allocation rule

and *payment rule* $(\mathcal{A}, \mathbf{p})$, in which \mathcal{A} is an allocation algorithm and $\mathbf{p} = (p^1, \dots, p^n)$ where each p^k (denoted by the payment rule for agent k) is a mapping from type profiles \mathbf{t} to $\mathbb{R}_+ \cup \{0\}$.

One can think of the interaction between strategic agents and a mechanism as following: Agents submit their *reported types* $\mathbf{s} = (s^1, \dots, s^n)$ and then the mechanism \mathcal{M} picks the outcome $o = \mathcal{A}(\mathbf{s})$ and charges each agent k with its payment $p^k(\mathbf{s})$. We also consider *interim allocation rule*, which is the allocation from the perspective of one agent when the other agent's reported types are drawn from their prior distribution. More concretely, we abuse notation and define $\mathcal{A}^k(s_k) \triangleq \mathcal{A}(s^k, \mathbf{t}^{-k})$ to be the distribution over outcomes induced by \mathcal{A} when agent k 's type is s^k and other agent types are drawn from \mathbf{F}^{-k} . Similarly, for agent k , we define *interim payment rule* $p^k(s^k) \triangleq \mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}^{-k}}[p^k(s^k, \mathbf{t}^{-k})]$ and *interim value* $v^k(s^k) \triangleq \mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}^{-k}}[v(s^k, \mathcal{A}^k(s^k, \mathbf{t}^{-k}))]$. In most parts of this article, we focus only on one agent, e.g., agent k , and we just work with the interim allocation algorithm $\mathcal{A}^k(\cdot)$. When it is clear from the context, we drop the agent's superscript, and therefore $\mathcal{A}(s)$ denotes the distribution over outcomes induced by $\mathcal{A}(s, \mathbf{t}^{-k})$ when $\mathbf{t}^{-k} \sim \mathbf{F}^{-k}$.

5.1.3 Bayesian and Dominant Strategy Truthfulness. We are only interested in designing mechanisms that are *interim truthful*, i.e., every agent bests off by reporting her true type assuming all other agent's reported types are drawn independently from their prior type distribution. More precisely, a mechanism \mathcal{M} is *Bayesian Incentive Compatible (BIC)* if for all agents k , and all types $s^k, t^k \in \mathcal{T}^k$,

$$\mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}^{-k}}[v(t^k, \mathcal{A}^k(t^k))] - p^k(t^k) \geq \mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}^{-k}}[v(t^k, \mathcal{A}^k(s^k))] - p^k(s^k). \quad (2)$$

As a stronger notion of truthfulness than Bayesian truthfulness, one can consider *dominant strategy truthfulness*. More precisely, a mechanism \mathcal{M} is *Dominant Strategy Incentive Compatible (DSIC)* if for all agents k , all types $s^k, t^k \in \mathcal{T}^k$ and all type profiles $\mathbf{t}^{-k} \in \mathcal{T}^{-k}$,

$$v(t^k, \mathcal{A}(\mathbf{t})) - p^k(\mathbf{t}) \geq v(t^k, \mathcal{A}(s^k, \mathbf{t}^{-k})) - p^k(s^k, \mathbf{t}^{-k}). \quad (3)$$

Moreover, an allocation algorithm \mathcal{A} is said to be BIC (or DSIC) if there exists a payment rule \mathbf{p} such that $M = (\mathcal{A}, \mathbf{p})$ is a BIC (or DSIC) mechanism. Throughout the article, we use the terms Bayesian (or dominant strategy) truthful and Bayesian (or dominant strategy) incentive compatible interchangeably. For randomized mechanisms, DSIC and BIC solution concepts are defined by considering expectation of utilities of agents over mechanism's internal randomness.

5.1.4 Social Welfare. We are considering mechanism design for maximizing *social welfare*, i.e., the sum of the utilities of agents and the mechanism designer. For quasi-linear agents, this quantity is in fact sum of the valuations of the agents under the outcome picked by the mechanism. For the allocation algorithm \mathcal{A} , we use the notation $\text{val}(\mathcal{A})$ for the expected welfare of this allocation and $\text{val}^k(\mathcal{A})$ for the expected value of agent k under this allocation, i.e., $\text{val}(\mathcal{A}) \triangleq \mathbf{E}_{\mathbf{t} \sim \mathbf{F}}[\sum_k v(t^k, \mathcal{A}(\mathbf{t}))]$ and $\text{val}^k(\mathcal{A}) \triangleq \mathbf{E}_{\mathbf{t} \sim \mathbf{F}}[v(t^k, \mathcal{A}(\mathbf{t}))]$.

5.2 Bayesian Black-box Reductions

A line of research initiated by Hartline and Lucier [15, 16] demonstrated that, for the welfare objective, Bayesian black-box reductions can exist. The constructed mechanism is expected to be an approximation scheme; for any ϵ the reduction gives a mechanism that is Bayesian incentive compatible (Definition 4.1) and obtains a welfare that is no smaller by an additive ϵ than the algorithm's welfare in expectation. More accurately, we define the following problem:

Definition 5.1 (BIC Black-box Reduction Problem). Given black-box oracle access to an allocation algorithm \mathcal{A} , simulate a Bayesian incentive compatible allocation algorithm $\tilde{\mathcal{A}}$ that approximately preserves welfare, i.e., for every agent a , $\text{val}^a(\tilde{\mathcal{A}}) \geq \text{val}^a(\mathcal{A}) - \epsilon$, and runs in time $\text{poly}(n, \frac{1}{\epsilon})$.

In this literature, Hartline and Lucier [15, 16] solve the case of single-dimensional agents and Hartline et al. [17, 18] solve the case of multi-dimensional agents with discrete type spaces. For the relaxation of the problem where only approximate incentive compatibility is required, Bei and Huang [7] solve the case of multi-dimensional agents with discrete type space, and Hartline et al. [17, 18] solve the general case by (1) achieving exact BIC for discrete type spaces and (2) achieving approximate BIC for general multi-dimensional type spaces. These reductions are approximation schemes that are polynomial in the number of agents, the desired approximation factor, and a measure of the size of the agents' type spaces (i.e., its dimension).

Notably, one could also consider approximately preserving objectives other than welfare. However, Chawla et al. [10] have shown that BIC black-box reductions for the makespan objective cannot be computationally efficient in general. As another important note, in the Bayesian setting, agents' types are drawn from a distribution. The original algorithm ideally obtains good welfare for types drawn from this distribution in expectation; although this assumption is not necessary for the reduction to work, the black-box reduction in algorithmic mechanism design makes more sense when the algorithm is assumed to obtain good welfare in such a Bayesian sense.

5.3 Surrogate Selection and the Replica-surrogate Matching

A main conclusion of the literature on Bayesian black-box reductions for mechanism design is that the multi-agent problem of reducing Bayesian mechanism design to algorithm design, itself, reduces to a single-agent problem of *surrogate selection*. Consider any agent in the original problem and the *induced algorithm* with the inputs from other agents. These inputs are hardcoded as random draws from their respective type distributions. The induced algorithm maps the type of this agent to a distribution over outcomes. If this distribution over outcomes is maximal-in-range, then there exists payments for which the induced algorithm is incentive compatible (Theorem 4.3). If not, the problem of surrogate selection is to map the type of the agent to an input to the algorithm to satisfy three properties:

- (a) The composition of surrogate selection and the induced algorithm is maximal-in-range,
- (b) The composition approximately preserves welfare,
- (c) The surrogate selection preserves the type distribution.

Condition (c), a.k.a. *stationarity*, implies that fixing the non-maximality-of-range of the algorithm for a particular agent does not affect the outcome for any other agents. With such an approach each agent's incentive problem can be resolved independently from that of other agents.

THEOREM 5.2 (HARTLINE ET AL. [18]). *The composition of an algorithm with a profile of surrogate selection rules, which maps the profile of agent types to an input to the algorithm, is Bayesian incentive compatible and approximately preserves the algorithms welfare (the loss in welfare is the sum of the losses in welfare of each surrogate selection rule).*

The surrogate selection rule of Hartline et al. [18] is based on setting up a matching problem between random types from the distribution (replicas) and the outcomes of the algorithm on random types from the distribution (surrogates). The true type of the agent is one of the replicas, and the surrogate selection rule outputs the surrogate to which this replica is matched. Given an induced allocation algorithm $\mathcal{A}(\cdot)$, assigning a replica r_i to a *surrogate outcome* $\mathcal{A}(s_j)$ —which basically is a distribution over possible outcomes in O that the induced algorithm produces for a surrogate s_j —produces a stochastic value equal to $v(r_i, o)$, where $o \sim \mathcal{A}(s_j)$. In the aforementioned matching problem, we think of expectations of these stochastic values, i.e., the quantities $\mathbf{E}_{o \sim \mathcal{A}(s_j)}[v(r_i, o)]$ for each (r_i, s_j) , as weights on the edges. Now, this approach addresses the three properties of surrogate selection rules as:

- (a) if the matching selected is maximal-in-range given the weights, then the composition of the surrogate selection rule with the induced algorithm is maximal-in-range,
- (b) the welfare of the matching is the welfare of the reduction and the maximum weighted matching approximates the welfare of the original algorithm, and
- (c) any maximal matching gives a stationary surrogate selection rule.

In fact, the main reason to consider a replica-surrogate matching rather than assigning the reported type to the maximum value surrogate outcome is to obtain both welfare preservation (when market size m is large enough) and stationarity property; see Hartline et al. [18] for more details. For a detailed discussion on why maximal-in-range matching will result in a BIC mechanism after composing the corresponding surrogate selection rule with the allocation algorithm, we refer the interested reader to look at Lemma A.1 and Lemma A.2 in Appendix A.

In this article, we consider a slight generalization of the surrogate selection rule in Hartline et al. [18], which is a family of surrogate selection rules based on *many-to-one matchings with budgets*. For reasons that will be clear in Section 5.4, this degree of freedom will critically help us to go beyond ideal computational model and obtain exact BIC black-box reductions in the expectations from samples computational model.

Definition 5.3. The *replica-surrogate matching* surrogate selection rule; for a k -to-1 matching algorithm M , an integer market size m , and budget k ; maps a type t to a surrogate type as follows:

- (1) Pick the real-agent index i^* uniformly at random from $\{1, \dots, km\}$.
- (2) Define the *replica type profile* \mathbf{r} , a km -tuple of types by setting $r_{i^*} = t$ and sampling the remaining $km - 1$ replica types \mathbf{r}_{-i^*} i.i.d. from the type distribution F .
- (3) Sample the *surrogate type profile* \mathbf{s} , an m -tuple of i.i.d. samples from the type distribution F .
- (4) Run matching algorithm M on the complete bipartite graph between replicas and surrogates.
- (5) Output the surrogate j^* that is matched to i^* .

The value that a replica obtains for the outcome that the induced algorithm produces for a surrogate is a random variable. The analysis of Hartline et al. [18] is based on the study of an ideal computational model where the value of any replica r_i for any surrogate outcome $\mathcal{A}(s_j)$, i.e., the quantity $\mathbf{E}_{o \sim \mathcal{A}(s_j)}[v(r_i, o)]$, is known exactly. In this computationally unrealistic model and with these values as weights, the maximum weight matching algorithm can be employed in the replica-surrogate matching surrogate selection rule above, and it results in a Bayesian incentive compatible mechanism. Hartline et al. [18] analyze the welfare of the resulting mechanism in the case where the budget is $k = 1$, prove that conditions (a)–(c) are satisfied, and give (polynomial) bounds on the size m that is necessary for the expected welfare of the mechanism to be an additive ϵ from that of the algorithm.

Remark. Given a BIC allocation algorithm $\tilde{\mathcal{A}}$ through a replica-surrogate matching surrogate selection, the payments that satisfy Bayesian incentive compatibility can be easily calculated with black-box access to $\tilde{\mathcal{A}}$ as implicit payments (Section 6).

If M is maximum matching, conditions (a)–(c) clearly continue to hold for our generalization to budget $k > 1$. Moreover, the welfare of the reduction will only weakly increase for $k > 1$.

LEMMA 5.4. *In the ideal computational model (where the value of a replica for being matched to a surrogate is given exactly) the per-replica welfare of the replica-surrogate maximum matching for $k = 1$ is no larger than the per-replica welfare of the replica-surrogate maximum matching for any budget $k > 1$.*

PROOF. Consider a non-optimal matching that groups replicas into k groups of size m and finds the optimal 1-to-1 matching between replicas in each group and the surrogates. As these are random ($k = 1$)–matchings, the expected welfare of each such matching is equal to the expected welfare of the ($k = 1$)–matching. These matchings combine to give a feasible matching between the mk replicas and m surrogates. Thus, the total expected welfare of the optimal k -to-1 matching between replicas and surrogates is at least k times the expected welfare of the ($k = 1$)–matching. Thus, the per-replica welfare, i.e., normalized by mk , is at least the per-replica welfare of the ($k = 1$)–matching. \square

Our main result in the remainder of this section is an approximation scheme for the ideal reduction of Hartline et al. [18]. We replace this ideal matching algorithm with an approximation scheme for the black-box model where replica values for surrogate outcomes can only be accessed by sampling, i.e., we only have sample access to random variables $v(r_i, o)$ for $o \sim \mathcal{A}(s_j)$. For any ϵ , we identify a $k > 1$ and a polynomial (in m and $1/\epsilon$) time k -to-1 matching algorithm for the black-box model and prove that the expected welfare of this matching algorithm (per-replica) is within an additive ϵ of the expected welfare (per-replica) of the maximum weighted matching in the ideal model with budget $k = 1$ analyzed by 18. The welfare of the ideal model is monotone non-decreasing in budget k due to Lemma 5.4; therefore, it will be sufficient to identify a polynomial time algorithm in the black-box model that has ϵ loss relative to the ideal model for that same budget k . Moreover, we show our algorithm produces a perfect (and so maximal) matching, and therefore the surrogate selection rule is stationary; and the algorithm is maximal-in-range for the true agent’s replica, and therefore the resulting mechanism is Bayesian incentive compatible.

5.4 Entropy Regularized Matching

The main idea in this section is to figure out the right maximal-in-range allocation for the replica-surrogate matching problem, so it approximates the maximum matching allocation by an additive ϵ loss in the per-replica welfare and also is implementable for the black-box model discussed previously. To this end, we define an *entropy regularized* bipartite matching problem and discuss its solution. While this solution cannot be implemented as it is for reasons that we will discuss later in this section, it is the key in having a polynomial time approximate scheme for the black-box model.

Consider a complete bipartite graph with km vertices on the left-hand side and m vertices on the right-hand side. We will refer to the left-hand side vertices as replicas and the right-hand side vertices as surrogates. Fix a replica type profile \mathbf{r} and a surrogate type profile \mathbf{s} . The weights on the edge between replica $i \in \{1, \dots, km\}$ and surrogate $j \in \{1, \dots, m\}$ will be denoted by $v_{i,j}$. In our application to the replica-surrogate matching defined in the previous section, the weights will be set to $v_{i,j} = \mathbf{E}_{o \sim \mathcal{A}(s_j)}[v(r_i, o)]$ for $(i, j) \in [km] \times [m]$.

Definition 5.5. For weights $\mathbf{v} = [v_{i,j}]_{(i,j) \in [km] \times [m]}$, the entropy regularized matching program for parameter $\delta > 0$ is:

$$\begin{aligned} \max_{\{x_{i,j}\}_{(i,j) \in [km] \times [m]}} \quad & \sum_{i,j} x_{i,j} v_{i,j} - \delta \sum_{i,j} x_{i,j} \log x_{i,j}, \\ \text{s.t.} \quad & \sum_i x_{i,j} \leq k & \forall j \in [m], \\ & \sum_j x_{i,j} \leq 1 & \forall i \in [km]. \end{aligned}$$

The optimal value of this program is denoted $\text{OPT}(\mathbf{v})$.

The dual variables for right-hand side constraints of the matching polytope can be interpreted as *prices* for the surrogate outcomes. Given prices, the *utility* of a replica for a surrogate outcome

given prices is the difference between the replica's value and the price. The following lemma, whose proof is a direct application of Karush-Kuhn-Tucker conditions, shows that for the right choice of dual variables, the maximizer of the entropy regularized matching program is given by exponential weights with weights equal to the utilities.

LEMMA 5.6. *For the optimal Lagrangian dual variables $\alpha^* \in \mathbb{R}^m$ for surrogate feasibility in the entropy regularized matching program (Definition 5.5), namely,*

$$\alpha^* = \operatorname{argmin}_{\alpha \geq 0} \max_{\mathbf{x}} \left\{ \mathcal{L}(\mathbf{x}, \alpha) : \sum_j x_{i,j} \leq 1, \forall i \right\},$$

where $\mathcal{L}(\mathbf{x}, \alpha) \triangleq \sum_{i,j} x_{i,j} v_{i,j} - \delta \sum_{i,j} x_{i,j} \log x_{i,j} + \sum_j \alpha_j (k - \sum_i x_{i,j})$ is the Lagrangian objective function; the optimal solution \mathbf{x}^* to the primal is given by exponential weights:

$$x_{i,j}^* = \frac{\exp\left(\frac{v_{i,j} - \alpha_j^*}{\delta}\right)}{\sum_{j'} \exp\left(\frac{v_{i,j'} - \alpha_{j'}^*}{\delta}\right)}, \quad \forall i, j.$$

Lemma 5.6 recasts the entropy regularized matching as, for each replica, sampling from the distribution of exponential weights. For any replica i and fixed dual variables α our Fast Exponential Bernoulli Race (Algorithm 3) gives a polynomial time algorithm for sampling from the distribution of exponential weights in the expectations from samples computational model.

LEMMA 5.7. *For replica i and any prices (dual variables) $\alpha \in [0, h]^m$, allocating a surrogate j drawn from the exponential weights distribution*

$$x_{i,j} = \frac{\exp\left(\frac{v_{i,j} - \alpha_j}{\delta}\right)}{\sum_{j'} \exp\left(\frac{v_{i,j'} - \alpha_{j'}}{\delta}\right)}, \quad \forall j \in [m], \quad (4)$$

is maximal-in-range for replica i , as defined in Definition 4.3, and this random surrogate j can be sampled with $O\left(\frac{h^4 m^2 \log(hm/\delta)}{\delta^4}\right)$ samples from replica-surrogate-outcome value distributions.

PROOF. To see that the distribution is maximal-in-range when assigning surrogate outcome j to replica i , consider the regularized welfare maximization

$$\operatorname{argmax}_{\mathbf{x}'} \sum_j v_{i,j} x'_j - \delta \sum_j x'_j \log x'_j - \sum_j \alpha_j x'_j$$

for replica i . By looking at the first-order conditions, similar to Lemma 5.6, it is easy to see that the exponential weight distribution in Equation (4) is the unique maximizer of this concave program.

To apply the Fast Exponential Bernoulli Race to the utilities, of the form $v_{i,j} - \alpha_j \in [-h, 1]$, we must first normalize them to be on the interval $[0, 1]$. This normalization is accomplished by adding h to the utilities (which has no effect on the exponential weights distribution, and therefore preserves being maximal-in-range) and then scaling by $1/(h+1)$. The scaling needs to be corrected by setting λ in the Fast Exponential Bernoulli Race (Algorithm 3) to $(h+1)/\delta$. The expected number of samples from the value distributions that are required by the algorithm, per Theorem 3.7, is $O(h^4 m^2 \log(hm/\delta) \delta^{-4})$. \square

If we knew the optimal Lagrangian variables α^* from Lemma 5.6, it would be sufficient to define the surrogate selection rule by simply sampling from the true agent i^* 's exponential weights distribution (which is polynomial time per Lemma 5.7). Notice that the wrong values of α correspond to violating primal constraints for the surrogates. Thus, the outcome from sampling from exponential weights for such α would not correspond to a matching, while remains to be maximal-in-range

for each replica. In the next section, we propose a polynomial time approximation scheme that outputs a matching that is maximal-in-range for each replica, and therefore for the true agent i^* , and at the same time approximates sampling from the correct α^* .

5.5 Online Entropy Regularized Matching

In this section, we reduce the entropy regularized matching problem to the problem of sampling from exponential weights (as described in Lemma 5.7) in a sequential fashion over all replicas. Although the actual problem is indeed an offline problem, we treat it as an *online problem* where replicas arrive online, but the ordering under which they arrive is in our control. This treatment helps us to preserve the maximal-in-range property of our assignment for each replica while guaranteeing primal feasibility and near-optimal objective value in the entropy regularized matching problem (and hence near-optimal social welfare for small enough δ).

Similar to Section 5.4, fix arbitrary profiles of replicas r and surrogates s . Now consider going over replicas r in a *random order*, over times $i = 1, \dots, km$, and assigning them to the surrogates by sampling from the exponential weights distribution as given by Lemma 5.7 with prices $\alpha^{(i)}$, $i = 1, \dots, km$ (we will detail later how to set these prices). The basic observation is that (near-optimal) approximate dual variables $\alpha^{(i)}$, $i = 1, \dots, km$ are sufficient for an online assignment of each replica to a surrogate via Lemma 5.7 to obtain (near-optimal) approximations to the optimum offline regularized matching. Moreover, such a sequential assignment will result in a maximal-in-range allocation for each replica.

How to construct a sequence of dual prices $\alpha^{(i)}$ for $i = 1, \dots, km$ —ideally in an online fashion—that can play the role of near-optimal approximations to the optimal dual prices α^* of the offline problem? How to preserve the feasibility of our assignment by respecting the matching constraints of the surrogate side? To address these questions, we propose a primal-dual algorithm by borrowing techniques used in designing online algorithms for stochastic online convex programming problems (Agrawal and Devanur [1], Chen and Wang [11]), and stochastic online packing problems (Agrawal et al. [2], Badanidiyuru et al. [6], Devanur et al. [12], Kesselheim et al. [24]).

Our primal-dual online algorithm (Algorithm 4, below) considers the replicas in a random order, updates the dual variables based on the current number of allocated replicas to each surrogate (*dual update* step), and allocates an available surrogate to each arriving replica by sampling from the exponential weights distribution as given by Lemma 5.7 with updated dual variables (*primal assignment* step). Under the hood, the dual update is essentially running a no-regret learning algorithm—such as exponential gradient ascent [30] (also known as multiplicative weights), or follow-the-perturbed-leader [22], or online mirror descent [9]—for a specific adversarial online linear optimization problem (which we explain later). Roughly speaking, this online learning algorithm tries to learn a dual assignment that fits the primal allocation the best in terms of dual complementary slackness, or equivalently tries to minimize the objective value of the Fenchel dual program of the primal entropy regularized matching problem (cf. Boyd et al. [8]). For the ease of exposition, we use exponential gradient ascent for our dual updates in Algorithm 4, but in principle it can be replaced by any online learning algorithm with the same regret guarantees.

Algorithm 4 is parameterized by δ , the scale of the regularizer; by η , the rate at which the algorithm learns the dual variables α ; and by scale parameter γ . For technical reasons, our algorithm uses scaled dual prices $\gamma\alpha^{(i)}$; we detail later why this modification is needed and how to set scale parameter γ .

The final algorithm needs to satisfy four properties to be useful as a surrogate selection rule in a polynomial time Bayesian incentive compatible black-box reduction. First, it needs to produce a maximal matching so the replica-surrogate matching surrogate selection rule is stationary, specifically via condition (c) in Section 5.3. It needs to be maximal-in-range for the real agent (replica

ALGORITHM 4: Online Entropy Regularized Matching Algorithm (with parameters $\delta, \eta, \gamma \in \mathbb{R}_+$)

-
- 1: **input:** sample access to replica-surrogate matching instance with expected values $\{v_{i,j}\}$ for replicas $i \in \{1, \dots, mk\}$ and surrogates $j \in \{1, \dots, m\}$.
 - 2: Shuffle the replicas by a uniform random permutation $\pi : [mk] \rightarrow [mk]$, hence indexed by $\pi(1), \dots, \pi(mk)$.
 - 3: **for all** $i \in \{1, \dots, km\}$ **do**
 - 4: Let k_j be the number of replicas previously matched to each surrogate j and $J = \{j : k_j < k\}$ the set of surrogates with availability remaining.
 - 5: Dual update: set $\alpha^{(i)}$ according to the exponential weights distribution with weights $\eta \cdot k_j$ for available surrogates $j \in J : \alpha_j^{(i)} = \exp(\eta \cdot k_j) / \sum_{j' \in J} \exp(\eta \cdot k_{j'})$.
 - 6: Primal assignment: By running the fast exponential Bernoulli race (Algorithm 3), match the arriving replica at time i (i.e., replica $\pi(i)$) to an available surrogate $j \in J$ drawn according to the exponential weights distribution with weights $(v_{\pi(i),j} - \gamma \alpha_j^{(i)}) / \delta$.
 - 7: **end for**
-

i^*). In fact, all replicas are treated symmetrically and allocated by sampling from an exponential weights distribution that is maximal-in-range via Lemma 5.7. Third, it needs to have good welfare compared to the ideal matching (no smaller than ϵ from optimal welfare in the ideal model). Fourth, its runtime needs to be polynomial. The first two properties are immediate and imply the theorem below. The last two properties are analyzed in the next section.

THEOREM 5.8. *The mechanism that maps types to surrogates via the replica-surrogate matching surrogate selection rule with the online entropy regularized matching algorithm (Algorithm 4) is Bayesian incentive compatible (truthful payments are computed implicitly from Theorem 4.3).*

5.6 Social Welfare Loss

We analyze the welfare loss of the online entropy regularized matching algorithm (Algorithm 4) with regularizer parameter δ , learning rate η , and scale parameter γ . During the analysis, we show how to set these parameters to guarantee the per-replica expected welfare loss is at most ϵ .

THEOREM 5.9. *There are parameter settings for online entropy regularized matching algorithm (Algorithm 4) for which (1) its per-replica expected welfare is within an additive ϵ of the optimal welfare of the replica surrogate matching, and (2) given oracle access to \mathcal{A} , the running time of this algorithm is polynomial in m and $1/\epsilon$.*

PROOF OVERVIEW. To prove Theorem 5.9, we consider the following three steps:

- *Step I:* We first analyze the performance of Algorithm 4 with learning rate η in the entropy regularized matching problem and argue that our online algorithm and the offline optimal entropy regularized matching algorithm have nearly the same (per-replica) objective value in the convex program, up to an additive loss of $O(\eta)$. We show this result holds if the scale parameter γ is set appropriately and k is large enough (still polynomial in m and $1/\eta$).
- *Step II:* It turns out that to obtain the result in Step I, we need to set γ to be a constant approximation to the k -fraction of the optimal objective value of the offline convex program in Definition 5.5, and also an overestimation. We then argue how to set γ to be such an approximation/estimation for the optimal objective value of our offline convex program with high probability, and with efficient sampling. We do this step while preserving incentive compatibility.

- *Step III:* We argue that for small enough regularizer parameter $\delta > 0$, the value of the convex objective of the offline optimal entropy regularized matching is nearly as large as the welfare of the offline optimal matching.
- *Step IV:* Finally, the proof of the theorem is finished by combining the above steps with the right choice of parameters δ and η (as functions of ϵ and m), and observing this choice guarantees (i) an additive per-replica social welfare loss of $O(\epsilon)$ with respect to any replica-surrogate k -to-1 matching, and (ii) polynomial in m and $1/\epsilon$ black-box oracle complexity (and also running time) for the final algorithm.

PROOF DETAILS OF THEOREM 5.9. We provide the details of the above four steps below.

Step I: Additive per-replica loss of the online entropy regularized matching algorithm. As before, fix arbitrary profiles of replicas \mathbf{r} and surrogates \mathbf{s} , and hence the replica-surrogate expected values \mathbf{v} . Recall that the expected values \mathbf{v} play the role of edge weights in our bipartite graph. Also, recall that $\text{OPT}(\mathbf{v})$ denotes the optimal objective value of the entropy regularized matching program. We now prove the following proposition:

PROPOSITION 5.10. *For a fixed regularizer parameter $\delta > 0$, learning rate $\eta > 0$, scale parameter $\gamma > 0$, budget $k \in \mathbb{N}$, and market size $m \in \mathbb{N}$ that satisfy*

$$k \geq \frac{m \log(m/\eta)}{\eta^2} \text{ and } \text{OPT}(\mathbf{v})/k \leq \gamma \leq O(1) \text{OPT}(\mathbf{v})/k,$$

the online entropy regularized matching algorithm (Algorithm 4) obtains an objective value within an additive $O(\eta)$ of the objective value of the optimal entropy regularized matching $\text{OPT}(\mathbf{v})$ (Definition 5.5).

We start by showing the following technical lemma, which is going to be useful in several steps of the proof of Proposition 5.10:

LEMMA 5.11. *Given vectors $\mathbf{z}_1, \dots, \mathbf{z}_T \in [0, 1]^d$, where $d, T \in \mathbb{N}$, and uniform random permutation $\pi : [T] \rightarrow [T]$ over $\{1, 2, \dots, T\}$, define $\mathbf{y}_t \triangleq \sum_{\tau=t+1}^T \mathbf{z}_{\pi(\tau)}/(T-t)$. Then:*

$$\sum_{t=1}^T \mathbf{E} \left[\max_{i \in [d]} |y_{t,i} - y_{0,i}| \right] \leq O \left(\sqrt{T(\log T + \log d)} \right).$$

PROOF. Fix $i \in [d]$ and consider the stochastic sequence $y_{t,i}$ for $t = 1, 2, \dots, T$. We have:

$$\mathbf{E} [y_{t,i} | \pi(1), \dots, \pi(t-1)] = \frac{\mathbf{E} \left[\sum_{\tau=t+1}^T z_{\pi(\tau),i} | \pi(1), \dots, \pi(t-1) \right]}{T-t} = \frac{\frac{T-t}{T-t+1} \sum_{\tau=t}^T z_{\pi(\tau),i}}{T-t} = y_{t-1,i}.$$

Therefore, $\{y_{t,i}\}$ is a martingale sequence with respect to random variables $\pi(1), \pi(2), \dots, \pi(T)$. Moreover, this martingale sequence has bounded difference, simply because

$$\begin{aligned} |y_{t,i} - y_{t-1,i}| &= \left| \frac{\sum_{\tau=t+1}^T z_{\tau,i}}{T-t} - \frac{\sum_{\tau=t}^T z_{\tau,i}}{T-t+1} \right| = \frac{|\sum_{\tau=t+1}^T z_{\tau,i} - (T-t)z_{t,i}|}{(T-t)(T-t+1)} \leq \frac{\sum_{\tau=t+1}^T |z_{\tau,i} - z_{t,i}|}{(T-t+1)(T-t)} \\ &\leq \frac{1}{T-t+1}. \end{aligned}$$

Let $c_t \triangleq 1/(T-t+1)$. Then by using Azuma-Hoeffding concentration bound for bounded difference martingales, for every $\delta > 0$, we have:

$$\Pr \{|y_{t,i} - y_{0,i}| > \delta\} \leq 2 \exp \left(-\frac{\delta^2}{\sum_{\tau=1}^t c_\tau^2} \right) = 2 \exp \left(-\frac{\delta^2}{\sum_{\tau=1}^t \frac{1}{(T-\tau+1)^2}} \right) \leq 2 \exp \left(-\frac{\delta^2(T-t)}{2} \right),$$

where the last inequality holds because

$$\sum_{\tau=1}^t \frac{1}{(T-\tau+1)^2} = \sum_{\tau=T-t+1}^T \frac{1}{\tau^2} \leq \int_{T-t}^T \frac{1}{x^2} dx = \frac{1}{T-t} - \frac{1}{T} < \frac{1}{T-t}.$$

By applying union bound, we have:

$$\Pr \left\{ \max_{i \in [d]} \{|y_{t,i} - y_{0,i}|\} > \delta \right\} \leq 2d \exp \left(-\frac{\delta^2(T-t)}{2} \right).$$

Therefore, by setting $\delta = \sqrt{\frac{2 \log(Td)}{T-t}}$, we have

$$\mathbf{E} \left[\max_{i \in [d]} |y_{t,i} - y_{0,i}| \right] \leq \delta + 2d \exp \left(-\frac{\delta^2(T-t)}{2} \right) = \sqrt{\frac{2 \log(Td)}{T-t}} + \frac{2}{T} = O \left(\sqrt{\frac{\log(Td)}{T-t}} \right).$$

Now, by summing the above upper-bound over $t = 1, 2, \dots, T$, we have

$$\sum_{t=1}^T \mathbf{E} \left[\max_{i \in [d]} |y_{t,i} - y_{0,i}| \right] \leq 1 + \sum_{t=1}^{T-1} \mathbf{E} \left[\max_{i \in [d]} |y_{t,i} - y_{0,i}| \right] \leq O \left(\sqrt{\log(Td)} \right) \sum_{t=1}^{T-1} \frac{1}{\sqrt{T-t}} = O \left(\sqrt{T \log(Td)} \right),$$

where the equation holds, because $\sum_{t=1}^{T-1} \frac{1}{\sqrt{T-t}} \leq \int_0^{T-1} \frac{1}{x^{1/2}} dx = O(\sqrt{T})$. \square

PROOF OF PROPOSITION 5.10. Let permutation $\pi : [km] \rightarrow [km]$ denote the replica arrival ordering, meaning that replica $r_{\pi(i)}$ arrives at time $i \in [km]$. π is a uniformly random permutation. Let $\mathbf{x} \in \{0, 1\}^{mk \times m}$ denote the actual allocation of Algorithm 4, that is, $x_{i,j} = 1$ if replica $r_{\pi(i)}$ is matched to surrogate s_j and $x_{i,j} = 0$ otherwise. We further use \bar{x} to denote the ‘‘conditional matching probabilities’’ of replicas to surrogates in Algorithm 4, that is,

$$\forall j \in J : \bar{x}_{i,j} = \frac{\exp \left(\frac{v_{\pi(i),j} - \gamma \alpha_j^{(i)}}{\delta} \right)}{\sum_{j' \in J} \exp \left(\frac{v_{\pi(i),j'} - \gamma \alpha_{j'}^{(i)}}{\delta} \right)},$$

$$\forall j \notin J : \bar{x}_{i,j} = 0.$$

Define stopping time τ to be the first time that one of the surrogates becomes unavailable (because all k copies are matched to previous replicas), i.e.,

$$\tau \triangleq \min \left\{ \{t \in [mk] : \exists j \text{ s.t. } \sum_{i=1}^t x_{i,j} > k\} \cup \{mk+1\} \right\}.$$

Notice that either $\tau - 1 = mk$ or there exists surrogate j such that $\sum_{i=1}^{\tau-1} x_{i,j} = k$. Moreover, define the following quantities for each $i \in [km]$:

$$\text{ALG}_i(\mathbf{v}) \triangleq \sum_{j \in [m]} v_{\pi(i),j} x_{i,j},$$

$$\overline{\text{ALG}}_i(\mathbf{v}) \triangleq \sum_{j \in [m]} v_{\pi(i),j} \bar{x}_{i,j} - \delta \sum_{j \in [m]} \bar{x}_{i,j} \log \bar{x}_{i,j}.$$

Note that $\text{ALG}_i(\mathbf{v})$ is the contribution of the algorithm at time i to the social welfare, and $\overline{\text{ALG}}_i(\mathbf{v})$ is the (fractional) contribution of (allocation probabilities of) the algorithm at time i to the convex objective of the entropy regularized matching problem. Likewise, let \mathbf{x}^* denote the fractional optimum solution of the offline convex optimization for entropy regularized matching, indexed in a

way that \mathbf{x}_i^* assigns replica $r_{\pi(i)}$ (fractionally) to surrogates. Let $\text{OPT}_i(\mathbf{v})$ denote the contribution of \mathbf{x}_i^* to the convex objective of entropy regularized matching for each $i \in [km]$, that is,

$$\text{OPT}_i(\mathbf{v}) \triangleq \sum_{j \in [m]} v_{\pi(i),j} x_{i,j}^* - \delta \sum_{j \in [m]} x_{i,j}^* \log x_{i,j}^*.$$

Note that the optimum objective value $\text{OPT}(\mathbf{v})$ of the entropy regularized matching problem (see Definition 5.5) is oblivious to the ordering induced by π and is equal to $\sum_{i=1}^{km} \text{OPT}_i(\mathbf{v})$. For simplicity of notation, we drop the input argument \mathbf{v} from $\text{OPT}(\mathbf{v})$, $\text{OPT}_i(\mathbf{v})$, $\text{ALG}_i(\mathbf{v})$, and $\overline{\text{ALG}}_i(\mathbf{v})$ in the rest of the proof.

Now consider times $i = 1, 2, \dots, \tau - 1$. At each time i , a new replica $r_{\pi(i)}$ arrives. For any given allocation $\mathbf{x}'_i = (x'_{i,1}, \dots, x'_{i,m})$ of replica $r_{\pi(i)}$ to surrogates and any given scaled prices/dual variables $\gamma \boldsymbol{\alpha}^{(i)}$ at time i , define the contribution of replica $r_{\pi(i)}$ to the Lagrangian objective of Lemma 5.6 as

$$\mathcal{L}^{(i)}(\mathbf{x}'_i, \gamma \boldsymbol{\alpha}^{(i)}) \triangleq \sum_{j \in [m]} v_{\pi(i),j} x'_{i,j} - \delta \sum_{j \in [m]} x'_{i,j} \log x'_{i,j} + \sum_{j \in [m]} \gamma \alpha_j^{(i)} \left(\frac{1}{m} - x'_{i,j} \right). \quad (5)$$

At each time $i = 1, 2, \dots, \tau - 1$, the difference between $\bar{\mathbf{x}}_i = (\bar{x}_{i,1}, \dots, \bar{x}_{i,m})$ picked by Algorithm 4 and $\mathbf{x}_i^* = (x_{i,1}^*, \dots, x_{i,m}^*)$ picked by the offline optimum is that the algorithm selects its matching (conditional) probabilities with respect to dual variables $\gamma \boldsymbol{\alpha}^{(i)}$, while the offline optimum selects its fractional matching with respect to the optimal dual variables $\boldsymbol{\alpha}^*$ (Lemma 5.6). In fact, we have:

$$\bar{\mathbf{x}}_i \in \operatorname{argmax}_{\mathbf{x}'_i \in \Delta^m} \mathcal{L}^{(i)}(\mathbf{x}'_i, \gamma \boldsymbol{\alpha}^{(i)}), \quad \mathbf{x}_i^* \in \operatorname{argmax}_{\mathbf{x}'_i \in \Delta^m} \mathcal{L}^{(i)}(\mathbf{x}'_i, \boldsymbol{\alpha}^*).$$

The optimality of $\bar{\mathbf{x}}_i$ for dual variables $\gamma \boldsymbol{\alpha}^{(i)}$ —combined with Equation (5)—implies

$$\overline{\text{ALG}}_i + \sum_{j \in [m]} \gamma \alpha_j^{(i)} \left(\frac{1}{m} - \bar{x}_{i,j} \right) \geq \text{OPT}_i + \sum_{j \in [m]} \gamma \alpha_j^{(i)} \left(\frac{1}{m} - x_{i,j}^* \right).$$

By rearranging the terms, we have

$$\begin{aligned} \overline{\text{ALG}}_i &\geq \text{OPT}_i + \gamma \boldsymbol{\alpha}^{(i)} \cdot \bar{\mathbf{x}}_i - \gamma \boldsymbol{\alpha}^{(i)} \cdot \mathbf{x}_i^* \\ &= \mathbf{E}[\text{OPT}_i] + \gamma \boldsymbol{\alpha}^{(i)} \cdot \bar{\mathbf{x}}_i - \gamma \boldsymbol{\alpha}^{(i)} \cdot \mathbf{E}[\mathbf{x}_i^*] + (\text{OPT}_i - \mathbf{E}[\text{OPT}_i]) - \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{x}_i^* - \mathbf{E}[\mathbf{x}_i^*]) \\ &\stackrel{(1)}{\geq} \mathbf{E}[\text{OPT}_i] + \gamma \boldsymbol{\alpha}^{(i)} \cdot \left(\bar{\mathbf{x}}_i - \frac{1}{m} \mathbf{1} \right) + (\text{OPT}_i - \mathbf{E}[\text{OPT}_i]) - \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{x}_i^* - \mathbf{E}[\mathbf{x}_i^*]) \\ &\stackrel{(2)}{=} \frac{\text{OPT}}{km} + \gamma \boldsymbol{\alpha}^{(i)} \cdot \left(\bar{\mathbf{x}}_i - \frac{1}{m} \mathbf{1} \right) + (\text{OPT}_i - \mathbf{E}[\text{OPT}_i]) - \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{x}_i^* - \mathbf{E}[\mathbf{x}_i^*]), \end{aligned}$$

where Inequality (1) holds because:

$$\forall j \in [m] : \mathbf{E}[x_{ij}^*] = \frac{\sum_{i' \in [km]} x_{\pi(i'),j}^*}{km} \leq \frac{k}{km} = \frac{1}{m},$$

and Equality (2) holds because:

$$\mathbf{E}[\text{OPT}_i] = \frac{\sum_{i' \in [km]} \text{OPT}_{\pi(i')}}{km} = \frac{\text{OPT}}{km}.$$

Now, suppose the observed history up to time i is denoted by \mathcal{H}_{i-1} . For each $i = 1, 2, \dots, \tau - 1$, consider a history path \mathcal{H}_{i-1} that leads to $i \leq \tau - 1$. By taking expectation conditioned on any such

history path \mathcal{H}_{i-1} , we have:

$$\begin{aligned}
\mathbf{E}[\overline{\text{ALG}}_i \mid \mathcal{H}_{i-1}] &\geq \frac{\text{OPT}}{km} + \gamma \mathbf{E} \left[\boldsymbol{\alpha}^{(i)} \cdot \left(\bar{\mathbf{x}}_i - \frac{1}{m} \mathbf{1} \right) \mid \mathcal{H}_{i-1} \right] + (\mathbf{E}[\text{OPT}_i \mid \mathcal{H}_{i-1}] - \mathbf{E}[\text{OPT}_i]) \\
&\quad - \gamma \mathbf{E} \left[\boldsymbol{\alpha}^{(i)} \cdot (\mathbf{x}_i^* - \mathbf{E}[\mathbf{x}_i^*]) \mid \mathcal{H}_{i-1} \right] \\
&\stackrel{(1)}{=} \frac{\text{OPT}}{km} + \gamma \boldsymbol{\alpha}^{(i)} \cdot \left(\mathbf{E}[\bar{\mathbf{x}}_i \mid \mathcal{H}_{i-1}] - \frac{1}{m} \mathbf{1} \right) + (\mathbf{E}[\text{OPT}_i \mid \mathcal{H}_{i-1}] - \mathbf{E}[\text{OPT}_i]) \\
&\quad - \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i^* \mid \mathcal{H}_{i-1}] - \mathbf{E}[\mathbf{x}_i^*]) \\
&\stackrel{(2)}{=} \frac{\text{OPT}}{km} + \gamma \boldsymbol{\alpha}^{(i)} \cdot \left(\mathbf{E}[\mathbf{x}_i \mid \mathcal{H}_{i-1}] - \frac{1}{m} \mathbf{1} \right) + (\mathbf{E}[\text{OPT}_i \mid \mathcal{H}_{i-1}] - \mathbf{E}[\text{OPT}_i]) \\
&\quad - \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i^* \mid \mathcal{H}_{i-1}] - \mathbf{E}[\mathbf{x}_i^*]) \\
&= \frac{\text{OPT}}{km} + \gamma \boldsymbol{\alpha}^{(i)} \cdot \left(\mathbf{x}_i - \frac{1}{m} \mathbf{1} \right) + (\mathbf{E}[\text{OPT}_i \mid \mathcal{H}_{i-1}] - \mathbf{E}[\text{OPT}_i]) \\
&\quad - \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i^* \mid \mathcal{H}_{i-1}] - \mathbf{E}[\mathbf{x}_i^*]) + \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i \mid \mathcal{H}_{i-1}] - \mathbf{x}_i),
\end{aligned}$$

where Equality (1) holds as $\boldsymbol{\alpha}^{(i)}$ is only a function of history up to time i , and Equality (2) holds as $\mathbf{E}[\bar{\mathbf{x}}_i \mid \mathcal{H}_{i-1}, \pi(i)] = \mathbf{E}[\mathbf{x}_i \mid \mathcal{H}_{i-1}, \pi(i)]$ for any $\pi(i)$, and therefore $\mathbf{E}[\bar{\mathbf{x}}_i \mid \mathcal{H}_{i-1}] = \mathbf{E}[\mathbf{x}_i \mid \mathcal{H}_{i-1}]$. To simplify the calculations, we introduce the following extra notations:

$$\begin{aligned}
O_i &\triangleq \mathbf{E}[\text{OPT}_i \mid \mathcal{H}_{i-1}] - \mathbf{E}[\text{OPT}_i], \\
Z_i &\triangleq \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i^* \mid \mathcal{H}_{i-1}] - \mathbf{E}[\mathbf{x}_i^*]), \\
L_i &\triangleq \gamma \boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i \mid \mathcal{H}_{i-1}] - \mathbf{x}_i).
\end{aligned}$$

By summing over times $i = 1, 2, \dots, \tau$, we have:

$$\sum_{i=1}^{\tau-1} \mathbf{E}[\overline{\text{ALG}}_i \mid \mathcal{H}_{i-1}] \geq \frac{\tau-1}{km} \text{OPT} + \gamma \sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}^{(i)}) - \sum_{i=1}^{\tau-1} |O_i| - \sum_{i=1}^{\tau-1} |Z_i| + \sum_{i=1}^{\tau-1} L_i, \quad (6)$$

where $g_i : [0, 1]^m \rightarrow \mathbb{R}$ is defined as follows for each $i = 1, 2, \dots, \tau - 1$:

$$g_i(\boldsymbol{\alpha}) \triangleq \boldsymbol{\alpha} \cdot \left(\mathbf{x}_i - \frac{1}{m} \mathbf{1} \right).$$

To bound the term $\sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}^{(i)})$ in Equation (6) from below, consider a full-information adversarial online linear maximization problem [19, 30] for rounds $i = 1, 2, \dots, \tau - 1$ ², where at each round the decision maker (player 1) chooses the dual vector $\boldsymbol{\alpha}^{(i)} \in \{\boldsymbol{\alpha} \in [0, 1]^m : \|\boldsymbol{\alpha}\|_1 \leq 1\}$, and an adaptive adversary (player 2) chooses the linear cost function $g_i(\boldsymbol{\alpha}) = \boldsymbol{\alpha} \cdot (\mathbf{x}_i - \frac{1}{m} \mathbf{1})$ defined above. For any given adversarial realization of random variables $\{\mathbf{x}_i\}$, which defines the strategies of player 2, the goal of player 1 is to produce a sequence $\boldsymbol{\alpha}^{(1)}, \boldsymbol{\alpha}^{(2)}, \dots, \boldsymbol{\alpha}^{(\tau-1)}$ that maximizes the linear objective function $\sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}^{(i)})$.

Now, consider ‘‘dual update’’ steps of Algorithm 4. These steps are essentially equivalent to player 1 running the exponential weight forecaster (also known as multiplicative weight updates) as a vanishing regret online learning algorithm in the aforementioned online linear maximization

²Note that τ is a random variable, and hence the number of rounds in this online linear optimization problem is stochastic.

problem. This algorithm, which is parametric with learning rate η , picks the sequence $\alpha^{(i)}$ for $i = 1, 2, \dots, \tau - 1$ to be the exponential weights distributions with weights $\eta \cdot k_j$ and guarantees the following regret bound [19]:

$$\sum_{i=1}^{\tau-1} g_i(\alpha^{(i)}) \geq (1 - \eta) \max_{\|\alpha\|_1 \leq 1, \alpha \geq 0} \sum_{i=1}^{\tau-1} g_i(\alpha) - \frac{\log m}{\eta} \geq (1 - \eta) \left(k - \frac{\tau - 1}{m} \right) - \frac{\log m}{\eta}, \quad (7)$$

where the last inequality holds because at the time $\tau - 1$, either there exists j such that $\sum_{i=1}^{\tau-1} x_{i,j} = k$, or $\tau - 1 = mk$ and all surrogate outcome budgets are exhausted. In the former case, we have

$$\max_{\|\alpha\|_1 \leq 1, \alpha \geq 0} \sum_{i=1}^{\tau-1} g_i(\alpha) \geq \sum_{i=1}^{\tau-1} g_i(\mathbf{e}_j) \geq k - \frac{\tau - 1}{m},$$

and in the latter case, we have

$$\max_{\|\alpha\|_1 \leq 1, \alpha \geq 0} \sum_{i=1}^{\tau-1} g_i(\alpha) \geq 0 \geq k - \frac{\tau - 1}{m}.$$

By applying the regret bound in Equation (7) to the RHS of the inequality in (6), we have:

$$\begin{aligned} \sum_{i=1}^{\tau-1} \mathbf{E}[\overline{\text{ALG}}_i \mid \mathcal{H}_{i-1}] &\geq \frac{\tau - 1}{km} \text{OPT} + \gamma(1 - \eta) \left(k - \frac{\tau - 1}{m} \right) - \gamma \frac{\log m}{\eta} - \sum_{i=1}^{\tau-1} |O_i| - \sum_{i=1}^{\tau-1} |Z_i| + \sum_{i=1}^{\tau-1} L_i \\ &\stackrel{(1)}{\geq} \frac{\tau - 1}{km} \text{OPT} + \text{OPT}(1 - \eta) \left(1 - \frac{\tau - 1}{km} \right) - \gamma \frac{\log m}{\eta} - \sum_{i=1}^{\tau-1} |O_i| - \sum_{i=1}^{\tau-1} |Z_i| + \sum_{i=1}^{\tau-1} L_i \\ &\stackrel{(2)}{\geq} \text{OPT} \left(1 - \eta - O(1) \frac{\log m}{k\eta} \right) - \sum_{i=1}^{\tau-1} |O_i| - \sum_{i=1}^{\tau-1} |Z_i| + \sum_{i=1}^{\tau-1} L_i \\ &\stackrel{(3)}{\geq} \text{OPT} - \eta \left(km + kmO \left(\frac{\log m}{m \log(m/\eta)} \right) \right) - \sum_{i=1}^{\tau-1} |O_i| - \sum_{i=1}^{\tau-1} |Z_i| + \sum_{i=1}^{\tau-1} L_i \\ &= \text{OPT} - O(\eta km) - \sum_{i=1}^{\tau-1} |O_i| - \sum_{i=1}^{\tau-1} |Z_i| + \sum_{i=1}^{\tau-1} L_i, \end{aligned}$$

where Inequality (1) above holds as $\gamma \geq \text{OPT}/k$, Inequality (2) above holds as $\gamma \leq O(1)\text{OPT}/k$, and Inequality (3) above holds because $k \geq m \log(m/\eta)/\eta^2$ and $\text{OPT}(\mathbf{v}) \leq km$. Further notice that for times $i = 1, 2, \dots, \tau - 1$, if we consider history paths \mathcal{H}_{i-1} in which $i \leq \tau - 1$, we have:

$$\begin{aligned} \mathbf{E}[\text{ALG}_i \mid \mathcal{H}_{i-1}, \pi(i)] &= \sum_{j \in [m]} \mathbf{E}[x_{i,j} \mid \mathcal{H}_{i-1}, \pi(i)] v_{\pi(i),j} \\ &= \sum_{j \in [m]} \mathbf{E}[\bar{x}_{i,j} \mid \mathcal{H}_{i-1}, \pi(i)] v_{\pi(i),j} \geq \mathbf{E}[\overline{\text{ALG}}_i \mid \mathcal{H}_{i-1}, \pi(i)]. \end{aligned}$$

Therefore, by taking expectation with respect to $\pi(i)$, we have $\mathbf{E}[\text{ALG}_i \mid \mathcal{H}_{i-1}] \geq \mathbf{E}[\overline{\text{ALG}}_i \mid \mathcal{H}_{i-1}]$ for $i = 1, 2, \dots, \tau - 1$. Denoting $\text{ALG} = \sum_{i=1}^{km} \text{ALG}_i$, we have:

$$\begin{aligned} \mathbf{E}[\text{ALG}] &\geq \mathbf{E}\left[\sum_{i=1}^{\tau-1} \text{ALG}_i\right] = \mathbf{E}\left[\sum_{i=1}^{\tau-1} \mathbf{E}[\text{ALG}_i \mid \mathcal{H}_{i-1}]\right] \geq \mathbf{E}\left[\sum_{i=1}^{\tau-1} \mathbf{E}[\overline{\text{ALG}}_i \mid \mathcal{H}_{i-1}]\right] \\ &\geq \text{OPT} - O(\eta km) - \mathbf{E}\left[\sum_{i=1}^{\tau-1} |O_i|\right] - \mathbf{E}\left[\sum_{i=1}^{\tau-1} |Z_i|\right] + \mathbf{E}\left[\sum_{i=1}^{\tau-1} L_i\right] \\ &\geq \text{OPT} - O(\eta km) - \mathbf{E}\left[\sum_{i=1}^{km} |O_i|\right] - \mathbf{E}\left[\sum_{i=1}^{km} |Z_i|\right] + \mathbf{E}\left[\sum_{i=1}^{\tau-1} L_i\right]. \end{aligned} \quad (8)$$

We now finish the proof by:

- (i) Showing $\mathbf{E}[\sum_{i=1}^{\tau-1} L_i] = 0$,
- (ii) Upper bounding $\mathbf{E}[\sum_{i=1}^{km} |O_i|]$ by $O(\sqrt{km \log(km)})$ and $\mathbf{E}[\sum_{i=1}^{km} |Z_i|]$ by $O(\gamma \sqrt{km \log(km)})$.

Assuming (i) and (ii) are done, and recalling that $k \geq m \log(m/\eta)/\eta^2$, we have:

$$\mathbf{E}\left[\sum_{i=1}^{km} |O_i|\right] \leq O\left(\sqrt{km \log(km)}\right) = kmO\left(\eta \sqrt{\frac{\log(m^2 \log(m/\eta)) - 2 \log(1/\eta)}{m^2 \log(m/\eta)}}\right) = O(\eta k). \quad (9)$$

Similarly, recalling that $\gamma \leq O(1)\text{OPT}(\mathbf{v})/k \leq O(m)$, we have:

$$\mathbf{E}\left[\sum_{i=1}^{km} |Z_i|\right] \leq O\left(\gamma \sqrt{km \log(km)}\right) = O(\eta km). \quad (10)$$

Combining the bounds (9) and (10) with $\mathbf{E}[\sum_{i=1}^{\tau-1} L_i] = 0$ and plugging them in (8) gives us the final result, i.e., $\mathbf{E}[\text{ALG}]/km \geq \text{OPT}/km - O(\eta)$.

To show (i), we simply use the fact that conditioned on any history path \mathcal{H}_{i-1} for which $\tau - 1 \geq i$, we have:

$$\mathbf{E}[L_i \mid \mathcal{H}_{i-1}] = \mathbf{E}\left[\alpha^{(i)} \cdot (\mathbf{E}[x_i \mid \mathcal{H}_{i-1}] - x_i) \mid \mathcal{H}_{i-1}\right] = \alpha^{(i)} \cdot (\mathbf{E}[x_i \mid \mathcal{H}_{i-1}] - \mathbf{E}[x_i \mid \mathcal{H}_{i-1}]) = 0.$$

Therefore, $\mathbf{E}[\sum_{i=1}^{\tau-1} L_i] = \sum_{i=1}^{km} \Pr[\tau - 1 \geq i] \mathbf{E}_{\mathcal{H}_{i-1}: \tau-1 \geq i}[\mathbf{E}[L_i \mid \mathcal{H}_{i-1}]] = 0$.

To show (ii), we use the technical Lemma 5.11:

- To upper bound $\mathbf{E}[\sum_{i=1}^{km} |O_i|]$, let $d = 1$, $T = km$, $z_i = \text{OPT}_{\pi^{-1}(i)}$ and $y_i = \mathbf{E}[\text{OPT}_{i+1} \mid \mathcal{H}_i]$ for $i = 1, 2, \dots, km$. Notice that

$$\begin{aligned} y_i &= \mathbf{E}[\text{OPT}_{i+1} \mid \mathcal{H}_i] = \frac{\sum_{i'=i+1}^{km} \text{OPT}_{i'}}{km - i} = \frac{\sum_{i'=i+1}^{km} z_{i'}}{km - i}, \\ y_0 &= \mathbf{E}[\text{OPT}_1] = \mathbf{E}[\text{OPT}_{i+1}] = \text{OPT}/km, \\ |O_i| &= |y_i - y_0| \end{aligned}$$

Thus, the conditions of Lemma 5.11 are satisfied and $\mathbf{E}[\sum_{i=1}^{km} |O_i|] \leq O(\sqrt{km \log(km)})$.

- To upper bound $\mathbf{E}[\sum_{i=1}^{km} |Z_i|]$, let $d = m$, $T = km$, $\mathbf{z}_i = \mathbf{x}_{\pi^{-1}(i)}^*$ and $\mathbf{y}_i = \mathbf{E}[\mathbf{x}_{i+1}^* \mid \mathcal{H}_i]$ for $i = 1, 2, \dots, km$. Again, notice that

$$\mathbf{y}_i = \mathbf{E}[\mathbf{x}_{i+1}^* \mid \mathcal{H}_i] = \frac{\sum_{i'=i+1}^{km} \mathbf{x}_{i'}^*}{km - i} = \frac{\sum_{i'=i+1}^{km} \mathbf{z}_{i'}}{km - i},$$

$$\mathbf{y}_0 = \mathbf{E}[\mathbf{x}_1^*] = \mathbf{E}[\mathbf{x}_{i+1}^*] = \sum_{i=1}^{km} \mathbf{x}_i^* / km.$$

Therefore, the conditions in the statement of Lemma 5.11 are satisfied. Moreover, note that

$$|Z_i| = \gamma |\alpha^{(i)} \cdot (\mathbf{y}_i - \mathbf{y}_0)| \leq \gamma \max_{j \in [m]} |y_{i,j} - y_{0,j}|,$$

where the inequality above holds as $\|\alpha^{(i)}\|_1 \leq 1$. Therefore,

$$\mathbf{E}\left[\sum_{i=1}^{km} |Z_i|\right] \leq O(\gamma \sqrt{km(\log(km) + \log(m))}) = O(\gamma \sqrt{km \log(km)}),$$

which concludes the proof. \square

Step II: parameter γ and approximating the offline optimal. Pre-setting γ to be an estimate of the optimal objective of the convex program in Definition 5.5 is necessary for guarantee of Algorithm 4 in Proposition 5.10. Also, γ should be set in a symmetric and incentive compatible way across replicas to preserve stationarity property. To this end, we look at an instance generated by an independent random draw of mk replicas (while fixing the surrogates). In such an instance, we estimate the expected values by sampling and taking the empirical mean for each edge in the replica-surrogate bipartite graph. We then solve the convex program exactly (which can be done in polytime using an efficient separation oracle). Obviously, this scheme is incentive compatible, as we do not even use the reported type of true agent in our calculation for γ , and it is symmetric across replicas. We now show how this approach leads to a constant approximation to the optimal value of the offline program in Definition 5.5 with high probability. Note that this scheme should be run offline as a *pre-processing* step before running our online Algorithm 4 used in Proposition 5.10.

PROPOSITION 5.12. *If $k = \Omega(\frac{\log(1/\epsilon')}{\delta^2 m (\log m)^2})$ for some $\epsilon' > 0$, then there exists a polynomial time algorithm that approximately calculates $\text{OPT}(\mathbf{v})/k$, that is, it outputs γ such that*

$$\text{OPT}(\mathbf{v})/k \leq \gamma \leq O(1) \text{OPT}(\mathbf{v})/k$$

with probability at least $1 - \epsilon'$. Moreover, this algorithm only needs polynomial in $m, k, 1/\delta$ and $1/\epsilon'$ samples to black-box allocation \mathcal{A} .

To formalize the approximation scheme and to prove Proposition 5.12, first fix the surrogate type profile \mathbf{s} . For a given replica profile \mathbf{r} and replica-surrogate edge (i, j) , let $v_{i,j}(r_i) = \mathbf{E}[v(r_i, \mathcal{A}(s_j))]$ and $\hat{v}_{i,j}(r_i)$ be the empirical mean of N samples of the random variable $v(r_i, \mathcal{A}(s_j))$. Suppose $\mathbf{v}(\mathbf{r})$ and $\hat{\mathbf{v}}(\mathbf{r})$ are the corresponding vectors of expected values and empirical means under replica profile \mathbf{r} . Now, draw \mathbf{r}' independently at random from the distribution of \mathbf{r} . We now show that $\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ is a constant approximation to $\text{OPT}(\mathbf{v}(\mathbf{r}))$ with high probability, and, therefore, we can use $\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ to set γ .

We prove this in two steps. In Lemma 5.13, we show for a given \mathbf{r}' , $\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ is a constant approximation to $\text{OPT}(\mathbf{v}(\mathbf{r}'))$ with high probability over randomness in $\{\mathcal{A}(s_j)\}$. Then, in Lemma 5.15, we show if \mathbf{r} and \mathbf{r}' are two random independent draws from replica profile distribution, then $\text{OPT}(\mathbf{v}(\mathbf{r}'))$ is a constant approximation to $\text{OPT}(\mathbf{v}(\mathbf{r}))$ with high probability over randomness in \mathbf{r} and \mathbf{r}' . These two pieces together prove our claim.

LEMMA 5.13. *If $N \geq \frac{\log(4m^2k/\epsilon')}{\delta^2(\log m)^2}$, then $1/2 \cdot \text{OPT}(\mathbf{v}(\mathbf{r}')) \leq \text{OPT}(\hat{\mathbf{v}}(\mathbf{r}')) \leq 2 \text{OPT}(\mathbf{v}(\mathbf{r}'))$ with probability at least $1 - \epsilon'/2$.*

PROOF. By using standard Chernoff-Hoeffding bound together with union bound, with probability at least $1 - 2m^2ke^{-\frac{\delta^2(\log m)^2 \cdot N}{2}} \geq 1 - \epsilon'/2$, we have

$$\forall (i, j) \in [km] \times [m] : |\hat{v}_{i,j}(r'_i) - v_{i,j}(r'_i)| \leq 1/2 \cdot \delta \log m.$$

Suppose \mathbf{x}^* is the optimal solution of the regularized matching convex program with values $\mathbf{v}(\mathbf{r}')$ and \mathbf{x}^{**} is the optimal solution with values $\hat{\mathbf{v}}(\mathbf{r}')$. Denoting the entropy function by $H(\cdot)$, we have:

$$\begin{aligned} \text{OPT}(\hat{\mathbf{v}}(\mathbf{r}')) &= \sum_i (\mathbf{x}_i^{**} \cdot \hat{\mathbf{v}}_i + \delta H(\mathbf{x}_i^{**})) \geq \sum_i (\mathbf{x}_i^* \cdot \hat{\mathbf{v}}_i + \delta H(\mathbf{x}_i^*)) \\ &\geq \sum_i (\mathbf{x}_i^* \cdot \mathbf{v}_i + \delta H(\mathbf{x}_i^*)) - \frac{\delta km \log m}{2} = \text{OPT}(\mathbf{v}(\mathbf{r}')) - \frac{\delta km \log m}{2} \geq 1/2 \cdot \text{OPT}(\mathbf{v}(\mathbf{r}')), \end{aligned}$$

where the last inequality holds, as $\text{OPT}(\mathbf{v}(\mathbf{r}'))$ is bounded below by the value of the uniform allocation, i.e., $\text{OPT}(\mathbf{v}(\mathbf{r}')) \geq \delta \cdot mk \log(m)$. Similarly, one can show $\text{OPT}(\mathbf{v}(\mathbf{r}')) \geq 1/2 \cdot \text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$, which completes the proof. \square

Before proving the second step, we prove the following lemma, which basically shows the optimal value of regularized matching $\text{OPT}(\mathbf{v}(\cdot))$ is a 1-Lipschitz multivariate function:

LEMMA 5.14. *For every $i \in [km]$, replica profile \mathbf{r} , and replica type r'_i , we have:*

$$|\text{OPT}(\mathbf{v}(r_i, r_{-i})) - \text{OPT}(\mathbf{v}(r'_i, r_{-i}))| \leq 1.$$

PROOF. Let \mathbf{x} and \mathbf{x}' be the optimal assignments in $\text{OPT}(\mathbf{v}(r_i, r_{-i}))$ and $\text{OPT}(\mathbf{v}(r'_i, r_{-i}))$, respectively. We have

$$\begin{aligned} \text{OPT}(\mathbf{v}(r_i, r_{-i})) &= \sum_l (\mathbf{x}_l \cdot \mathbf{v}_l(r_l) + H(\mathbf{x}_l)) \geq \sum_{l \neq i} (\mathbf{x}'_l \cdot \mathbf{v}_l(r_l) + H(\mathbf{x}'_l)) + \mathbf{x}'_i \cdot \mathbf{v}_i(r_i) + H(\mathbf{x}'_i) \\ &\geq \sum_{l \neq i} (\mathbf{x}'_l \cdot \mathbf{v}_l(r_l) + H(\mathbf{x}'_l)) + \mathbf{x}'_i \cdot \mathbf{v}_i(r'_i) + H(\mathbf{x}'_i) - 1 = \text{OPT}(\mathbf{v}(r'_i, r_{-i})) - 1, \end{aligned}$$

where the last inequality holds, because $\mathbf{x}'_i \cdot (\mathbf{v}_i(r'_i) - \mathbf{v}_i(r_i)) \leq 1$. Similarly, $\text{OPT}(\mathbf{v}(r'_i, r_{-i})) \geq \text{OPT}(\mathbf{v}(r_i, r_{-i})) - 1$ by switching the roles of r_i and r'_i . \square

LEMMA 5.15. *If $k \geq \frac{32 \log(8/\epsilon')}{\delta^2 m (\log m)^2}$, then $1/2 \cdot \text{OPT}(\mathbf{v}(\mathbf{r})) \leq \text{OPT}(\mathbf{v}(\mathbf{r}')) \leq 3/2 \cdot \text{OPT}(\mathbf{v}(\mathbf{r}))$ with probability at least $1 - \epsilon'/2$.*

PROOF. The lemma can be proved directly by McDiarmid's inequality [26] and using the 1-Lipschitzness proved in Lemma 5.14. For completeness, we prove it here from first principles. We start by defining the following Doob martingale sequence [27], where (conditional) expectations are taken over the randomness in the replica profile \mathbf{r} :

$$\begin{aligned} X_0 &= \mathbf{E}[\text{OPT}(\mathbf{v}(\mathbf{r}))], \\ X_n &= \mathbf{E}[\text{OPT}(\mathbf{v}(\mathbf{r})) | r_1, \dots, r_n], n = 1, 2, \dots, km. \end{aligned}$$

It is easy to check that $\mathbf{E}[X_n | r_1, \dots, r_{n-1}] = X_{n-1}$, and therefore $\{X_n\}$ forms a martingale sequence with respect to $\{r_n\}$. Moreover, $|X_n - X_{n-1}| \leq 1$ because of Lemma 5.14. Now, by using Azuma-Hoeffding bound for martingales, we have

$$\Pr \left\{ |X_{km} - X_0| \geq \frac{\delta km \log(m)}{4} \right\} \leq 2e^{-\frac{km \delta^2 (\log m)^2}{32}},$$

and therefore w.p. at least $1 - 2e^{-\frac{km\delta^2(\log m)^2}{32}}$, we have $|\text{OPT}(\mathbf{v}(\mathbf{r})) - \mathbb{E}[\text{OPT}(\mathbf{v}(\mathbf{r}))]| \leq \frac{\delta km \log(m)}{4}$. Similarly, w.p. at least $1 - 2e^{-\frac{km\delta^2(\log m)^2}{32}}$, we have $|\text{OPT}(\mathbf{v}(\mathbf{r}')) - \mathbb{E}[\text{OPT}(\mathbf{v}(\mathbf{r}))]| \leq \frac{\delta km \log(m)}{4}$. Therefore, w.p. at least $1 - 4e^{-\frac{km\delta^2(\log m)^2}{32}}$, we have $|\text{OPT}(\mathbf{v}(\mathbf{r})) - \text{OPT}(\mathbf{v}(\mathbf{r}'))| \leq \frac{\delta km \log(m)}{2}$. By using the lower-bound of $\delta km \log(m)$ for $\text{OPT}(\mathbf{v}(\mathbf{r}))$ (due to uniform assignment), we conclude that with probability at least $1 - 4e^{-\frac{km\delta^2(\log m)^2}{32}} \geq 1 - \epsilon'/2$, we have the following, as desired:

$$1/2 \cdot \text{OPT}(\mathbf{v}(\mathbf{r})) \leq \text{OPT}(\mathbf{v}(\mathbf{r}')) \leq 3/2 \cdot \text{OPT}(\mathbf{v}(\mathbf{r})). \quad \square$$

By putting Lemma 5.13 and Lemma 5.15 together, we immediately get the following corollary:

COROLLARY 5.16. *If $N \geq \frac{\log(4m^2k/\epsilon')}{\delta^2(\log m)^2}$ and $k \geq \frac{32 \log(8/\epsilon')}{\delta^2 m (\log m)^2}$, then $\gamma = \frac{4}{k} \text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ satisfies*

$$\text{OPT}(\mathbf{v}(\mathbf{r}))/k \leq \gamma \leq 12 \cdot \text{OPT}(\mathbf{v}(\mathbf{r}))/k,$$

with probability at least $1 - \eta$.

We conclude the above discussion as the proof of Proposition 5.12 is immediate from Corollary 5.16.

Step 3: convex objective of optimal entropy regularized matching vs. welfare of optimal matching. The last ingredient we need in the proof is the following lemma:

LEMMA 5.17. *With parameter $\delta \geq 0$ the per-replica convex objective value of the optimal entropy regularized matching is within an additive $\delta \log m$ of the welfare of the optimal matching.*

PROOF. The entropy $-\sum_{i,j} x_{i,j} \log x_{i,j}$ is non-negative and maximized with $x_{i,j} = 1/m$. The maximum value of the entropy term is thus $\delta mk \log m$. The optimal convex objective value of the entropy regularized matching exceeds that of the optimal matching; thus, it is within an additive $\delta mk \log m$ of the welfare of the optimal (unregularized) matching. As a result, the per-replica convex objective value of the optimal entropy regularized matching is within $\delta \log m$ of the per-replica welfare of the optimal matching. \square

Step 4: putting all the pieces together. We conclude the section by combining Propositions 5.12, 5.10, and Lemma 5.17 to prove the main theorem.

PROOF OF THEOREM 5.9. Let $\delta = \frac{\epsilon}{3} \cdot \frac{1}{\log m}$ and $\eta = \frac{\epsilon}{3} \cdot \frac{1}{c}$, where c is a constant such that the per-replica welfare of Algorithm 4 is within an additive $c \cdot \eta$ of the per-replica offline optimal objective value of the entropy regularized matching problem when estimation γ is set appropriately (Proposition 5.10). Moreover, let $k = \frac{m \log(m/\eta)}{\eta^2} = \Theta\left(\frac{m \log(m/\epsilon)}{\epsilon^2}\right)$, to satisfy the required condition in Proposition 5.10. Finally, set $\epsilon' = \frac{\epsilon}{3}$ and the number of samples N in the scheme described in Corollary 5.16 to $N = \Theta\left(\frac{\log(m^2k/\epsilon')}{\delta^2(\log m)^2}\right) = \Theta\left(\frac{\log(m/\epsilon)}{\epsilon^2(\log m)^2}\right)$, so γ is the appropriate estimator used in Proposition 5.10 with probability $1 - \frac{\epsilon}{3}$.

The expected per-replica welfare of Algorithm 4 is within an additive $c\eta = \epsilon/3$ of the per-replica optimal convex objective value of the entropy regularized matching, with probability at least $1 - \epsilon$, due to Propositions 5.10 and 5.12. This probability is over the internal randomness of the samples used in the estimation scheme for γ in Corollary 5.16. The per-replica optimal objective value of the entropy regularized matching is at most 1. Therefore, the overall expected per-replica welfare of Algorithm 4 is within an additive $\epsilon/3 + \epsilon' = 2\epsilon/3$ of the per-replica optimal convex objective value of the entropy regularized matching. Following Lemma 5.17, the per-replica optimal value of the entropy regularized matching is within an additive $\delta \log m = \epsilon/3$ of the per-replica welfare

of the optimal matching, and therefore the expected per-replica welfare of Algorithm 4 is within an additive $2\epsilon/3 + \epsilon/3 = \epsilon$ of the per-replica welfare of the optimal matching.

Finally, due to Proposition 5.12 and the fact that k is polynomial in m and $1/\epsilon$, the algorithm's running time is polynomial in m and $1/\epsilon$. \square

5.7 The End-to-End BIC Black-box Reduction

We now summarize the proposed BIC black-box reduction. We incorporate our surrogate selection rule (by using Algorithm 4 as the matching algorithm in Definition 5.3) in the reduction under ideal-model proposed in Hartline et al. [18] and we set the market size parameter m accordingly to maintain the welfare preservation property of this reduction.

Definition 5.18 (Hartline et al. [18]). The doubling dimension of a metric space is the smallest constant Δ such that every bounded subset S can be partitioned into at most 2^Δ subsets, each having diameter at most half of the diameter of S .

We now use the following theorem in Reference [18], which states the welfare preservation of the maximum weight replica-surrogate matching in the ideal model if m is large enough:

THEOREM 5.19 (HARTLINE ET AL. [18]). *For any agent with type space \mathcal{T} that has doubling dimension $\Delta \geq 2$, if*

$$m \geq \frac{1}{2\epsilon^{\Delta+1}},$$

then the expected per-replica welfare of the maximum matching in the ideal model of Hartline et al. [18] with load $k = 1$ is within an additive 2ϵ of the expected welfare of allocation \mathcal{A} for that agent.

We now have the following immediate corollary by combining Theorem 5.9 with Theorem 5.19:

COROLLARY 5.20 (BIC BLACK-BOX REDUCTION). *If the market size parameter m is set to $\lceil \frac{1}{2\epsilon^{\Delta+1}} \rceil$, and the parameters of Algorithm 4 are set as stated in Theorem 5.9, then the composition of surrogate selection rule defined by Algorithm 4 with the allocation \mathcal{A} is (1) a BIC mechanism, (2) the expected welfare is within an additive 3ϵ of the expected welfare of \mathcal{A} for each agent, and (3) its running time is polynomial in n and $1/\epsilon$ given access to black-box oracle \mathcal{A} .³*

6 IMPLICIT PAYMENT COMPUTATIONS

In this section, we describe one standard reduction for computing implicit payments in our general setting, given access to a BIC allocation algorithm $\tilde{\mathcal{A}}$: a multi-parameter counterpart of the single-parameter payment computation procedure used for example by Archer et al. [3], Hartline and Lucier [15], which makes $n + 1$ calls to $\tilde{\mathcal{A}}$, thus incurring a factor $n + 1$ overhead in running time. A different implicit payment computation procedure, described in Babaioff et al. [4, 5], avoids this overhead by calling $\tilde{\mathcal{A}}$ only once in expectation, but incurs a $1 - \epsilon$ loss in expected welfare and potentially makes payments of magnitude $\Theta(1/\epsilon)$ from the mechanism to the agents.

The implicit payment computation procedure assumes that the agents' type spaces $(\mathcal{T}^k)_{k \in [n]}$ are *star-convex at 0*, meaning that for any agent k , any type $t^k \in \mathcal{T}^k$, and any scalar $\lambda \in [0, 1]$, there is another type $\lambda t^k \in \mathcal{T}^k$ with the property that $v(\lambda t^k, o) = \lambda v(t^k, o)$ for every $o \in \mathcal{O}$. (The assumption is without loss of generality, as argued in the next paragraph.) The implicit payment computation procedure, applied to type profile \mathbf{t} , samples $\lambda \in [0, 1]$ uniformly at random and computes outcomes $o^0 \triangleq \tilde{\mathcal{A}}(\mathbf{t})$ as well as $o^k \triangleq \tilde{\mathcal{A}}(\lambda t^k, \mathbf{t}^{-k})$ for all $k \in [n]$. The payment charged to

³Our result obviously holds when the doubling dimensions of type spaces are considered to be constant. For arbitrary large-dimensional type spaces, the running time is polynomial in n and $1/\epsilon^\Delta$.

agent k is $v(t^k, o^0) - v(t^k, o^k)$. Note that, in expectation, agent k pays

$$p^k(\mathbf{t}) = v(t^k, \tilde{\mathcal{A}}(\mathbf{t})) - \int_0^1 v(t^k, \tilde{\mathcal{A}}(\lambda t^k, \mathbf{t}^{-k})) d\lambda,$$

in accordance with the payment identity for multi-parameter BIC mechanisms when type spaces are star-convex at 0; see Babaioff et al. [4] for a discussion of this payment identity.

Finally, let us justify the assumption that \mathcal{T}^k is star-convex for all k . This assumption is without loss of generality for the allocation algorithms $\tilde{\mathcal{A}}$ that arise from the RSM reduction, because we can enlarge the type space \mathcal{T}^k if necessary by adjoining types of the form λt^k with $t^k \in \mathcal{T}^k$ and $0 \leq \lambda < 1$. Although the output of the original allocation algorithm \mathcal{A} may be undefined when its input type profile includes one of these artificially adjoined types, the RSM reduction never inputs such a type into \mathcal{A} . It only calls \mathcal{A} on profiles of surrogate types sampled from the type-profile distribution F , whose support excludes the artificially adjoined types. Thus, even when the input to $\tilde{\mathcal{A}}$ includes an artificially adjoined type λt^k , it occurs as one of the replicas in the reduction. The behavior of algorithm $\tilde{\mathcal{A}}$ remains well defined in this case, because replicas are only used as inputs to the valuation function $v(r_i, o_j)$, whose output is well defined even when $r_i = \lambda t^k$ for $\lambda < 1$.

7 SUMMARY AND FUTURE DIRECTIONS

In this article, we investigated the question of designing Bayesian incentive compatible black-box reductions in mechanism design. We provided a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design for welfare maximization problems. Unlike prior results, our reduction achieves exact incentive compatibility for problems with multi-dimensional and continuous type spaces. We showed how to employ and generalize the computational model in the literature on Bernoulli Factories. In particular, we considered a generalization that we called the expectations from samples computational model, in which a problem instance is specified by a function mapping the expected values of a set of input distributions to a distribution over outcomes. The challenge is to give a polynomial time algorithm that exactly samples from the distribution over outcomes given only sample access to the input distributions. In this model, we gave a polynomial time algorithm for the function given by exponential weights: Expected values of the input distributions correspond to the weights of alternatives, and we wish to select an alternative with probability proportional to an exponential function of its weight. As we showed, this algorithm is the key ingredient in designing an incentive compatible mechanism for bipartite matching, which can be used to make the approximately incentive compatible reduction of Hartline et al. [18] exactly incentive compatible.

While the existence of such a reduction is good news for Bayesian mechanism design, there are limitations that are mostly unavoidable.

- *Beyond expected social welfare.* It is tempting to try converting an arbitrary algorithm for an optimization problem into a computationally efficient Bayesian truthful mechanism. Interestingly, this is not possible for all optimization objectives. In particular, Chawla et al. [10] show that no black-box reduction is possible for the objective of makespan, even if we only require Bayesian truthfulness and an average-case performance guarantee. This precludes extending our result beyond the expected-welfare objective in a general fashion.
- *Exponential dependence on dimension.* Notably, our reduction is a fully polynomial time approximation scheme to the reduction of Hartline et al. [18]. However, the running time of Hartline et al. [18] has exponential dependence on Δ . Therefore, our reduction also suffers from the same exponential dependence. Intuitively, this seems to be unavoidable for

reductions that can only access the type space by sampling and can only access the outcome space by calling the allocation function on sampled type profiles.

We conclude with some open questions. The first natural question, directly related to the second limitation above, is to determine whether or not the exponential dependence on Δ in the black-box reduction is unavoidable. Are there black-box reductions whose running time exhibits a milder dependence on the structure of the type space? Another interesting question is to find more connections between Bayesian mechanism design and the expectations from samples computational model. Finally, one might be interested in a generalization of Bernoulli race to more interesting combinatorial settings, e.g., can one sample a base of a matroid given access to marginal coins so the sampling procedure satisfies the marginals *exactly*? It would also be interesting to see if these tools result in simpler black-box reductions.

APPENDIX

A SURROGATE SELECTION AND BIC REDUCTION—FURTHER DETAILS

LEMMA A.1. *If matching algorithm $M(\mathbf{r}, \mathbf{s})$ produces a perfect k -to-1 matching for the instance in Definition 5.3, then its corresponding surrogate selection rule, denoted by Γ^M , is stationary.*

PROOF. Proof of Lemma A.1. Each surrogate s_j is an i.i.d. sample from F . Moreover, by the principle of deferred decisions the index i^* (the real agent's index in the replica type profile) is a uniform random index in $[mk]$, even after fixing the matching. Since this choice of replica is uniform in $[mk]$ and M is a perfect k -to-1 matching, the selection of surrogate outcome is uniform in $[m]$, and therefore the selection of surrogate type associated with this outcome is also uniform in $[m]$. As a result, the output distribution of the selected surrogate type is F . \square

LEMMA A.2. *If $M(\mathbf{r}, \mathbf{s})$ is a feasible replica-surrogate k -to-1 matching and is a truthful allocation rule (in expectation over allocation's random coins) for all replicas (i.e., assuming each replica is a rational agent, no replica has any incentive to misreport), then the composition of Γ^M and interim allocation algorithm $\mathcal{A}(\cdot)$ forms a BIC allocation algorithm for the original mechanism design problem.*

PROOF. Proof of Lemma A.2. Each replica-agent $i \in [mk]$ (including the real agent i^*) bests off by reporting her true replica type under some proper payments. Now, consider an agent in the original mechanism design problem with true type t . For any given surrogate type profile \mathbf{s} , using the Γ^M -reduction the agent receives the same outcome distribution as the one he gets matched to in M in a Bayesian sense, simply because of stationary property of Γ^M (Lemma A.1). As allocation M is incentive compatible, this agent does not benefit from misreporting her true type as long as the value he receives for reporting t' is $v(t, \mathcal{A}(\Gamma^M(t')))$. Therefore, conditioning on \mathbf{s} and non-real replicas in \mathbf{r} , the final allocation is BIC from the perspective of this agent. The lemma then follows by averaging over the random choice of \mathbf{s} and non-real agent replicas in \mathbf{r} . \square

ACKNOWLEDGMENTS

The authors would like to thank Nikhil Devanur, Shipra Agrawal and Pooya Jalaly for helpful discussions and comments on various parts of the article. We want to emphasize that an early conference version of this work was presented at the Annual ACM Symposium on Theory of Computing (STOC'17) [13].

REFERENCES

- [1] Shipra Agrawal and Nikhil R. Devanur. 2015. Fast algorithms for online stochastic convex programming. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1405–1424.

- [2] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. 2009. A dynamic near-optimal algorithm for online linear programming. *arXiv preprint arXiv:0911.2974* (2009).
- [3] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. 2004. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Math.* 1, 2 (2004), 129–150.
- [4] Moshe Babaioff, Robert Kleinberg, and Aleksandrs Slivkins. 2013. Multi-parameter mechanisms with implicit payment computation. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. 35–52.
- [5] Moshe Babaioff, Robert Kleinberg, and Aleksandrs Slivkins. 2015. Truthful mechanisms with implicit payment computation. *J. ACM* 62, 2 (May 2015).
- [6] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. 2013. Bandits with knapsacks. In *Proceedings of the IEEE 54th Symposium on Foundations of Computer Science (FOCS'13)*. IEEE, 207–216.
- [7] Xiaohui Bei and Zhiyi Huang. 2011. Bayesian incentive compatibility via fractional assignments. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 720–733.
- [8] Stephen Boyd and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge University Press.
- [9] Sébastien Chubek. 2014. Convex optimization: Algorithms and complexity. *arXiv preprint arXiv:1405.4980* (2014).
- [10] Shuchi Chawla, Nicole Immorlica, and Brendan Lucier. 2012. On the limits of black-box reductions in mechanism design. In *Proceedings of the 44th ACM Symposium on Theory of Computing*. ACM, 435–448.
- [11] Xiao Alison Chen and Zizhuo Wang. 2015. A dynamic learning algorithm for online matching problems with concave returns. *Eur. J. Op. Res.* 247, 2 (2015), 379–388.
- [12] Nikhil R. Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A. Wilkens. 2011. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. ACM, 29–38.
- [13] Shaddin Dughmi, Jason D. Hartline, Robert Kleinberg, and Rad Niazadeh. 2017. Bernoulli factories and black-box reductions in mechanism design. In *Proceedings of the 49th ACM SIGACT Symposium on Theory of Computing*. 158–169.
- [14] Theodore Groves. 1973. Incentives in teams. *Econometrica: J. Economet. Soc.* 41, 4 (1973), 617–631.
- [15] Jason D. Hartline and Brendan Lucier. 2010. Bayesian algorithmic mechanism design. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*. ACM, 301–310.
- [16] Jason D. Hartline and Brendan Lucier. 2015. Non-optimal mechanism design. *Amer. Econ. Rev.* 105, 10 (2015), 3102–3124.
- [17] Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. 2011. Bayesian incentive compatibility via matchings. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*.
- [18] Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. 2015. Bayesian incentive compatibility via matchings. *Games. Econ. Behav.* 92, C (2015), 401–429.
- [19] Elad Hazan. 2016. Introduction to online convex optimization. *Found. Trends Optim.* 2, 3–4 (2016), 157–325.
- [20] Zhiyi Huang and Sampath Kannan. 2012. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. In *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*. IEEE, 140–149.
- [21] Mark Huber. 2015. Optimal linear Bernoulli factories for small mean problems. *CoRR* abs/1507.00843 (2015).
- [22] Adam Kalai and Santosh Vempala. 2005. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.* 71, 3 (2005), 291–307.
- [23] M. S. Keane and George L. O'Brien. 1994. A Bernoulli factory. *ACM Trans. Model. Comput. Simul.* 4, 2 (1994), 213–219.
- [24] Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. 2014. Primal beats dual on online packing LPs in the random-order model. In *Proceedings of the 46th ACM Symposium on Theory of Computing*. ACM, 303–312.
- [25] Krzysztof Latuszyński. 2010. The Bernoulli factory, its extensions and applications. In *Proceedings of the International Workshop for Applied Public Key Infrastructure*. 1–5.
- [26] Colin McDiarmid. 1989. On the method of bounded differences. *Surv. Combinat.* 141, 1 (1989), 148–188.
- [27] Rajeev Motwani and Prabhakar Raghavan. 2010. *Randomized Algorithms*. Chapman & Hall/CRC.
- [28] Șerban Nacu and Yuval Peres. 2005. Fast simulation of new coins from old. *Ann. Appl. Prob.* 15, 1A (2005), 93–115.
- [29] Noam Nisan and Amir Ronen. 2001. Algorithmic mechanism design. *Games Econ. Behav.* 35, 1-2 (2001), 166–196.
- [30] Shai Shalev-Shwartz et al. 2012. Online learning and online convex optimization. *Found. Trends® Mach. Learn.* 4, 2 (2012), 107–194.

Received May 2018; revised June 2020; accepted November 2020

Copyright of Journal of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.