Coordination Complexity: Small Information Coordinating Large Populations

[Extended Abstract]^{*}

Rachel Cummings[†]
Caltech

Katrina Ligett[‡]
Caltech

Jaikumar Radhakrishnan
Tata Institute of Fundamental
Research

Aaron Roth[§]
University of Pennsylvania

Zhiwei Steven Wu[¶]
University of Pennsylvania

ABSTRACT

We study a quantity that we call *coordination complexity*. In a distributed optimization problem, the information defining a problem instance is distributed among n parties, who need to each choose an action, which jointly will form a solution to the optimization problem. The coordination complexity represents the minimal amount of information that a centralized coordinator, who has full knowledge of the problem instance, needs to broadcast in order to coordinate the n parties to play a nearly optimal solution.

We show that upper bounds on the coordination complexity of a problem imply the existence of good jointly differentially private algorithms for solving that problem, which in turn are known to upper bound the price of anarchy in certain games with dynamically changing populations.

We show several results. We fully characterize the coordination complexity for the problem of computing a many-to-one matching in a bipartite graph by giving almost matching lower and upper bounds. Our upper bound in fact extends much more generally, to the problem of solving a linearly separable convex program. We also give a different upper

bound technique, which we use to bound the coordination complexity of coordinating a Nash equilibrium in a routing game, and of computing a stable matching.

CCS Concepts

•Theory of computation → Communication complexity; Approximation algorithms analysis; Algorithmic game theory;

Keywords

coordination complexity, privacy

1. INTRODUCTION

In this paper, we study a quantity which we call *coordination complexity*. This quantity measures the amount of information that a centralized coordinator needs to broadcast in order to coordinate n parties, each with only local information about a problem instance, to jointly implement a globally optimal solution. Unlike in *communication complexity*, there is no need for the communication protocols in our setting to derive the optimal solution starting with nothing but local information, nor even *verify* that a proposed solution is optimal (as is the goal in non-deterministic communication complexity). Instead, in our setting, there is a central coordinator who already has complete knowledge of the problem instance — and hence also of the optimal solution. His goal is simply to publish a concise message to guide the n parties making up the problem instance to coordinate on the desired solution — ideally using fewer bits than would be (trivially) needed to simply publish the optimal solution itself.¹

Aside from its intrinsic interest, our motivation for studying this quantity is two-fold. First, as we show, problems with low coordination complexity also have good protocols for implementing nearly optimal solutions under the constraint of *joint differential privacy* [DMNS06, KPRU14] — i.e. protocols that allow the joint implementation of a nearly

^{*}A full version of this paper is available on arXiv at <http://arxiv.org/abs/1508.03735>

[†]Supported by Simons Award for Graduate Students in Theoretical Computer Science and NSF CNS-1254169.

[‡]Supported in part by NSF grant CNS-1254169, NSF grant CNS-1518941 US-Israel Binational Science Foundation grant 2012348, the Charles Lee Powell Foundation, a Google Faculty Research Award, an Okawa Foundation Research Grant, a Microsoft Faculty Fellowship.

[§]Supported in part by NSF Grant CCF-1101389, an NSF CAREER award, and an Alfred P. Sloan Foundation Fellowship.

[¶]Supported in part by NSF Grant CCF-1101389.

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 ACM 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.

ITCS'16, January 14–16, 2016, Cambridge, MA, USA

© 2016 ACM. ISBN 978-1-4503-4057-1/16/01...\$15.00

DOI: <http://dx.doi.org/10.1145/2840728.2840767>

¹Within our framework of coordination complexity, we assume that the players are not strategic — they will faithfully follow the coordination protocol upon observing the message broadcast by the coordinator. We do study the interface between the coordination complexity and the strategic variants of some problems in Section 5.

optimal solution in a manner such that no coalition of parties can learn much about the portion of the instance known by any party not in the coalition. The existence of jointly differentially private protocols in turn have recently been shown to imply a low “price of anarchy” for no-regret players in the strategic variant of the optimization problem when the game in question is smooth – even when the population is dynamically changing [LST15]. Hence, as a result of the connection we develop in this paper, in order to show dynamic price of anarchy bounds of the sort given in [LST15], it is sufficient to show that the game in question has low coordination complexity, without needing to directly develop and analyze differentially private algorithms. Using this connection we also derive new results for what can be implemented under the constraint of pure joint differential privacy – results that were previously only known subject to approximate joint differential privacy.

Second, coordination complexity is a stylized measure of the power of concise broadcasts (e.g. prices in the setting of allocation problems, or congestion information in the setting of routing problems) to coordinate populations in the absence of any interaction.² Here we note that prices seem to coordinate markets, despite the fact that individuals do not actually participate in any kind of interactive “Walrasian mechanism” of the sort that would be needed to compute the allocation itself, in addition to the prices (see e.g. [KC82, DNO14]). Indeed, prices alone are generally not sufficient to coordinate high welfare allocations because prices on their own can induce a large number of indifferences that might need to be resolved in a coordinated way – and hence Walrasian equilibria are defined not just as vectors of equilibrium prices, but as vectors of prices paired with optimal allocations. Publishing a Walrasian equilibrium would be a trivial solution in our setting, because it involves communicating the entire solution that we wish to coordinate – the optimal allocation. Nevertheless, we show that the coordination complexity of the allocation problem is – up to log factors – equal to the number of types of goods in a commodity market. This is the same as what would be needed to communicate prices (indeed, our solution can be viewed as communicating prices in a slightly different, “regularized” market), and can be substantially smaller than what would be needed to communicate the optimal allocation itself.

1.1 Our Results and Techniques

In our model (which we formally define in Section 2), a *problem instance* D is defined by an n -tuple from some abstract domain $\mathcal{X}$: $D \in \mathcal{X}^n$. We write $D = (D^{(1)}, \dots, D^{(n)})$ to denote the fact that the information defining the problem instance is partitioned among n agents, and each agent i knows only his own part $D^{(i)}$. The solution space is also a product space: $\mathcal{A}^n$, and each agent i can choose a single *action* $a_i \in \mathcal{A}$ – the choices of all of the agents jointly form a solution $a = (a_1, \dots, a_n)$. The *coordinator* knows the entire problem instance D , and publishes a signal $\sigma(D) \in \{0, 1\}^\ell$. Each agent then chooses an action $a_i := \pi(D^{(i)}, \sigma(D))$ based only on the coordinator’s signal and her own part of the problem instance. The jointly induced solution $a = (a_1, \dots, a_n)$

²Of course, the connection here is in a stylized model – in a market, there is not in fact any party with complete information of the problem instance – but the market is nevertheless encoding good “distributional information” about the population of buyers likely to arrive.

is the output of the interaction. The pair of functions σ, π jointly form a protocol, and ℓ , the length of the coordinator’s signal is the coordination complexity of the protocol. The coordination complexity of a problem is the minimal coordination complexity of any protocol solving the problem.

A canonical example to keep in mind is many-to-one matchings: Here, a problem instance is defined by a bipartite graph between n agents and k types of goods. Each good j has a supply s_j , and the goal is to find a maximum cardinality matching such that no agent is matched to more than one good, and no good j is matched to more than s_j agents. Here, the portion of the instance known to agent i is the set of goods adjacent to agent i – but nothing about the goods adjacent to other agents. Note that describing a matching requires $\Omega(n \log k)$ bits, which is the trivial upper bound on the coordination complexity for this problem. For this problem, we show nearly matching upper and lower bounds: no protocol with coordination complexity $o(k)$ can guarantee a constant approximation to the optimal solution, whereas there is a protocol with coordination complexity $O(k \log n)$ that can obtain a $(1 + o(1))$ -approximation to the optimal solution. Our upper bound in fact extends much more generally, to any problem that can be written down as a convex program whose objective and constraints are linearly separable between agents’ data.

The idea of the upper bound is to broadcast a portion of the optimal dual solution to the convex program – one dual variable for every constraint that is defined by the data of multiple agents (there is no need to publish the dual variables corresponding to constraints that depend only on the data of a single agent). For the many-to-one matching problem, these dual variables correspond to “prices” – one for each of the k types of goods. This idea on its own does not work, however, because a dual optimal solution to a convex program is not generally sufficient to specify the primal optimal solution. When specialized to the case of matchings, this is because optimal “market clearing prices” can induce a large number of indifferences among goods for each of the n agents, and these indifferences might need to be broken in a coordinated way to induce an optimal matching. To solve this problem, we instead release the dual variables corresponding to a slightly different convex program, in which a *strongly convex regularizer* has been added to the objective. The effect of the strongly convex regularizer is that the optimal dual solution now uniquely specifies the optimal primal solution – although now the optimal primal solution to a modified problem. The rest of our approach deals with trading off the weight of the regularizer with the number of bits needed to approximately specify each of the dual variables, and the error of the regularized optimal solution relative to the optimal solution to the original problem.

We also give several other positive results, based on a different technique: broadcasting the truncated transcript of a process known to converge to a solution of interest. Using this technique, we give low coordination complexity protocols for the problem of coordinating on an equilibrium in a routing game, and for the problem of coordinating on a *stable* many-to-one matching.

Finally, we show that problems that have both low sensitivity objectives (as all of the problems we study do) and low coordination complexity also have good *jointly differentially private* protocols. Using the results of [LST15], this also shows a bound on the price of anarchy of the strategic

variant of these problems, whenever they are *smooth games*, which holds even under a dynamically changing population.

1.2 Related Work

Our model of coordination complexity is related to, but distinct from, the well-studied notion of communication complexity – see [KN97] for a textbook introduction. While both complexity notions measure the number of bits that must be transmitted among decentralized parties to reach a particular outcome, they differ in the initial endowment of information, as well as in the requirements of each player to know the final outcome. In communication complexity, the information describing the problem instance is fully distributed, and communication is necessary for all parties to know the outcome. Coordination complexity in contrast assumes the existence of a coordinator who knows the entire problem instance, and must broadcast information to the players which will allow them to each compute their part of the output – there is no need for any of the parties to know the entire output. More similar to our setting is *non-deterministic communication complexity*, in which we may imagine that there is an oracle who knows the inputs of all players and broadcasts a message (perhaps partially) describing a solution together with a certificate that allows the parties to verify the optimality of the solution. In contrast, in our model of coordination complexity, the coordinator does not need to provide any certificate allowing parties to verify that the coordinated solution is optimal (indeed, each party need not have any information about the portion of the solution proposed to other parties).

The informational requirements of coordinating matchings has a long history of study in economics, and has recently gained attention in theoretical computer science. Hayek’s classic paper [Hay45] conjectured that Walrasian price mechanisms, which coordinate matchings via a tâtonnement process that updates market-clearing prices based on demand, are “informationally efficient,” in that they verify optimal allocations with the least amount of information. This was later formalized by [Hur60] and [MR74] in specific settings of interest, using an informational metric that measured smooth real-valued communication. Nisan and Segal study the communication complexity of matchings using the tools of communication complexity as developed in computer science, and show that any communication protocol that determines an optimal allocation must also determine supporting prices [NS06]. Recently, [DNO14] and [ANRW15] studied the problem of computing an optimal matching through the lens of interactive communication complexity, showing that interactive protocols can have significantly lower communication complexity than non-interactive ones. Note that the communication complexity bounds given in these papers are always larger than the description length of the matching itself – in contrast, here when we study coordination complexity, nontrivial bounds must not just be smaller than the input, but must also be smaller than the size of the optimal matching.

Finally, there are two papers that study a very similar setting to ours, although they obtain rather different results.³ Calsamiglia [Cal84] studies a real-valued communication model in which a central coordinator with full knowledge of the instance needs to broadcast a concise signal to

coordinate an allocation in an exchange market—see [Seg06] for context on how this result fits into the economic literature on communication complexity. Deng, Papadimitriou, and Safra also study a similar model in Section 4 of [DPS02], which they call “Market Communication”. Despite the similarity in models, the results of both [Cal84] and [DPS02] stand in sharp contrast to ours—they both give *lower bounds*, showing that the amount of communication necessary needs to grow linearly with the number of buyers n , while we give upper bounds showing that it is necessary to grow only with the number of different types of goods k . Calsamiglia does not allow approximation in his model, which is necessary for our results. Deng, Papadimitriou, and Safra allow for approximation, but study an instance of a problem that cannot be expressed as a linearly separable convex program, which shows that structure of the sort that we use is necessary.

A line of work [KPRU14, RR14, CKRW14, HHR⁺14, HHRW14, RRW15] has studied protocols for implementing outcomes in various settings under the constraint of joint differential privacy [DMNS06, KPRU14], which allows n parties to jointly implement some solution while ensuring that no coalition of parties can learn much about the input of any party outside the coalition. Most (but not all) of these algorithms are actually private coordination protocols of the sort we study here, in which the algorithm can be viewed as a coordinator who is constrained to broadcast a private signal. These jointly private algorithms are not constrained to transmit a short signal – and indeed, the private signals can sometimes be verbose. But as we show, problems with low coordination complexity also have good jointly differentially private algorithms, which was one of our original motivations for studying this quantity.

Lykouris, Syrgkanis, and Tardos [LST15] show that the existence of a jointly differentially private algorithm for solving an optimization problem implies that the strategic variant of the problem has a low “price of anarchy” for learning agents, even in dynamic settings, in which player types change over time, as long as the game is smooth. Because we show in Section 5 that any problem with a low sensitivity objective and low coordination complexity has a good jointly differentially private algorithm, using the results of [LST15], to prove a bound on the price of anarchy in a smooth dynamic game, we show it suffices to bound the coordination complexity of the game.

2. PRELIMINARIES

A *coordination problem* is defined by a set of n agents, a data domain $\mathcal{X}$, an action range $\mathcal{A}$, and a social objective function $S: \mathcal{X}^n \times \mathcal{A}^n \rightarrow \mathbb{R}$. An *instance* of a coordination problem consists of a set of n elements from the data domain: $D = (D^{(1)}, \dots, D^{(n)}) \in \mathcal{X}^n$. Each agent i has knowledge only of $D^{(i)}$, his own portion of the problem instance, and the goal is for a centralized *coordinator* to broadcast a concise message to the agents to allow them to arrive at a solution $a = (a_1, \dots, a_n) \in \mathcal{A}^n$ that approximately maximizes the objective function $S(D, a)$.

A *coordination protocol* consists of two functions, an *encoding function* $\sigma: \mathcal{X}^n \rightarrow \{0, 1\}^*$ and a *decoding function* $\pi: \mathcal{X} \times \{0, 1\}^* \rightarrow \mathcal{A}$. A coordination protocol (σ, π) proceeds in two stages:

- First the coordinator broadcasts the message $\sigma(D)$ to all agents using the encoding function.

³We thank Ilya Segal for pointing out [Cal84] to us, and thank Sepehr Assadi for pointing out Section 4 of [DPS02].

- Then each agent selects an action a_i on the basis of her own portion of the problem instance and the broadcast message, using the decoding function: $a_i := \pi(D^{(i)}, \sigma(D))$.

Both functions σ and π may be randomized. The approximation ratio of a protocol is the ratio of the optimal objective value to the expected objective value of the solution induced by the protocol, in the worst case over problem instances.

DEFINITION 1 (APPROXIMATION RATIO). A coordination protocol (σ, π) obtains a ρ approximation to a problem if:

$$\max_{D \in \mathcal{X}^n} \frac{\text{OPT}(D)}{\mathbb{E}_{a_1, \dots, a_n} [S(D, a)]} \leq \rho$$

where each $a_i = \pi(D^{(i)}, \sigma(D))$, and the expectation is taken over the randomness of σ and π .

The coordination complexity of a protocol is the maximum number of bits the encoding function broadcasts, in the worst case over problem instances.

DEFINITION 2 (COORDINATION COMPLEXITY). A coordination protocol (σ, π) has coordination complexity ℓ if:

$$\max_{D \in \mathcal{X}^n} |\sigma(D)| = \ell.$$

The coordination complexity of obtaining a ρ approximation to a problem is the minimum value of the coordination complexity of all protocols (σ, π) that obtain a ρ approximation to the problem.

We conclude by making several observations about coordination protocols. First, as we have defined them, they are *non-interactive* – the coordinator first broadcasts a signal, and then the agents respond. This is without loss of generality, since the coordinator has full knowledge of the problem instance. Any interactive protocol could be reduced at no additional communication cost to a non-interactive protocol, simply by having the coordinator publish the transcript that would have arisen from the interactive protocol. This is in contrast to the setting of communication complexity, in which interactive protocols can be more powerful than non-interactive protocols (and makes it easier to prove lower bounds for coordination complexity).

Second, the coordination complexity of a problem is trivially upper bounded both by the description length of the problem instance (as is communication complexity), and by the description length of the problem's optimal solution (unlike in non-deterministic communication complexity, there is no need to pair the optimal solution with a certificate allowing individual agents to verify it). Hence, non-trivial bounds will be asymptotically smaller than both of these quantities.

Bipartite Matching.

The primary coordination problem we study in this paper is the *bipartite matching problem*. In this problem, there is a bipartite graph $G = (V, W, E)$, in which every node in V is associated with a player and every node in W represents a good. Each player i 's private data is the set of edges incident to her node – i.e. $D^{(i)} = \{j : (i, j) \in E\}$. We study two variants of this problem. In the *one-to-one* matching problem, W represents a set of distinct goods, and the goal is to coordinate a maximum cardinality matching $E' \subseteq E$ such that for every $i \in V$, $|\{j \in W : (i, j) \in E'\}| \leq 1$

and for every $j \in W$, $|\{i \in V : (i, j) \in E'\}| \leq 1$. In the *many-to-one* matching problem, W represents a set of k commodities j , each with a supply b_j . The goal is to coordinate a maximum cardinality many-to-one matching $E' \subseteq E$ such that for every $i \in V$, $|\{j \in W : (i, j) \in E'\}| \leq 1$ and for every $j \in W$, $|\{i \in V : (i, j) \in E'\}| \leq b_j$. The social objective in this setting is the welfare or the cardinality of the matching, and we will use $\text{OPT}(G)$ to denote the optimal welfare objective.

Note that the resulting solution might not be feasible since the players' demands are not always satisfied. We need to make sure that we are not over-counting when measuring the welfare. In one-to-one matchings, if more than one players select a good, only the first player is matched to it. In many-to-one matchings, if more than b_j players select a good of type j , only the first b_j players are matched the good j .

Notation.

We use $\|\cdot\|$ to denote the ℓ_2 norm, and more generally use $\|\cdot\|_p$ to denote the ℓ_p norm.

3. LOWER BOUNDS FOR MATCHINGS

In this section, we present lower bounds on the coordination complexity of bipartite matching problems. As a building block, we prove a lower bound for the one-to-one matching problem on a bipartite graph with n vertices on each side, showing an $\Omega(n)$ lower bound – i.e. that no substantial improvement on the trivial solution is possible. We then extend this lower bound to the problem of many-to-one matchings, in which there are n agents who must be matched to k goods (each good can be matched to many agents, up to its supply). Here, we show an $\Omega(k)$ lower bound. In the next section, we give a nearly matching upper bound, which substantially improves over the trivial solution.

3.1 A Variant of the Index Function Problem

Before we present our lower bound, we introduce a variant of the *random index function* problem [KNR99], which will be useful for our proof.

MULTIPLE-INDEX.

There are two players Alice and Bob. Alice receives as input a sequence of t pairs, $I = \langle (S_i, u_i) : i = 1, 2, \dots, t \rangle$, where the S_i are disjoint sets each with k elements, and u_i is uniformly distributed in S_i . Based on her input Alice sends Bob a message $M(I)$. Bob then receives (S_j, j) , where j is chosen from $[t]$. Bob must determine u_j ; let his output be $B(S_j, j, M(I)) \in S_j$. We say that the protocol succeeds if $B(S_j, j, M(I)) = u_j$. Let $\ell(t, k, p)$ be the minimum number of bits (for the worst input) that Alice must send in order for Bob to succeed with probability at least p .

Note that if Bob guesses randomly, then the protocol already succeeds with probability $p = 1/k$. The following result shows that any significant improvement over this trivial probability of success will require Alice to send Bob a long message. See the full version for a full proof.

LEMMA 1. For $p \geq 1/k$, we have $\ell(t, k, p) \geq (8 \log e)t(p - 1/k)^2$.

3.2 Lower Bound for One-to-One Matchings

We will first focus on the lower bound on one-to-one matching and show the following.

THEOREM 1. *Suppose the coordination protocol Π for one-to-one matching guarantees an approximation ratio of ρ in expectation. Then, the coordinator of Π must broadcast $\Omega(n/\rho^4)$ bits on problem instances of size n (in the worst case).*

Fix the protocol Π . We will extract a two-party communication protocol for the **MULTIPLE-INDEX** problem from Π , and use the above lemma. As a first step for our lower bound proof, we will consider the following random graph construction process **RanG**.

Random Graph Construction RanG(ρ, n):.

Let $\kappa = \frac{n}{8\rho}$ and $A = \frac{n}{16\rho^2}$. Consider the following random bipartite graph G with vertex set (V, W) such that $|V| = |W| = n$.

- Randomly generate an ordering $w_1, w_2, \dots, w_n$ of W (all $n!$ orderings being equally likely), and partition W as $W_1 \cup W_2$ such that $W_1 = \{w_1, w_2, \dots, w_\kappa\}$, and $W_2 = \{w_{\kappa+1}, w_{\kappa+2}, \dots, w_n\}$.
- Similarly, randomly generate an ordering $v_1, v_2, \dots, v_n$ of V , and partition V into n/A blocks, $B_1, B_2, \dots, B_{n/A}$ (each with A vertices), where

$$B_j := \{v_i : (j-1)A + 1 \leq i \leq jA\}.$$

- Connect B_j and W as follows. First, we describe the connections between V and W_1 . The neighbourhoods of the vertices in each B_j will be disjoint: we partition W_1 into equal-sized disjoint sets $(T_v : v \in B_j)$, and let the neighbours of $v \in B_j$ be exactly the 2ρ vertices in T_v .
- In addition, assign each vertex in v one neighbor in W_2 , by connecting V with W_2 in round-robin fashion — connect vertex v_i to vertex w_j , where $j = (\kappa + i \bmod (n - \kappa))$.

Before we prove Theorem 1, let us first observe that a graph generated by **RanG** always has a matching with high welfare.

LEMMA 2. *Each graph G generated by the above process **RanG**($\cdot, n$) has optimal welfare $\text{OPT}(G) \geq \frac{7n}{8}$.*

PROOF PROOF OF THEOREM 1. Let Π be a coordination protocol with coordination complexity ℓ and approximation ratio ρ . This means that on a graph instance generated by **RanG**, the parties can coordinate on a matching with expected weight at least $\frac{7n}{8\rho}$. Since $|W_1| = \frac{n}{8\rho}$, we know in expectation at least $\frac{7n}{8\rho} - \frac{n}{8\rho} = \frac{3n}{4\rho}$ of the vertices in W_2 are matched. Let α_v be the probability that in Π , agent v picks her neighbor in W_2 . Then, by linearity of expectation, $\sum_{v \in V} \alpha_v \geq \frac{3n}{4}$, that is, $\mathbb{E}_{v \in V}[\alpha_v] \geq \frac{3}{4}$. Hence, there must be some block B_j such that $\mathbb{E}_{v \in B_j}[\alpha_v] \geq \frac{3}{4}$. We will now restrict attention to the block B_j and consider the following instance of **MULTIPLE-INDEX**: for each $v \in B_j$, set $S_v = N(v)$ —the neighborhood of vertex v , and let u_v be the vertex unique vertex in $S_v \cap W_2$. Since the message broadcast by the coordination protocol allows the players to identify the special element with average success probability

of $\frac{3}{4\rho}$, by Lemma 1 the length of message

$$\begin{aligned} \ell &\geq (8 \log e) |B_j| \left(\frac{3}{4\rho} - \frac{1}{2\rho+1} \right)^2 \\ &\geq (8 \log e) \left(\frac{n}{16\rho^2} \right) \left(\frac{3}{4\rho} - \frac{1}{2\rho+1} \right)^2 \geq \Omega\left(\frac{n}{\rho^4}\right), \end{aligned}$$

which completes the proof. $\square$

3.3 Lower Bound for Many-to-One Matchings

Finally, we give the following lower bound on coordination complexity for many-to-one matchings. The lower bound relies on the result from Section 3.2—we show that any coordination protocol for many-to-one matchings can also be reduced to a protocol for one-to-one matchings, and so the lower bound in Section 3.2 can be extended to give a lower bound for the many-to-one setting.

For our lower bound instance, we consider bipartite graphs $G = (V, E)$ such that the vertices in V represent n different players and W represent a set of k goods j , each with a supply b .

THEOREM 2. *Suppose that there exists a coordination protocol for many-to-one matchings that guarantees an approximation ratio of ρ in expectation. Then such a protocol has coordination complexity of $\Omega(k/\rho^4)$.*

We will start by considering a one-to-one matching instance generated by **RanG** with k vertices on each side of the graph $G' = (V', W', E')$. By Lemma 2, the optimal matching of G' has size $\text{OPT}' \geq \frac{7k}{8}$.

Now we will turn this into an instance of a many-to-one matching problem: make b copies of each vertex in V' to obtain vertex set V , and set $W := W'$ such that the supply of each good j is b ; then for an edge (v', w') in the original graph, connect all copies of v' to w' in the new graph. This gives a bipartite graph $G = (V, W, E)$. The following claim is straightforward.

CLAIM 1. *The new graph G has a matching of size at least $(b \text{OPT}')$.*

Now suppose that we could coordinate the players in V to obtain a matching M^* of size $\frac{b \text{OPT}'}{\rho}$ in G . Then with a simple sampling procedure, we can extract a high cardinality matching for the original graph: for each vertex in $v' \in V'$, sample one of the b copies of v' in G uniformly at random along with its incident matched edge. If two vertices in V' are connected to the same type of good in W' , break ties arbitrarily and keep only one of the edges.

LEMMA 3. *The sampled matching in G' has expected size at least $\frac{\text{OPT}'}{3\rho}$.*

We will defer the proof to the full version. We now have all the pieces to prove Theorem 2.

PROOF PROOF OF THEOREM 2. Suppose that there exists a coordination protocol (σ, π) for many-to-one matchings with a guaranteed approximation ratio of ρ . By the result of Lemma 3, we know that this coordination protocol for one-to-one matchings has an approximation ratio most $O(\rho)$. By the lower bound in Theorem 1, we know that the length of $\sigma(G)$ is at least $\Omega(k/\rho^4)$. $\square$

4. COORDINATION PROTOCOL FOR LINEARLY SEPARABLE CONVEX PROGRAMS

In this section, we give a coordination protocol for problems which can be expressed as linearly separable convex programs, with coordination complexity scaling only with the number of constraints that bind between agents (so called *coupling constraints*, defined below). In the next section, we show how to specialize this protocol to the special case of many-to-one matchings, which gives coordination complexity nearly matching our lower bound.

DEFINITION 3. *A linearly separable convex optimization problem consists of n players and for each player i ,*

- *a compact and bounded convex feasible set $\mathcal{F}^{(i)} \subseteq \{x^{(i)} \in \mathbb{R}^l \mid \|x^{(i)}\| \leq 1\}$,*
- *a concave objective and 1-Lipschitz function $v^{(i)}: \mathcal{F}^{(i)} \rightarrow \mathbb{R}$ such that $v^{(i)}(\mathbf{0}) = 0$,*
- *and k convex constraint function $c_j^{(i)}: \mathcal{F}^{(i)} \rightarrow [0, 1]$ (indexed by $j = 1, \dots, k$).*

The convex optimization problem is:

$$\begin{aligned} & \max_x \sum_{i=1}^n v^{(i)}(x^{(i)}) \\ & \text{subject to } \sum_{i=1}^n c_j^{(i)}(x^{(i)}) \leq b_j \quad \text{for } j = 1, \dots, k \\ & \quad x^{(i)} \in \mathcal{F}^{(i)} \quad \text{for } i = 1, \dots, n \end{aligned}$$

where each player i controls the block of decision variable $x^{(i)}$.

Viewed as a coordination problem, the data held by each agent i is $D^{(i)} = \{\mathcal{F}^{(i)}, v^{(i)}, c_1^{(i)}, \dots, c_k^{(i)}\}$, his action range is $\mathcal{A}_i = \mathcal{F}^{(i)}$, and the social objective function is S the objective of the convex program. We will call the first set of constraints the *coupling constraints*, and the second set of constraints the *personal constraints*.

We will denote the product of the personal constraints by $\mathcal{F} = \mathcal{F}^{(1)} \times \dots \times \mathcal{F}^{(n)}$, the objective function by $v(x)$, and the optimal value by OPT . In this notation we can write the problem as

$$\max_{x \in \mathcal{F} \text{ and } \sum_{i=1}^n c_j^{(i)}(x^{(i)}) \leq b_j \text{ for all } j} v(x).$$

Note that here the problem is constrained both by the personal constraints $\mathcal{F}$ and by the coupling constraints. We will assume the problem above is feasible and our goal is coordinate the players to play an aggregate solution $x = (x^{(i)})_{i \in [n]}$ that is approximately feasible and optimal. Our solution consists of two steps:

1. We will first introduce a regularization term $\eta\|x\|^2$ to our objective function, and coordinate the players to maximize the regularized objective. The purpose of adding this regularization term is to make the objective function *strongly concave*, which will cause it to have the property that an optimal dual solution will uniquely specify the optimal primal solution.

2. Then we will show that the resulting optimal solution to the regularized problem is close to being optimal for the original (unregularized) problem. The weight of the regularization has to be traded off against the bit precision to which we need to communicate the optimal dual variables.

4.1 Coordination through Regularization

In the first step, we add a small regularization term to our original objective function. Consider the following convex optimization problem:

$$\max_{x \in \mathcal{F} \text{ and } \sum_{i=1}^n c_j^{(i)}(x^{(i)}) \leq b_j \text{ for all } j} v'(x) = \sum_{i=1}^n v^{(i)}(x^{(i)}) - \frac{\eta}{2}\|x\|^2$$

CLAIM 2. *The objective function v' is η -strongly concave.*

To solve the convex program, we will work with the partial *Lagrangian* $\mathcal{L}(x, \lambda)$, which results from bringing only the coupling constraints into the objective via Lagrangian dual variables, but leaving the personal constraints to continue to constrain the primal feasible region:

$$\sum_{i=1}^n \left(v^{(i)}(x^{(i)}) - \frac{\eta}{2}\|x^{(i)}\|^2 - \sum_{j=1}^k \lambda_j \left(\sum_{i=1}^n c_j^{(i)}(x^{(i)}) - b_j \right) \right)$$

Let OPT' denote the optimum of the convex program, and by *strong duality* we have

$$\max_{x \in \mathcal{F}} \min_{\lambda \in \mathbb{R}_+^k} \mathcal{L}(x, \lambda) = \min_{\lambda \in \mathbb{R}_+^k} \max_{x \in \mathcal{F}} \mathcal{L}(x, \lambda) = \text{OPT}'$$

Fixing the optimal dual variables, λ , the optimal primal solution y satisfies

$$y = \operatorname{argmax}_{x \in \mathcal{F}} \mathcal{L}(x, \lambda)$$

Note that the result of moving the coupling constraints into the Lagrangian is that we can now write the primal optimization problem over a feasible region defined only by the personal constraints. Because of this fact, and because the Lagrangian objective is linearly separable across players, given λ , each player's portion of the solution $y^{(i)}$ is

$$\operatorname{argmax}_{x \in \mathcal{F}_i} \left[v^{(i)}(x^{(i)}) - \frac{\eta}{2}\|x^{(i)}\|^2 - \sum_{j=1}^k \lambda_j c_j^{(i)}(x^{(i)}) \right]. \quad (1)$$

Thus, if the argmax were unique, this means that the optimal dual variables λ would be sufficient to coordinate each of the parties to find their portion of the optimal solution, without the need for further communication (the problem, in general, is that the argmax need not be unique, and ties may need to be broken in a coordinated fashion). However, because we have added a strongly concave regularizer, the argmax is unique in our setting:

CLAIM 3. *The solution to*

$$\operatorname{argmax}_{x \in \mathcal{F}_i} v^{(i)}(x^{(i)}) - \frac{\eta}{2}\|x^{(i)}\|^2 - \sum_{j=1}^k \lambda_j \left(\sum_{i=1}^n c_j^{(i)}(x^{(i)}) \right)$$

is unique.

This gives rise to our simple coordination mechanism **ReC**. The mechanism first computes the optimal dual variables in our regularized partial Lagrangian problem, rounds them to

finite precision, and then publishes these variables. Then each individual player finds her part of the near optimal solution by performing the optimization in Equation (1). The details are in Algorithm 4.1.

Algorithm 1 Coordination Protocol for Linearly Separable Convex Programs $\text{ReC}(\eta, \varepsilon)$

Input: a linearly separable convex program instance I , regularization parameter η and target accuracy ε

Initialize: $\alpha = \frac{\eta\varepsilon^2}{4\sqrt{nk}}$

Modify the objective of I into

$$\max_{x \in \mathcal{F}} v(x) - \frac{\eta \|x\|^2}{2}$$

Compute the optimal dual solution $\lambda^\bullet$ for the modified convex program

Round each coordinate of $\lambda^\bullet$ into a multiple of $\alpha/\sqrt{k}$ and obtain $\hat{\lambda}$

Broadcast the rounded dual solution $\hat{\lambda}$. To decode, each player i computes $\hat{x}^{(i)}(\lambda)$:

$$\arg\max_{x \in \mathcal{F}_i} \left[v^{(i)}(x^{(i)}) - \frac{\eta}{2} \|x^{(i)}\|^2 - \sum_{j=1}^k \hat{\lambda}_j c_j^{(i)}(x^{(i)}) \right]$$

Next we show that the resulting solution is close to the optimal solution of the regularized convex program (i.e. that we do not lose much by truncating the dual variables to have finite bit precision). Let $(x^\bullet, \lambda^\bullet)$ be an optimal primal-dual pair for the regularized convex program. Note that since the objective of the program is strongly concave, $x^\bullet$ is unique. First, we will show that if the broadcast dual vector $\hat{\lambda}$ is close to an optimal dual solution $\lambda^\bullet$, the resulting solution $\hat{x}$ will also be close to the optimal primal solution $x^\bullet$.

LEMMA 4. Suppose we have a dual vector $\hat{\lambda}$ such that $\|\lambda^\bullet - \hat{\lambda}\| \leq \alpha$. Let $\hat{x} = \arg\max_{x \in \mathcal{F}} \mathcal{L}(x, \hat{\lambda})$, then

$$\|\hat{x} - x^\bullet\| \leq \frac{2\sqrt{\alpha}(nk)^{1/4}}{\sqrt{\eta}}$$

The proof relies on some basic properties of the Lagrangian and strong concavity and is deferred to the full version.

LEMMA 5. The coordination mechanism ReC instantiated with regularization parameter η and target accuracy parameter ε will coordinate the players to play a solution $\hat{x}$ that satisfies $\|\hat{x} - x^\bullet\| \leq \varepsilon$, and has a coordination complexity of $O(k \log(nk/\eta\varepsilon))$.

PROOF. Note that $\alpha = \frac{\eta\varepsilon^2}{4\sqrt{nk}}$, and the mechanism rounds each coordinate of the optimal dual solution $\lambda^\bullet$ to a multiple of $\alpha/\sqrt{k}$, so the approximate dual vector $\hat{\lambda}$ can be specified with $O(k \log(\sqrt{k}/\alpha))$ bits.

Since for each coordinate j , $|\lambda_j^\bullet - \hat{\lambda}_j| \leq \alpha/\sqrt{k}$, we also have that $\|\lambda^\bullet - \hat{\lambda}\| \leq \alpha$. By Lemma 4, we know that $\|\hat{x} - x^\bullet\| \leq \varepsilon$. $\square$

4.2 Approximate Feasibility and Optimality

Now we carry out the second step to show that if we choose the regularization parameter η carefully, the solution resulting from the coordination mechanism above is both approximately feasible and optimal. Let $x^\bullet$ denote the optimal

solution of the original convex program, $x^\bullet$ denote the optimal solution of the regularized convex program, and $\hat{x}(\eta)$ denote the solution resulting from the coordination mechanism when we use parameter η .

As an intermediate step, we will first bound the objective difference between $x^\bullet$ and x^* .

LEMMA 6. For any choice of η , $v(x^*) - v(x^\bullet) \leq \frac{\eta n}{2}$.

Next we bound the objective difference between $\hat{x}$ and $x^\bullet$ using Lipschitzness.

LEMMA 7. Suppose that $\|\hat{x} - x^\bullet\| \leq \varepsilon$, then

$$v(x^\bullet) - v(\hat{x}) \leq n\varepsilon.$$

THEOREM 3. The coordination mechanism $\text{ReC}(\eta, \varepsilon)$ coordinates the players to play a joint solution $\hat{x}$ that satisfies

$$v(\hat{x}) \geq \text{OPT} - n(\varepsilon + \eta) \quad \text{and} \quad \min_{x \in \mathcal{F}} \|x - \hat{x}\| \leq \varepsilon$$

and has coordination complexity of $O(k \log(nk/\eta\varepsilon))$.

PROOF. Follows easily from the previous lemmas. $\square$

4.3 Application to Many-to-One Matchings

Next we show a simple instantiation of our coordination mechanism for linearly separable convex programs to give a coordination complexity upper bound for many-to-one matchings. First, let's consider the following linear program formulation of the matching problem.

$$\max_x \sum_{i=1}^n \sum_{j=1}^k v_{i,j} x_{i,j} \tag{2}$$

$$\text{subject to } \sum_{i=1}^n x_{i,j} \leq b_j \quad \text{for } j = 1, \dots, k \tag{3}$$

$$\sum_{j=1}^k x_{i,j} \leq 1 \quad \text{for } i = 1, \dots, n \tag{4}$$

$$x_{i,j} \geq 0 \quad \text{for } i = 1, \dots, n \text{ and } j = 1, \dots, k \tag{5}$$

Observe that the matching linear program is an example of a linearly separable convex program as defined in Definition 3. Each player i has valuation $v_{i,j} \in \{0, 1\}$ for each type of good j and controls the decision variables $\{x_{i,j}\}_{j=1}^k$. Each supply constraint in Equation (3) corresponds to a coupling constraint, and constraints in both Equation (4) and Equation (5) are personal constraints.

A nice property about the matching linear program is that any extreme point is integral. However, this structure no longer holds if we add a regularization term to the welfare objective, so the resulting solution $\hat{x}$ resulting from the coordination mechanism will be fractional. To obtain an integral solution, we can simply use independent rounding, which does not require any further coordination. In order to obtain an integral solution, each player i will take their portion of the fractional solution $(\hat{x}_{i,j})_{j=1}^k$ and will independently sample a good by selecting each good j with probability $\hat{x}_{i,j}$. We will continue to use similar notation: let $v(\cdot)$ denote the welfare objective in the linear program, let x^* be the optimal solution for the matching linear program with welfare OPT , $\hat{x}$ be the optimal solution for the regularized program with welfare $\hat{V}$, x' be the rounded solution of $\hat{x}$, and let $\mathcal{F}$ denote the feasible region defined by all the constraints of Equation (3) in the linear program. The following lemma bounds the loss of welfare due to rounding.

LEMMA 8. Let $\beta \in (0, 1)$. Then with probability at least $1 - \beta$, the rounded solution x' satisfies

$$v(x') \geq \left(1 - \frac{\log(2/\beta)}{\sqrt{\widehat{V}}}\right) \widehat{V}$$

Now we look at approximate feasibility of x' .

LEMMA 9. Suppose that $\min_{x \in \mathcal{F}} \|x - \widehat{x}\| \leq \varepsilon$. Then with probability $1 - \beta$, x' satisfies

$$\sum_{j=1}^k \left(\sum_{i=1}^n x'_{i,j} - b_j \right)_+ \leq \sqrt{3k \log(k/\beta) \widehat{V}} + \sqrt{nk\varepsilon}$$

Observe that since this is a packing linear program, if desired, it is easy to obtain exact feasibility by simply scaling down the supply constraints: this transfers the approximation factor in the feasibility bound to the become an approximation factor in the objective.

Lastly, we are ready to establish the welfare guarantee for the rounded solution. Since the solution we obtain might slightly violate the feasibility constraints, we want to make sure we are not over-counting. If more than b_j parties select a particular good of type j , we only count the first b_j parties to select it when measuring our welfare guarantee.

THEOREM 4. There exists a coordination protocol with coordination complexity $O(k \log(nk))$ such that the parties coordinate on a matching x' with total weight:

$$\sum_{j=1}^k \min \left\{ \sum_{i=1}^n v_{i,j} x'_{i,j}, b_j \right\} \geq \left(1 - O \left(\frac{\sqrt{k} \log(k/\beta)}{\sqrt{\text{OPT}}} \right) \right) \text{OPT}$$

as long as $\text{OPT} \geq 1$.

Observe that in the setting of many-to-one matchings, when the supply of each good is $s_j \gg 1$, we expect that $\text{OPT} \gg k$, and hence in this setting, the above theorem guarantees a solution with weight $(1 - o(1)) \text{OPT}$.

5. INTERFACE WITH PRIVACY AND EFFICIENCY IN GAMES

In this section, we explain a simple implication of our results: Problems that have low sensitivity objectives (i.e. problems such that one party's data and action do not substantially affect the objective value) and low coordination complexity also have good algorithms for solving them subject to *joint differential privacy*. When the strategic variant of the optimization problem is a smooth game, they also have good welfare properties for no-regret players, even when agent types are dynamically changing.

5.1 Privacy Background

A database $D \in \mathcal{X}^n$ is an n -tuple of private records, each from one of n agents. Two databases D, D' are i -neighbors if they differ only in their i -th index: that is, if $D_j = D'_j$ for all $j \neq i$. If two databases D and D' are i -neighbors for some i , we say that they are *neighboring databases*. We write $D \sim D'$ to denote that D and D' are neighboring. We will be interested in randomized algorithms that take a database as input, and output an element from some abstract range $\mathcal{R}$.

DEFINITION 4 ([DMNS06]). A mechanism $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{R}$ is (ε, δ) -differentially private if for every pair of neighboring databases $D, D' \in \mathcal{X}^n$ and for every subset of outputs $S \subseteq \mathcal{R}$,

$$\Pr[\mathcal{M}(D) \in S] \leq \exp(\varepsilon) \Pr[\mathcal{M}(D') \in S] + \delta.$$

For the class of problems we consider, elements in both the domain and the range of the mechanism are partitioned into n components, one for each player. In this setting, *joint differential privacy* [KPRU14] is a more natural constraint: For all i , the *joint* distribution on outputs given to players $j \neq i$ is differentially private in the input of player i . Given a vector $x = (x_1, \dots, x_n)$, we write $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ to denote the vector of length $(n - 1)$ which contains all coordinates of x except the i -th coordinate.

DEFINITION 5 ([KPRU14]). A mechanism $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{R}^n$ is (ε, δ) -jointly differentially private if for every i , for every pair of i -neighbors $D, D' \in \mathcal{X}^n$, and for every subset of outputs $S \subseteq \mathcal{R}^{n-1}$,

$$\Pr[\mathcal{M}(D)_{-i} \in S] \leq \exp(\varepsilon) \Pr[\mathcal{M}(D')_{-i} \in S] + \delta.$$

If $\delta = 0$, we say that $\mathcal{M}$ is ε -differentially private. The case of $\delta > 0$ is sometimes referred to as *approximate differential privacy*.

Note that this is still a very strong privacy guarantee; the mechanism preserves the privacy of any player i against arbitrary coalitions of other players. It only weakens the constraint of differential privacy by allowing player i 's output to depend arbitrarily on her *own* input.

An important class of jointly differentially private algorithms – particularly amenable to our purposes – are those that work in the so-called *billboard model*. Algorithms in the billboard model compute a differentially private signal as a function of the input database; then each player i 's portion of the output is computed as a function only of this private signal and the private data of player i . The following lemma shows that algorithms operating in the billboard model satisfy joint differential privacy.

LEMMA 10 ([HHR⁺14]). Suppose $\mathcal{M}: \mathcal{X}^n \rightarrow \mathcal{R}$ is (ε, δ) -differentially private. Consider any set of functions $f_i: \mathcal{R} \times \mathcal{R} \rightarrow \mathcal{R}'$. Then the mechanism $\mathcal{M}'$ that outputs to each player i : $f_i(D_i, \mathcal{M}(D))$ is (ε, δ) -jointly differentially private.

Note the similarity between algorithms operating in the billboard model and coordination complexity protocols: a signal is computed by a central party, and then the action of each agent is a function only of this signal and of their own portion of the problem instance. Thus, the following lemma is immediate:

LEMMA 11. A coordination protocol (σ, π) satisfies (ε, δ) -joint differential privacy if the coordinator's encoding function σ satisfies (ε, δ) -differential privacy.

5.2 A Generic Private Coordination Protocol

Next, we give a general way to convert any coordination protocol to a jointly differentially private algorithm – and the lower the coordination complexity of the protocol, the better the utility guarantee of the private algorithm. The

tool we use is the *exponential mechanism* of [MT07], one of the most basic tools in differential privacy. To formally define this mechanism, we consider some arbitrary range $\mathcal{R}$ and some quality score function $q: \mathcal{X}^n \times \mathcal{R} \rightarrow \mathbb{R}$, which maps database-output pairs to quality scores.

DEFINITION 6 (THE EXPONENTIAL MECHANISM [MT07]). *The exponential mechanism $\mathcal{M}_E(D, q, \mathcal{R}, \varepsilon)$ selects and outputs an element $r \in \mathcal{R}$ with probability proportional to*

$$\exp\left(\frac{\varepsilon q(D, r)}{2\Delta(q)}\right),$$

where

$$\Delta(q) \equiv \max_{D, D' \in \mathcal{X}^n, D \sim D'} |q(D) - q(D')|.$$

McSherry and Talwar showed that the exponential mechanism is private and with high probability selects an outcome with high quality.

THEOREM 5 ([MT07]). *The exponential mechanism satisfies $(\varepsilon, 0)$ -differential privacy, and for any $D \in \mathcal{X}^n$ it outputs an outcome $r \in \mathcal{R}$ that satisfies*

$$q(D, r) \geq \max_{r'} q(D, r') - \frac{2\Delta(q)(\log(|\mathcal{R}|/\beta))}{\varepsilon}$$

with probability at least $1 - \beta$.

Using the exponential mechanism, we can take any coordination protocol (σ, π) , and construct a jointly differentially private coordination protocol (σ', π) with the same coordination complexity, and almost the same approximation factor. The idea is to construct a differentially private encoding function σ' that selects from the message space of σ using the exponential mechanism. Without loss of generality, we assume that the social objective function S has low-sensitivity:

$$\max_{i \in [n], a \in \mathcal{A}^n, D \sim D', a'_i \in \mathcal{A}_i} |S(D, a) - S(D', (a'_i, a_{-i}))| \leq 1.^4$$

Now we present the private protocol as follows:

Algorithm 2   Jointly   private   algorithm
PriCoor $((\sigma, \pi), q, \varepsilon, D)$

Input: A coordination protocol (σ, π) , objective function f , and input instance D

Let $\mathcal{R} = \{\sigma(D') \mid D' \in \mathcal{X}^n\}$ be the space of all possible messages in the range of σ

Let quality function q be defined as $q(D, r) = \mathbb{E}_\pi[S(D, (\pi(r, D^{(i)}))_{i \in [n]})] \forall D \in \mathcal{X}^n, r \in \mathcal{R}$

Let $\sigma'(D) = \mathcal{M}_E(D, q, \mathcal{R})$ be the message selected by the exponential mechanism

Output $a = (\pi(\sigma'(D), D^{(i)}))_{i=1}^n$

LEMMA 12. *Suppose that (σ, π) has coordination complexity ℓ and approximation ratio ρ for the objective f . Then the algorithm **PriCoor** $((\sigma, \pi), f, \varepsilon, D)$ satisfies $(\varepsilon, 0)$ -joint differential privacy, and with probability at least $1 - \beta$, the resulting action profile a satisfies*

$$\mathbb{E}[S(D, a)] \geq \frac{\text{OPT}(D)}{\rho} - \frac{2(\ell + \log(1/\beta))}{\varepsilon},$$

⁴We can always obtain this condition by scaling. It is already satisfied in the matching problem.

where the expectation is taken over the internal randomness of the encoding function σ' and decoding function π .

5.3 Efficiency in Games with Dynamic Population

Now we briefly discuss a connection between coordination complexity and the efficiency in games with dynamic population, which leverages the connection to joint differential privacy discovered by [LST15]. We briefly introduce the model in [LST15], but the discussion will necessarily be lacking in detail – see [LST15] for a formal treatment.

Let G be an n -player normal form *stage game*. We consider this game played repeatedly with a changing population of players over T rounds. Each player i has an action set $\mathcal{A}_i$, type $D^{(i)}$, and a utility function $u(D^{(i)}, a) = u_i(a)$. For concreteness, we can think about allocation games defined by auction rules M , which take as input an action profile and output an allocation $X_i(a)$ and a payment $P_i(a)$ for each player. Players have quasi-linear utility $u_i(a) = v(D^{(i)}, X_i(a)) - P_i(a) = v_i(X_i(a)) - P_i(a)$, where $v_i: \mathcal{A}^n \rightarrow [0, 1]$ denotes the valuation of player i over the allocation. In these games, a natural objective function is social welfare: $S(D, a) = \sum_{i=1}^n v_i(X_i(a))$. We write $\text{OPT}(D) = \max_{a \in \mathcal{A}^n} S(D, a)$ to denote the optimal welfare with respect to an instance D .

In the model of [LST15], after each round, every player independently exits with some probability p . Whenever a player leaves the game, she is replaced a new player, whose type is chosen adversarially. We will write D^t to denote the game instance, and a^t to denote the action profile played at round t . Lastly, we also assume that each player in the game is a *no-regret learner* and plays some *adaptive learning algorithm*.⁵

The main result of [LST15] is that the existence of jointly differentially private algorithms that find action profiles approximately optimizing the welfare in a game implies that when the dynamically changing game is played by no-regret players, their average welfare is high.

THEOREM 6 (COROLLARY 5.2 OF [LST15]). *Consider a mechanism with dynamic population (M, T, p) , such that the stage mechanism M is allocation based (λ, μ) -smooth and $T \geq 1/p$. Assume that there exists an (ε, δ) -joint differentially private allocation algorithm $X^\bullet: \mathcal{X}^n \rightarrow \mathcal{A}^n$ such that for any input instance $D \in \mathcal{X}^n$ it computes a feasible outcome that is ρ -approximately optimal*

$$\mathbb{E}[S(D, X^\bullet(D))] \geq \text{OPT}(D)/\rho.$$

If all players use adaptive learning in the repeated mechanism, then the overall welfare satisfies

$$\sum_t \mathbb{E}[S(D^t, a^t)] \geq \frac{\lambda}{\rho \max\{1, \mu\}} \sum_t \mathbb{E}[\text{OPT}(D^t)] - \frac{nT}{\max\{1, \mu\}} \sqrt{2p(1 + n(\varepsilon + \delta)) \ln(NT)},$$

where $N = \max_i |\mathcal{A}_i|$.

We show the following result connecting low coordination complexity with the welfare guarantee in games with dynamic population. (See the full version for details).

⁵For more details of adaptive learning algorithms and adaptive regret, see [HS07].

LEMMA 13. Consider a mechanism with dynamic population (M, T, p) such that the stage mechanism M is allocation based (λ, μ) -smooth and $T \geq 1/p$. Assume there is a coordination protocol (σ, π) with coordination complexity ℓ and approximation ratio ρ for the corresponding welfare maximization problem.

Then if all players use adaptive learning in the repeated mechanism, the average welfare satisfies

$$\sum_t \mathbb{E}[S(D^t, a^t)] \geq \frac{\lambda}{\rho \max\{1, \mu\}} \sum_t \mathbb{E}[\text{OPT}(D^t)] - \frac{T}{\max\{1, \mu\}} \cdot \inf_{\varepsilon > 0} \left\{ n\sqrt{4pn\varepsilon \ln(NT)} + \frac{\lambda}{\rho} \frac{2(\ell + \log(n))}{\varepsilon} \right\}.$$

6. FURTHER RESULTS

In the full version, we also present coordination protocols for coordinating a Nash equilibrium in atomic routing games and many-to-one stable matchings.

References

- [ANRW15] Noga Alon, Noam Nisan, Ran Raz, and Omri Weinstein. Welfare maximization with limited interaction. *CoRR*, abs/1504.01780, 2015.
- [Cal84] Xavier Calsamiglia. *Informational requirements of parametric resource allocation processes*. Asociación Sudeuropea de Economía Teórica, 1984.
- [CKRW14] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. Privacy and truthful equilibrium selection for aggregative games. *CoRR*, abs/1407.7740, 2014.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference, TCC '06*, pages 265–284, 2006.
- [DNO14] Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *Proceedings of the 46th ACM Symposium on Theory of Computing, STOC '14*, pages 233–242, 2014.
- [DPS02] Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. On the complexity of equilibria. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 67–71. ACM, 2002.
- [Hay45] Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.
- [HHR⁺14] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *Proceedings of the 46th ACM Symposium on Theory of Computing, STOC '14*, pages 21–30, 2014.
- [HHRW14] Justin Hsu, Zhiyi Huang, Aaron Roth, and Zhiwei Steven Wu. Jointly private convex programming. *CoRR*, abs/1411.0998, 2014.
- [HS07] Elad Hazan and C. Seshadhri. Adaptive algorithms for online decision problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(088), 2007.
- [Hur60] Leonid Hurwicz. *Optimality and Informational Efficiency in Resource Allocation Processes*. Mathematical Methods in the Social Sciences. Stanford University Press, 1960.
- [KC82] Alexander S Kelso and Vincent P Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, pages 1483–1504, 1982.
- [KN97] Eyal Kushilevitz and Noam Nisan. Communication complexity. *Advances in Computers*, 44:331–360, 1997.
- [KNR99] Ilan Kremer, Noam Nisan, and Dana Ron. On randomized one-round communication complexity. *Computational Complexity*, 8(1):21–49, 1999.
- [KPRU14] Michael Kearns, Mallesh M. Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: incentives and privacy. In *Proceedings of the 5th Innovations in Theoretical Computer Science, ITCS '14*, pages 403–410, 2014.
- [LST15] Thodoris Lykouris, Vasilis Syrgkanis, and Éva Tardos. Learning and efficiency in games with dynamic population. *CoRR*, abs/1505.00391, 2015.
- [MR74] Kenneth Mount and Stanley Reiter. The information size of message spaces. *Journal of Economic Theory*, 8(2):161–192, June 1974.
- [MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS '07*, pages 94–103, 2007.
- [NS06] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.
- [RR14] Ryan M. Rogers and Aaron Roth. Asymptotically truthful equilibrium selection in large congestion games. In *Proceedings of the 16th ACM Conference on Economics and Computation, EC '14*, pages 771–782, 2014.
- [RRUW15] Ryan M. Rogers, Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. Inducing approximately optimal flow using truthful mediators. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 471–488, 2015.
- [Seg06] Ilya Segal. Communication in economic mechanisms. *ECONOMETRIC SOCIETY MONOGRAPHS*, 41:222, 2006.