



## 1. Introduction

In an epidemic, healthcare resources (e.g., ventilators, ICU beds and healthcare workers) are scarce and, what is worse, limited resources may not be used efficiently. The surging demand for care often leads to healthcare congestion, which can have significant welfare consequences. As we have seen in the COVID-19 outbreak, government policies to fight the epidemic are targeted mainly at reducing hospital congestion (i.e., by “flattening the curve”).

Another distinctive feature of an epidemic is the shortage of information. Information about the novel virus, such as its modes of transmission and the severity of the disease, is difficult to obtain. At the early stage of an epidemic outbreak, the general populace relies on public authorities for that information. Existing studies confirm that public information disclosure about the severity of the disease and recommendations for social distancing have been effective in shaping individuals’ beliefs and behaviors during COVID-19 outbreak, thereby determining the ultimate welfare consequences of the epidemic.<sup>1</sup>

What role does information disclosure play in healthcare congestion? Is full disclosure always the best disclosure policy in terms of protecting the interests of the general public? If not, what would be the benefits of coarse information disclosure, and what is the optimal choice for government authorities?

This paper addresses these questions. We construct a stylized model to understand the reasons for the congestion in the healthcare system and its welfare consequences. In our model, the healthcare system has a capacity constraint and can offer only a limited number of hospital admissions. However, a large number of agents are susceptible to the virus but may not know if they are infected and, even if they are, whether the infection will develop into serious symptoms that necessitate hospitalization. If that happens, then without hospital care, the agent will bear a cost. This common cost  $s$ , referred to as the *severity of the disease*, represents the underlying *state* of the economy.<sup>2</sup> This cost can be avoided only if the agent goes to the hospital and is admitted for care. Uncertain about how severe the disease could be, each agent chooses whether to visit the hospital based on the common belief about the disease’s severity and the private likelihood that hospitalization is needed.

Congestion occurs if the total number of hospital visits exceeds the healthcare system’s fixed capacity. At the early stage of an epidemic, due to the lack of knowledge about the novel virus, the healthcare sector does not have an effective means to pre-screen visitors. We assume that, in congestion, a hospital randomly admits the visitors up to its capacity. A critical feature of the congestion game is that one agent’s hospital visit imposes congestion externalities on others, because that visit lowers others’ chances of admission during congestion.

We solve for the unique equilibrium, in which each agent chooses to pay a fixed cost to visit the hospital if and only if the private likelihood that hospitalization will be needed surpasses a threshold. This threshold is determined by the expected severity level,  $\tilde{s}$ . When the disease is expected to be more severe, the threshold decreases and more agents choose to visit the hospital. Therefore, congestion occurs only if  $\tilde{s}$  exceeds a cutoff level,  $s_0$ , which we refer to as the *exhaustion level*. Whenever congestion occurs, those unlikely to need hospitalization get access to hospital care, while some agents that are likely to need hospitalization are not admitted. This misallocation of healthcare resources, together with the wasted visiting costs, defines the efficiency loss associated with congestion.

<sup>1</sup> For details, see, among others, Simonov et al. (2020), Bursztyn et al. (2020), and Allcott et al. (2020)).

<sup>2</sup> One interpretation of the disease’s severity is the expected mortality rate caused by this infectious disease.

Next, we introduce a principal who can commit to a public disclosure policy and whose goal is to maximize social welfare. Such a policy determines the expected severity level  $\tilde{s}$  for any true disease severity  $s$ . Clearly, congestion externality is the underlying reason for such a policy intervention. If there were no congestion, then private decision making would be socially efficient, and, thus, it would be optimal for the principal to adopt a full disclosure policy.

In the presence of congestion externality, an information disclosure policy improves total welfare only if it can alleviate congestion. For true severity level  $s > s_0$ , if the disclosure policy is able to induce a posterior belief, under which the expected severity level is exactly  $s_0$ , then the congestion externalities would be corrected, and the ex-post efficiency would be achieved without congestion. However, rational agents would form that expectation only if the principal also reports  $s_0$  at some severity levels lower than the exhaustion level  $s_0$ . The principal can always accomplish this by committing to the disclosure policy that reports  $s_0$  to the agents whenever the true severity level  $s$  appears in an interval  $[s_-, s_+]$  around  $s_0$  whose conditional mean is  $s_0$ .

The welfare gain from pooling states in  $[s_-, s_+]$  stems from avoiding congestion when  $s \in (s_0, s_+]$ . On the other hand, for  $s \in [s_-, s_0)$ , this exaggerates the disease severity and, therefore, distorts the agents' decision making. However, such distortion never results in congestion. We prove that the welfare losses caused by this distortion are dominated by the welfare gain from avoiding congestion, provided that the interval  $[s_-, s_+]$  is properly chosen. As such, full disclosure is never an optimal policy.

We fully characterize an optimal disclosure policy, assuming that the distribution of the likelihood of agents' needs for hospitalization has an increasing hazard rate. This optimal policy pushes the above-discussed intuition to the limit by identifying the largest interval  $[s_-, s_+]$  that fully exploits the benefit from avoiding congestion. We can interpret this optimal disclosure policy as a simple *censorship rule*. This policy censors all the states  $s$  between  $s_-$  and  $s_+$ , while fully revealing all other states. Depending on the parameters, it may be a middle censorship rule that censors only a strict intermediate range of states, or an upper censorship rule that censors all states above  $s_-$ .<sup>3</sup> Under the optimal policy, congestion never occurs when information is censored. For any true severity level  $s \in (s_0, s_+]$ , since the censorship rule induces a posterior belief  $\tilde{s} = s_0$ , it enables the healthcare system to run efficiently at its full capacity without congestion. Further, we show that this middle censorship rule is essentially the unique optimal disclosure policy.<sup>4</sup>

In the ongoing COVID outbreak, hospital congestion significantly threatens the efficiency of the healthcare system; meanwhile, the general public demands information transparency. Our model demonstrates that disclosure policy, if properly designed, can be an effective tool to alleviate healthcare congestion. From an ex-ante perspective, not revealing the disease's true severity, but pooling the severity levels around the exhaustion level, can in fact work to protect the interests of all agents. To examine the robustness of this simple idea, we further extend our benchmark model to investigate the cases with partially effective pre-screening and with limited information leakage. Our analyses confirm that this insight is largely robust to these considerations.

<sup>3</sup> The lower bound of the censoring range  $s_-$  is always strictly higher than the lowest possible severity level 0.

<sup>4</sup> Other optimal policies differ only in the states in which visiting the hospital is a strictly dominated strategy, as the cost is strictly higher than the benefit, regardless of what others do. Pooling those states together does not change the allocation of healthcare resources or ex-ante welfare. The middle censorship policy is, in fact, the most informative of all optimal policies. See Proposition 4 for details.

*Related literature* This paper contributes to a broader literature on policy interventions in epidemics. A recent growing economics literature looks at policy analysis within the SIR or SEIR framework (e.g., Acemoglu et al. (2020), Atkeson (2020), Alvarez et al. (2020), Eichenbaum et al. (2021), Fajgelbaum et al. (2021) and Jones et al. (2021)). This strand of the literature emphasizes the impact of lockdown and quarantine policies on contagion dynamics. The focus of our paper is different. We construct a model to rationalize hospital congestion and link that to agents' information and beliefs. In this sense, we provide a micro-foundation for the congestion externality usually taken as given in macroeconomic models.<sup>5</sup> In addition, we investigate how an information disclosure policy can be adopted to reduce healthcare congestion.

Our paper is also related to the literature on how to use limited or costly testing to extract as much information (regarding infection) as possible. For example, Deb et al. (2020) consider the method of targeted testing and targeted transfers to compensate for limited testing; Ely et al. (2021) study the optimal allocation of tests with different sensitivities and specificities among agents with differing infection risk; Lipnowski and Ravid (2021) investigate the optimal design for testing a pooled sample. In an extension of our baseline model, we consider pre-screening, which may include, but is not limited to, testing. Pre-screening produces useful results to distinguish between agents with different needs for hospitalization and to inform hospital admission decisions. To focus our attention on the disclosure policy about the common disease severity, we take pre-screening as exogenous and discuss how its precision can change the optimal disclosure policy.<sup>6</sup>

In terms of methodology, we apply the information design approach, initiated by Rayo and Segal (2010) and Kamenica and Gentzkow (2011), to the healthcare congestion problem in epidemics. The information design problem we consider features a group of agents with differing needs for hospitalization, who play a congestion game with strategic substitution, and a benevolent principal who chooses the information structure on a continuous state space to address congestion externality and restore efficiency.<sup>7</sup> Under mild conditions, we clearly characterize the properties of the principal's objective function, which endogenously arises from the healthcare congestion game. Given these properties, we apply the method in Dworczak and Martini (2019) and identify the optimal disclosure policy.

Disclosure policies involving various kinds of censorship rules have proven to be optimal in different economic contexts.<sup>8</sup> For example, in a recent work, Kolotilin et al. (2019) show that an upper censorship rule is optimal for all distributions of state if and only if the designer's indirect utility function is convex below some threshold and concave above that threshold, i.e., *S*-shaped. In our analysis, we find that the principal's indirect utility function is a piecewise convex function with a kink, i.e., *W*-shaped. Therefore, middle censorship can arise as an optimal policy.

*Structure* The rest of the paper is organized as follows. Section 2 presents the benchmark model of healthcare congestion. In Section 3, we set up the information disclosure problem and solve for the optimal disclosure policy. Section 4 discusses the robustness of our results, and Section 5

<sup>5</sup> See, for example, Berger et al. (2020) and Jones et al. (2021).

<sup>6</sup> See Section 4.1 for details.

<sup>7</sup> Another related paper is Das et al. (2017). In a quite different setup, they investigate the optimal information structure in traffic networks to reduce road congestion.

<sup>8</sup> For example, see Alonso and Câmara (2016) for voting, Goldstein and Leitner (2018) for banks' risk sharing, Gehlbach and Sonin (2014) and Ginzburg (2019) for media control, Inostroza and Pavan (2020) for regime change, and Zapechelnyuk (2020) for quality control.

concludes the paper. Proofs are relegated to the Appendix. All missing proofs can be found in the online appendix.

## 2. A simple model of healthcare congestion

In this section, we develop a stylized model of healthcare congestion. We model scarce public healthcare resources as the limited number of treatments and intensive care the healthcare system can provide. We focus on the strategic behavior of demanding public healthcare resource and link that to individual beliefs about the disease. This simple model helps us understand the causes of healthcare congestion, as well as its welfare consequences. It serves as the baseline for the information design problem in the next section.

### 2.1. Benchmark model

There is a unit mass of risk-neutral agents, indexed by  $i \in [0, 1]$ . During an epidemic, they are susceptible to the spread of a virus. The agents worry that a possible infection can lead to serious symptoms, requiring the treatment and intensive care provided by a hospital. Each agent decides whether or not to visit the hospital. Let  $a_i \in [0, 1]$  denote the probability that agent  $i$  visits a hospital, and  $n = \int_{i \in [0, 1]} a_i di$  denote the total number of hospital visits.

*Limited capacity and congestion* The capacity of the healthcare system, such as the number of doctors, ICU beds, and ventilators, is limited. We assume that the healthcare system can, at most, provide hospital care to  $\bar{n} \in (0, 1)$  visitors. When the total number of agents seeking admission is no greater than this capacity (i.e.,  $n \leq \bar{n}$ ), everyone will be admitted to the hospital. Only if an agent is admitted will they receive a comprehensive diagnosis about whether they are infected and if so, whether they will experience serious enough symptoms that require hospital care. When the total number of visitors exceeds hospital capacity (i.e.,  $n > \bar{n}$ ), congestion occurs. In this case, visitors will be randomly admitted, and, thus, each one will be admitted with probability  $\frac{\bar{n}}{n}$ .

We make this assumption to match the observation that, at the early stage of the epidemic, there is very little precise knowledge about the novel virus. Due to this lack of knowledge, the hospital does not have the reliable means to decide which patients will soon experience serious or even life-threatening symptoms.<sup>9</sup> In this case, hospital admissions are on a first-come, first-served basis, which can be taken as random admissions in our static setting.<sup>10</sup> Therefore, given the total number of visitors  $n$ , we can concisely write the probability of admission as

$$p(n, \bar{n}) = \min \left\{ 1, \frac{\bar{n}}{n} \right\}.$$

*Agent's type* Agents are uncertain about whether they are infected and if they are, whether the infection will soon escalate to serious symptoms. Specifically, each agent  $i$  has private information about the likelihood that they require hospital care. This likelihood, or the agent's *type*, is

<sup>9</sup> It is worth noting that, as we focus on the strategic hospital visits, we do not consider agents with life-threatening symptoms. For these agents, going to the hospital is a non-strategic choice, which never depends on hospital congestion or their beliefs about the severity of the disease. Moreover, the hospital will be able to identify those with life-threatening symptoms and give them priority for admission.

<sup>10</sup> In Section 4.1, we extend the benchmark model to consider a partially effective pre-screening protocol that is adopted for admission.

denoted by  $q_i \in [0, 1]$ .<sup>11</sup> Let  $H : [0, 1] \rightarrow [0, 1]$  be the cumulative distribution function of this  $q_i$ . We assume that it is common across all agents. By the law of large numbers, the distribution of this likelihood across the population is also  $H$ . Throughout the paper, we assume that  $H$  has a continuous and strictly positive density,  $h$ .

*Severity of the disease* We use  $s \in \mathcal{S} \equiv [0, \bar{s}]$  to denote the *severity of the disease*; that is, the payoff loss to any agent when they experience serious symptoms without hospitalization. Equivalently,  $s$  measures the welfare loss when hospital care is needed but that demand is not met.<sup>12</sup> The value of  $s$  is common to all the agents. Given the novelty of the virus, agents do not have precise information about the value of  $s$ . Rather, they rely on public disclosure to learn about  $s$ , and have a posterior belief that is characterized by the expected severity level  $\tilde{s} \in \mathcal{S}$ . In this section, we solve the agents' problem for any given expected severity level  $\tilde{s}$ . We will explore the optimal disclosure policy about the severity  $s$  in the next section.

*Payoffs* For a hospital visit, each agent must pay a fixed cost  $c$ . The cost can be interpreted, for instance, as the opportunity cost of the time spent on one's visit, the costs associated with transportation and making an appointment, and so on. For simplicity, we assume that this cost is identical across all agents.

In addition, each agent's payoff depends on whether he indeed needs hospitalization and whether he receives hospital care if he needs it. If an agent does not need hospitalization ex-post, their payoff, net of the visiting cost, is normalized to 0. This is independent of their decision to visit the hospital.<sup>13</sup> If an agent is infected and has serious symptoms that necessitate hospitalization ex-post, they will receive such care only if they choose to visit the hospital and are admitted. In this case, they will receive intensive care and necessary treatment so that the welfare loss  $s$  will not materialize. For that reason, their payoff, net of the visiting cost, is also 0.<sup>14</sup> If an agent indeed needs

TF.2313cD.0005Tc(nason,)Tj/F11Tf2.86010TD0Tc()Tj/F31Tf.2620TD.0005Tc(needs)are

Table 1  
Ex-post payoff.

		Hospital care is not needed	Hospital care is needed
not visit the hospital		0	$-s$
visit	admitted	$-c$	$-c$
	not admitted	$-c$	$-c - s$

precisely because of the limited capacity of the healthcare system. Consequently, agents' decisions about hospital visits are *strategic substitutes*. Conditional on there being congestion, an agent's incentive to go to a hospital decreases when others are more likely to go.

## 2.2. Equilibrium

All agents simultaneously and independently decide whether to visit the hospital, given their type,  $q_i$ , and the public belief,  $\tilde{s}$ . A strategy for agent  $i$  is a mapping  $a_i(\cdot; \tilde{s}) : [0, 1] \rightarrow [0, 1]$ . A Bayesian Nash equilibrium requires that (i) given the total number of hospital visits  $n(\tilde{s})$ , the hospital visit decision of every type of agent will be optimal:

$$a_i^*(q_i; \tilde{s}) \in \arg \max_{a_i \in [0, 1]} v_i(a_i; q_i, \tilde{s}, n(\tilde{s})), \quad \forall q_i \in [0, 1]; \quad (2)$$

and (ii) agents' decisions in turn determine the total number of hospital visits:

$$n(\tilde{s}) = \int_0^1 \left( \int_0^1 a_i^*(q; \tilde{s}) dH(q) \right) di. \quad (3)$$

The following proposition fully characterizes the unique Bayesian Nash equilibrium.

**Proposition 1.** *For any expected severity level  $\tilde{s} > 0$ , there is a unique Bayesian Nash equilibrium. In this equilibrium, agents play a symmetric cutoff strategy:*

$$a^*(q_i; \tilde{s}) = \begin{cases} 1, & \text{if } q_i \geq H^{-1}(1 - n(\tilde{s})), \\ 0, & \text{if } q_i < H^{-1}(1 - n(\tilde{s})), \end{cases} \quad (4)$$

where the total number of hospital visits  $n(\tilde{s})$  satisfies

$$n(\tilde{s}) = \begin{cases} 0, & \text{if } 0 \leq \tilde{s} \leq c, \\ 1 - H\left(\frac{c}{\tilde{s}}\right), & \text{if } c < \tilde{s} \leq \frac{c}{H^{-1}(1-\bar{n})}, \end{cases} \quad (5)$$

and is the unique solution to

$$n(\tilde{s}) = 1 - H\left(\frac{cn(\tilde{s})}{\tilde{s}\bar{n}}\right) \quad (6)$$

if  $\tilde{s} > \frac{c}{H^{-1}(1-\bar{n})}$ . Congestion occurs if and only if  $\tilde{s} > \frac{c}{H^{-1}(1-\bar{n})}$ .

Given the strategies taken by others and, thus, the total number of visits  $n(\tilde{s})$ , an agent with type  $q_i$  visits the hospital if their expected payoff from visiting exceeds that from not visiting

—i.e.,  $-c - q_i(1 - p(n(\tilde{s}), \bar{n}))\tilde{s} \geq -q_i\tilde{s}$ , or, equivalently,  $q_i \geq \frac{c}{p(n(\tilde{s}), \bar{n})\tilde{s}} \equiv \beta(\tilde{s})$ . This, in turn, implies that the total number of visits is

$$n(\tilde{s}) = 1 - H(\beta(\tilde{s})) = 1 - H\left(\frac{c}{p(n(\tilde{s}), \bar{n})\tilde{s}}\right). \quad (7)$$

From this equation, we can write the cutoff  $\beta(\tilde{s})$  as  $H^{-1}(1 - n(\tilde{s}))$ , as stated in (4). Equations (5) and (6) give a full characterization of the solution  $n(\tilde{s})$  to (7).

To understand (5) and (6), it is helpful to think about the extreme case with  $\bar{n} = 1$  as a reference. In this case, hospital congestion never happens. An agent with type  $q_i$  visits the hospital only if their expected payoff from visiting exceeds that from not visiting:  $-c \geq -q_i\tilde{s}$ . When  $\tilde{s} \leq c$ , the severity of the disease is very mild in expectation so that no one chooses to visit. Otherwise, when  $\tilde{s} > c$ , visiting is optimal only if the likelihood that hospital care is needed exceeds  $\frac{c}{\tilde{s}}$ . In this case, the total number of visits is  $1 - H(\frac{c}{\tilde{s}})$ .

Now, consider the case in which there is a capacity constraint  $\bar{n} < 1$ . When  $1 - H(\frac{c}{\tilde{s}}) \leq \bar{n}$ , or, equivalently,  $\tilde{s} \leq \frac{c}{H^{-1}(1-\bar{n})}$ , each agent simply behaves the same as they would in the case with unlimited capacity, because congestion will not occur given that all others behave in this way. Therefore, the equilibrium number of visits is  $n(\tilde{s}) = 0$  if  $\tilde{s} \leq c$ , and  $n(\tilde{s}) = 1 - H(\frac{c}{\tilde{s}})$  if  $c < \tilde{s} \leq \frac{c}{H^{-1}(1-\bar{n})}$ , as (5) claims.

However, when  $1 - H(\frac{c}{\tilde{s}}) > \bar{n}$ , or, equivalently,  $\tilde{s} > \frac{c}{H^{-1}(1-\bar{n})}$ , the situation becomes different. If the agents still behave as if there were unlimited capacity, then the total number of visitors would exceed  $\bar{n}$ . In this case, congestion would take place, and each visitor's probability of admission would be strictly lower than 1. This, in turn, lowers each agent's incentive to visit. Indeed, an agent with type  $q_i$  will visit the hospital if  $q_i > \frac{cn(\tilde{s})}{\tilde{s}\bar{n}}$  since the probability of admission now is only  $\frac{\bar{n}}{n(\tilde{s})}$ . Therefore, the equilibrium total number of visitors must satisfy (6).

In summary, the equilibrium number of hospital visits is increasing as the public belief increases. In equilibrium, congestion occurs only when the agents' belief exceeds  $s_0 \equiv \frac{c}{H^{-1}(1-\bar{n})}$ . We refer to  $s_0$  as the *exhaustion level*. This is the public belief, under which all health care resources are exhausted in equilibrium.

### 2.3. Total welfare

Next, to understand the welfare consequences of congestion, we calculate the equilibrium total welfare based on Proposition 1. This is the aggregated equilibrium payoff of the population.

**Lemma 1.** Suppose



**Proposition 2.** *For any true severity level  $s$ , the expected severity level that achieves ex-post efficiency is*

$$\tilde{s} = \begin{cases} \text{any value in } [0, c], & \text{if } s \leq c; \\ s, & \text{if } s \in (c, s_0]; \\ s_0, & \text{if } s > s_0. \end{cases}$$

The result should be intuitive. When the disease is very mild —i.e.,  $s \leq c$ — the cost of a hospital visit outweighs the benefit. Thus, ex-post efficiency requires that there is no hospital visit. This outcome can be achieved under any  $\tilde{s} \leq c$  since  $n(\tilde{s}) = 0$  according to Proposition 1.

When  $s > c$ , in the absence of limited hospital capacity, those agents with type  $q_i \geq \frac{c}{s}$  should go to the hospital since their benefit  $q_i s$  from doing so outweighs the cost  $c$ . In the presence of capacity constraint, there are two cases. First, if  $s \in (c, s_0]$

We assume that it is common knowledge that the severity level,  $s$ , is distributed according to a continuous and strictly increasing cumulative distribution function,  $G$ , over the interval  $[0, \bar{s}]$ . We also assume that  $\bar{s} > s_0$ , to avoid the trivial case of no congestion.<sup>16</sup> A disclosure policy is a random variable that is arbitrarily correlated with the true severity level,  $s$ . Every disclosure policy induces a joint distribution of the true severity level,  $s$ , and the public belief,  $\tilde{s}$ , which is the posterior mean. This joint distribution, in turn, determines the ex-ante expected total welfare

$$\begin{aligned}\tilde{\mathbb{E}}U(\tilde{s}, s) &= \tilde{\mathbb{E}}\left[-s\mathbb{E}q - cn(\tilde{s}) + s \min\{n(\tilde{s}), \bar{n}\}\mathbb{E}(q|q \geq \beta(\tilde{s}))\right] \\ &= -\tilde{\mathbb{E}}s\mathbb{E}q + \tilde{\mathbb{E}}\left[-cn(\tilde{s}) + \tilde{s} \min\{n(\tilde{s}), \bar{n}\}\mathbb{E}(q|q \geq \beta(\tilde{s}))\right],\end{aligned}\quad (9)$$

where  $\tilde{\mathbb{E}}$  is with respect to the joint distribution of  $s$  and  $\tilde{s}$ . The second equality comes from the fact that  $\tilde{\mathbb{E}}(s|\tilde{s}) = \tilde{s}$ . The first term in (9),  $-\tilde{\mathbb{E}}s\mathbb{E}q$ , is the expected welfare loss caused by the disease if there were no healthcare system. This is determined solely by  $G$  and is independent of the disclosure policy. For the second term, let

$$V(\tilde{s}) \equiv -cn(\tilde{s}) + \tilde{s} \min\{n(\tilde{s}), \bar{n}\}\mathbb{E}(q|q \geq \beta(\tilde{s})). \quad (10)$$

This represents the *conditional value of the healthcare system*, given any public belief  $\tilde{s}$ . The principal's problem is then to design a disclosure policy to maximize  $\tilde{\mathbb{E}}V(\tilde{s})$ , which depends only on the *marginal distribution* of the public belief  $\tilde{s}$ . It is well known that there exists a disclosure policy that induces a distribution  $\tilde{G}$  of the public belief  $\tilde{s}$  if and only if  $\tilde{G}$  is a mean-preserving contraction (PC hereafter) of distribution  $G$ .<sup>17</sup> For instance,  $G$  is an PC of itself, which is the distribution of  $\tilde{s}$  under full information disclosure. We can then formulate the principal's problem as choosing an PC of distribution  $G$  that maximizes the expected welfare.

$$\tilde{G}^*(\tilde{s}) = \begin{cases} G(\tilde{s}), & \text{if } 0 \leq \tilde{s} < s_-, \\ G(s_-), & \text{if } s_- \leq \tilde{s} < s_0, \\ G(s_+), & \text{if } s_0 \leq \tilde{s} < s_+, \\ G(\tilde{s}), & \text{if } s_+ \leq \tilde{s} \leq \bar{s}. \end{cases} \quad (12)$$

With the increasing hazard rate assumption on  $h$ , Proposition 3 identifies a range  $[s_-, s_+]$  of severity levels around the exhaustion level,  $s_0$ , such that simply pooling the severity levels in this range and fully revealing them outside this range is a principal's optimal disclosure policy. This optimal policy can be interpreted as a *censorship rule*, which censors only states in  $[s_-, s_+]$ . If  $s_+ < \bar{s}$ , it is a *middle censorship rule*, which censors an intermediate range of true severity levels and fully reveals both very low and very high states. If  $s_+ = \bar{s}$ , it is an *upper censorship rule*, which censors all states above a cutoff  $s_- > 0$  and reveals only the states below  $s_-$ . Under this disclosure policy, the agents' public belief after observing the censored message is precisely the exhaustion level,  $s_0$ , so that the healthcare system runs at exactly its full capacity, avoiding congestion when  $s \in (s_0, s_+]$ .

The proof of Proposition 3 is built on Theorem 1 in Dworczak and Martini (2019). To get the underlying intuition, it is crucial to understand how the conditional value of the healthcare system,  $V(\tilde{s})$ , changes as the public belief,  $\tilde{s}$ , varies. Using Proposition 1, we can rewrite the conditional value of the healthcare system in a more explicit form

$$V(\tilde{s}) = \begin{cases} 0, & \text{if } \tilde{s} \in [0, s_-), \\ \dots, & \text{if } \tilde{s} \in [s_-, s_0), \\ \dots, & \text{if } \tilde{s} \in [s_0, s_+), \\ \dots, & \text{if } \tilde{s} \in [s_+, \bar{s}]. \end{cases}$$

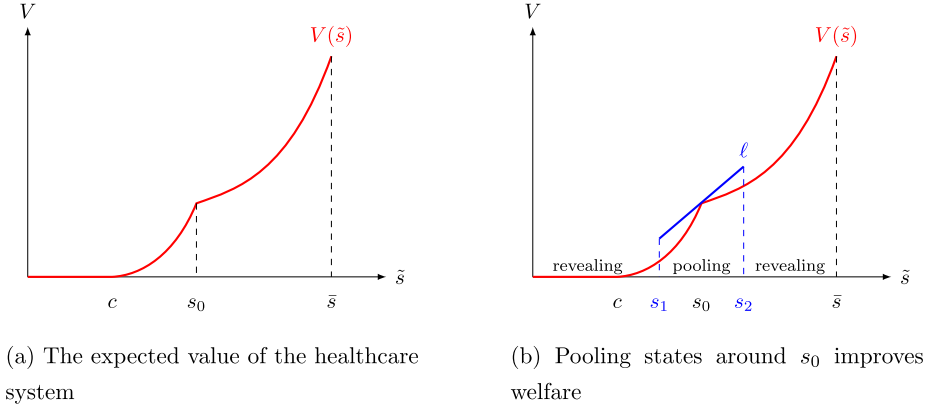


Fig. 1. Full information is not optimal.

(b) of Fig. 1 for an illustration. Moreover, the interval  $[s_1, s_2]$  can be properly chosen so that  $\mathbb{E}_G(s|s_1 \leq s \leq s_2) = s_0$ . Then,


$$\int_{s_1}^{s_2} V(s_0) dG(\tilde{s}) - \int_{s_1}^{s_2} V(\tilde{s}) dG(\tilde{s}) = \int_{s_1}^{s_2} \ell(\tilde{s}) dG(\tilde{s}) - \int_{s_1}^{s_2} V(\tilde{s}) dG(\tilde{s}) > 0.$$

It is easy to see that the left-hand side measures the welfare difference between two information disclosure policies: one is the policy that censors states in  $[s_1, s_2]$  and fully reveals otherwise; and the other is full disclosure. The difference being strictly positive means that this censorship rule does strictly better than full disclosure.


In view of Proposition 2, such a censorship rule has two effects on welfare. On the one hand, when  $s \in (s_0, s_2]$ , the censorship rule induces a public belief,  $s_0$ , under which the healthcare capacity is fully used without congestion. That leads to a welfare gain and achieves ex-post efficiency. On the other hand, when  $s \in [s_1, s_0)$ , the censorship rule exaggerates the severity level. Under the public belief  $s_0 > s$ , some agents will make unnecessary hospital visits, which results in a welfare loss. Our analysis demonstrates that, when the interval  $[s_1, s_2]$  is properly chosen, the welfare gain from avoiding congestion for  $s \in (s_0, s_2]$  will be of first order, which dominates the welfare loss induced by unnecessary hospital visits for  $s \in [s_1, s_0)$ . After all, the welfare gain is achieved by eliminating both misallocation and unnecessary visits, while the welfare loss occurs only due to some unnecessary visits, which do not lead to congestion or any misallocation.

Along the same line of reasoning, panel (b) of Fig. 1 also suggests that the total welfare can be further improved by extending the censoring range. Intuitively, we can find the largest interval  $[s_-, s_+]$  around  $s_0$  that: i) satisfies  $\mathbb{E}_G(s|s_- \leq s \leq s_+) = s_0$ ; and ii) admits a straight line over  $[s_-, s_+]$  that coincides with  $V$  at  $s_0$  and is above  $V$  elsewhere. Under the additional assumption that the density,  $h$ , has an increasing hazard rate, the second part of Proposition 3 shows that pooling the states in this  $[s_-, s_+]$  interval and fully revealing otherwise is not only welfare improving compared to full disclosure, but is also an optimal disclosure policy.

Fig. 2 gives a graphical illustration of the two possible cases of the censoring interval  $[s_-, s_+]$ . In panel (a), this interval is in the middle of the interval  $[0, \bar{s}]$ , while, in panel (b), this interval contains all states above  $s_-$ . The corresponding disclosure policies are middle censorship and



(a) Middle censorship



pooling states around the exhaustion level can be welfare improving. In this section, we relax two important assumptions in our benchmark — random admission and information monopoly — one at a time, to discuss the robustness of our results.

#### 4.1. Pre-screening

We have shown that hospital congestion leads to welfare loss, which leaves room for the principal to strategically disclose information. One of the major sources of welfare loss with congestion is the misallocation of scarce healthcare resources, due to the fact that the hospital cannot identify the patients who are the most likely to need care. Intuitively, in congestion, any method that helps the hospital screen patients and, thus, allocate limited healthcare resources more efficiently can be welfare improving. But as long as such a screening method can not completely eliminate the possibility of congestion, the principal can still improve welfare by manipulating information.

In this section, we confirm this intuition by extending our benchmark model to incorporate some effective pre-screening. It is reasonable to believe that as an epidemic evolves, the health-care sector learns more about the nature of the virus and gains a better understanding of the factors (e.g., different comorbidities) that can predict serious symptoms and the need for hospital care. Moreover, other protocols (e.g., a disease test) are developed to guide the hospital admissions.

We model pre-screening as a diagnostic evaluation that produces either a positive or negative result. We consider the case in which there is no false negative, so a patient who needs hospital care definitely tests positive.<sup>23</sup> However, an agent who does not need hospital care obtains a negative result with probability  $\gamma \in [0, 1]$ . Parameter  $\gamma$  measures the precision of the pre-screening. The larger the  $\gamma$ , the more precise the pre-screening. The previous benchmark corresponds to  $\gamma = 0$ . When  $\gamma = 1$ , we have perfect screening.

Because a negative result perfectly reveals no need for hospital care, the hospital admits only those who obtain a positive result. If the total number of visitors who test positive does not exceed hospital capacity  $\bar{n}$ , then they are all admitted. Otherwise, the hospital randomly admits those with positive results up to its capacity, as in the benchmark. Given no false negative, if  $Q \subset [0, 1]$  agents go to the hospital, the admission probability of an agent who needs hospital care is then

$$\left\{ 1, \frac{\bar{n}}{\int_Q (q + (1 - \gamma)(1 - q))h(q) dq} \right\}.$$

Clearly, this probability is weakly increasing in  $\gamma$  because, when the pre-screening is more precise, it is more likely to screen those who do not need hospital care.

We assume that the total demand for hospital care in the society exceeds the hospital capacity, i.e.,  $\mathbb{E}q > \bar{n}$ . Under this assumption, congestion is possible regardless of  $\gamma$ . For instance, if all agents choose to go to the hospital, the total number of positive results will exceed  $\bar{n}$  even with perfect pre-screening.<sup>24</sup>

<sup>23</sup> Such a specification makes this extension more aligned with the benchmark and facilitates the comparison. The main insight that full information disclosure is not optimal as long as pre-screening does not completely rule out the possibility of hospital congestion still holds if we incorporate the false negative.

<sup>24</sup> If  $\mathbb{E}q \leq \bar{n}$ , there exists  $\underline{\gamma} \in (0, 1]$  such that  $\int_0^1 (q + (1 - \underline{\gamma})(1 - q))h(q) dq = \bar{n}$ . Then, when the pre-screening is sufficiently precise — i.e.,  $\gamma \geq \underline{\gamma}$  — congestion is completely eliminated.

As in the benchmark, after the agents form public belief  $\tilde{s}$  about the disease's severity level, they simultaneously and independently decide whether to go to the hospital. Qualitatively similar to the benchmark, there is a unique Bayesian Nash equilibrium for each  $\gamma$ , in which agents use symmetric cutoff strategies.<sup>25</sup> The equilibria for different  $\gamma$  differ quantitatively in the exhaustion level  $s_{0,\gamma}$ . When the pre-screening is completely uninformative—i.e.,  $\gamma = 0$ — $s_{0,0}$  is exactly the benchmark exhaustion level  $s_0$ . As  $\gamma$  increases, the exhaustion level  $s_{0,\gamma}$  also strictly increases. This is intuitive, as with a larger  $\gamma$ , when making admission decisions, the hospital is more capable of identifying who needs hospital care. If the same set of agents choose to visit the hospital, the pool of agents from which the hospital considers admissions becomes strictly smaller, and, therefore, the capacity will not be exhausted under the public belief,  $\tilde{s} = s_{0,\gamma}$ .

Based on the equilibrium analysis, we can then investigate the principal's information disclosure, as in the benchmark. The next proposition summarizes the results. Let  $\bar{\gamma} \equiv \inf\{\gamma \in [0, 1] | s_{0,\gamma} \geq \bar{s}\}$  with the convention that  $\bar{\gamma} = +\infty$  if the set is empty.

**Proposition 5.** *For any  $\gamma < \bar{\gamma}$ , full information disclosure is not optimal. If in addition  $h$  is differentiable and has an increasing hazard rate, then censoring a certain interval  $[s_{-,\gamma}, s_{+,\gamma}]$  around  $s_{0,\gamma}$  is optimal. For any  $\gamma \geq \bar{\gamma}$ , full information disclosure is optimal.*

To understand how the optimal disclosure policy depends on the precision  $\gamma$ , recall that we assumed that  $s_0 < \bar{s}$ , or, equivalently,  $s_{0,0} < \bar{s}$ , in the benchmark model. As the exhaustion level,  $s_{0,\gamma}$ , is increasing in  $\gamma$ , this assumption does not guarantee  $s_{0,\gamma} < \bar{s}$  for all  $\gamma \in [0, 1]$ . Thus, we identify  $\bar{\gamma}$  such that  $s_{0,\gamma} < \bar{s}$  for any  $\gamma < \bar{\gamma}$ . In such cases, congestion still occurs in equilibrium under full disclosure. Therefore, for any  $\gamma < \bar{\gamma}$ , the optimal disclosure policy is qualitatively similar to the one in the benchmark model: censoring an interval around the exhaustion level  $s_{0,\gamma}$  whose conditional mean is exactly  $s_{0,\gamma}$ . In contrast, when  $\gamma \geq \bar{\gamma}$ , we have  $s_{0,\gamma} \geq \bar{s}$ . This means that congestion never occurs in equilibrium under full information disclosure. Consequently, full information disclosure is optimal according to Proposition 2.

The critical precision level,  $\bar{\gamma}$ , by definition, depends on the upper bound of the disease severity  $\bar{s}$ . If  $\bar{s}$  is sufficiently high, then  $\bar{\gamma} = +\infty$ , which means that  $s_{0,\gamma} < \bar{s}$  for all  $\gamma \in [0, 1]$ . In this case, it is worth noting that the corresponding censorship policy is optimal even under perfect pre-screening. This is the case because, although healthcare resources are always allocated efficiently due to perfect pre-screening, hospital congestion still leads to welfare loss due to unnecessary hospital visits, which can be mitigated through strategic disclosure. In this sense, even perfect pre-screening cannot completely replace the role of strategic disclosure in alleviating congestion.

#### 4.2. Information leakage

We now consider the possibility that the information about the true severity level of the disease is leaked to a fraction of the population. This could happen due to, for example, limited access to experts' opinions or restricted media coverage. Suppose that  $\tau \in (0, 1)$  fraction of agents are always informed of the true severity level  $s$ . The other  $1 - \tau$  fraction of agents remains uninformed as they are in the benchmark. Whether informed or not, each agent's probability of needing hospital care is still distributed according to  $H$ . Therefore, the overall type distributions

<sup>25</sup> We provide a detailed analysis in Section E in the online appendix.

among these two groups of agents are identically  $H$ . As before, we consider the case in which the principal's information disclosure is public. Under such a policy, both the informed and uninformed agents observe the public signal.<sup>26</sup> We also assume that there is no communication between these two groups of agents. This environment is common knowledge among all agents. The uninformed agents are aware that some fraction of the population has better information than they do, and the informed agents know that the uninformed agents observe only the public signal.

Obviously, the model reduces to the benchmark at the extreme case  $\tau = 0$ . At the other extreme,  $\tau = 1$ , all agents are fully informed, and, clearly, disclosure policy can never be effective. For any intermediate case  $\tau \in (0, 1)$ , the asymmetric information between these two groups of agents poses a challenge to our analysis. In the current setting, it is the uninformed agents' *posterior belief*, not just the *posterior mean* as in the benchmark, that matters for their equilibrium behavior. This is because that the uninformed agents must form a belief not only about the severity level, but also about the behavior of the informed agents. Nonetheless, the next proposition states that one of the key insights of Proposition 3 still carries over provided there is only a small fraction of informed agents.

**Proposition 6.** *There exists  $\bar{\tau} \in (0, 1)$  such that full information disclosure is not optimal for any  $\tau \leq \bar{\tau}$ . In particular, censoring a certain interval of the true severity levels around  $s_0$  does strictly better than full information disclosure.*

Proposition 3 shows that without information leakage, pooling a certain interval around  $s_0$  can eliminate congestion and improve social welfare. Under such a policy, the public belief is  $s_0$  after the censored signal, and each agent goes to the hospital if their type  $q_i \geq \frac{c}{s_0}$ . In the presence of information leakage, by pooling a properly designed interval around  $s_0$ , we show that it is still possible to induce the same *behavior* of the uninformed agents — each of them goes to the hospital if their type  $q_i \geq \frac{c}{s_0}$  after the censored message.<sup>27</sup> Proposition 6 verifies that such a disclosure policy is indeed welfare improving compared to full information disclosure, when the fraction of the informed agents is not too big.

Admittedly, the effectiveness of such a policy is diminishing as the fraction of informed agents becomes larger. On the one hand, the existence of the informed agents reduces the population whose behavior can be manipulated through strategic disclosure. On the other hand, and perhaps more importantly, such censorship has a counter effect on the informed agents' behavior. When the true state is above  $s_0$ , knowing that the uninformed agents have a lower incentive to visit the hospital given the censored signal, the informed agents will actually have a greater incentive to visit due to strategic substitution. Therefore, it is intuitive that such a censorship policy will no longer be effective when  $\tau$  is sufficiently large.

## 5. Concluding remarks

A main theme of economics concerns the efficient allocation of scarce resources. The epidemic outbreak, as a public health crisis, features a situation in which the free market with pricing competition cannot guarantee efficient resource allocation. At the same time, central planning

<sup>26</sup> For our result in Proposition 6, it does not matter whether the informed agents observe the signal or not.

<sup>27</sup> The uninformed agents' posterior mean may not be  $s_0$  after the censored message. See Section F.2 in the online appendix for the detailed construction of such a policy.



cannot achieve efficiency since the social planner does not have enough information about individuals' demand. In such an environment, congestion can pose a critical threat to economic efficiency.

In this paper, we construct a model to study the healthcare congestion problem. Information disclosure policy is found to be effective in mitigating the congestion and improving the efficiency of the healthcare system. In particular, pooling the states around the exhaustion level and fully revealing other states always dominates the policy of full disclosure, and, under mild conditions, it prevails as the optimal disclosure policy. We further show that this insight is robust to considering partially effective pre-screening and limited information leakage.

On the applied side, we restrict our attention to disclosing information about the disease severity and its effect on agents' hospital visits. Our results rely on the assumption that all agents are fully rational and Bayesian, which, admittedly, may not hold true during an epidemic. In addition, there are other dimensions of information that can be critical to both agents' strategic decision making and social welfare in epidemics. For instance, disclosing information about the infectiousness of the virus may change social-distancing choices, and revealing effectiveness and safety information about vaccines may matter for decisions of whether to be vaccinated. Further, through these channels, information may endogenously shape the distribution of agents' needs for hospitalization, which is taken as exogenous in our paper. On the theory side, the information design problem in our main model naturally boils down to a linear persuasion problem. However, we believe that the key insight gained from this study — pooling states around the exhaustion level can alleviate congestion and improve efficiency — should hold in a more general congestion setting, as is suggested by our analysis in Section 4.2, where the problem is no longer linear. We believe these are promising avenues for future research.

## Appendix A

**Proof of Proposition 1.** Suppose the strategy profile  $\{a_i^*(\cdot; \tilde{s})\}_{i=1}$  is a Bayesian Nash equilibrium. Consider agent  $i$  with probability of infection  $q_i$ . Agent  $i$ 's expected payoff from his visiting decision  $a_i$  is

$$v_i(a_i; q_i, \tilde{s}, n(\tilde{s})) = a_i(-c - q_i(1 - p(n(\tilde{s}), \bar{n}))\tilde{s}) + (1 - a_i)(-q_i\tilde{s}),$$

where  $n(\tilde{s})$ , which is the total number of hospital visits, satisfies (3). Therefore, we have

$$a_i^*(q_i; \tilde{s}) = \begin{cases} 1, & \text{if } q_i > \frac{c}{p(n(\tilde{s}), \bar{n})\tilde{s}}, \\ 0, & \text{if } q_i < \frac{c}{p(n(\tilde{s}), \bar{n})\tilde{s}}. \end{cases} \quad (14)$$

Since  $i$  is arbitrary, we immediately know that the equilibrium is symmetric and every agent uses the same strategy as in (14). By (3), we know that  $n(\tilde{s})$  satisfies

$$n(\tilde{s}) = 1 - H\left(\frac{c}{p(n(\tilde{s}), \bar{n})\tilde{s}}\right).$$

It remains to show that this equation has a unique solution  $n(\tilde{s})$  that is given by (5) or (6), depending on  $\tilde{s}$ . We leave the details to the online appendix. See Lemma A.2.  $\square$

**Proof of Lemma 1.** By Proposition 1, the aggregate welfare is

$$\begin{aligned} U(\tilde{s}, s) &= \int_{q < \beta(\tilde{s})} (-qs) dH(q) + \int_{q > \beta(\tilde{s})} (-c - q(1 - p(n(\tilde{s}), \bar{n}))s) dH(q) \\ &= -s\mathbb{E}q - c(1 - H(\beta(\tilde{s}))) + p(n(\tilde{s}), \bar{n})s \int_{q > \beta(\tilde{s})} q dH(q) \\ &= -s\mathbb{E}q - cn(\tilde{s}) + s \min\{n(\tilde{s}), \bar{n}\} \mathbb{E}(q | q > \beta(\tilde{s})), \end{aligned}$$

where the last equality comes from  $n(\tilde{s}) = 1 - H(\beta(\tilde{s}))$ .  $\square$

**Proof of Proposition 2.** We first characterize the ex-post efficient allocation and then show that it is achieved in equilibrium by a certain public belief  $\tilde{s}$ . Suppose the true severity level is  $s$ . Consider the planner's problem of deciding who should go to the hospital to maximize the ex-post total welfare. It is

$$\begin{aligned} \max_{a: [0, 1] \rightarrow \{0, 1\}} & \int_0^1 a(q)(-c - q(1 - p(n, \bar{n}))s) + (1 - a(q))(-qs) dH(q), \\ \text{s.t. } n &= \int_0^1 a(q) dH(q). \end{aligned}$$

It is straightforward to see that the optimal allocation must take a cut-off form  $a(q) = 1$  if and only if  $q \geq q^*$  for some  $q^* \in [0, 1]$ . Restricted to such a cut-off form, the planner's problem can be written as

$$\max_{q^* \in [0, 1]} -s\mathbb{E}q + \int_{q^*}^1 (-c + q \min\{1, \frac{\bar{n}}{1 - H(q^*)}\}s) dH(q).$$

If  $\bar{n} < 1 - H(q^*)$ , or equivalently  $q^* < H^{-1}(1 - \bar{n})$ , the objective function becomes

$$-s\mathbb{E}q - c(1 - H(q^*)) + \mathbb{E}(q | q \geq q^*),$$

which is clearly strictly increasing in  $q^*$ . Therefore,  $q^* < H^{-1}(1 - \bar{n})$  is never optimal. Thus, the planner's problem is equivalent to

$$\max_{q^* \in [H^{-1}(1 - \bar{n}), 1]} -s\mathbb{E}q + \int_{q^*}^1 (-c + qs) dH(q).$$

The solution to this problem is  $q^* = \max\{H^{-1}(1 - \bar{n}), \frac{c}{s}\}$ , or equivalently,

$$q^* = \begin{cases} 1, & \text{if } s \leq c, \\ \frac{c}{s}, & \text{if } c < s \leq s_0, \\ H^{-1}(1 - \bar{n}), & \text{if } s > s_0, \end{cases}$$

where, recall,  $s_0 = \frac{c}{H^{-1}(1 - \bar{n})}$  is the exhaustion level.

By Theorem 1, we can see that the equilibrium cut-off is

$$\beta(\tilde{s}) = \begin{cases} 1, & \text{if } \tilde{s} \leq c, \\ \frac{c}{\tilde{s}}, & \text{if } c < \tilde{s} \leq s_0. \end{cases}$$

Therefore, the equation  $\beta(\tilde{s}) = q^*$  has a solution

$$\tilde{s} = \begin{cases} \text{any value in } [0, c], & \text{if } s \leq c, \\ s, & \text{if } c < s \leq s_0, \end{cases}$$

For the first equality, notice that  $G$  and  $\tilde{G}^*$  coincide over  $[0, s_-] \cup [s_+, \bar{s}]$  and that  $W$  is linear over  $[s_-, s_+]$  by construction. Thus,

$$\begin{aligned} \int_0^{\bar{s}} W(\tilde{s}) dG(\tilde{s}) - \int_0^{\bar{s}} W(\tilde{s}) d\tilde{G}^*(\tilde{s}) &= \int_{s_-}^{s_+} W(\tilde{s}) dG(\tilde{s}) - \int_{s_-}^{s_+} W(\tilde{s}) d\tilde{G}^*(\tilde{s}) \\ &= (G(s_+) - G(s_-))W(s_0) - (\tilde{G}^*(s_+) - \tilde{G}^*(s_-))W(s_0) = 0. \end{aligned}$$

For the second equality, notice that  $\text{supp} \tilde{G}^* = [0, s_-] \cup \{s_0\} \cup [s_+, \bar{s}]$  if  $s_+ < \bar{s}$  and  $\text{supp} \tilde{G}^* = [0, s_-] \cup \{s_0\}$  if  $s_+ = \bar{s}$ . In either case, we have  $V = W$  over this set by the construction of  $W$ . This completes the proof.  $\square$

## Appendix. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jet.2022.105469>.

## References

- Acemoglu, Daron, Chernozhukov, Victor, Werning, Iván, Whinston, Michael D., 2020. A multi-risk SIR model with optimally targeted lockdown. NBER Working Paper (2020.02).
- Allcott, Hunt, Boxell, Levi, Conway, Jacob, Gentzkow, Matthew, Thaler, Michael, Yang, David, 2020. Polarization and public health: partisan differences in social distancing during the coronavirus pandemic. *J. Public Econ.* 191, 104254.
- Alonso, Ricardo, Câmara, Odilon, 2016. Persuading voters. *Am. Econ. Rev.* 106 (11), 3590–3605.
- Alvarez, Fernando E., Argente, David, Lippi, F515.892Tm(,)Tj/F11Tf.33060TD()Tj/F61Tf.21680TD()Tj/Fg1Tf7.9701007.970110

- Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. *Am. Econ. Rev.* 101 (6), 2590–2615.
- Kolotilin, Anton, 2018. Optimal information disclosure: a linear programming approach. *Theor. Econ.* 13 (2), 607–635.
- Kolotilin, Anton, Mylovanov, Tymofiy, Zapechelnyuk, Andriy, 2019. *Censorship as Optimal Persuasion*. University of New South Wales.
- Lipnowski, Elliot, Ravid, Doron, 2021. Pooled testing for quarantine decisions. *J. Econ. Theory* 198, 105372.
- Rayo, Luis, Segal, Ilya, 2010. Optimal information disclosure. *J. Polit. Econ.* 118 (5), 949–987.
- Simonov, Andrey, Sacher, Szymon K., Dubé, Jean-Pierre H., Biswas, Shirsho, 2020. The persuasive effect of fox news: non-compliance with social distancing during the covid-19 pandemic. NBER Working Paper, (w27237).
- Zapechelnyuk, Andriy, 2020. Optimal quality certification. *Am. Econ. Rev.: Insights* 2 (2), 161–176.