

An Imperfect Test for a Virus Can Be Worse Than No Test at All

Mark Whitmeyer*

January 17, 2021

Abstract

This note studies the effect of the availability of a test for a virus on the public health of a population. It is shown by example that the existence of a freely available and moderately informative test for a virus may lower society's welfare in comparison to the case where no test exists or access to the test is restricted. In this setting, any test provided to any subset of agents who would find it optimal not to isolate absent the test improves welfare.

Keywords: COVID-19; Coronavirus; Information Design; Group Testing

JEL Classifications: D82; D83; C72; I18

1 Introduction

In 2020, the virus COVID-19 swept through the globe. One particular difficulty presented by the virus is that infected individuals may be virtually asymptomatic and carry the virus without knowing it. Moreover, in the early stages of the pandemic, there were well-publicized shortages of

*Hausdorff Center for Mathematics & Institute for Microeconomics, University of Bonn

Email: mark.whitmeyer@gmail.com. I thank Scott Delhommer, Rosemary Hopcroft, Vasudha Jain, Joseph Whitmeyer, and Thomas Wiseman for their feedback. This paper has also benefited from the comments of two anonymous referees and the editor, Davide Dragone. This work was generously funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy-GZ 2047/1, Projekt-ID 390685813.

medical tests for the virus, which made it impossible to test everyone, even everyone with symptoms. To get around this, various solutions were proposed, including testing people in groups (Gollier and Gossner (2020)) and testing for inconclusive symptoms like a high temperature.¹

This purpose of this note is to present a simple example that illustrates that the existence (and availability) of a moderately informative test can actually lower social welfare in comparison to the scenario when no such test exists (or is available). This analysis, thus,

1. Supports the regulation of costless tests;
2. Provides a word of caution against moderately informative tests; and
3. Emphasizes the importance of the choice as to whom should be tested.

In the model we explore, there is a heterogeneous population of agents with different exposure likelihoods. Each agent has a simple decision: whether to stay home and isolate (or self-quarantine) or refrain from isolating and instead go out. An agent incurs a reward from not isolating, but possibly suffers a cost as well—she encounters others if out and neither wishes to become infected (if she is not infected) nor wishes to infect others (if she is infected). Thus, the prevalence of infected people who are not isolating is an endogenous equilibrium object determined by the individual isolation decisions of the agents.

Each agent's welfare is affected by the decisions of others, and crucially, a change in the infection rate of those out and about affects all of those who are not isolating. Accordingly, the mechanism that generates the possible welfare loss due to a test is the worsening of the participant pool. A false negative from a less than fully informative test can encourage risky people to refrain from isolating, increasing the chances of disease transmission. Although the information provided is itself valuable, the gain in welfare as a result of this information is outweighed by the increased disease prevalence among those an agent encounters.

The frequency of false negatives for the COVID-19 virus is well-documented. Indeed, the most common test for the virus, the reverse transcriptase polymerase chain reaction (RT-PCR) test has false negative rates at initial presentation of symptoms that range from 30 to 40 percent (Yang et al. (2020); Fang et al. (2020); and Ai et al. (2020)). Moreover, these rates can vary considerably

¹See, e.g., this April 4th CNBC article.

depending on the time since exposure (Kucirka et al. (2020)). CT scans may be more effective (Fang et al. (2020); Ai et al. (2020); and Caruso et al. (2020)), but even those may have high false negative rates in the first few days following the onset of symptoms (Kanne et al. (2020)).

In this note, the agents are completely rational, yet can be made worse off by the availability of an imperfect test. This is because of the strategic nature of the societal interaction. If this were merely a decision problem for each agent, any test would increase welfare.² Moreover, it is understood that an imperfect test can encourage sub-optimal behavior if people misunderstand or if people do not realize that the test is flawed. Here we discover that an imperfect test can be detrimental *even when its quality is common knowledge and when agents have no behavioral biases*.

Although we find that there are some tests and testing protocols that can lower welfare, Proposition 2.2 reveals that *any* test, provided it is given to people who have a (relatively) low likelihood of exposure—those who would refrain from isolating absent a test—is welfare improving, since it both provides those agents with more information but also ensures that the pool of agents participating in society improves. That is not to say that that group of people is the optimal group to test, merely that such a protocol, according to this framework, cannot reduce welfare.

The last statement leads us to the following caveat: in this paper, we do not attempt to characterize optimal testing protocols, nor do we provide a thorough cost-benefit analysis of the participation/social-distancing trade-off. These are both worthwhile concerns, yet the goal of this paper is more modest. Instead, we merely wish to expose a counter-intuitive aspect to testing, one that has seemingly been heretofore unmentioned and overlooked.

2 The Model

Let us consider the following formal model. There is a population that consists of a continuum of agents with measure 1. The Bernoulli random variable, Θ , corresponds to the infection status of an agent, where $\Theta = 1$ denotes that an agent is infected and $\Theta = 0$ denotes that an agent is not infected. The only source of heterogeneity is an agent's prior exposure to the disease, which we term her type. That is, an agent's type is her likelihood of infection, $\mu \in [0, 1]$, where

²See e.g. Ramsey (1990); Blackwell (1951), who (among many others) establish that the well-known result that in decision problems, the value of information is always positive.

$\mu := \mathbb{P}(\Theta = 1)$. We impose that the population distribution of types has an atomless³ cumulative distribution function F with support on $[0, 1]$.

Each agent has a simple choice: either isolate (action I) or participate (action P). If an agent isolates, she obtains a payoff normalized to 0. If an agent participates, then with probability τ she interacts with someone who is infected, where τ is the average infection likelihood of those participating. Hence, τ is the chance that a agent who is not isolating encounters an infected person. It is important to keep in mind that τ is an endogenous equilibrium object, determined by the isolation decisions of the agents in the populace.

If an agent of type μ participates then her payoff is

$$u(P; \mu, \tau) = A - B\mu(1 - \tau) - C(1 - \mu)\tau$$

where $A \geq 0$ is her reward from participating, $C \geq 0$ is her loss from becoming infected, and $B \geq 0$ is her loss from infecting someone else. Note that, in contrast to τ , A , B , and C are exogenous parameters.

The solution concept that we use is Nash Equilibrium: given the actions of the other agents, no agent has a (unilateral) profitable deviation. Due to the linearity of the payoff from P in an agent's type, any equilibrium must be of a particular cut-off form: all agents whose types are above (or below) a certain threshold will isolate, and all those whose types are below (or above) that threshold will not. Because there is a coordination-like aspect of the game, there may exist two kinds of equilibria; both those in which high types isolate and low types participate, but also the inverse. Throughout we restrict attention to the first class of equilibria, since it seems to correspond closest to people's behavior during the pandemic—by and large, sick people are encouraged to stay home, not go out. Moreover, the condition below ensures that such an equilibrium exists.

We begin by looking at the case in which there is no test, so that we can subsequently compare welfare to the case in which a test is available. We impose the following condition on the parameters:

Condition 2.1. There exists some type $\hat{\mu}$ such that

³This assumption is not significant, but eases notation.

$$(i) \quad \begin{aligned} A + ((B + C) \hat{\tau} - B) \mu - C \hat{\tau} &\geq 0, \quad \text{for all } \mu \leq \hat{\mu}, \quad \text{and} \\ A + ((B + C) \hat{\tau} - B) \mu - C \hat{\tau} &\leq 0, \quad \text{for all } \mu \geq \hat{\mu}, \end{aligned}$$

where

$$\hat{\tau} = \tau(\hat{\mu}) := \frac{\int_0^{\hat{\mu}} x dF(x)}{F(\hat{\mu})}$$

and

$$(ii) \quad \hat{\mu} \leq \frac{C}{B + C}$$

$\hat{\tau}$ is a conditional expectation: it is the average infected likelihood of those participating (those whose type, μ , is less than the cutoff type, $\hat{\mu}$). Inequality (ii) ensures that in this equilibrium, the payoff of each type $\mu \in [0, \hat{\mu}]$ is decreasing in τ , which is realistic: the welfare of the participants gets worse as it becomes more likely that they encounter infected individuals.

A necessary condition for Condition 2.1 is that

$$(B + C) \hat{\tau} \leq B$$

That is, given $\hat{\tau}$, an agent's participation payoff is decreasing in her own type. If Condition 2.1 holds, then trivially there exists an equilibrium in which all types $\mu < \hat{\mu}$ participate and all types $\mu > \hat{\mu}$ isolate (indeed the first part of the condition is necessary and sufficient for such an equilibrium to exist). Denote the aggregate payoff from this equilibrium by W , *viz.*,

$$W := \int_0^{\hat{\mu}} \{A + ((B + C) \hat{\tau} - B) x - C \hat{\tau}\} dF(x)$$

Clearly, this is an exceedingly simple model: the scenario is static and there are no benefits to testing other than to guide agents' isolation decisions. For instance, there is no contact tracing in this model, nor are there benefits to society from obtaining statistics about the spread, morbidity, or mortality of the disease. These are all important considerations for determining optimal testing policies.

Furthermore, outside of their infection likelihoods, agents in this model are homogeneous. We do not distinguish between essential and inessential workers, say, which is another vital component of a thorough cost-benefit analysis of testing. The model's homogeneity in this dimension also does not allow us to tackle the subtle issue that, in reality, not all agents with the same infection likelihood are the same. Perhaps some high-likelihood agents are travelers who

have recently returned from a virus hot-spot, whereas others are doctors, who are likely to be sick by virtue of their occupation. It is easy to see how distinguishing between these sorts of agents could be very important when choosing whom to test.⁴

Yet another simplification is that the agents who participate interact with each other randomly. They cannot choose the types with whom they interact, and neither the encounter rate of an agent nor whom an agent meets is affected by her infection probability. Again, this is perhaps unrealistic—a doctor is more likely to come into contact with infected individuals than a traveller who has returned home from a hot-spot, yet our model does not allow for this distinction.

Naturally, a model intended to guide policy directly should include many, if not all, of these details that are conspicuously absent from this paper. However, as is noted above, the intent of this paper is not prescriptive. Our goal is to discover and understand the counterintuitive notion that testing may lower welfare. The model is kept deliberately simple in order to clearly illustrate the intuition of the result. We should expect similar incentives to be present in a more complicated model, but they might be hidden or obfuscated by other factors.

2.1 Testing Participants Cannot Hurt

Now let us introduce testing to the scenario and derive our first result. Formally, a test, π , is a stochastic map:

$$\pi : \{0, 1\} \rightarrow \Delta(S)$$

where S is some (compact) set of signal realizations. In the example that we explore later on, we assume that S consists merely of two signal realizations—a positive result and a negative result—but for now, we need not make such a restriction.

Then,

Proposition 2.2. *Let Condition 2.1 hold. Then, any test, π , given to any subset of types in the interval $[0, \hat{\mu}]$ begets an equilibrium that yields society a payoff that is (at least weakly) greater than W .*

Proof. It suffices to show that the payoff of each type $\mu \in [0, \hat{\mu}]$ (weakly) increases in expectation. Moreover, recall that the second part of Condition 2.1 implies that the payoff of each type $\mu \in [0, \hat{\mu}]$

⁴I thank an anonymous referee for highlighting this aspect of optimal testing.

is decreasing in τ . Hence, if the pool of agents in society improves (τ decreases) the payoffs of those agents increase.

If the testing protocol begets an equilibrium in which the new average likelihood of participants (τ^*) is equal to $\hat{\tau}$, then in expectation, the payoff of each type who is tested must weakly improve (this follows from Ramsey (1990) and Blackwell (1951)). Naturally, this payoff is further improved if $\tau^* \leq \hat{\tau}$ and so we need only establish that there is an equilibrium in which all types $\mu \in [0, \mu^*]$, with $\mu^* \geq \hat{\mu}$ participate, and $\tau^* \leq \hat{\tau}$.

Let G be the new distribution over types (beliefs about infection likelihood) as a result of the test. Note that by definition,⁵ G is a mean-preserving spread of F .⁶ There are two cases to consider: Case I, where $G(\hat{\mu}) = F(\hat{\mu})$; and Case II, where $G(\hat{\mu}) < F(\hat{\mu})$. The first case trivially yields the desired result (the pool of participants and the participation decisions of the agents are unchanged).

Let us consider Case II and search for an equilibrium in which all types $\mu \in [0, \mu^*]$ participate, where $\mu^* \geq \hat{\mu}$. It is straightforward to see⁷ that

$$\hat{\tau} > \frac{\int_0^{\hat{\mu}} x dG(x)}{G(\hat{\mu})}$$

Define

$$\tau(\mu) := \frac{\int_0^{\mu} x dG(x)}{G(\mu)}$$

which is obviously increasing in μ . Accordingly, since we have assumed that Condition 2.1 holds, there must exist some $\mu' > \hat{\mu}$ such that $\tau(\mu') = \hat{\tau}$. Clearly,

$$\varphi(\mu) := A + ((B + C) \tau(\mu) - B) \mu - C \tau(\mu)$$

is continuous in μ . Evidently, $\varphi(\hat{\mu})$ is positive and $\varphi(\mu')$ is negative, so by the intermediate value theorem there exists some $\mu^* \in (\hat{\mu}, \mu')$, with $\varphi(\mu^*) = 0$ and $\tau(\mu^*) < \hat{\tau}$. Consequently, there exists an equilibrium in which all types $\mu \in [0, \mu^*]$ participate, where $\mu^* > \hat{\mu}$ and $\tau^* := \tau(\mu^*) < \hat{\tau}$.

⁵Indeed, this is precisely the definition of a mean-preserving spread introduced in Rothschild and Stiglitz (1970).

⁶Moreover, since F is atomless, G will not have atoms except possibly at 0 and 1. The proof proceeds with the implicit assumption that G does not have an atom on 1, but that is just to ease notation and simplify the arguments.

If G does have an atom there, the result continues to hold.

⁷Refer to Lemma A.1 in the Appendix.




Figure 1: Example 2.1.1 Splitting of Beliefs (Types)




Figure 2: Example 2.1.1 Old (F , Dashed Lines) and New (G , Solid Lines) Distributions of Types

Note that because belief is a martingale, for each agent that is tested, there must be some test result that ensures that her posterior type (remember, this is her posterior belief about her infection likelihood following a test result) is weakly less than $\hat{\mu}$ and hence μ^* . ■

2.1.1 An Example

To add intuition to the proof, let us briefly explore an example. Namely, let F be the uniform distribution on $[0, 1]$ and the parameters be such that $\hat{\mu} = \frac{1}{2}$. Consider a binary test, π , that is given to types $\mu \in [\frac{1}{4}, \frac{1}{2}]$, with

$$\pi(+|1) = \pi(-|0) = \frac{2}{3}$$

where the set of test outcomes is $S = \{+, -\}$. Accordingly, each type in $[\frac{1}{4}, \frac{1}{2}]$ will be “split” into two new types: with probability $\frac{\mu+1}{3}$, outcome + will realize and the posterior belief (new type) of agent μ will be $\frac{2\mu}{1+\mu} \in [\frac{2}{5}, \frac{2}{3}]$; and with probability $\frac{2-\mu}{3}$, outcome – will realize and the posterior belief (new type) of agent μ will be $\frac{\mu}{2-\mu} \in [\frac{1}{7}, \frac{1}{3}]$. This splitting is depicted in Figure 1.

Furthermore, Figure 2 depicts the new cdf of types, G . Also depicted are $\int_0^\mu G(x)dx$ and $\int_0^\mu F(x)dx$, which illustrates that G is a mean-preserving spread of F .

3 How Tests Can Lower Welfare

Let us turn our attention to a simplified version of the previous section’s model. We reduce the model as follows: now, the population is inhabited by just two types of agents, ω_H (high likelihood) and ω_L (low likelihood). Proportion $q = 3/4$ of the population are ω_L .

High likelihood agents, ω_H , are infected with (prior) probability μ_H , where

$$\mu_H := \mathbb{P}(\Theta = 1 | \omega_H) = 5/8$$

Likewise, μ_L denotes the prior probability that type ω_L is infected:

$$\mu_L := \mathbb{P}(\Theta = 1 | \omega_L) = 1/8$$

Recall that an agent may choose either to isolate (I) or not (P). We impose the following values for the parameters: A , her payoff from participation, equals $5/4$; B , her penalty from infecting someone uninfected, is 2 ; and C , her penalty from becoming infected, is 4 . Thus, if she has a belief μ that she is infected, her payoff from not isolating (P) is

$$u(P; \mu, \tau) = \frac{5}{4} - 2\mu(1 - \tau) - 4(1 - \mu)\tau$$

which simplifies to

$$u(P; \mu, \tau) = \frac{5}{4} - 4\tau - (2 - 6\tau)\mu \tag{1}$$

With no testing, the unique equilibrium is that in which agents of type ω_H isolate and agents of type ω_L participate. Since only agents of type ω_L are participating, the likelihood that an agent who is not isolating is infected is merely the likelihood that an agent of type ω_L is infected; *viz.*, $\tau = \mu_L = 1/8$. On path, an agent of type ω_L obtains a payoff of $19/32$, which is bigger than 0 , her

payoff from isolating. Should an agent of type ω_H deviate and participate, she would obtain a payoff of $-1/32$, less than her isolation payoff of 0.

In fact, an agent of type ω_H would always prefer to isolate (since μ_H is so high), unless τ were precisely 0, which could never happen at equilibrium. Consequently, this equilibrium is unique.

Without testing, the aggregate welfare for society, V , is

$$V = q \left(\frac{5}{4} - 4\tau - (2 - 6\tau)\mu_L \right) = \frac{3}{4} \left(\frac{19}{32} \right) = \frac{57}{128} \approx .45$$

3.1 After the Introduction of a Test

We introduce a binary test to the scenario. The set of signal realizations is $S := \{+, -\}$, corresponding to a positive test and a negative test, respectively. Consequently, π can be written in terms of the variable p , where

$$p := \pi(-|1), \quad \text{and} \quad \pi(-|0) = 1$$

The situation can be conveniently described by the following joint distribution of S and Θ for an agent with prior μ_i , $i = H, L$:

$\Theta \setminus S$	-	+	$\mathbb{P}(\Theta)$
1	$\mu_i p$	$\mu_i (1 - p)$	μ_i
0	$(1 - \mu_i)$	0	$1 - \mu_i$
$\mathbb{P}(S)$	$\mu_i p + (1 - \mu_i)$	$\mu_i (1 - p)$	1

Note that in this model there are no false positives. This is done for expositional convenience and to allow us to focus on the effect of the high false negative rates noted in the introduction.

Next, we look at the Nash Equilibria of the participation game with testing. As in Section 2, we focus on equilibria in which agents who think their infection likelihood is low participate and agents who think their infection likelihood is high do not. In this example, absent a test, this is the only equilibrium that exists, and we suppose that this is still the equilibrium selected following a test. To put another way, we assume that introduction of the test does not qualitatively alter the equilibrium selected to an (ostensibly less realistic) equilibrium in which low types isolate and high types participate.

We consider three cases: in the first, only agents of type ω_H have access to the test.

3.1.1 Case 1 (Testing only High Likelihood)

Here we suppose that only agents of type ω_H may take the test. There are three regions of false negative probabilities, p , each of which beget a different equilibrium. If p is sufficiently low ($p \leq .73$), then agents of type ω_L participate and agents of type ω_H participate if and only if they get a negative test result. On the other hand, if p is in an intermediate range ($.9 \geq p \geq .73$), then such an equilibrium no longer exists. All agents of type ω_L continue to participate, but now only a fraction of the types ω_H who have received negative results participate. Finally, if p is too high, only agents of type ω_L participate.

We start by calculating the values of p such that at equilibrium agents of type ω_L participate and agents of type ω_H participate if and only if they get a negative test result.

The crucial variable is τ , the likelihood of encountering an infected agent while participating. Using the law of total probability, it is

$$\tau_p^H = \frac{q\mu_L + (1-q)\mu_H p}{q + (1-q)(\mu_H p + 1 - \mu_H)} = \frac{5p + 3}{5p + 27}$$

Using Bayes' law, the probability that type ω_H is infected after a negative test is

$$\mu_H^- = \frac{\mu_H p}{\mu_H p + 1 - \mu_H} = \frac{5p}{5p + 3} \quad (2)$$

Then, using Expression 1, agents of type ω_H will participate after – if and only if

$$\frac{5}{4} - 4\tau_p^H - (2 - 6\tau_p^H)\mu_H^- \geq 0 \quad (3)$$

which simplifies to $p \leq \frac{9}{5} - \frac{12}{5\sqrt{5}} \approx .73$.

It is easy to verify that agents of type ω_L prefer to participate for this range of p . Moreover, clearly, agents of type ω_H have no profitable deviation to action P after +, since they are sure that they are infected.

Should Inequality 3 fail to hold, there is no longer an equilibrium in which all agents of type ω_H who have obtained a negative test participate. However, for moderate p , there is also no equilibrium in which *none* of those agents participate, since the resulting τ would be low enough

to entice participation. Instead, fraction σ of the agents of type ω_H who have seen a negative test participate.

Observe that $\mu_{\bar{H}}$ is the same as above, but the likelihood of encountering an infected agent is shaped by σ . This is

$$\tau(\sigma) = \frac{q\mu_L + (1-q)\sigma\mu_H p}{q + (1-q)\sigma(\mu_H p + 1 - \mu_H)} = \frac{5p\sigma + 3}{5p\sigma + 3\sigma + 24} \quad (4)$$

Because some of the agents of type ω_H isolate after a negative result and others do not, we need them to be indifferent as to whether they participate after -. Thus, using Expression 1,

$$\frac{5}{4} - 4\tau(\sigma^*) - (2 - 6\tau(\sigma^*))\mu_{\bar{H}} = 0$$

Substituting in for $\tau(\sigma^*)$, we obtain

$$\sigma^* = \frac{24(10p - 9)}{5(25p^2 - 42p + 9)} \quad (5)$$

which is feasible (lies in the interval $[0, 1]$) provided $.73 \leq p \leq .9$. Substituting σ^* into Equation 4 we obtain the equilibrium infection likelihood

$$\tau(\sigma^*) = \frac{15(1-p)}{48-40p}$$

If $p \geq .9$, then the only equilibrium yields the same payoff as the scenario without testing. The test is too uninformative to persuade any agents of type ω_H and so as in the case without testing, only agents of type ω_L participate, yielding a payoff of .45.

We finish the analysis of Case 1 by inspecting aggregate welfare as a function of the probability p , V^H (H for ‘‘High likelihood’’), the details of whose derivation we leave to Appendix A.1:

$$V^H = \begin{cases} \frac{125p^2 - 1530p + 1917}{128(5p + 27)}, & 0 \leq p \leq \frac{9}{5} - \frac{12}{5\sqrt{5}} \\ -\frac{105p - 9}{640p - 768}, & \frac{9}{5} - \frac{12}{5\sqrt{5}} \leq p \leq \frac{9}{10} \\ \frac{57}{128}, & \frac{9}{10} \leq p \leq 1 \end{cases}$$

Comparing this to V , aggregate welfare when there are no tests, we see that $V^H < V = 57/128$ for $p \in (.21, .9)$, and $V^H = V$ for $p \in [.9, 1]$. Moreover, for all $p \in [.73, .9]$, there is a Pareto decrease in welfare: agents of type ω_H are no better off and agents of type ω_L are strictly worse off.




Figure 3: Society's Welfare in Cases 1 (V^H , Solid Line), 2 (V^L , Solid Line), 3 (V^T , Dashed Line), & With No Testing (V , Dashed Line)

3.1.2 Case 2 (Testing only Low Likelihood)

What if those tests from Case 1 were instead given to agents of type ω_L ? To ensure a fair comparison, suppose that there are only enough tests to serve measure 1/4 of agents (so only 1/3 of agents of type ω_L get tested). At equilibrium, untested agents of type ω_L participate, tested agents of type ω_L participate if and only if they have a negative test result, and agents of type ω_H do not participate.

Now,

$$\tau_p^L = \frac{\frac{2q}{3} + \frac{q}{3}p}{\frac{2q}{3} + \frac{q}{3}(\mu_L p + 1 - \mu_L)} \mu_L = \frac{p + 2}{p + 23}$$

It is obvious that agents of type ω_L participate after a negative test since they would even with a perfectly uninformative test ($p = 1$). Likewise, they do not participate after a positive test since they are sure that they are infected. Agents of type ω_H do not participate since τ_p^L is strictly greater than 0.

Consequently, aggregate welfare as a function of p is V^L (L for "Low likelihood"):

$$\begin{aligned} V^L &= \frac{2q}{3} \left(\frac{5}{4} - 4\tau_p^L - (2 - 6\tau_p^L) \mu_L \right) + \frac{q}{3} (\mu_L p + 1 - \mu_L) \left(\frac{5}{4} - 4\tau_p^L - (2 - 6\tau_p^L) \mu_L \right) \\ &= \frac{5p^2 - 274p + 1637}{128(p + 23)} \end{aligned}$$

where μ_L^- , the probability that type ω_L , is infected after a negative test is obtained using Bayes' law. Evidently, $V^L \geq V$ for all p .

This case yields precisely the result that we uncovered in Section 2.1. Namely, any test applied to any subset of the population whose members are participating in the absence of a test must improve society's welfare.

3.1.3 Case 3 (Testing Everyone)

What if both types of agents have access to the test? Perhaps unsurprisingly, this case is qualitatively identical to Case 1. For a sufficiently low p ($p \leq .76$), agents of type ω_L participate and agents of type ω_H participate if and only if they get a negative test result. If p is in an intermediate range ($.91 \geq p \geq .76$), then all agents of type ω_L continue to participate, and only a fraction of the types ω_H who have received negative results participate. If p is too high, only agents of type ω_L participate.

Leaving its derivation to Appendix A.1—since it is identical to the work for Case 1 *mutatis mutandis*—aggregate welfare as a function of p is V^T :

$$V^T = \begin{cases} \frac{5p^2 - 42p + 45}{16(p+3)}, & 0 \leq p \leq \frac{39 - 6\sqrt{11}}{25} \\ \frac{3p}{24 - 20p}, & \frac{39 - 6\sqrt{11}}{25} \leq p \leq \frac{69 - 2\sqrt{534}}{25} \\ \frac{15p^2 - 294p + 735}{128p + 896}, & \frac{69 - 2\sqrt{534}}{25} \leq p \leq 1 \end{cases}$$

Evidently, $V^T \leq V$ for $p \in [.51, .9]$. As in Case 1, there is an interval of p values, $[.76, .9]$, that begets a Pareto decrease in the equilibrium welfare for society. Curiously, Case 3 illustrates that an extremely uninformative test can be strictly welfare improving. It provides some information to agents of type ω_L but is insufficiently informative to persuade any agents of type ω_H to participate and so does not worsen the participant pool.

Unsurprisingly, welfare for society is higher when everyone can get tested (Case 3), than when only agents of type ω_H can get tested (Case 1). In both cases, an increase in τ drives the decrease in welfare, but in the Case 3, the effect of this increase in τ is not as pernicious, due to the information acquired by the low likelihood types.

A graph of society's welfare as a function of p when there is no testing and in each of the three cases is depicted in Figure 3.

Finally, recall that in the example, the test does not generate false positives. How does the false positive rate affect welfare? As it turns out, the effect of the false positive rate is also ambiguous. It is possible that a high false positive rate could be better for society due to its effect on the isolation decisions of agents of type ω_H . Naturally, there are also regions of the parameters in which welfare is decreasing in the false positive rate, which is driven by the intrinsic value of information.

4 Brief Discussion

Determining the optimal way to guide a society through a pandemic and calculating optimal testing protocols is a challenging task. Such an undertaking is made even more difficult by the myriad of constraints that must be satisfied—logistical, resource, cognitive, and political, to name a few. Here we bypass such concerns and merely explore a counter-intuitive aspect of testing engendered by the interconnected nature of society.

Namely, we find that the introduction of a moderately informative test can make society worse off and can even result in a Pareto decrease in welfare. This is because the test worsens the pool of participants—it makes it more likely that agents encounter infected agents who, thanks to the test, are sufficiently confident that they are not infected due to a false negative produced by the test.

On the other hand, as Proposition 2.2 illustrates, every test is welfare improving, so long as it is given (or available) only to agents who participate in equilibrium in the absence of the test. Such a protocol guarantees that the pool of participants in society improves and provides valuable information to agents, to boot.

As we have stressed throughout, this paper does not prescribe optimal testing protocols. There are a number of contemporaneous papers that do, under a variety of different assumptions and set-ups. Kasy and Teytelboym (2020) allow for perfectly informative (but costly) tests, and investigate the dynamic problem of how to test individuals in order to both inform quarantine protocols and also learn about the virus’s prevalence rate. Ely et al. (2020) assume a finite quantity of heterogeneous tests and study how to allocate such tests to heterogeneous agents. Lipnowski and Ravid (2020) explore a similar problem, but look at how to *pool* tests optimally.

Deb et al. (2020), in turn, allow for targeted testing *and* transfers to agents.

Crucially, in comparison to this paper, none of those papers contain a strategic interaction between the agents in the population. In each such paper, agents and planners are merely tasked with decision problems and so the value of information is positive. The important decision then, is the choice of whom to test or how to test, given limited resources or costly tests. As we encounter here, introducing strategic concerns adds a novel wrinkle, one that may be important to keep in mind when determining policy.

Our discovery that information may be detrimental is related to the phenomenon discovered by Kremer (1996), who shows that increased prevalence of the AIDS virus may worsen the pool of available partners due to the fatalism of high-activity people. In this note, the welfare loss is also due to a worsened pool of participants, which effect is driven by not by fatalism but by the increased (rational) confidence of high-likelihood individuals consequent to negative test results.

A number of other recent papers have pointed out other counterintuitive incentives in models of epidemics and disease transmission. Those include Talamàs and Vohra (2020), who show that the introduction of a moderately effective vaccine can result in Pareto losses in welfare for society; and Heinsalu (2020), who illustrates that increasing the infection risk early in a pandemic may be optimal. However, the mechanism behind the result in Talamàs and Vohra (2020) is completely different to that in this paper. There, the authors assume a network structure and allow agents to choose with whom they match. In contrast to this paper, fixing the partnership structure of their model or imposing that matching is random ensures that even a partially-effective vaccine is a Pareto improvement to welfare. In their paper, it is only when agents may choose with whom to match that imperfect vaccines can have a destabilizing effect on society's partnership network, begetting a denser (and hence worse) network, and thereby lowering welfare.

With that in mind, it would be interesting to investigate whether the results of this paper extend to a finite population with a network structure. Likewise, it is unclear how allowing agents to make more sophisticated decisions—letting them choose not just whether to isolate but also with whom to interact—would affect the findings. Given the results of Talamàs and Vohra (2020), it is reasonable to suspect that a similar destabilizing effect could manifest as the result of an imperfect test.

References

- Tau Ai, Zhenlu Yang, Hongyan Hou, Chenao Zhan, Chong Chen, Wenzhi Lv, Qian Tao, Ziyong Sun, and Liming Xia. Correlation of chest ct and rt-pcr testing in coronavirus disease 2019 (covid-19) in china: A report of 1014 cases. *Radiology*, 2020.
- David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley, Calif., 1951. University of California Press.
- Damiano Caruso, Marta Zerunian, Michela Polici, Francesco Pucciarelli, Tiziano Polidori, Carlotta Rucci, Gisella Guido, Benedetta Bracci, Chiara de Dominicis, and Andrea Laghi. Chest ct features of covid-19 in rome, italy. *Radiology*, 2020.
- Rahul Deb, Mallesh Pai, Akhil Vohra, and Rakesh Vohra. Testing alone is insufficient. *Mimeo*, May 2020.
- Jeffrey Ely, Andrea Galeotti, and Jakub Steiner. Optimal test allocation. *Mimeo*, May 2020.
- Yicheng Fang, Huangqi Zhang, Jicheng Xie, Minjie Lin, Lingjun Ying, Peipei Pang, and Wenbin Ji. Sensitivity of chest ct for covid-19: Comparison to rt-pcr. *Radiology*, 2020.
- Christian Gollier and Olivier Gossner. Group testing against covid-19. *Mimeo*, March 2020.
- Sander Heinsalu. Infection arbitrage. *Mimeo*, April 2020.
- Jeffrey Kanne, Brent Little, Jonathan Chung, Brett Elicker, and Loren Ketai. Essentials for radiologists on covid-19: An update—radiology scientific expert panel. *Radiology*, 2020.
- Maximilian Kasy and Alexander Teytelboym. Adaptive targeted infectious disease testing. *Oxford Review of Economic Policy*, 36:S77–S93, 06 2020.
- Michael Kremer. Integrating Behavioral Choice into Epidemiological Models of AIDS. *The Quarterly Journal of Economics*, 111(2):549–573, 05 1996.

Lauren M Kucirka, Stephen A Lauer, Oliver Laeyendecker, Denali Boon, and Justin Lessler. Variation in false-negative rate of reverse transcriptase polymerase chain reaction–based sars-cov-2 tests by time since exposure. *Annals of Internal Medicine*, 2020.

Elliot Lipnowski and Doron Ravid. Group testing for quarantine decisions. *Mimeo*, July 2020.

F. P. Ramsey. Weight or the value of knowledge. *The British Journal for the Philosophy of Science*, 41(1):1–4, 1990.

Michael Rothschild and Joseph E Stiglitz. Increasing risk: I. a definition. *Journal of Economic Theory*, 2(3):225 – 243, 1970.

Eduard Talamàs and Rakesh Vohra. Free and perfectly safe but only partially effective vaccines can harm everyone. *Games and Economic Behavior*, Forthcoming, 2020.

Yang Yang, Minghui Yang, Chenguang Shen, Fuxiang Wang, Jing Yuan, Jinxiu Li, Mingxia Zhang, Zhaoqin Wang, Li Xing, Jinli Wei, Ling Peng, Gary Wong, Haixia Zheng, Mingfeng Liao, Kai Feng, Jianming Li, Qianting Yang, Juanjuan Zhao, Zheng Zhang, Lei Liu, and Yingxia Liu. Evaluating the accuracy of different respiratory specimens in the laboratory diagnosis and monitoring the viral shedding of 2019-ncov infections. *medRxiv*, 2020. URL <https://www.medrxiv.org/content/early/2020/02/17/2020.02.11.20021493>.

A Omitted Results and Derivations

Lemma A.1.

$$\hat{\tau} = \frac{\int_0^{\hat{\mu}} x dF(x)}{F(\hat{\mu})} > \frac{\int_0^{\hat{\mu}} x dG(x)}{G(\hat{\mu})}$$

Proof. Directly, using integration by parts

$$\frac{\int_0^{\hat{\mu}} x dF(x)}{F(\hat{\mu})} = \hat{\mu} - \frac{\int_0^{\hat{\mu}} F(x) dx}{F(\hat{\mu})} > \hat{\mu} - \frac{\int_0^{\hat{\mu}} F(x) dx}{G(\hat{\mu})} > \hat{\mu} - \frac{\int_0^{\hat{\mu}} G(x) dx}{G(\hat{\mu})} = \frac{\int_0^{\hat{\mu}} x dG(x)}{G(\hat{\mu})}$$

which used the facts that $F(\hat{\mu}) > G(\hat{\mu})$ and G is a mean preserving spread of F , of which one definition is that

$$\int_0^{\mu} G(x) dx \geq \int_0^{\mu} F(x) dx$$

for all $\mu \in [0, 1]$ (an equivalent statement is that F second order stochastically dominates G). ■

A.1 Sections 3.1.1 and 3.1.3 Aggregate Welfare Derivations

Case 2 (Section 3.1.2) is omitted, since the lone derivation may be found in the text.

A.1.1 Case 1 (Testing only High Likelihood)

For $0 \leq p \leq \frac{9}{5} - \frac{12}{5\sqrt{5}} \approx .73$, aggregate welfare, V^H , is

$$\begin{aligned} V^H &= q \left(\frac{5}{4} - 4\tau_p^H - (2 - 6\tau_p^H) \mu_L \right) + (1 - q) (\mu_H p + 1 - \mu_H) \left(\frac{5}{4} - 4\tau_p^H - (2 - 6\tau_p^H) \mu_H^- \right) \\ &= \frac{125p^2 - 1530p + 1917}{128(5p + 27)} \end{aligned}$$

For $\frac{9}{10} \geq p \geq \frac{9}{5} - \frac{12}{5\sqrt{5}}$, since each high type, ω_H is receiving a payoff of 0, we have

$$V^H = q \left(\frac{5}{4} - 4\tau(\sigma^*) - (2 - 6\tau(\sigma^*)) \mu_L \right) = -\frac{105p - 9}{640p - 768}$$

If $1 \geq p \geq \frac{9}{10}$, aggregate welfare is the same as in the case without a test; namely, $V^H = \frac{57}{128} \approx .45$.

A.1.2 Case 3 (Testing Everyone)

First, let $0 \leq p \leq \frac{39-6\sqrt{11}}{25} \approx .76$. Let us calculate τ_p^T . Directly,

$$\tau_p^T = p \frac{q\mu_L + (1 - q)\mu_H}{q(\mu_L p + 1 - \mu_L) + (1 - q)(\mu_H p + 1 - \mu_H)} = \frac{p}{p + 3}$$

Then, aggregate welfare, V^T , is

$$\begin{aligned} V^T &= q(\mu_L p + 1 - \mu_L) \left(\frac{5}{4} - 4\tau_p^T - (2 - 6\tau_p^T) \mu_L^- \right) \\ &\quad + (1 - q)(\mu_H p + 1 - \mu_H) \left(\frac{5}{4} - 4\tau_p^T - (2 - 6\tau_p^T) \mu_H^- \right) \end{aligned}$$

which simplifies to

$$V^T = \frac{5p^2 - 42p + 45}{16(p + 3)}$$

For $\frac{39-6\sqrt{11}}{25} \leq p \leq \frac{69-2\sqrt{534}}{25} \approx .91$, we have

$$V^T = q(\mu_L p + 1 - \mu_L) \left(\frac{5}{4} - 4\tau(\sigma^*) - (2 - 6\tau(\sigma^*)) \mu_L^- \right) = \frac{3p}{24 - 20p}$$

If $1 \geq p \geq \frac{69-2\sqrt{534}}{25}$, aggregate welfare is

$$V^T = q(\mu_L p + 1 - \mu_L) \left(\frac{5}{4} - 4\tau^- - (2 - 6\tau^-) \mu_L^- \right) = \frac{15p^2 - 294p + 735}{128p + 896}$$

where we used the fact that

$$\tau^- = \mu_L^- = \frac{p}{7 + p}$$