

# Social Distancing via Social Scheduling

Deepesh Kumar Lall  
Oracle India  
Bengaluru, India  
deepeshsunny8@gmail.com

Garima Shakya  
Kyushu University  
Fukuoka, Japan  
garima@inf.kyushu-u.ac.jp

Swaprava Nath  
Indian Institute of Technology  
Bombay  
Mumbai, India  
swaprava@cse.iitb.ac.in

## ABSTRACT

Motivated by the need for *social distancing* during a pandemic, we consider an approach to schedule the visitors of a facility (e.g., a general store). Our algorithms take input from the citizens and schedule the store’s discrete time-slots based on their importance in visiting the facility. We consider *indivisible* customer job requests that take single or multiple slots to complete. The salient properties of our approach are: it (a) ensures social distancing by ensuring a maximum population in a given time-slot at the facility, (b) prioritizes individuals based on the importance of the jobs, (c) maintains truthfulness of the reported importance by adding a *cooling-off* period after their allocated time-slot, during which the individual cannot re-access the same facility, (d) guarantees voluntary participation of the citizens, and yet (e) is computationally tractable. The mechanisms we propose are prior-free. The problem is NP-complete for indivisible *multi-slot* jobs, and we provide a polynomial-time mechanism that is truthful, individually rational, and approximately optimal. Experiments with data collected from a store show that visitors with more important (single-slot) jobs are allocated more preferred slots, which comes at the cost of a longer cooling-off period and significantly reduces social congestion. For the multi-slot jobs, our mechanism yields reasonable approximation while reducing the computation time significantly. While our solutions are primarily motivated by the ongoing raging pandemic, our formulation naturally applies to a broad range of scheduling settings.

## KEYWORDS

social distancing; social scheduling; strong polynomial mechanism design; poly-time approximation

### ACM Reference Format:

Deepesh Kumar Lall, Garima Shakya, and Swaprava Nath. 2023. Social Distancing via Social Scheduling. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 11 pages.

## 1 INTRODUCTION

Pandemics show the limits of pharmaceutical interventions (e.g., vaccines). Infectious diseases have appeared multiple times in the history of the human race (Spanish flu, Ebola, SARS, COVID-19, etc.), and vaccine development has taken different approaches. However, the unique first defense had always been a non-pharmaceutical intervention called *social distancing*, a term that was added to all major English dictionaries in 2020. From the times of the Spanish

flu (1918) [21] to the recent COVID-19 [36], it has been proven to be the most effective early solution.

The method of social distancing, however, has evolved. From the mass exodus of the population from the afflicted areas or forcing ships to anchor for months before entering a port in the early 20th century, we can now leverage communication and AI technologies to efficiently execute social distancing without disrupting human habitation or occupation. What we need is to *socially schedule* these individuals to visit the public facilities and prevent them from overcrowding. This is the approach we consider in this paper.

In our setting, each customer has an infinite queue of jobs that have different importances and lengths (both are privately known only to the customer). However, the customers are *myopic*, i.e., worry only about the last unprocessed job.<sup>1</sup> They experience a better value if the job is assigned their preferred slots, but also a *disutility* to wait before submitting their next job for allocation. All jobs are *indivisible*, i.e., has to be completed once started. In this paper, we consider *two* settings: (i) all jobs are of single time-slot length, (ii) different jobs are of different integral time-slot lengths.

Though cast in the context of social scheduling for pandemics, a similar problem arises in general scheduling settings, e.g., scheduling traffic in freeways or multi-ownership computational jobs in a single-core processor. Since all such settings have multiple agents competing for a common resource and the importance of the jobs is private, the solutions involving truthful revelation in a computationally tractable manner also apply to those settings.

This paper introduces a novel approach to pandemic containment using the mechanism design that reduces congestion in facilities, satisfies various desirable theoretical properties, and exhibits fair performance in practice. The following section details the contributions of this paper.

## 1.1 Brief Problem Description and Contributions

The opening hours of a facility are divided into *periods* (e.g., a day), each of which has multiple *slots* (e.g., every hour when it is open). The customers have an unlimited number of outstanding jobs to be processed in a sequence at the facility, and they report the valuations of the immediate unprocessed job.<sup>2</sup> A valuation  $v_{ij}$  denotes agent  $i$ ’s importance for that job if it starts in slot  $j$ . Since this information is private to agent  $i$ , a mechanism needs to elicit this truthfully. In a setting where the agents’ preferences are private, if the mechanism has no additional structures (e.g., transfer of individual payoff), only *dictatorial* mechanisms (where a pre-selected agent’s favorite

<sup>1</sup>A large number of studies in behavioral science point to such myopic behavior [23].

<sup>2</sup>A typical shopper knows that she needs to visit a store many times but precisely knows the importance of the immediate visit.

*Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

outcome is always selected) are truthful [29, Thm 7.2]. Therefore, the use of transfers in some form is inevitable to ensure the truthfulness of the agents. However, for pandemic containment, the use of money for scheduling citizens is unethical and illegal in certain countries. Hence, we use *time-delay* as a replacement for money. Waiting time is often seen as a resource that individuals agree to trade with [25]. Our scheduling approach will work in all places where payment can be replaced with a time-delay. Quite naturally, an agent prefers to have a more valuable slot assigned to her with less time-delay. We model the agents' payoffs using the well-known quasi-linear payoff model [32, Chap 10]. This competitive scenario induces a *game* where agents' actions are to report the valuations. The agents may *overstate* (or *understate*) their actual valuations.

The contributions of this paper can be summarized into the following *four* major points:

- ▷ We show that the problem of maximizing the *social welfare* (sum of the agents' valuations) of the slot allocation is computationally easy to solve for jobs with single-slot length (Theorems 5.1 and 5.5) despite it being a combinatorial optimization problem.
- ▷ The single-slot case has the advantage that the delay (cooling-off time) can be calculated via the Vickrey-Clarke-Groves (VCG) mechanism [11, 18, 35] that ensures truthfulness and participation of the agents.
- ▷ Our main contribution is in the multi-slot jobs. We show that the welfare-maximizing allocation of multi-slot jobs is computationally hard (Theorem 6.4). We propose a polynomial time mechanism (Theorem 6.8) which ensures participation (Theorem 6.6), truthfulness (Theorem 6.7), and is approximately optimal (Theorem 6.9).
- ▷ Our real and synthetic data experiments (Section 7) show that visitors with more important jobs are allocated more preferred slots, which comes at the cost of a longer delay to re-access the store. We show that social distancing is significantly improved using users' visit data from a store (Section 7.1). For the multi-slot jobs, our approximately optimal mechanism provides a reasonable approximation at a much reduced computational cost in practice (Section 7.2).

The mechanisms presented in this paper are *prior-free*, i.e., they do not depend on the probabilistic information of the valuations.

## 1.2 Related Work

The literature on pandemic control has widely documented the quantitative benefits of social distancing [13, 15, 33]. The benefits of social distancing depend on the extent to which the citizens follow it. An extensive part of the recent research related to social distancing during COVID-19 aims to understand the relationship of social distancing with different public policies and other factors [1, 10, 16, 27, 28].

Studies related to the behavior of individuals during the pandemic show that despite the infection probability not decreasing, during the equilibrium social distancing phase, individuals gradually reduce their social distancing efforts [8, 34].

In a slightly different strand of literature, the problem of resource allocation with monetary transfers has been addressed to ensure truthfulness, e.g., in the context of jobshop scheduling [20], which

is also close to our work. Lavi and Nisan [24] study online supply curves based auction of *identical divisible goods* that ensures truthfulness. In this paper, we consider an offline allocation problem but a comparatively more complex one (multiple resources and indivisible tasks of different length). Chen et al. [9] propose a truthful approximate mechanism for online allocation of job to machines where the job can resume or restart once preempted. We provide a comparatively better approximation ratio for efficiency, albeit in an offline setting. The other related line of work involves designing incentives in queueing problems with *specific cost structures* that aim to find an efficient allocation truthfully and also ensures budget balance [4, 14, 26], while our model can admit costs of *any* structure.

In the context of job scheduling without money, Koutsoupias [22] studies the allocation of independent tasks to machines. Every machine reports the time it takes to execute each task and the mechanism provides an approximation to the minimum *makespan* in a truthful manner without money for one task—which can be repeated for multiple tasks. In this paper, we maximize the *sum of the visitors' valuations* which are independent of the length of the job and provide an approximate mechanism for multiple tasks maintaining the slot-capacity. Braverman et al. [6] study a similar problem of the allocation of medical treatments at hospitals that have differential costs to patients and the patients value the hospitals differently. The waiting time before being admitted to the hospital helps to get a stable matching. However, the value of the agents do not change over slots and hence is different from our setting.

## 2 BASIC SETUP AND SINGLE-SLOT JOBS

Define  $N := \{1, \dots, n\}$  as the set of agents trying to access a facility  $\mathcal{F}$ . Time is divided into *periods*, and each period is further divided into *slots*. The set of slots is denoted by  $M := \{1, \dots, m\}$ . Every slot has a maximum capacity of  $k$ , which is decided by the region's social distancing norm based on the size of the facility.<sup>3</sup> A *central planner* (e.g., an AI app) allocates these slots to the agents, maintaining the capacity constraint. Every agent  $i$  has a *cardinal* preference  $v_{ij} \in \mathbb{R}_{\geq 0}$  (called the agent's *valuation*) if her immediate unprocessed job is allocated slot  $j$ . The valuation implicitly reflects the importance of visiting the facility for that agent. The valuation vector of  $i$  is represented by  $v_i = (v_{ij}, j \in M) \in \mathbb{R}_{\geq 0}^m$ . In this paper, we consider different facilities independently. The joint facility-slot allocation problem can be modeled as a similar problem with the additional constraint that the same slot cannot be allocated to the same agent at different facilities. We leave the detailed analysis for it as future work.

The planner decides the allocation, which can be represented as a matrix  $A = [a_{ij}]$ , where  $a_{ij} = 1$  if agent  $i$  is allotted slot  $j$ , and zero otherwise. We assume that every agent can be assigned at most one slot in a period, and the total number of agents assigned to each slot is at most  $k$ . We denote the slot assigned to  $i$  by  $a_i^*$ . The planner also decides a delay  $d = (d_i, i \in N)$ , where  $d_i$  is the time-delay (in the same unit as the valuation) of agent  $i$  before which she cannot make another request to the system. The net payoff of an agent is

<sup>3</sup>The analysis and results will follow even if the capacity  $k_j$  varies with the slots  $j \in M$ .

assumed to follow a standard *quasi-linear form* [32], which implies that every agent wants a more valued slot to be assigned to her and also wants to wait less.

$$u_i((A, d), v_i) = v_i(A) - d_i, \text{ where } v_i(A) = v_{ia_i^*}. \quad (1)$$

Denote the set of all allocations by  $\mathcal{A}$ . The delays  $d_i \in \mathbb{R}_{\geq 0}, \forall i \in N$ . The planner does not know the valuations of the agents. Therefore he needs the agents to report their valuations to decide the allocation and the delay. This leaves the opportunity for an agent to misrepresent her true valuation. To distinguish, we use  $v_{ij}$  for the true valuation and  $\hat{v}_{ij}$  for reported valuations. In the first part of this paper, we will consider single-slot jobs and use the shorthand  $v = (v_i)_{i \in N}$  to denote the true valuation profile represented as an  $m \times n$  real matrix, and  $\hat{v}$  to denote the reported valuation profile. The notation  $v_{-i}$  denotes the valuation profile of the agents except  $i$ . Therefore, the planner's decision problem is formally captured by the following function.

*Definition 2.1 (Social Scheduling Function (SSF)).* A social scheduling function (SSF) is a mapping  $f : \mathbb{R}^{m \times n} \rightarrow \mathcal{A} \times \mathbb{R}^n$  that maps the reported valuations to an allocation and delay for every agent. Hence,  $f(\hat{v}) = (A(\hat{v}), d(\hat{v}))$ , where  $A$  is the allocation and  $d$  is the delay function.<sup>4</sup>

### 3 PRELIMINARY DEFINITIONS

In this section, we formally define a few desirable properties that a social scheduling function should satisfy. The properties address the issues of prioritization, truthfulness, voluntary participation, and computational complexity.

The first property ensures that the allocation is *efficient* in each period, i.e., maximizes the sum of the valuations of all the agents.

*Definition 3.1 (Efficient Per Period (EPP)).* An SSF  $f$  is *efficient per period (EPP)* if, at every period, it chooses an allocation  $A^*$  that maximizes the sum of the valuations of all the agents. Formally, if  $f(\cdot) = (A^*(\cdot), d(\cdot))$ , then

$$A^*(v) \in \operatorname{argmax}_{A \in \mathcal{A}} \sum_{i \in N} \sum_{j \in M} v_{ij} a_{ij}. \quad (2)$$

However, since the planner can only access the reported values  $\hat{v}_i$ 's, which can be different from the true  $v_i$ 's, the reported values must indeed be the true values. The following property ensures that the agents are incentivized to 'truthfully' reveal this information *irrespective of the reports of the other agents*.

*Definition 3.2 (Per Period Dominant Strategy Truthful).* An SSF  $f(\cdot) = (A(\cdot), d(\cdot))$  is *truthful in dominant strategies per period* if for every  $v_i, \hat{v}_i, \hat{v}_{-i}$ , and  $i \in N$

$$v_i(A(v_i, \hat{v}_{-i})) - d_i(v_i, \hat{v}_{-i}) \geq v_i(A(\hat{v}_i, \hat{v}_{-i})) - d_i(\hat{v}_i, \hat{v}_{-i}).$$

The next property ensures that it is always weakly beneficial for every rational agent to participate in such a mechanism.

*Definition 3.3 (Individual Rationality).* An SSF  $f(\cdot) = (A(\cdot), d(\cdot))$  is *individually rational* if for every  $v$ , and  $i \in N$

$$v_i(A(v)) - d_i(v) \geq 0.$$

<sup>4</sup>We overload the notation  $A$  and  $d$  to denote both functions and values of those functions since their use will be clear from the context.

Large facilities with a large number of high-capacity slots lead to an exponential increase in the size of the set  $\mathcal{A}$ . This largeness of  $\mathcal{A}$  makes it challenging to find a solution quickly. In a practical setting, where the allocations and delays need to be decided before every period, it is desirable to have an SSF that is computable in a time polynomial in  $n$  and  $m$  so that it finishes the computation in a time negligible to the time duration of the period. We consider algorithms that are *strongly polynomial* [17]. An SSF is strongly polynomial-time computable if there exists an algorithm that computes it in a time strongly polynomial in  $n$  and  $m$ , irrespective of the size of the actual data, such as the value of the  $v_i$ s or  $k$ .

### 4 PERIODIC MECHANISMS

We consider mechanisms that run at every period of this social scheduling problem. The agents report their valuations at the beginning of every period. The planner decides the schedules and delays.<sup>5</sup> Since the agents have the opportunity to overstate their importance to get a better slot allotted to them, our approach that uses the ideas of mechanism design [5] to this social scheduling problem is useful. We use the delay as a surrogate for transferable utility among the agents to satisfy several desirable properties.

For the single-slot job setup, the delays of agents are computed via the standard VCG payment rule, and we call the allocation and delay together as the mechanism **VCG-T** (VCG with Time delays). **VCG-T** mechanism is as follows.

*Description of VCG-T.* The SSF needs to decide on the allocation  $A$  and the delay  $d$ . **VCG-T** computes the allocation as follows.

$$\begin{aligned} & \operatorname{argmax}_A \sum_{j \in M} \sum_{i \in N} v_{ij} a_{ij} \\ & \text{s.t. } \sum_{j \in M} a_{ij} \leq 1, \forall i \in N; \sum_{i \in N} a_{ij} \leq k, \forall j \in M \\ & a_{ij} \geq 0, \forall i \in N, j \in M. \end{aligned} \quad (3)$$

This is an LP relaxation of the actual allocation problem, which allows  $a_{ij}$ s to be only in  $\{0, 1\}$ . We will show that this is without loss of optimality since the solution to LP (3) will always be integral and coincide with the solution of the corresponding IP.

The delays of agents are computed via the standard VCG payment rule. Denote the optimal allocation of LP (3) by  $A^*(v)$ . Also, denote the allocation given by LP (3) when agent  $i$  is removed from the system by  $A_{-i}^*(v_{-i})$ . For agent  $i$ , the delay is given by,

$$d_i := \sum_{\ell \in N \setminus \{i\}} v_{\ell}(A_{-i}^*) - \sum_{\ell \in N \setminus \{i\}} v_{\ell}(A^*). \quad (4)$$

The mechanism in every period is described in Section 4.

<sup>5</sup>For mechanisms that consider the dynamic extension of such allocation problems with finite or infinite horizon [3, e.g.], (a) the designer needs to know the transition probabilities, (b) equilibrium guarantees are weaker, and (c) are computationally expensive. These factors made us restrict our attention to periodic mechanisms.

**Algorithm 1** VCG-T in every period

- 1: **Input:** for every agent  $i \in N$ , the value  $\hat{v}_i$
- 2: compute  $A^*(\hat{v})$  (Equation (3)) as the allocation
- 3: charge a delay  $d_i(\hat{v})$  (Equation (4)) to every  $i \in N$  for which they cannot access the scheduling mechanism again
- 4: **Output:**  $A^*(\hat{v})$  and  $d(\hat{v})$

Since VCG-T uses the VCG payment expression to compute the time delay and because the allocated slots are *goods* to the agents, the following two facts follow from the known properties of the VCG mechanism. For completeness, the proof for these facts is provided in the Appendix B.

FACT 1. VCG-T is per period dominant strategy truthful.

FACT 2. VCG-T is individually rational for every agent.

In the next section, we show that the VCG-T mechanism for the single-slot job is per period dominant strategy truthful, individually rational, and runs in strongly polynomial time.

## 5 VCG-T FOR SINGLE-SLOT JOBS

We first show that the allocation given by VCG-T indeed maximizes per-period social welfare.

**THEOREM 5.1.** *The allocation of VCG-T given by LP (3) always gives integral solutions.*

**PROOF.** Consider the vector  $x^T = (a_{11}, \dots, a_{1m}, \dots, a_{n1}, \dots, a_{nm})$ , i.e., the rows of  $A$  linearized as a vector. We can write the constraints of LP (3) in using a  $(n+m) \times nm$  constraint matrix, s.t.,

$$\begin{pmatrix} 1 & \dots & 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & \dots & 0 & 1 & \dots & 1 & 0 & \dots & 0 \\ & & & & \dots & & & & \\ 1 & \dots & 0 & 1 & \dots & 0 & 1 & \dots & 0 \\ 0 & 1 & \dots & 0 & 1 & \dots & 0 & 1 & 0 \\ & & & & \dots & & & & \end{pmatrix} x \leq \begin{pmatrix} 1 \\ \vdots \\ \vdots \\ k \\ \vdots \\ \vdots \end{pmatrix}$$

Denote the matrix on the LHS by  $C$ . The first  $n$  and the next  $m$  rows correspond to the first and second set of constraints of LP (3), respectively. We show that  $C$  is totally unimodular (TU), which is sufficient to conclude that LP (3) has integral solutions. We use the Ghouila-Houri characterization [7] to prove that  $C$  is TU. According to this characterization, a  $p \times q$  matrix  $C$  is TU if and only if each set  $R \subseteq \{1, 2, \dots, p\}$  can be partitioned into two sets  $R_1$  and  $R_2$ , such that,  $\sum_{i \in R_1} a_{ij} - \sum_{i \in R_2} a_{ij} \in \{1, 0, -1\}$ , for  $j = 1, 2, \dots, q$ . Note that, in  $C$  every column has two 1's, one in the first  $n$  rows and one in the next  $m$  rows. Pick any subset  $R$  of the rows, construct the  $R_1$  to be the rows that come from the first  $n$  rows, and  $R_2$  to be the rows that come from the last  $m$  rows (one of these partitions can be empty). Clearly, the difference in each column of the rows  $R$  will lie in  $\{1, 0, -1\}$ . Hence proved.  $\square$

The result above shows that the optimal solution of LP (3) is an optimal solution of the corresponding integer program that maximizes the per-period social welfare. Hence, we conclude the following.

**COROLLARY 5.2.** VCG-T is EPP.

Even though the LP formulation of VCG-T is without loss of optimality, in general, LPs can be weakly polynomial, i.e., the space used by the algorithm may not be bounded by a polynomial in the size of the input. However, we show that an even stronger result holds for VCG-T. The forthcoming results show that the allocation and delays of VCG-T are strongly polynomial. To show this, we will first reduce the allocation problem (LP (3)) to a minimum weight  $b$ -matching problem, which is known to be strongly polynomial [31].

Consider an edge-weighted bipartite graph  $(N, M, E)$ , where  $N$  and  $M$  are the agent set and set of slots, respectively. The set  $E$  denotes the edges  $(i, j)$  with weights  $-v_{ij}$ . The matching constraints are given by  $b_i = 1, \forall i \in N$ , and  $b_j = k, \forall j \in M$ .

**LEMMA 5.3.** *Let  $E^* \subseteq E$  be a perfect  $b$ -matching in  $(N, M, E)$  and  $A^* = [a_{ij}^*]_{i \in N, j \in M}$  be an allocation where  $a_{ij}^* = 1 \Leftrightarrow (i, j) \in E^*$ . The matching  $E^*$  is a minimum weight perfect  $b$ -matching iff  $A^*$  is an optimal solution to LP (3).*

**PROOF.** We prove this via contradiction. Suppose  $A^*$  is not an optimal solution to LP (3), i.e., there exists  $A'$  which satisfies the constraints and yet gives a larger value to the objective function than that of  $A$ . Hence,  $\sum_{j \in M} \sum_{i \in N} v_{ij} a'_{ij} > \sum_{j \in M} \sum_{i \in N} v_{ij} a^*_{ij}$ . Consider the edge set  $E'$  corresponding to  $A'$ . This is a perfect  $b$ -matching since  $A'$  satisfies the constraints of LP (3), and  $E'$  gives a lower weight than  $E^*$ , which proves that  $E^*$  is not the minimum weight perfect  $b$ -matching. The implications can be reversed to obtain the other direction of the proof.  $\square$

Note that the delays are calculated by solving an equivalent LP like LP (3) with one less agent. Therefore, each of these LPs is strongly polynomial, and the planner needs to solve  $n$  of them. The computation of each delay requires the addition of  $2(n-1)$  terms and one subtraction. Hence, the number of computations is polynomial in the number of numbers in the input instance, and the space required is polynomial in the input size. Therefore we conclude the following.

**COROLLARY 5.4.** *The computation of the delays in VCG-T is strongly polynomial.*

Combining Theorem 5.1, Lemma 5.3, and Corollary 5.4, we get the following result.

**THEOREM 5.5.** *VCG-T provides a combinatorial, strongly polynomial algorithm for computing a social schedule and delays.*

The main contribution of this paper is to schedule the multi-slot jobs that are relatively difficult to schedule. We present our results for multi-slot jobs in the next section.

## 6 MULTI-SLOT JOBS

In this section, we consider jobs with different lengths, i.e., for agent  $i$ , the job may be of length  $l_i \geq 1$ . Since the job is *indivisible*, the entire length  $l_i$  of the job requires contiguous slots for execution within the period. For example, an individual may visit a facility (e.g., a shopping mall) for quick shopping, which may take a shorter duration, or for dining, which may take longer. However, all these

jobs are indivisible, and the allocation needs to provide contiguous time-slots to that agent. The agents report the valuations and lengths of their jobs. We show that the optimal allocation problem in such a setting can be computationally intractable. The notation is mildly updated as follows to accommodate the multi-slot jobs.

Each agent  $i$  gets a valuation  $v_{ij}$  for her last unprocessed job if her job begins at slot  $j$  and has a length  $l_i$ . The value of the job is zero if (a) it starts at any of the last  $(l_i - 1)$  time-slots of the period (since it cannot finish within the period), and (b) if the job is unallocated.

A matrix  $V$  consists of the agents' reported valuations, and  $L$  consists of the lengths of agents' jobs. Allocation is given by the matrix  $\mathbf{A} = [\mathbf{a}_{ij}]$ , where  $\mathbf{a}_{ij} = 1$  if agent  $i$ 's job starts at slot  $j$ , else  $\mathbf{a}_{ij} = 0$ , and  $\mathbf{a}_i$  represents the slot allocation vector for agent  $i$ . Keeping all other notations as before, we define the **MIA** problem as follows.

**Definition 6.1 (Multi-slot Indivisible jobs Allocation problem (MIA)).** : Given  $(N, M, V, L, k)$ , find an allocation  $\mathbf{A}$ , such that  $\sum_{i \in N} \sum_{j \in M} v_{ij}(\mathbf{a}_{ij})$  is maximum, subject to the constraints that the total number of jobs allocated in a slot does not exceed the capacity of the slot, and each job  $i$  is assigned to at most  $l_i$  contiguous slots. Mathematically, **MIA** is given by the following integer program (IP).

$$\begin{aligned} & \operatorname{argmax}_{\mathbf{A}} \sum_{i \in N} \sum_{j \in M} v_{ij} \mathbf{a}_{ij} \\ \text{s.t.} & \sum_{i \in N} \sum_{\substack{p \in M \\ j \in [p, p+l_i-1]}} \mathbf{a}_{ip} \leq k, \forall j \in M, \\ & \sum_{j \in M} \mathbf{a}_{ij} \leq 1, \forall i \in N; \mathbf{a}_{ij} \in \{0, 1\}, \forall i \in N, \forall j \in M \end{aligned} \quad (5)$$

The first set of inequalities ensures that the number of jobs processed in a slot does not exceed the slot capacity  $k$ . We sum over every job  $i \in N$  and check if it is under execution at  $j$  for every  $j \in M$ . A job  $i$  is under execution at slot  $j$  if allocated at a slot  $p$  s.t.  $j \leq p + l_i - 1$ . The second set of inequalities ensures that no job is allocated more than once.

We show that **MIA** is computationally intractable by performing a polynomial reduction from the Multi-Unit Combinatorial Auction (**MUCA**), which is NP-complete.

**Description of MUCA:** Consider a multiset  $\mathcal{M} = (\mathcal{G}, y)$ , where  $\mathcal{G} = \{1, 2, 3, \dots, g\}$  is a set of goods and  $y$  is a function,  $y : \mathcal{G} \rightarrow \mathbb{Z}_{\geq 0}$  representing the multiplicity or the number of available units of the elements of  $\mathcal{G}$  in  $\mathcal{M}$ . Each agent  $i \in N = \{1, 2, \dots, n\}$  is a multi-minded bidder, which means  $i$  has a positive valuation  $w_i(\cdot)$  for multiple bundles of available goods. We call the set of bundles for which agent  $i$  has a positive valuation to be the *demand set* of  $i$ , represented by  $\mathcal{D}_i$ . The valuation function is such that,  $w_i(q) \in \mathbb{R}_{\geq 0}, \forall q \in \mathcal{D}_i$ . We use the following notation  $\mathcal{D} = [\mathcal{D}_i]_{i \in N}$  and,  $W_i = (w_i(q))_{q \in \mathcal{D}_i}$ ,  $W = [W_i]_{i \in N}$ . This paper assumes that every agent demands at most one unit of every good. With this assumption, an allocation of a bundle of goods to the agents is represented as a matrix  $\mathcal{S} = [\mathbf{s}_{iq}]$ , where  $\mathbf{s}_{iq} = 1$ , if the bundle  $q \in \mathcal{D}_i$  is allocated to  $i$ , else  $\mathbf{s}_{iq} = 0$ . For an allocation  $\mathcal{S}$ , every agent  $i$  gets a valuation,  $w_i(\mathcal{S}) = \sum_{q \in \mathcal{D}_i} w_i(q) \mathbf{s}_{iq}$ , otherwise  $w_i(\mathcal{S}) = 0$ . The formal definition is as follows.

**Definition 6.2 (Multi-Unit Combinatorial Auction (MUCA)).** Given  $(N, \mathcal{M}, W, \mathcal{D})$ , find an allocation  $\mathcal{S}$  of goods to the agents such that  $\sum_{i \in N} \sum_{q \in \mathcal{D}_i} w_i(q) \mathbf{s}_{iq}$  is maximum, and the total units of good  $j \in \mathcal{G}$  allocated to the agents does not exceed  $j$ 's availability  $y(j)$ . Every agent  $i$  is assigned at most one of the demanded bundles from  $\mathcal{D}_i$ . Mathematically, **MUCA** is given by the following integer program (IP):

$$\begin{aligned} & \operatorname{argmax}_{\mathcal{S}} \sum_{i \in N} \sum_{q \in \mathcal{D}_i} w_i(q) \mathbf{s}_{iq} \\ \text{s.t.} & \sum_{i \in N} \sum_{\substack{q \in \mathcal{D}_i \\ j \in q}} \mathbf{s}_{iq} \leq y(j), \forall j \in \mathcal{G} \\ & \sum_{q \in \mathcal{D}_i} \mathbf{s}_{iq} \leq 1, \forall i \in N; \mathbf{s}_{iq} \in \{0, 1\}, \forall i \in N, \forall q \in \mathcal{D}_i \end{aligned} \quad (6)$$

The reduction of **MIA** to **MUCA** proceeds as follows. For a given instance  $(N, M, V, L, k)$  of **MIA**, construct an instance of **MUCA**  $(N, \mathcal{M}, W, \mathcal{D})$  problem such that the set of agents  $N$  is  $N$ , the set of goods  $\mathcal{G}$  is the set of the slots  $M$  within the period, where  $y(j) = k, \forall j \in M$ . For every  $i \in N$ , the demand set  $\mathcal{D}_i$  consists of  $(m - l_i + 1)$  distinct bundles. Each of the bundles in  $\mathcal{D}_i$  is of size  $l_i$  and consists of  $l_i$  contiguous slots. We denote a bundle as  $q_j$  if it contains  $l_i$  contiguous slots starting from slot  $j$ , and  $w_i(q_j)$  is equal to  $v_{ij}$  (the valuation  $i \in N$  gets if her job starts at slot  $j \in M$ ). The above construction is done in polynomial steps of the input size. We construct a solution of **MIA** from a solution of **MUCA** in the following way: for every  $q_j \in \mathcal{D}_i$  and  $i \in N$ , if  $\mathbf{s}_{iq_j} = 1$  then,  $\mathbf{a}_{ij} = 1$  for every  $i \in N$  and  $j \in M$ . Similarly, we construct a solution of **MUCA** from a solution **MIA** in the following way: if  $\mathbf{a}_{ij} = 1$  for  $i \in N$  and  $j \in M$  then,  $\mathbf{s}_{iq_j} = 1$  for every  $q_j \in \mathcal{D}_i$  and  $i \in N$ . The following lemma shows that an optimal solution of **MIA** is an optimal solution of **MUCA** and vice-versa.

**LEMMA 6.3.** Let  $\mathcal{S}^*$  is a solution for **MUCA** for a multiset of goods  $\mathcal{M}$ , and  $\mathbf{A}^*$  is such that,  $\mathbf{a}_{ij}^* = 1$  for  $i \in N$  and  $j \in M$ , if and only if  $\mathbf{s}_{iq_j}^* = 1$  in  $\mathcal{S}^*$  for  $i \in N$  and  $q_j \in \mathcal{D}_i$ , then  $\mathbf{A}^*$  is an optimal solution for **MIA** if and only if  $\mathcal{S}^*$  is an optimal solution for **MUCA**.

Since **MUCA** is NP-complete [12, 30], using Lemma 6.3, we get the following theorem.

**THEOREM 6.4.** **MIA** is NP-complete.

However, it is possible to find an approximately efficient allocation in polynomial time that is truthful and individually rational. To find that, we leverage the approximation algorithm of **MUCA** due to Bartal et al. [2, Theorem 3]. Using Lemma 6.3 and the following few results, we prove that there exists a polynomial time truthful mechanism (**MIA Approximation Algorithm** or **MAA**) to achieve  $O(km^{\frac{1}{k-2}})$  approximation to the optimal solution of **MIA**.

The operational principle of **MAA** is a sequential dictatorship, where the sequence is an arbitrary order (WLOG  $1, 2, \dots, n$ ) of the agents and is independent of the information submitted by them. The mechanism comes with a price<sup>6</sup> vector, which is updated while iterating over the agents in the sequence. We use a superscript  $i$  to denote the price faced by the agent  $i$  for slot  $j$ ,  $P_j^i$  when  $i$ 's

<sup>6</sup>The terms *price* and *delay* are equivalent in the rest of the paper.

**Algorithm 2** MAA in every period

---

```

1: Procedure MAA( $N, M, V, L, k$ )
2:  $b \leftarrow \arg \max_i \max_{i \in N, j \in M} v_{ij}$ ;  $v_{\max} \leftarrow \max_{i \in N, j \in M} v_{ij}$ ; //  $b$  is the agent
   with highest valuation ( $v_{\max}$ ) for any slot.//
    $r \leftarrow (6m(k-1))^{\frac{1}{k-2}}$ 
3:  $Q^0 \leftarrow [0, 0, \dots, m \text{ times}]$ 
4:  $\mathbf{a}_{bs'}^* = 1$ , where  $s' \leftarrow \operatorname{argmax}_{j \in M} (v_{bj})$ ; // Allocate the highest
   valued slot for  $b$ .//
    $\mathbf{a}_{bj}^* = 0, \forall j \in M \setminus \{s'\}$ ;  $\mathcal{P}_b = v_{\max}^{-b}$ 
5: for  $i = \{1, 2, \dots, n\}$  and  $i \neq b$  do
6:   for  $j = \{1, 2, \dots, m\}$  do
7:      $P_j^i \leftarrow \pi_0 \cdot r^{Q_j^{i-1}}$ 
8:      $\mathbf{a}_{is}^* = 1$ , where  $s \leftarrow \mathbf{max}(P^i, v_i)$  (Equation (7))
9:      $\mathbf{a}_{ij}^* = 0, \forall j \in M \setminus \{s\}$ 
10:     $\mathcal{P}_i \leftarrow \sum_{j=s}^{s+l_i-1} P_j^i$ 
11:    for  $j = \{1, 2, \dots, m\}$  do
12:      if  $j \in [s, s+l_i-1]$  then
13:         $Q_j^i \leftarrow Q_j^{i-1} + 1$ 
14:      else
15:         $Q_j^i \leftarrow Q_j^{i-1}$ 
16: return  $\mathbf{a}^*, \mathcal{P}$ 

```

---

turn comes. Hence,  $P^i = [P_j^i]_{j \in M}$  denotes the price vector seen by  $i$ . The mechanism also uses a function  $\mathbf{max}$  that returns the slot  $s$  that maximizes agent  $i$ 's utility given her valuation vector  $v_i$  and the price vector  $P^i$  when her job of length  $l_i$  starts from slot  $s$ . Mathematically,  $\mathbf{max}$  is defined as follows:

$$\mathbf{max}(v_i, P^i) = \operatorname{argmax}_{s \in M} (v_{is} - \sum_{j \in [s, s+l_i-1]} P_j^i) \quad (7)$$

MAA maintains a vector  $Q^i = [Q_j^i]_{j \in M}$ , where  $Q_j^i$  denotes the current allocated population of the slot  $j$  after allocating slots to  $i$ . First, MAA picks the agent  $b$  that has the maximum valuation  $v_{\max}$  for any slot, and initializes  $Q_j^0 = 0, \forall j \in M$ . The initial price of every slot is set to  $\pi_0 := \frac{v_{\max}}{6m(k-1)}$  and a constant factor  $r = (6m(k-1))^{\frac{1}{k-2}}$  is defined. Consider an arbitrary order (WLOG  $(1, 2, \dots, n)$ ) of the agents. For each agent  $i \neq b$  in sequence, the price  $P_j^i$  is computed using  $Q_j^{i-1}$ ,  $r$ , and  $\pi_0$  such that the prices of the slots increase by a multiplicative factor of a suitable exponent of  $r$  such that the prices for more congested slots are higher. Then the highest utility-deriving slot  $s$  to start  $i$ 's job is found using  $\mathbf{max}$ , and the corresponding allocation vector for  $i$  is represented as  $\mathbf{a}_i^*$ , where  $\mathbf{a}_{is}^* = 1, \mathbf{a}_{ij}^* = 0, \forall j \neq s$ . The total price (or delay) charged to  $i$  is denoted by  $\mathcal{P}_i$  and is equal to  $\sum_{j \in [s, s+l_i-1]} P_j^i$ . The vector  $Q^i$  is updated after allocating slots to agent  $i$ . The agent  $b$  gets her most valued starting slot and pays the maximum valuation among all other agents and slots, represented by  $v_{\max}^{-b}$ .

An important feature of Algorithm 2 is that it *does not* explicitly check the capacity constraint. However, we show that the choices of  $\pi_0$  and  $r$  implicitly maintain that in the following result. The units of slot  $j$  after Algorithm 2 executes that are occupied by agents except  $b$  are  $Q_j^*$ .

LEMMA 6.5. *Let  $\pi_0, r, \delta > 0$  be such that  $\pi_0 r^\delta \geq v_{\max}$ , then  $Q_j^* \leq \delta + 1$ . This implies that MAA maintains the capacity constraints for each slot for  $\delta = k - 2$ .*

PROOF. Assume for contradiction that  $Q_j^* > \delta + 1$  and let  $i$  be the first customer due to which this contradiction takes place for some slot  $j$ , i.e.,  $Q_j^i > \delta + 1$ . Since each customer does not get more than one unit of any slot, then it must be that  $Q_j^{i-1} > \delta$ . Hence, for slot  $j$ , the following holds:  $P_j^i > \pi_0 r^\delta \geq v_{\max} \geq \max_{j \in M} v_{ij}$ . This makes  $i$ 's total price for  $l_i$  contiguous slots, including slot  $j$  to be more than her corresponding total valuation for those slots. This contradicts the definition of  $\mathbf{max}$  since the utility becomes negative for agent  $i$ .  $\square$

The allocation to  $b$  is at most *one* unit from each slot. With carefully choosing  $\pi_0$  and  $r$ , we bound the units of any slot allocated to all the other agents,  $Q_j^*$  to  $(k-1)$ , maintaining the possibility of the maximum use of every slot.

Since  $\mathbf{max}$  allocates a slot only if that allocation increases the agent's utility, the following result holds.

THEOREM 6.6. *MAA is individually rational.*

Next, we show that misreporting the private information  $(v_i, l_i)$  is never beneficial for any agent  $i$ .

THEOREM 6.7. *In MAA, reporting  $v_i$  and  $l_i$  truthfully in every period is a dominant strategy for all  $i \in N$ .*

PROOF. For the agent  $b$ , we see that the utility is that of a second price auction and is independent of its length report. Since for the second price auction revealing valuation truthfully is a dominant strategy, therefore, truthfully revealing valuation and length is a dominant strategy for  $b$ .

For the other agents, note that MAA considers the agents sequentially and allocates the utility-maximizing available slots in their turn. The order of the agents in MAA is independent of the valuations and lengths of the jobs. Consider agent  $i$ . When her turn comes, the mechanism picks the slots that give the maximum difference between the valuation of  $i$  for those slots and the current prices of those slots. Note that, the prices of those allocated slots are not dependent on the valuation or length reported by agent  $i$  (rather, it is dependent on the reports of the previous agents in the sequence and  $b$ ), and the mechanism allocates her the optimal set of slots. Hence, by misreporting the valuation  $v_i$ , agent  $i$  can either continue to get the same slots or get a worse set of slots w.r.t. her true valuation. Hence, there is no incentive for  $i$  to misreport her valuation.

*Misreporting length:* If  $\hat{l}_i < l_i$ , MAA allocates only  $\hat{l}_i$  number of contiguous slots to  $i$  (which can have zero value as  $\hat{l}_i$  is not sufficient for completion of her job) and  $i$  can get a negative payoff as she has to pay  $\mathcal{P}_i$  (which is non-negative). If  $\hat{l}_i > l_i$ , then MAA allocates more slots than  $i$  actually needs. This allocation does not increase agent  $i$ 's valuation but increases the price since now she will be charged for  $\hat{l}_i$  slots which is larger than the true length.

We get the claim by combining the above two arguments that hold for all  $i \in N$  irrespective of the reports of the other agents.  $\square$

To find the best slot for an agent, **max** checks the feasibility constraints and computes the allocation considering the slots’ valuation and current price. As there are  $(m - l_i + 1)$  possible allocations, **max** requires at most  $O(m)$  time for every agent  $i$ . Therefore, the following result on the complexity of **MAA** holds.

**THEOREM 6.8.** *MAA has time complexity  $O(mn)$ .*

We need a few more results to show the approximation factor of **MAA**. Due to paucity of space, we present those results (similar to the ones by Bartal et al. [2], modified according to **MIA**) in the full version of the paper.<sup>7</sup> These results help us prove the main theorem of this section.

**THEOREM 6.9.** *There exists a polynomial time ( $O(mn)$ ), incentive compatible, and individually rational mechanism to achieve  $O(km^{\frac{1}{k-2}})$  approximation to the optimal solution for **MIA**.*

The result above shows the existence of an approximately efficient mechanism that satisfies the other three desirable properties. The question of finding a lower bound on the approximation ratio remains open.

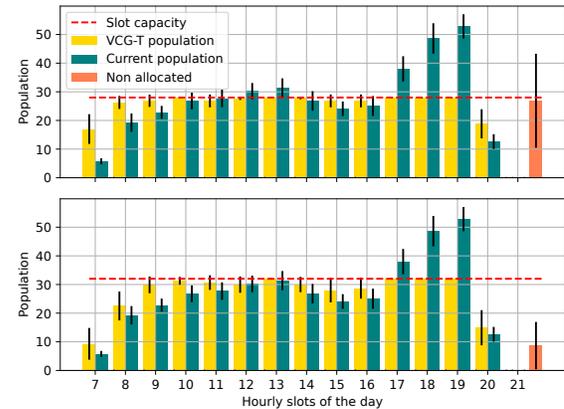
## 7 EXPERIMENTS

While the mechanisms presented in this paper satisfy several desirable properties of a social scheduling mechanism for indivisible single and multiple-slot jobs, its prioritizing profile for different classes of importance, costs of prioritization, and reduction in social congestion are not theoretically captured. In this section, we investigate these properties using real and synthetic datasets. The real dataset we collected from a store gave us only the checkout times. In the absence of the length information of the visits, we resorted to the single-slot job model (with hourly slots) and tested the performance of **VCG-T** on this data (§7.1). For multi-slot jobs, we simulated **MAA** to find the suboptimality and the reduction in the running time (§7.2). For these experiments, we consider three discrete levels of valuations of the agents denoted by 3, 2, and 1, which can be interpreted as **high**, **medium**, and **low** respectively. The numbers represent the agent’s valuation if they are allocated their most preferred slot. We used Gurobi [19] under an academic license for all experiments.<sup>8</sup>

### 7.1 Reduction in the social congestion

We consider real data of customer footfall in a general store that we have collected from the store (the dataset will be made publicly available post-publication). The dataset contains the customers’ hourly checkout (billing) time from 7 AM to 9 PM (opening hours of the store) for the whole month of July 2020. Since the dataset was anonymized for customer identification, we have assumed that the billing timestamps are unique users for a day. Given the store size, around 32 people an hour should be a fair capacity to maintain social distance. However, the data shows that the monthly average during the periods 5-6 PM, 6-7 PM, and 7-8 PM were 38.00, 48.63, and 52.83, respectively. Interestingly, the monthly average of the population in an hour is 26.5, which is well within the safety limits.

Therefore, this dataset works as a perfect example of where users can benefit significantly from social scheduling.



**Figure 1: Social congestion reduction, slot capacities 28 (top) and 32 (bottom).**

We divide the store opening hours into 14 hourly slots between 7 AM to 9 PM. For each day, the slots are sorted in decreasing order of footfall. We fix this order of the slots as the preference order of each agent for that day. The valuations of every customer for her most preferred slot is drawn from a distribution  $\{\mathbf{high}:0.1, \mathbf{medium}:0.3, \mathbf{low}:0.6\}$ . The valuations for the other slots are assumed to decrease with a multiplicative factor  $\delta = 0.8$  in the order of the slot preference, i.e., the valuation for the  $t$ -th most preferred slot of a **medium** agent will be  $2\delta^{t-1}$ . In this experiment, if a user is not allocated a slot on a certain day, she is given the option to update her importance and preferences for the next day. For the experiments, we assume that the user increases the importance by one level, e.g., **low** becomes **medium**, and keeps the slot preferences the same. After three consecutive days, if an agent is not allocated, she is considered ‘non-allocated,’ and alternative arrangements (e.g., home delivery) are made. Figure 1 shows the comparison of the average current population with that allocated by **VCG-T** for slot capacities of 28 (above) and 32 (below). The figures also show the daily non-allocated population in red. Each plot in this section shows the average values with a 95% confidence interval. The plots show the trade-off between better social distancing (lower slot capacity) and its cost (non-allocation). However, in both these cases, social congestion is reduced by approximately 50% during rush hours.

**VCG-T** also prioritizes the jobs at a cost. Figure 2 shows the allocated slot preference and the delays for the three different classes of valuations for slot capacity 30. It shows that a higher valuation comes with a higher delay.

### 7.2 Suboptimality vs Complexity Reduction (MAA)

The sub-optimality of **MAA** (Algorithm 2) was obtained for a worst-case scenario in §6. Here we investigate the sub-optimality of **MAA** and the amount of time it reduces w.r.t. a brute-force algorithm that finds the optimal allocation of the slots. The top plot of Figure 3 shows the percentage reduction  $((t_{OPT} - t_{MAA})/t_{OPT})$  in the running

<sup>7</sup>[https://www.cse.iitb.ac.in/~swaprava/papers/social\\_distancing\\_AAMAS23.pdf](https://www.cse.iitb.ac.in/~swaprava/papers/social_distancing_AAMAS23.pdf)

<sup>8</sup>The codes and collected dataset can be found here: [https://www.cse.iitb.ac.in/~swaprava/papers/codes\\_and\\_data.zip](https://www.cse.iitb.ac.in/~swaprava/papers/codes_and_data.zip).

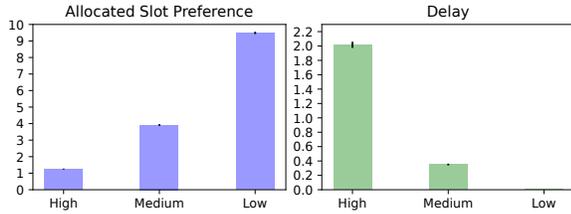


Figure 2: Priority and delay trade-off of VCG-T.

time of MAA compared to the optimal mechanism, where  $t_{OPT}$  and  $t_{MAA}$  are the running times of the optimal and MAA mechanisms, respectively. The bottom plot shows the ratio of the optimal social welfare to the welfare yielded by MAA. For each agent, a slot preference order is generated uniformly at random from the set of all feasible preference orders over the slots. Valuations for the most preferred slot are generated from a uniform distribution over {high, medium, low} and for other slots is assumed to decrease with a multiplicative factor  $\delta = 0.5$ , same as that in §7.1. The length of jobs is generated randomly with uniform distribution in the range 1 to  $m$ . The experiment is run with  $n = 6$ ,  $k = 5$ , and  $m$  varying from 3 to 8. For each value of  $m$ , the experiment is repeated for 100 valuation matrix and length vector pairs. The parameters  $m, n, k$  are chosen such that the optimal mechanism is computable in a reasonable time, yet the experiment yields an insightful result. We see that MAA reduces the running time by more than 99.5% and yields an approximation of roughly 1.75 on average.

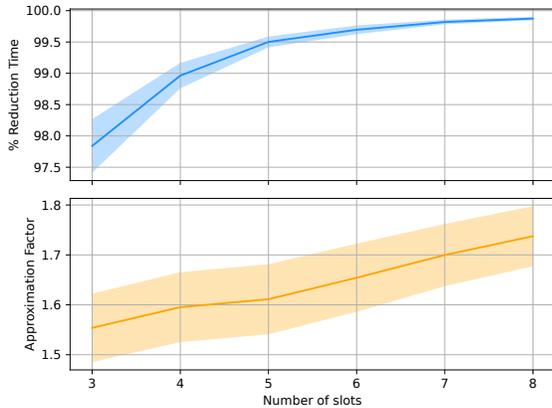


Figure 3: Running time and approximation factor trade-off for MAA.

### 7.3 Prioritizing profile and its cost

In this section, we investigate what the typical priority slots allotted to an agent of a specific class in VCG-T are. The top plot of Figure 4 shows the agents’ allocated slot preferences (mean with one standard deviation) versus the population ( $n$ ) plot where  $m = 5$ ,  $k = 4$ , and  $\delta = 0.65$ . The importance of an agent is picked uniformly at random. Values of  $n$  vary between 2 to  $1.1mk$  in steps of 1 (for a population beyond  $mk$ , some agents have to be dropped). The experiment is repeated 100 times for every  $n$ . The plot shows that the higher the importance, the lower the allocated slot preference for the agents, which is desirable.

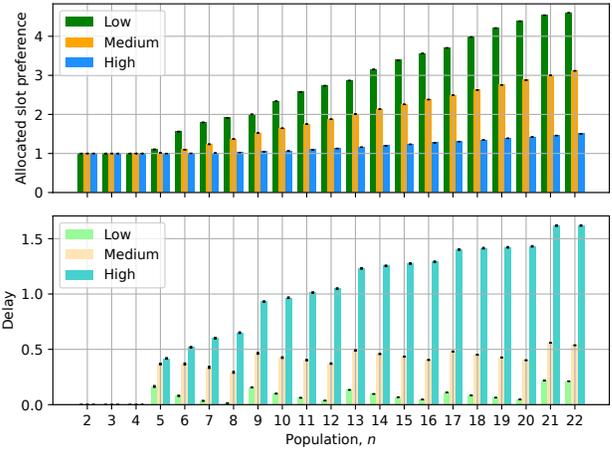


Figure 4: Priority-delay trade-off for VCG-T.

However, VCG-T does the prioritized allocations of the agents at the cost of their delays. The bottom plot of Figure 4 shows the corresponding delays decided by VCG-T for each of these three classes. The plot shows that an early slot allocation of an agent because of her importance also comes with a longer delay and shows the trade-off between these two decisions.

## 8 CONCLUSION

This paper presents results addressing some unique approaches toward social distancing that distinguish them from the extant literature. We use time delay as a method to incentivize individuals to report their true importance of a store visit. To the best of our knowledge, there is no work on the social distancing problem through scheduling the visitors efficiently and maintaining economic properties. Most works (and apps) monitor and enforce or roster individuals to maintain social distancing. The first-come-first-served type apps do not consider the importance of a visit. Hence, our work fills this critical gap in the social distancing literature by providing computationally efficient scheduling schemes, truthfully eliciting information on the importance of a visit, and prioritizing more important jobs. In other words, this approach solves the problem at the source, even before people leave for the store. Our algorithm shows how such approaches can efficiently reduce the stress on the stores by scheduling individuals on real-world data. There are some natural extensions possible for this work that do not immediately follow from the results and warrant a separate investigation. In the future, we like to extend this for dynamic arrivals and departures and multiple stores simultaneously.

## 9 ACKNOWLEDGMENTS

Nath is supported by a MATRICS grant (MTR/2021/000367) from Science and Engineering Research Board, Govt. of India, a Technocraft Center for Applied Artificial Intelligence fund (DO/2021-TCAI002-009), an Indian Institute of Technology Bombay fund (2022058).

## REFERENCES

- [1] R Arazi and A Feigl. 2021. Discontinuous transitions of social distancing in the SIR model. *Physica A: Statistical Mechanics and its Applications* 566 (2021), 125632.
- [2] Yair Bartal, Rica Gonen, and Noam Nisan. 2003. Incentive Compatible Multi Unit Combinatorial Auctions. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge* (University of Indiana, Indiana) (TARK '03). Association for Computing Machinery, New York, NY, USA, 72–87. <https://doi.org/10.1145/846241.846250>
- [3] Dirk Bergemann and Juuso Välimäki. 2010. The Dynamic Pivot Mechanism. *Econometrica* 78 (2010), 771–789.
- [4] Francis Bloch. 2017. Second-best mechanisms in queuing problems without transfers: The role of random priorities. *Mathematical Social Sciences* 90 (2017), 73–79.
- [5] Tilman Börgers. 2015. *An introduction to the theory of mechanism design*. Oxford University Press, USA.
- [6] Mark Braverman, Jing Chen, and Sampath Kannan. 2016. Optimal provision-after-wait in healthcare. *Mathematics of Operations Research* 41, 1 (2016), 352–376.
- [7] Paul Camion. 1965. Characterization of totally unimodular matrices. *Proc. Amer. Math. Soc.* 16, 5 (1965), 1068–1073.
- [8] Ruggiero Cavallo. 2021. Social Distancing Equilibrium. (2021). <http://www.eecs.harvard.edu/~cavallo/wp/cavallo-socdisteq.pdf> Unpublished.
- [9] Xujin Chen, Xiaodong Hu, Tie-Yan Liu, Weidong Ma, Tao Qin, Pingzhong Tang, Changjun Wang, and Bo Zheng. 2016. Efficient mechanism design for online scheduling. *Journal of Artificial Intelligence Research* 56 (2016), 429–461.
- [10] Samuel Cho. 2020. Mean-field game analysis of SIR model with social distancing. *arXiv preprint arXiv:2005.06758* (2020).
- [11] E. H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11 (1971), 17–33.
- [12] Peter Cramton, Yoav Shoham, Richard Steinberg, et al. 2004. *Combinatorial auctions*. Technical Report. University of Maryland, Department of Economics-Peter Cramton.
- [13] Min W Fong, Huizhi Gao, Jessica Y Wong, Jingyi Xiao, Eunice YC Shiu, Sukhyun Ryu, and Benjamin J Cowling. 2020. Nonpharmaceutical measures for pandemic influenza in nonhealthcare settings—social distancing measures. *Emerging infectious diseases* 26, 5 (2020), 976.
- [14] Sambuddha Ghosh, Yan Long, and Manipushpak Mitra. 2020. Prior-free online mechanisms for queueing with arrivals. *Economic Theory* (2020), 1–30.
- [15] Robert J Glass, Laura M Glass, Walter E Beyeler, and H Jason Min. 2006. Targeted social distancing designs for pandemic influenza. *Emerging infectious diseases* 12, 11 (2006), 1671.
- [16] Marko Gosak, Moritz UG Kraemer, Heinrich H Nax, Matjaž Perc, and Bary SR Pradelski. 2021. Endogenous social distancing and its underappreciated impact on the epidemic curve. *Scientific reports* 11, 1 (2021), 1–10.
- [17] Martin Grötschel, László Lovász, and Alexander Schrijver. 1993. Complexity, oracles, and numerical computation. In *Geometric Algorithms and Combinatorial Optimization*. Springer, 21–45.
- [18] T. Groves. 1973. Incentives in Teams. *Econometrica* 41 (1973), 617–631.
- [19] LLC Gurobi Optimization. 2020. Gurobi Optimizer Reference Manual. <http://www.gurobi.com>
- [20] Mohammad T. Hajiaghayi. 2005. Online Auctions with Re-Usable Goods. In *Proceedings of the 6th ACM Conference on Electronic Commerce* (Vancouver, BC, Canada) (EC '05). Association for Computing Machinery, New York, NY, USA, 165–174. <https://doi.org/10.1145/1064009.1064027>
- [21] Richard J Hatchett, Carter E Mecher, and Marc Lipsitch. 2007. Public health interventions and epidemic intensity during the 1918 influenza pandemic. *Proceedings of the National Academy of Sciences* 104, 18 (2007), 7582–7587.
- [22] Elias Koutsoupias. 2014. Scheduling without payments. *Theory of Computing Systems* 54, 3 (2014), 375–387.
- [23] Thomas Langer and Martin Weber. 2005. Myopic prospect theory vs. myopic loss aversion: how general is the phenomenon? *Journal of Economic Behavior & Organization* 56, 1 (2005), 25–38.
- [24] Ron Lavi and Noam Nisan. 2004. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science* 310, 1 (2004), 159–180. [https://doi.org/10.1016/S0304-3975\(03\)00391-8](https://doi.org/10.1016/S0304-3975(03)00391-8)
- [25] France Leclerc, Bernd H Schmitt, and Laurette Dube. 1995. Waiting time and decision making: Is time like money? *Journal of Consumer Research* 22, 1 (1995), 110–119.
- [26] Manipushpak Mitra. 2001. Mechanism design in queueing problems. *Economic Theory* 17, 2 (2001), 277–305.
- [27] Balázs Pejó and Gergely Biczók. 2020. Corona Games: Masks, Social Distancing and Mechanism Design. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Modeling and Understanding the Spread of COVID-19*. 24–31.
- [28] Timothy C Reluga. 2010. Game theory of social distancing in response to an epidemic. *PLoS Comput Biol* 6, 5 (2010), e1000793.
- [29] Kevin Roberts. 1979. The characterization of implementable choice rules. *Aggregation and revelation of preferences* 12, 2 (1979), 321–348.
- [30] Michael H Rothkopf, Aleksandar Pekeć, and Ronald M Harstad. 1998. Computationally manageable combinatorial auctions. *Management science* 44, 8 (1998), 1131–1147.
- [31] Alexander Schrijver. 2003. *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. Springer Science & Business Media.
- [32] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511811654>
- [33] Linda Thunström, Stephen C Newbold, David Finnoff, Madison Ashworth, and Jason F Shogren. 2020. The benefits and costs of using social distancing to flatten the curve for COVID-19. *Journal of Benefit-Cost Analysis* (2020), 1–27.
- [34] F.M.O Toxvaerd. 2020. *Equilibrium Social Distancing*. Cambridge Working Papers in Economics 2021. Faculty of Economics, University of Cambridge. <https://ideas.repec.org/p/cam/camdae/2021.html>
- [35] W. Vickrey. 1961. Counter Speculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 1 (1961), 8–37.
- [36] Annelies Wilder-Smith and David O Freedman. 2020. Isolation, quarantine, social distancing and community containment: pivotal role for old-style public health measures in the novel coronavirus (2019-nCoV) outbreak. *Journal of travel medicine* (2020).