

# Optimal Information Provision for Strategic Hybrid Workers

Sohil Shah<sup>1\*</sup>, Saurabh Amin<sup>1</sup> and Patrick Jaillet<sup>1</sup>

**Abstract**—We study the problem of information provision by a strategic central planner who can publicly signal about an uncertain infectious risk parameter. Signalling leads to an updated public belief over the parameter, and agents then make equilibrium choices on whether to work remotely or in-person. The planner maintains a set of desirable outcomes for each realization of the uncertain parameter and seeks to maximize the probability that agents choose an acceptable outcome for the true parameter. We distinguish between stateless and stateful objectives. In the former, the set of desirable outcomes does not change as a function of the risk parameter, whereas in the latter it does. For stateless objectives, we reduce the problem to maximizing the probability of inducing mean beliefs that lie in intervals computable from the set of desirable outcomes. We derive the optimal signalling mechanism and show that it partitions the parameter domain into at most two intervals with the signals generated according to an interval-specific distribution. For the stateful case, we consider a practically relevant situation in which the planner seeks to enforce in-person work capacity limits that progressively get more stringent as the risk parameter increases. We show that the optimal signalling mechanism in this case can be obtained by solving a linear program. We numerically verify the improvement in achieving desirable outcomes using our information design relative to the no-information and full-information benchmarks.

## I. INTRODUCTION

### A. Motivation

The COVID-19 pandemic sparked tremendous interest in practical tools to mitigate disease spread [1], [2], [3]. One can distinguish between two types of non-pharmaceutical interventions central planners use: (i) hard and (ii) soft.

*Hard interventions* are rigorously enforced measures, such as lockdowns or mask mandates, that prohibit certain actions that risk infectious spread. Balancing these interventions with their negative economic impact has recently been an active area of research [4], [5], [6]. These measures are particularly effective at containing disease spread in the early phases of a pandemic when cures and vaccines are unavailable [7]. However, these measures alone are not viable in the long-term as they tend to be economically and socially costly.

On the other hand, *soft interventions* aim to influence agents to choose less risky actions. Soft interventions are often much less costly to implement than hard interventions, but their design is a nascent area of study [8], [9], [10]. Our paper is motivated by the impact strategic provision of information has on public health. The basic idea is that a central planner can generate stochastic, possibly coarse

signals into the true value of an uncertain parameter that factors into agent utilities. The signals reshape agents' beliefs over this parameter and thereby influence their actions to induce a preferable outcome. In practice, signals correspond to the planner's reporting on the virulence of a disease or censoring data on the prevalence of the disease.

Specifically, we focus on how to optimally provision information to strategic hybrid workers. A workforce has  $K$  groups of non-atomic agents, each deriving value from working in-person. However, there is a stochastic cost associated with becoming infected that increases with the number of in-person contacts and an unknown continuous stochastic parameter,  $\theta \in \Theta$ , measuring the disease's infectious risk. The central planner has imperfect information about  $\theta$ , but can publicly generate signals into the true value of  $\theta$  using an information mechanism. On observing the signal, agents' public belief over  $\theta$  is updated and they simultaneously make equilibrium choices on where to work. The planner is strategic in that it values a set of outcomes as desirable based on the true value of  $\theta$  and some prescribed public health criteria. The novelty of our information provision problem (see Section II) is that these criteria only need to be described by a desirable set of outcomes for each parameter value. We distinguish between stateless and stateful objectives; in the former the set of desirable outcomes does not change as a function of the risk parameter, whereas in the latter it does. The central planner's objective is to maximize the probability that a desirable outcome is achieved. We investigate how to optimally provision information and evaluate its effectiveness to maximize the planner's objective.

In Section III, we derive the optimal signalling mechanisms for stateless objectives. We first characterize the equilibrium that agents will achieve as a function of their beliefs over  $\theta$  in Proposition 1. We observe that the set of possible equilibria is restricted to a one-dimensional manifold in the  $K$ -dimensional simplex. The intersection of this manifold and the desirable set of outcomes is precisely the set of outcomes the planner can hope to induce via information provision. We characterize the mapping between posterior mean beliefs over  $\theta$  and the achieved equilibrium in Lemma 1, and then we derive the optimal mechanism in Theorem 1. We find that the signalling depends on the position of the prior mean belief over  $\theta$  relative to the subset of posterior mean beliefs in  $\Theta$  that induce desirable outcomes. The mechanism uses at most two signals and partitions the parameter domain into at most two intervals with the signals generated according to an interval-specific distribution. This result adds to a growing literature on the optimality of these *monotone partitional signalling* mechanisms [11], [12], [13].

\*This material is based upon work supported in part by the NSF Graduate Research Fellowship under Grant No. 1745302, the C3.ai Digital Transformation Institute, and AFOSR grant FA9550-19-1-0263.

<sup>1</sup>Laboratory for Information and Decision Systems, Cambridge, MA 02139, USA {sshah95,amins,jaillet}@mit.edu

We next consider stateful outcome sets that enforce in-person capacity limits that progressively get more stringent as  $\theta$  increases in Section IV. For simplicity, we consider  $\Theta$  to be discrete and identify a linear program that solves for an optimal signalling mechanism (Proposition 2).

In Section V, we present a numerical comparison of our signalling mechanisms against two benchmarks: (i) revealing the true  $\theta$  (full information) and (ii) revealing nothing (no information). Our signalling mechanisms improve on these benchmarks for both stateless and stateful capacity objectives. Due to space limitation, all proofs are deferred to [14].

### B. Related Literature

Information design has been active area of research in the economics community spawned by the seminal results shown in [15] (see [16], [17], [18] for useful reviews).

Our contribution adds to the literature on information provision for multiple agents over a continuous unknown parameter. The authors of [19] show the equivalence between signalling mechanisms over continuous parameters and mean-preserving contractions of the parameter's prior distribution. The authors in [20] use this insight to characterize an optimal recommendation system for a discrete number of agents that are uncertain about a system state and must reveal a privately held type. Their results characterize how the designer must choose his recommended actions to maximize his quasilinear objective over these actions while retaining incentive compatibility. We instead focus on a setting with non-atomic agents and a broader class of objectives that need only be represented by sets of outcomes.

With regard to the design of soft interventions to mitigate disease spread, the authors of [9] consider the optimal design of rotation schemes where in-person workers alternate on some schedule. Related to information provision, [10] presents an activity-based model where the principal informs agents of the infection rate and identifies when full disclosure maximizes society's expected welfare. Most closely related to our work, [8] also investigates how central planners can optimally provision information about unknown stochastic risk during a pandemic. In [8], the central planner seeks to globally optimize over an objective function that takes a weighted average of functional measures of the economic health and infectious health. In contrast, our setup allows for a larger class of objectives that can incorporate other practically relevant factors such as peer effects or strict capacity constraints. Our results also go beyond a binary distribution over the infectious risk by considering the unknown parameter as being drawn from any bounded, continuous distribution. There is also significant recent literature on infection dynamics in mean-field games [21], [22], however our paper focuses on information design after a stochastic representation of infectious risk has been obtained. Finally, having more value towards the latter phases of a pandemic, our insights into soft intervention design should be considered in complement to the existing work studying the design of hard interventions to curtail initial infectious spread [4], [5], [6], [23].

## II. MODEL AND PROBLEM FORMULATION

We consider a model where strategic hybrid workers (agents) choose whether or not to work at a shared workspace amidst a pandemic of a communicable disease. Agents face uncertainty over the infectiousness of the disease (parameter), denoted by  $\theta$ , and thus over their likelihood of incurring the cost of becoming ill. In Sec. II-A, we describe agents' utilities and their dependence on the unknown parameter.

A central planner overseeing the agents maintains preferences over the aggregate outcomes of agents. The preferences may or may not depend on the true parameter. For a set of desirable agent outcomes for each value of the parameter, the planner seeks to maximize the probability that the equilibrium outcome lies in the set corresponding to the true parameter. Despite facing the same uncertainty over the parameter as the agents, the planner seeks to implement a mechanism that stochastically generates a publicly disclosed signal into the parameter's true value. The signal is revealed before agents make their choices, thereby influencing public belief over the parameter and affecting the agents' equilibrium choices. In Sec. II-B, we formalize this class of signalling mechanisms and identify the planner's problem of designing the optimal mechanism.

### A. Agent Behavior and Infectious Risk

We consider a unit mass of risk-neutral, Bayesian-rational agents. Each agent simultaneously commits to either working in-person at a shared workspace ( $\ell_O$ ) or working remotely ( $\ell_R$ ). Each agent belongs to a group  $k \in [K] := \{1, \dots, K\}$ . We represent the size of each group using the mass vector  $\mathbf{x} := (x_1, \dots, x_K) \in \mathbb{R}_+^K$  with the mass of agents belonging to each group  $k$ ,  $x_k$ , being common knowledge. Agents' choices result in a *remote mass* vector  $\mathbf{y} := (y_1, \dots, y_K)$ , where  $0 \leq y_k \leq x_k$  is the mass of agents from group  $k$  who elected to work remotely. The *in-person mass* vector is then  $\mathbf{x} - \mathbf{y}$ , with  $x_k - y_k \geq 0$  agents from group  $k$  working in-person. Letting  $\|\cdot\|$  correspond to the  $L^1$ -norm, note that  $\|\mathbf{x}\| = 1 = \|\mathbf{y}\| + \|\mathbf{x} - \mathbf{y}\|$ .

Any agent in any group  $k$  who works remotely faces no infectious risk, but also receives no benefit from the work environment. Consequently, their utility  $u_k(\ell_R, \mathbf{y}; \theta) = 0$  independent of  $\theta$  and  $\mathbf{y}$ .

Agents from group  $k$  that work in-person, however, each receive a commonly-known benefit  $v_k$  where  $v_k > 0$ . The benefit encapsulates the personal value agents derive from working in-person net any travel costs. Without loss of generality, we assume that  $v_k$  are distinct and we index the groups so that  $v_k$  is strictly decreasing in  $k$ .

Agents that work in-person face both an uncertain risk of contracting the disease per infectious contact and also face uncertain costs in the event they do become ill. Both the transmissivity and severity of the disease are natural sources of public uncertainty as they cannot be easily estimated by individual agents. The overall risk agents can expect to incur directly depends on the in-person mass in a known way, as the increased frequency of contacts and likelihood of infected individuals being present can be easily estimated.

Since agents are risk neutral, we capture the product of the uncertain risk per contact and uncertain costs into an uncertain *infectiousness parameter*  $\theta \in \Theta := [0, M]$  for some  $M \in \mathbb{R}_+$ . Critically, agents cannot perceive the value of the parameter  $\theta$  that was realized for this specific disease, but both the agents and planner all commonly believe a prior distribution  $F$  over  $\Theta$  from which  $\theta \sim F$ .

We model the (expected) stochastic infectious cost  $\beta(\theta, \mathbf{y})$  to be linear in  $\theta$  and decreasing in the remote mass (increasing for in-person mass). Specifically, for commonly-known, positive, differentiable functions  $c_1, c_2 : [0, 1] \rightarrow \mathbb{R}_+$  with  $c_1$  decreasing,  $c_2$  non-increasing and  $c_1(1) = c_2(1) = 0$ :

$$\beta(\theta, \mathbf{y}) := \theta c_1(\|\mathbf{y}\|) + c_2(\|\mathbf{y}\|) \quad (1)$$

We further detail this form of the infectious cost in [14]. Given the actions of all agents  $\mathbf{y}$  and the true parameter  $\theta$ , the utility of agents of group  $k$  working in-person is:

$$u_k(\ell_O, \mathbf{y}; \theta) = v_k - \beta(\theta, \mathbf{y}) \quad (2)$$

### B. Signalling and Central Planner Objective

The planner can implement a *signalling mechanism* that publicly provisions information into the true realization of  $\theta$  to all agents. Specifically, we assume that the planner publicly commits to and discloses a mechanism  $\pi = \langle \{g_\theta\}_{\theta \in \Theta}, \mathcal{S} \rangle$  where  $\{g_\theta\}_{\theta \in \Theta}$  are a set of probability distributions with each  $g_\theta$  distributed over the set of *signals*  $\mathcal{S}$ . If the true realization of the infectious risk is  $\hat{\theta} \in \Theta$ , then a signal  $s \in \mathcal{S}$  is randomly generated with probability  $g_{\hat{\theta}}(s)$ . The planner does not observe the true parameter value before committing to  $\pi$ . Thus, the planner cannot influence the mechanism after the true parameter is realized.

Prior to choosing their actions, all agents view  $s \in \mathcal{S}$  and symmetrically update their belief over  $\theta$  to a posterior distribution  $F_s$  where for all  $t$ :

$$F_s(t) = \mathbb{P}[\theta \leq t | s] = \frac{\int_0^t g_\theta(s) dF(\theta)}{\int_0^M g_\theta(s) dF(\theta)} \quad (3)$$

For any given  $\pi$ , it is useful to represent the corresponding probability with which each signal  $s \in \mathcal{S}$  is generated and the posterior mean of  $\theta$  upon viewing each signal  $s \in \mathcal{S}$ . If agents' strategies only depend on their mean beliefs over  $\Theta$  as in our case, this representation is equivalent to a *direct mechanism* where the planner performs the update on the agents' behalf and simply shares the updated posterior mean with all agents. We denote such a mechanism by the set of tuples  $\mathcal{T}_\pi = \{(q_s, \theta_s)\}_{s \in \mathcal{S}}$  where:

$$q_s := \int_0^M g_\theta(s) dF(\theta) \quad [\text{signal probability}] \quad (4)$$

$$\theta_s := \frac{\int_0^M \theta g_\theta(s) dF(\theta)}{\int_0^M g_\theta(s) dF(\theta)} \quad [\text{posterior mean}] \quad (5)$$

A particularly relevant subclass of mechanisms is one in which  $\Theta$  is partitioned into intervals and the parameters in each interval collectively map to a single interval-specific signal. Namely, we call a mechanism  $\pi$  *monotone partitional*

if there exists a finite partition of  $\Theta$ ,  $\mathcal{P} := \{\Theta_i\}_{i=1}^m = \{[t_{i-1}, t_i]\}_{i=1}^m$  for some  $m$  with  $t_0 = 0, t_m = M$  and some increasing sequence  $\{t_i\}_{i=1}^{m-1} \subseteq \Theta$  such that  $\mathcal{S} = [m]$  and for all  $\theta$ , the distribution  $g_\theta$  is such that  $g_\theta(s) = \mathbb{I}\{\theta \in [t_{s-1}, t_s]\}$ . Observe from (4) and (5) that for these mechanisms  $q_s = F(t_s) - F(t_{s-1})$  and  $\theta_s = \frac{\int_{t_{s-1}}^{t_s} \theta dF(\theta)}{q_s}$ . Signalling mechanisms of this form are prominently featured in the literature on optimal information provision for continuous parameters [11], [13], [20], [24].

Each instance of our game for a particular initial mass vector  $\mathbf{x}$  and public belief  $F$  over  $\Theta$  can be expressed in normal form  $\Gamma(\mathbf{x}, F) = \langle K, \mathbf{x}, (A_k)_{k \in [K]}, (u_k)_{k \in [K]} \rangle$ :

$K, \mathbf{x}$ :  $K$  groups of agents,  $x_k$  agents in each group

$A_k$ : group  $k$ 's agents' strategies,  $A_k = \{\ell_O, \ell_R\}, \forall k \in [K]$

$u_k$ : group  $k$ 's agents' utility functions  $u_k(a, \mathbf{y}; \theta)$  for strategy  $a \in A_k$ , aggregate strategy  $\mathbf{y}$  and parameter  $\theta$

To implement the mechanism we consider the notion of a (Nash) equilibrium where no agents have an incentive to deviate from where they decided to work. In particular, agents of group  $k$  at an equilibrium  $\mathbf{y}^*$  should always be choosing an action that generates utility equal to the best alternative of working remotely or in-person,  $\arg \max_{\ell \in \{\ell_O, \ell_R\}} \mathbb{E}[u_k(\ell, \mathbf{y}^*; \theta)]$ . In equilibrium, a group's agents spread across both actions if and only if this group's agents are indifferent between  $\ell_O$  and  $\ell_R$ . The following definition exploits this insight:

**Definition 1.** *The remote mass vector  $\mathbf{y}^*$  is a Nash equilibrium of our game  $\Gamma(\mathbf{x}, F)$  if and only if both:*

- (i)  $\mathbb{E}_{\theta \sim F}[u_k(\ell_O, \mathbf{y}^*; \theta)] > 0 \implies y_k^* = 0$
- (ii)  $\mathbb{E}_{\theta \sim F}[u_k(\ell_O, \mathbf{y}^*; \theta)] < 0 \implies y_k^* = x_k$

Once signal  $s \in \mathcal{S}$  is publicly revealed by the mechanism  $\pi$ , a signal-specific public belief  $F_s$  over the risk parameter is induced (see (3)) and we denote the equilibrium remote mass that agents choose in response by  $\mathbf{y}_\pi^*(s) \in \mathbb{R}_+^K$ .

In our model, the planner seeks to maximize the probability that the equilibrium that is achieved by the agents belongs to a set of desirable outcomes  $\mathcal{Y}(\theta) \subset \mathbb{R}^K$  that depends on the true parameter  $\theta$ . The definition of this set can arise out of practically relevant notions like capping the stochastic infection cost (i.e. the set  $\mathcal{Y}(\theta) := \{\mathbf{y} : \beta(\theta, \mathbf{y}) \leq \epsilon\}$ ) or ensuring a certain surplus level is maintained (i.e. the set  $\mathcal{Y}(\theta) := \{\mathbf{y} : \sum_{i=1}^K v_i(x_i - y_i) - (1 - \|\mathbf{y}\|)\beta(\theta, \mathbf{y}) \geq \epsilon\}$ ).

We focus explicitly on two classes of objectives. The first is what we denote as *stateless* objectives where  $\mathcal{Y}(\theta)$  does not depend on  $\theta$ . Hence, the objective set can be represented by a set  $\mathcal{Y}$  independent of  $\theta$ . The second is a *stateful* objective wherein we restrict the shapes of the objective sets to be nested as  $\theta$  increases. Particularly, we focus on a *scaled capacity* objective where  $\mathcal{Y}(\theta) = \{\mathbf{y} : \|\mathbf{y}\| \geq b(\theta)\}$  for some increasing function  $b(\cdot)$ . From a practical viewpoint, this stateful objective mimics capacity restrictions that are imposed on the shared workspace which are gradually scaled back as the infectiousness decreases.

We can express the (expected) objective  $V(\pi)$  of the

planner using  $\pi$  by:

$$V(\pi) = \mathbb{P}\{\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\theta)\} \quad (6)$$

For stateless objectives,  $V(\pi) = \mathbb{P}\{\mathbf{y}_\pi^*(s) \in \mathcal{Y}\}$  and the planner seeks to design  $\pi^*$  so that:

$$\pi^* \in \arg \max_{\pi: \langle \{g_\theta\}_{\theta \in \Theta}, \mathcal{S} \rangle} V(\pi) \quad (7)$$

### III. STATELESS INFORMATION DESIGN

In this section, we focus on designing  $\pi^*$  for a *stateless* planner objective  $\mathcal{Y}$ . In Sec. III-A, we compute the exact form of the equilibrium  $\mathbf{y}_\pi^*(s)$  as a function of agent's posterior mean beliefs  $\theta_s$ . We then derive the structure of the optimal mechanism used by the planner in Sec. III-B.

#### A. Equilibrium Characterization

Recall from (2) that agents' utilities have a linear dependence on  $\theta$ . Since agents are risk-neutral, it is clear from Definition 1 that the equilibrium outcomes  $\mathbf{y}_\pi^*(s)$  should only vary as a function of the posterior means  $\theta_s$ . Observe that agents that work in-person incur the same expected costs but may derive different benefits, so the agents deriving the highest value from working in-person should do so in an equilibrium. The following proposition formalizes this insight and characterizes the mass vector  $\mathbf{y}_\pi^*(s)$  and total in-person mass  $m(\theta_s) := 1 - \|\mathbf{y}_\pi^*(s)\|$  in equilibrium.

**Proposition 1.** *Let  $X_j := \sum_{i=1}^j x_i$  for all  $j \in \{0, 1, \dots, K\}$ , and  $v(u) := v_j$  if  $X_{j-1} \leq u < X_j$  for all  $u \in [0, 1]$ . For any signal  $s \in \mathcal{S}$  under mechanism  $\pi$ , the equilibrium in-person mass  $m(\theta_s) = \sup\{u : v(u) \geq c_1(1-u)\theta_s + c_2(1-u)\}$  and:*

$$(\mathbf{y}_\pi^*(s))_k = \begin{cases} 0 & \text{if } X_k \leq m(\theta_s) \\ m(\theta_s) - X_{k-1} & \text{if } X_{k-1} < m(\theta_s) < X_k \\ x_k & \text{if } X_{k-1} \geq m(\theta_s) \end{cases} \quad \square$$

Observe that the non-increasing function  $v(u)$  corresponds to the highest benefit amongst the mass of  $1-u$  agents with the smallest benefits  $v_i$ . Intuitively, the proposition sets the in-person mass  $m(\theta_s)$  by having as many agents work in-person as possible until the remaining agents find the infectious cost too high.

Proposition 1 also identifies a threshold-based structure that the equilibrium must obey. Specifically, it identifies a critical group  $k^*(s) = \inf\{k : X_k > m(\theta_s)\}$  where agents in groups  $k$  for  $k < k^*(s)$  all work in-person and agents in groups  $k$  for  $k > k^*(s)$  all work remotely. Hence, only a subset of all possible remote mass vectors  $\mathbf{y}$  are achievable in equilibrium. In fact, for any  $\pi$ , any equilibrium  $\mathbf{y}_\pi^*(s)$  must lie on a one-dimensional manifold  $\mathcal{Z}(\mathbf{x}) := \{c\mathbf{e}_k + \sum_{i=k+1}^K x_i \mathbf{e}_i | k \in [K], c \in [0, x_k]\}$  where  $\mathbf{e}_i \in \mathbb{R}^K$  with  $(\mathbf{e}_i)_k = 1$  if  $i = k$  else 0; see Fig. 1.

Moreover, as the posterior mean belief of the infectiousness  $\theta_s$  increases, agents are more cautious to choose in-person work because the infectious cost increases on expectation. The following lemma establishes that the equilibrium in-person mass is monotone and smooth in  $\theta_s$ .




Fig. 1: Example of  $\mathcal{Z}(\mathbf{x})$  (solid curve in blue) when  $K = 2$ .

**Lemma 1.**  *$m(\theta_s)$  is non-increasing, bounded and continuous in  $\theta_s$ .*  $\square$

Using Proposition 1 and Lemma 1, we conclude that as the posterior mean  $\theta_s$  increases, the equilibrium mass vector moves further from the initial point  $\mathbf{x}$  and closer towards the origin along  $\mathcal{Z}(\mathbf{x})$ . As a result, the design of the optimal signalling mechanism can be viewed as the problem of *maximizing the probability that posterior means* generated by the mechanism correspond to points  $\mathbf{y} \in \mathcal{Z}(\mathbf{x})$  that the planner deems desirable (i.e.  $\mathbf{y} \in \mathcal{Y}$ ).

#### B. Optimal Information Design

We now use Proposition 1 to solve the planner's problem of optimizing the signalling mechanism  $\pi$ . The manifold  $\mathcal{Z}(\mathbf{x})$  contains all possible equilibria  $\mathbf{y}_\pi^*(s)$  that could possibly be achieved for any  $\pi$  and any realized signal  $s$ . However, the manifold need not intersect with the goal set  $\mathcal{Y}$  at all, and in this case, no  $\pi$  can induce an equilibrium  $\mathbf{y}_\pi^*(s) \in \mathcal{Y}$ . Thus, we need to exploit the geometric relationship between  $\mathcal{Y}$  and  $\mathcal{Z}(\mathbf{x})$  to solve this design problem.

More generally, in relation to well-behaved sets  $\mathcal{Y}$ , our next result shows that the equilibria in  $\mathcal{Z}(\mathbf{x}) \cap \mathcal{Y}$  corresponds to a finite number of intervals of in-person mass. Specifically, we make the following assumption:

**Assumption A1.**  *$\mathcal{Y}$  is closed and convex.*

Moreover, we denote  $k(\mathbf{x}, u) = \min\{j > 0 : X_j > u\}$  as the lowest group that has not completely chosen in-person work after the mass of  $u$  agents with highest benefit choose in-person work. Likewise,  $z(\mathbf{x}, u)$  corresponds to the point on  $\mathcal{Z}(\mathbf{x})$  after  $u$  agents choose in-person work or equivalently  $z(\mathbf{x}, u) = (u - \sum_{k=1}^{k(\mathbf{x}, u)-1} x_k) \mathbf{e}_{k(\mathbf{x}, u)} + \sum_{k=k(\mathbf{x}, u)+1}^K x_k \mathbf{e}_k$ . Observe that  $z(\mathbf{x}, 1 - \|\mathbf{y}_\pi^*(s)\|) = \mathbf{y}_\pi^*(s)$  for all  $i$ .

**Lemma 2.** *Under Assumption A1, there exists  $\tilde{K} \leq K$  closed and disjoint intervals  $\Omega_i = [\omega_i^1, \omega_i^2] \subset [0, 1]$  for all  $i \in [\tilde{K}]$  with  $\mathcal{Y} \cap \mathcal{Z}(\mathbf{x}) = \cup_{i=1}^{\tilde{K}} \{z(\mathbf{x}, u) : u \in \Omega_i\}$ .*  $\square$

Observe that if  $\mathcal{Y} \cap \{z(\mathbf{x}, u) : X_{k-1} \leq u \leq X_k\}$  is empty, then  $\omega_k^1 > \omega_k^2$ . Trivially, notice that if  $\mathcal{Y}$  is contained in the interior of the positive orthant of  $\mathbb{R}^K$ , then there is at most only one non-empty segment of intersection  $\Omega_1$  since  $\{z(\mathbf{x}, u) : X_{k-1} \leq u \leq X_k\} \cap \text{int}(\mathbb{R}_+^K) = \emptyset$  for all  $k > 1$ . We note that Lemma 2 is generalizable beyond Assumption A1. Namely, if  $\mathcal{Y}$  is nonconvex, our results are valid as long as the  $\mathcal{Y} \cap \mathcal{Z}(\mathbf{x})$  is able to be represented by a finite number of intersecting line segments  $\Omega_i$ .

Since  $\mathbf{y}_\pi^*(s) \in \mathcal{Y}$  if and only if  $\mathbf{y}_\pi^*(s) \in \mathcal{Y} \cap \mathcal{Z}(\mathbf{x})$ , we can re-express (7) using Lemma 2 to instead depend on whether the remote masses lie in one of the intervals  $\Omega_i$ :

$$\begin{aligned} V^* &= \max_{\pi: \{\{g_\theta\}_{\theta \in \Theta}, \mathcal{S}\}} \mathbb{P}\{\mathbf{y}_\pi^*(s) \in \mathcal{Y} \cap \mathcal{Z}(\mathbf{x})\} \\ &= \max_{\pi: \{\{g_\theta\}_{\theta \in \Theta}, \mathcal{S}\}} \sum_{k=1}^{\tilde{K}} \mathbb{P}\{1 - \|\mathbf{y}_\pi^*(s)\| \in \Omega_i\} \\ &= \max_{\pi: \{\{g_\theta\}_{\theta \in \Theta}, \mathcal{S}\}} \sum_{k=1}^{\tilde{K}} \mathbb{P}\{\omega_k^1 \leq m(\theta_s) \leq \omega_k^2\} \quad (8) \end{aligned}$$

From Lemma 1, we note that since  $m$  is monotone, bounded and continuous, the preimage of this mapping over the closed interval  $[\omega_k^1, \omega_k^2]$  is another closed interval  $[\underline{\theta}_k, \bar{\theta}_k]$ , where  $m(\theta_s) \in \Omega_i$  if and only if  $\theta_s \in [\underline{\theta}_k, \bar{\theta}_k]$ . Without loss of generality, since the intervals  $\Omega_i$  are disjoint, we let  $\underline{\theta}_k$  be increasing in  $k$  which also implies that  $\bar{\theta}_k$  is increasing in  $k$ . Exploiting this fact and using a mechanism  $\pi$ 's corresponding direct mechanism  $\mathcal{T}_\pi$  as computed in (4) and (5), we can express the planner's program as:

$$V^* = \max_{\mathcal{T}_\pi: \{(q_s, \theta_s)\}_{s \in \mathcal{S}}} \sum_{k=1}^{\tilde{K}} q_s \mathbb{I}\{\underline{\theta}_k \leq \theta_s \leq \bar{\theta}_k\} \quad (9)$$

We make some immediate observations about this result. Firstly, observe that if the objective set  $\mathcal{Y}$  does not intersect with the manifold  $\mathcal{Z}(\mathbf{x})$ , then all the  $\Omega_i$  are empty and hence the objective in (8) is necessarily equal to 0. By contrast, if  $\mathbb{E}_{\theta \sim F}[\theta] := \mu \in [\underline{\theta}_l, \bar{\theta}_l]$  for some  $l$ , observe that a non-informative mechanism  $\pi$  has the property that  $\mathcal{S}$  is a singleton and  $\mathcal{T}_\pi = \{(1, \mu)\}$ . Hence, the planner achieves the maximal objective of 1 with the non-informative mechanism. However, if  $\mu \notin [\underline{\theta}_l, \bar{\theta}_l]$  for any  $l$ , then the non-informative mechanism achieves an objective of 0. Consequently, we find it useful to classify distributions  $F$  based on the positioning of the prior mean  $\mu$  relative to the intervals  $\{[\underline{\theta}_k, \bar{\theta}_k]\}_{k \in [\tilde{K}]}$  and solve for the optimal mechanism in each regime.

In order to characterize these regimes, consider the following increasing functions  $\underline{s}(t) = \mathbb{E}[\theta | \theta \leq t]$  and  $\bar{s}(t) = \mathbb{E}[\theta | \theta \geq t]$  which are computable from  $F$ . Then, we can define  $\delta(\alpha, \lambda, t) := \frac{\underline{s}(t)\lambda F(t) + \bar{s}(t)\alpha(1-F(t))}{\lambda F(t) + \alpha(1-F(t))}$  to represent the posterior mean belief when receiving a signal that is generated with probability  $\lambda$  for  $\theta \leq t$  and with probability  $\alpha$  for  $\theta > t$ . The range of  $\delta(\cdot, \cdot, \cdot)$  depends on the prior distribution  $F$  and describes the posterior mean beliefs that can be induced using a simple class of signalling mechanisms. We now introduce the following regimes and subregimes on the prior distribution  $F$ :

- (R1):  $\exists k$  such that  $\mu \in [\underline{\theta}_k, \bar{\theta}_k]$
- (R2):  $\exists k$  such that  $\mu \in (\bar{\theta}_k, \underline{\theta}_{k+1})$ 
  - (R2a):  $\exists l, k \in [\tilde{K}], t \in \Theta, \alpha, \lambda \in [0, 1]$   
s.t.  $\delta(\alpha, \lambda, t) \in [\underline{\theta}_k, \bar{\theta}_k]$   
and  $\delta(1 - \alpha, 1 - \lambda, t) \in [\underline{\theta}_l, \bar{\theta}_l]$
- (R3):  $\mu > \bar{\theta}_{\tilde{K}}$

$$(R4): \mu < \underline{\theta}_1$$

The above regimes are illustrated in Fig. 2. (R1) captures the aforementioned case where the prior distribution's mean  $\mu$  already lies in an interval. Both (R3) and (R4) capture distributions where  $\mu$  lies outside the convex hull of the intervals. (R2) considers when  $\mu$  does not lie in any interval, but falls in the gap between two intervals. Intuitively, in this case the planner seeks to spread mean beliefs outwards from the gap. However, this information design is not feasible for distributions  $F$  that are tightly concentrated in this gap and  $\delta(\cdot, \cdot, \cdot)$  will not take on many values outside this gap. Consequently, we focus on a more likely case – subregime (R2a) – where  $F$  has sufficient probability mass outside the gap between two intervals surrounding  $\mu$ . Before solving for an optimal mechanism in each regime, note that any direct mechanism  $\mathcal{T}_\pi$  that is a solution to (9) is not unique in general. This is immediate as any signal  $s \in \mathcal{S}$  can be forked into two signals  $s_1, s_2$  uniformly at random. This induces symmetric posterior means  $\theta_{s_1} = \theta_{s_2} = \theta_s$  with probability  $\frac{q_s}{2}$  each and hence achieves the same objective  $V^*$ . To eliminate such multiplicity, we provide the following lemma: it shows that we need not consider mechanisms that use a set of signals  $\mathcal{S}$  of size  $|\mathcal{S}| > \tilde{K} + 1$  since we require at most one signal to correspond to a posterior mean in each of the  $\tilde{K}$  intervals and one extra signal whose posterior mean lies outside these sets.

**Lemma 3.** *There exists an optimal mechanism  $\pi^* = \langle \{g_\theta^*\}_{\theta \in \Theta}, [\tilde{K} + 1] \rangle$  where  $V(\pi^*) = V^*$  such that  $\forall k \in [\tilde{K}]$ , both  $\theta_k \in [\underline{\theta}_k, \bar{\theta}_k]$  and  $\theta_{\tilde{K}+1} \notin [\underline{\theta}_k, \bar{\theta}_k]$ .  $\square$*

The insight given by Lemma 3 consequently reduces the search over all possible mechanisms  $\pi$  to those that use at most  $\tilde{K} + 1$  signals and can be represented by a  $\mathcal{T}_\pi$  with at most  $\tilde{K} + 1$  tuples. Using this reduction, we can without loss of optimality restrict the problem (9) to simpler signalling mechanisms  $\mathcal{T}_\pi$  with less than  $\tilde{K} + 1$  tuples. Notice that if we only specified  $\{(q_s, \theta_s)\}_{s \in [\tilde{K}+1]}$  instead of deriving it from a particular  $\pi$ , we might be optimizing over tuples that cannot actually be implemented by a mechanism  $\pi$ . To ensure implementability, we require that the distribution over the posterior means  $G$  with cumulative distribution  $G(t) = \sum_{i=1}^{\tilde{K}+1} q_i \mathbb{I}\{\theta_i \leq t\}$  is a mean-preserving contraction of  $F$ . The equivalence between the distributions  $G$  that are *mean-preserving contractions* of  $F$  and the distributions over posterior means that are implementable through a mechanism has been well-established [19], [25]. Formally,  $G \preceq F$  is a mean-preserving contraction if and only if  $\forall t \in \Theta$ ,  $\int_0^t F(t) dt \leq \int_0^t G(t) dt$ .

We now exploit the above reduction and solve for the optimal mechanism that solves (9). Define the increasing function  $h(\theta) := \sup\{s : \int_0^s F^{-1}(t) dt \leq s\theta\}$ . The following theorem characterizes the optimal objective and optimal mechanism across R1, R2a, R3 and R4 (we discuss the optimal direct signalling mechanism for R2 in the supplementary material [14]).




Fig. 2: The pdf for  $\theta$  and the intervals  $[\underline{\theta}_k, \bar{\theta}_k]$  (blue) with the solution for the four cases from Theorem 1 marked by the induced posterior means  $\theta_1, \theta_2$  (star). The posteriors are induced by mapping the probability mass to signals as marked (green, violet, yellow).

**Theorem 1.** *The optimal objective  $V^*$  and a corresponding optimal mechanism can be computed as follows for each regime:*

(R1):  $V^* = 1$  and  $\pi^*$  is monotone partitional with  $t_0 = 0$  and  $t_1 = M$

(R2a):  $V^* = 1$  and  $\pi^*$  has  $\mathcal{S} = \{1, 2\}$ ,  $g_\theta(1) = \lambda \mathbb{I}\{\theta \in [0, t]\} + \alpha \mathbb{I}\{\theta \in (t, M]\}$  and  $g_\theta(1) = (1 - \lambda) \mathbb{I}\{\theta \in [0, t]\} + (1 - \alpha) \mathbb{I}\{\theta \in (t, M]\}$ .


(R3):  $V^* = q_1^* := \min\{h(\bar{\theta}_K), \frac{M - \mu}{M - \bar{\theta}_K}\}$  and  $\pi^*$  is monotone partitional with  $t_0 = 0$ ,  $t_1 = F^{-1}(q_1^*)$  and  $t_2 = M$ .

(R4):  $V^* = 1 - q_2^*$  where  $q_2^* := \inf\{q \geq \frac{\theta_1 - \mu}{\theta_1} : q \leq h(\theta_1 - \frac{\theta_1 - \mu}{\theta_1})\}$  and  $\pi^*$  is monotone partitional with  $t_0 = 0$ ,  $t_1 = F^{-1}(q_2^*)$  and  $t_2 = M$ .  $\square$

Notice that in subcase (R1), an optimal mechanism is trivially monotone partitional with a single signal; in Figure 2, the entire probability mass over  $\Theta$  maps to a single signal. In (R3) and (R4) an optimal mechanism is also monotone partitional with two signals, where a threshold is specified by the theorem and the  $\Theta$  is partitioned into a “low” and “high” interval, where each interval corresponds to a signal. In (R2b), we create two signals and identify a threshold  $t$  where we split  $\Theta$  into two intervals at that threshold. In each interval, we then reveal the first signal with probability  $\alpha$  if  $\theta \leq t$  and with probability  $\lambda$  if  $\theta > t$ . Otherwise, we reveal the second signal. Consequently, the signals are generated according to an interval-specific distribution over just two signals. The analysis of these cases show that simple signalling mechanisms that use at most two signals are sufficient to achieve the optimum.

#### IV. STATEFUL INFORMATION DESIGN

In this section, we consider a stateful scaled capacity objective  $\{\mathcal{Y}(\theta)\}_{\theta \in \Theta}$  that is defined by a collection of sets  $\mathcal{Y}(\theta) = \{\mathbf{y} : \|\mathbf{y}\| \geq b(\theta)\}$  for some increasing function  $b(\cdot)$  in the domain  $\Theta$ . Such sets are practically relevant because they model a capacity limit on in-person work that gets



(a) Goal sets for stateful objective (b) A more general nested structure of goal sets.

Fig. 3

more stringent as the infectious risk increases. In this stateful setting, we cannot reduce the search over  $\pi$  by searching over the corresponding direct mechanisms  $\mathcal{T}_\pi$ . This is because a posterior mean  $\theta_s$  will induce a specific equilibrium outcome  $\mathbf{y}_\pi^*(s)$  but we cannot evaluate whether  $\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\theta)$  if we only know  $\theta_s$  and not the true  $\theta$ .

To proceed with analyzing the stateful case, for some  $N > 0$  we restrict  $\Theta := \{\nu_1, \nu_2, \dots, \nu_N\}$  with  $0 < \nu_1 < \dots < \nu_N$ . We denote  $\mathbb{P}[\theta = \nu_i] = p_i$  for all  $i \in [N]$ . To simplify notation, we also denote  $b(\nu_i) = b_i$  for all  $i \in [N]$  where  $b_i$  is strictly increasing in  $i$ ; see Fig. 3(a). Using (6), we can express the objective of the planner by:

$$V^* = \max_{\{g_\nu\}_{\nu \in \Theta, \mathcal{S}}} \mathbb{P}\{\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\nu)\} \\ = \max_{\{g_\nu\}_{\nu \in \Theta, \mathcal{S}}} \sum_{j=1}^N \sum_{s \in \mathcal{S}} p_j g_\nu(s) \mathbb{I}\{\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\nu_j)\} \quad (10)$$

As with the stateless design,  $\mathbf{y}_\pi^*(s)$  only depends on the induced posterior mean  $\theta_s$  and  $\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\nu_j)$  if and only if  $\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\nu_j) \cap \mathcal{Z}(\mathbf{x})$ . The latter follows if and only if  $b_j \leq \|\mathbf{y}_\pi^*(s)\| \leq 1$  or  $0 \leq m(\theta_s) \leq 1 - b_j$ . Using Proposition 1 and Lemma 1, observe that this corresponds to  $\theta_s \in [\gamma_j, \infty)$  for some constants  $\gamma_j$  that are strictly increasing in  $j$ . For completeness, we denote  $\gamma_0 = 0$  and  $\gamma_{N+1} = \infty$ . We can then reformulate (10) as follows:

$$V^* = \max_{\{g_\nu\}_{\nu \in \Theta, \mathcal{S}}} \sum_{j=1}^N \sum_{s \in \mathcal{S}} p_j g_\nu(s) \mathbb{I}\{\gamma_j \leq \theta_s < \infty\} \quad (11)$$

We highlight that despite these simplifications, our analysis is extensible to other practically relevant objectives aside from the scaled capacity objective. Suppose that the desirable sets satisfy Assumption A1 and are also (i) nested with  $\mathcal{Y}(\nu_{i+1}) \subset \mathcal{Y}(\nu_i)$  for all  $i \in [N - 1]$ , and (ii) half-bounded with  $\mathbf{x} \in \mathcal{Y}(\nu_i)$  for all initial mass vectors  $\mathbf{x}$  and all  $i \in [N]$ ; see Fig. 3(b). Then we arrive at an identical formulation to (11) since the minimal posterior mean beliefs for the agent outcomes to enter each  $\mathcal{Y}(\nu_j)$  is an increasing sequence  $\gamma_j$ .

Using a similar reduction as in stateless design, we can restrict our mechanism to induce a single posterior mean  $\theta_s$  to lie in each interval  $[\gamma_{s-1}, \gamma_s)$  for  $s \in \mathcal{S} := [N + 1]$ . These  $\theta_s$  consequently induce equilibrium outcomes that lie in  $\mathcal{Y}(\nu_j)$  if and only if  $\theta_s \geq \gamma_j$  or equivalently  $s \geq j + 1$ . This

implies that an optimal mechanism can be described by at most  $N + 1$  signals. In the proposition below, we provide a linear program (LP) to compute an optimal signalling mechanism for the stateful scaled capacity objective. We first introduce the LP:

$$\begin{aligned}
V^* &= \max_{\{z_{ji}\}} \sum_{j=1}^N \sum_{i=j+1}^{N+1} z_{ji} \\
\text{s.t.} \quad &\sum_{i=1}^{N+1} z_{ji} = p_j, & \forall j \in [N] \\
&\gamma_{i-1} \sum_{j=1}^N z_{ji} \leq \sum_{j=1}^N \nu_j z_{ji}, & \forall i \in [N+1] \\
&\sum_{j=1}^N \nu_j z_{ji} \leq \gamma_i \sum_{j=1}^N z_{ji}, & \forall i \in [N+1] \\
&0 \leq z_{ji} \leq p_j, & \forall i \in [N+1], j \in [N]
\end{aligned}$$

**Proposition 2.** *The optimal value  $V^*$  to the above LP is equal to  $\max_{\pi, \{\{g_\theta\}_{\theta \in \Theta, \mathcal{S}}\}} V(\pi)$  for a stateful scaling capacity objective as reduced in (11). For an optimal solution  $\{z_{ji}^*\}_{j \in [N], i \in [N+1]}$  of the LP, an optimal signalling mechanism is given by  $\pi^* = \langle \{g_\nu\}_{\nu \in \Theta, \mathcal{S}} \rangle$  where  $\mathcal{S} := [N+1]$  and  $g_{\nu_j}(i) = \frac{z_{ji}^*}{p_j}$  for all  $i \in \mathcal{S}, j \in [N]$ .  $\square$*

The optimal objective  $V^*$  and an optimal signalling mechanism  $\pi^*$  are hence computable in polynomial time.

If the  $b_j$  and hence  $\mathcal{Y}(\nu_j)$  were all identical, our setting becomes the stateless setting solved by Theorem 1. In this special case, all  $\nu_j$  correspond to a single  $\gamma = \gamma_j$  and only two posterior means  $\theta_1 < \gamma \leq \theta_2$  are required. If  $\mu = \sum_{j \in [N]} p_j \nu_j \geq \gamma$  then this corresponds to regime R1 and the non-informative mechanism is optimal. If  $\mu < \gamma$  this corresponds to regime R4 and likewise we recover a discretized interval-based signalling mechanism over  $\mathcal{S} = \{1, 2\}$  by solving the LP. In particular, we can verify that there exists a  $\hat{j} \in [N]$  such that for all  $j < \hat{j}$ ,  $g_{\nu_j}(1) = 1$  and for all  $j > \hat{j}$ ,  $g_{\nu_j}(2) = 1$ .

Finally, for any mechanism  $\pi$ , the probability of complying with the capacity limit conditioned on  $\theta = \nu_j$  can be represented by  $V_j(\pi) := \mathbb{P}\{\mathbf{y}_\pi^*(s) \in \mathcal{Y}(\theta) | \theta = \nu_j\}$ . Observe that  $V(\pi) = \sum_{j=1}^N p_j V_j(\pi)$  is a weighted average of the compliance for each possible realization of  $\theta$ . However, practically planners may face significantly more adverse risks of being overcapacitated as the infectiousness increases, and weight compliance heterogeneously across realizations of  $\theta$ . For this slightly more general objective, an optimal signalling mechanism for a different weighted sum  $\sum_{j=1}^N \alpha_j V_j(\pi)$  for  $\alpha_j \geq 0$  can likewise be solved using the LP by replacing the objective with  $\max_{\{z_{ji}\}} \sum_{i=j+1}^{N+1} \frac{\alpha_j}{p_j} z_{ji}$ .

## V. NUMERICAL EXPERIMENTS

We now present numerical experiments to illustrate the benefit of our optimally designed signalling mechanisms relative to no information and full information benchmarks. We first consider a stateless case where the in-person safe capacity the planner seeks to implement does not depend on

Probability of Compliance vs. Capacity

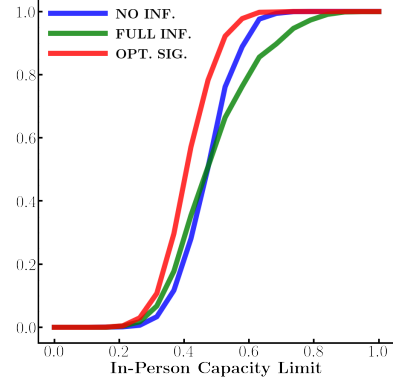


Fig. 4: Comparison between capacity compliance across non-informative, fully-informative and optimal signalling.

$\theta$ . We vary the capacity limit to demonstrate how the relative improvement with regard to each benchmark changes.

We implement an experiment where  $\mathcal{Y} = \{\mathbf{y} : \|\mathbf{y}\| \geq b\}$  and we vary the constants  $b$  between 0 and 1 in increments of  $\frac{1}{20}$ . The number of groups is fixed at  $K = 10$  and for each  $b$  we run 10,000 simulations by drawing  $\mathbf{x} \sim Unif(L_1)$ ,  $v_i \sim Unif[0, 10]$ ,  $\theta \sim Unif[5, 20]$  ( $\mu = \frac{25}{2}$ ). We evaluate the optimal mechanism given by Theorem 1 against two benchmarks. The first benchmark is the non-informative mechanism where agents receive no added information and consequently have a posterior mean belief of  $\mu$  over  $\Theta$ . The second benchmark is the fully-informative mechanism, where the true value of  $\theta$  is directly revealed so agents maintain posterior mean beliefs  $\theta'$  where  $\theta'$  is sampled from  $Unif[5, 20]$ . The resulting objective that measures the probability of compliance is averaged across simulations for each  $b$  and plotted in Figure 4. We indeed obtain that the optimal signalling mechanism strictly improves relative to both benchmarks. The improvement reduces as  $b \nearrow 1$  as the intersection between the outcomes achievable in equilibrium and the outcomes the planner deems acceptable reduces to the empty set. Naturally, the mechanism's ability to influence agents to a set of outcomes decreases as the set of outcomes becomes smaller. Likewise, as  $b \searrow 0$ , the set of acceptable outcomes grows to encompass all outcomes. Hence, the mechanism approaches full compliance, but so do the benchmark mechanisms. Consequently, the relative value of our optimally-designed signalling mechanism decreases as the set of outcomes the planner is willing to accept grows.

For the stateful capacity case, we highlight a specific instance where we scale-down the in-person capacity limit as the infectious risk increases. Consequently, the minimal posterior mean to induce outcomes in  $\mathcal{Y}(\nu_j)$  are represented by  $\gamma_j$  that are increasing in  $j$ . To gain intuition, we choose the specific set of tuples for  $N = 3$  where  $\{p_i\}_{i \in \{1,2,3\}} = (0.3, 0.3, 0.4)$ ,  $\{\nu_i\}_{i \in \{1,2,3\}} = (0.4, 0.6, 1.0)$ ,  $\{\gamma_i\}_{i \in \{1,2,3\}} = (0.5, 0.9, 1.2)$  and compare the performance of the optimal mechanism against the same two benchmarks as in the stateless case. In Table 1, we tabulate each information mechanism's performance. Observe that under no information, agents' mean belief over  $\theta$  is

Information Scheme	$V(\pi)$	$V_1(\pi)$	$V_2(\pi)$	$V_3(\pi)$
No Information	0.300	1	0	0
Full Information	0.000	0	0	0
Stateful Design	0.425	1	0.417	0

Table 1: Comparison of the overall probability of compliance and conditional probabilities of compliance for non-informative, fully-informative and stateful optimal signalling.

$\mu = \sum_{j=1}^N p_j \nu_j = 0.7$ . Since agents receive no signal, the equilibrium outcome is independent of  $\theta$  and agents only comply if  $\gamma_j < \mu$  which happens with probability  $p_1 = 0.3$ . Under full information, agents perfectly observe  $\theta$ , so the outcome is compliant with the limit if and only if  $\theta_j \geq \gamma_j$ . This occurs with probability 0 here. Thus, the stateful design significantly improves the probability of compliance relative to both benchmarks. We also observe that the stateful design does no worse in this setting than either benchmark in any of the conditional compliance probabilities  $V_j(\pi)$ .

## VI. CONCLUDING REMARKS

In this paper, we introduced a model to study information provision for strategic hybrid workers. Our model captures two key features: (a) an objective that aims to maximize the probability that the equilibrium outcome lies in a particular set which may or may not be state-dependent; (b) heterogeneous workers making strategic decisions to trade-off in-person work and infectious risk. We provided a complete description of the equilibria of the game and derived the optimal signalling mechanisms for the planner. Future work may focus on extending the presented results to designers using any smooth stateful objective functions. Preliminary investigation suggests approximate optimal design is achieved using analogous methods. Characterizing how the design is affected by agents with heterogeneous prior beliefs or alternative information sources is also a problem of significant interest.

## REFERENCES

- [1] T. Ji, H.-L. Chen, J. Xu, L.-N. Wu, J.-J. Li, K. Chen, and G. Qin, "Lockdown Contained the Spread of 2019 Novel Coronavirus Disease in Huangshi City, China: Early Epidemiological Findings," *Clinical Infectious Diseases*, vol. 71, no. 6, pp. 1454–1460, Sep. 2020.
- [2] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Analysis and Control of Epidemics: A Survey of Spreading Processes on Complex Networks," *IEEE Control Systems Magazine*, vol. 36, no. 1, pp. 26–46, Feb. 2016, conference Name: IEEE Control Systems Magazine.
- [3] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, "An Efficient Curing Policy for Epidemics on Graphs," *IEEE Transactions on Network Science and Engineering*, vol. 1, no. 2, pp. 67–75, Jul. 2014, conference Name: IEEE Transactions on Network Science and Engineering.
- [4] D. Acemoglu, V. Chernozhukov, I. Werning, and M. D. Whinston, "Optimal Targeted Lockdowns in a Multi-Group SIR Model," National Bureau of Economic Research, Working Paper 27102, May 2020, series: Working Paper Series.
- [5] D. Acemoglu, A. Fallah, A. Giometto, D. Huttenlocher, A. Ozdaglar, F. Parise, and S. Pattathil, "Optimal adaptive testing for epidemic control: combining molecular and serology tests," *arXiv:2101.00773 [physics, q-bio]*, Jan. 2021, arXiv: 2101.00773.
- [6] J. R. Birge, O. Candogan, and Y. Feng, "Controlling Epidemic Spread: Reducing Economic Losses with Targeted Closures," Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 3590621, May 2020.

- [7] R. M. Anderson, H. Heesterbeek, D. Klinkenberg, and T. D. Hollingsworth, "How will country-based mitigation measures influence the course of the COVID-19 epidemic?" *The Lancet*, vol. 395, no. 10228, pp. 931–934, Mar. 2020, publisher: Elsevier.
- [8] F. de Véricourt, H. Gurkan, and S. Wang, "Informing the Public About a Pandemic," *Management Science*, vol. 67, no. 10, pp. 6350–6357, Oct. 2021, publisher: INFORMS.
- [9] J. Ely, A. Galeotti, and J. Steiner, "Rotation as Contagion Mitigation," *Management Science*, vol. 67, no. 5, pp. 3117–3126, May 2021, publisher: INFORMS.
- [10] A. Hernandez-Chanto, C. Oyarzun, and J. Hedlund, "Contagion Management through Information Disclosure," Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 3988157, Dec. 2021.
- [11] P. Dworzak and G. Martini, "The simple economics of optimal persuasion," *Journal of Political Economy*, vol. 127, no. 5, pp. 1993–2048, 2019, publisher: The University of Chicago Press Chicago, IL.
- [12] A. Kolotilin, "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, vol. 13, no. 2, pp. 607–635, 2018, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE1805>.
- [13] Y. Guo and E. Shmaya, "The Interval Structure of Optimal Disclosure," *Econometrica*, vol. 87, no. 2, pp. 653–675, 2019, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15668>.
- [14] S. Shah, S. Amin, and P. Jaillet, "Optimal information provision for strategic hybrid workers," 2021, <https://sshah95.github.io/assets/pdf/inforprovisioncdc2022.pdf>. LIDS Technical Report.
- [15] E. Kamenica and M. Gentzkow, "Bayesian Persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, Oct. 2011.
- [16] O. Candogan, "Information Design in Operations," in *Pushing the Boundaries: Frontiers in Impactful OR/OM Research*, ser. INFORMS TutORials in Operations Research. INFORMS, Nov. 2020, pp. 176–201, section: 8.
- [17] D. Bergemann and S. Morris, "Information Design: A Unified Perspective," *Journal of Economic Literature*, vol. 57, no. 1, pp. 44–95, Mar. 2019.
- [18] E. Kamenica, "Bayesian Persuasion and Information Design," *Annual Review of Economics*, vol. 11, no. 1, pp. 249–272, 2019, eprint: <https://doi.org/10.1146/annurev-economics-080218-025739>.
- [19] M. Gentzkow and E. Kamenica, "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, vol. 106, no. 5, pp. 597–601, May 2016.
- [20] O. Candogan and P. Strack, "Optimal Disclosure of Information to Privately Informed Agents," Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 3773326, Jan. 2021.
- [21] S. Y. Olmez, S. Aggarwal, J. W. Kim, E. Miehling, T. Başar, M. West, and P. G. Mehta, "Modeling Presymptomatic Spread in Epidemics via Mean-Field Games," Nov. 2021.
- [22] D. La Torre, D. Liuzzi, R. Muggistro, and S. Marsiglio, "Mobility Choices and Strategic Interactions in a Two-Group Macroeconomic–Epidemiological Model," *Dynamic Games and Applications*, vol. 12, no. 1, pp. 110–132, Mar. 2022.
- [23] L. Cianfanelli, F. Parise, D. Acemoglu, G. Como, and A. Ozdaglar, "Lockdown interventions in SIR model: Is the reproduction number the right control variable?" *arXiv:2112.06546 [physics]*, Dec. 2021, arXiv: 2112.06546.
- [24] M. Ivanov, "Optimal Signals in Bayesian Persuasion Mechanisms," 2015.
- [25] D. Blackwell and M. Girshick, "Theory of Games and Statistical Decisions." New York: John Wiley and Sons, 1954.