Optimal Test Allocation*

Jeffrey Ely
Northwestern U

Andrea Galeotti
LBS

Ole Jann
Cerge-Ei

Jakub Steiner
Cerge-Ei and Zurich U

December 1, 2020

Abstract

A health authority chooses a binary action for each of several individuals that differ in their pretest probabilities of being infectious and in the additive losses associated with two types of decision errors. The authority is endowed with a portfolio of tests that differ in their sensitivities and specificities. We derive a simple necessary condition for optimality of test allocation. In special cases, precision parameters of the allocated test are monotone in the individuals' types. We characterize the marginal benefit of a test, provide an algorithmic solution for the test-allocation problem and consider the benefits of confirmatory testing.

1 Introduction

During the COVID-19 pandemic, health authorities around the world make millions of decisions based on whether they believe a patient is infectious or not. They do so with the help of a diverse portfolio of tests that vary in their accuracy and availability. Consequently, fighting the pandemic requires deciding who gets which test.

When faced with this decision, medical professionals and guidance documents usually rely on the concepts of "positive predictive value" (PPV) and "negative predictive value" (NPV). These express the posterior probabilities that a person is infectious (or not) after a

*We benefited from comments of Olivier Gossner, Pavel Kocourek and Flavio Toxvaerd. This work was supported by the ERC grant 770652 (Jann and Steiner), Experentia Foundation (Steiner), and ERC grant 724356 (Andrea Galeotti).

positive (or negative) test result. Targeting high PPV and NPV is a natural consequence of the desire to minimize the number of incorrect decisions based on test results. But when tests are scarce, such as during the COVID-19 pandemic, high PPV and NPV can at most be achieved for a small minority of individuals. Hence, obtaining precise information about one individual has an opportunity cost, as it means that less accurate information will have to be used on other people.

In this paper, we derive the optimal test allocation that takes into account (i) the information value of a test for a given individual and (ii) the opportunity cost of not using the test on other individuals. We show how simple monotonicity conditions simplify the problem in natural applications, and characterize the marginal benefit of a test expanding the available tests' portfolio.

To illustrate our model, consider the following problem: A public health authority has to make decisions about 2,000 individuals, of whom half are symptomatic patients and half are asymptomatic members of the general public. We assume that the authority ascribes a prior infection probability of 60% and 5% to these groups, respectively. For each individual, the authority decides whether to quarantine them or not. For the sake of simplicity, we assume in this introduction that errors are equally costly (so that quarantine is chosen if and only if posterior infection probability exceeds 50%). In the absence of testing, this means that all symptomatic patients and none of the general public get quarantined, for a combined expected number of $0.4 \times 1000 + 0.05 \times 1000 = 450$ errors.

Now assume that the health authority has 1,000 each of the polymerase chain reaction (PCR) and antigen (Ag) tests. Both these tests detect the presence of the virus but they differ in their sensitivities - probabilities that they identify an infectious person - and specificities - probabilities of correctly indicating that a person is not infectious. Let us assume that the PCR test has a sensitivity of 97% and a specificity of 100%, the Ag test a sensitivity of 90% and a specificity of 99%.[1]

The medical profession is well aware that screening of a low-prevalence group with an unspecific test leads to low PPV and possible overdiagnosis; see e.g. the discussion of mammography in Maxim et al. (2014). A decision maker who targets satisfactory PPV (so that most positives are true positives) would use Ag tests on the symptomatic and PCR tests on the asymptomatic group.[2] Then every single positive PCR test is correct, as are more than 99% of all positive Ag tests. This produces an expected number of $600 \times 0.1 + 400 \times 0.01 + 50 \times 0.03 = 65.5$ errors.

[1]These values are illustrative. Sensitivity for PCR tests varies from 68% to 100%. For antigen tests, the variance in sensitivity is larger (0% to 94%); average specificity is 99.5%. (Dinnes et al., 2020)

[2]E.g., WHO (2020): "Ag-RDTs are not recommended for routine surveillance purposes or case management in this setting. Positive test results would likely be false positives. Molecular testing is preferred."

But now consider what happens if the health authority uses PCR tests on symptomatic patients and screens the asymptomatic population with Ag tests instead. This provides a much lower PPV for Ag tests: Almost one out of every five positive Ag tests is now wrong. But overall, this allocation leads to a lower expected number of errors, $600 \times 0.03 + 50 \times 0.1 + 950 \times 0.01 = 32.5$.

Why is that? The desire to achieve a high PPV when testing asymptomatic people leads to accurate decisions about them, but ignores the opportunity cost of not using the more sensitive PCR tests on the symptomatic population. The PCR test is more accurate in both dimensions but its comparative advantage relative to Ag test is in finding infections, and it is therefore better used on the group where infections are more likely – even if that means making more mistakes in dealing with the asymptomatic group.

In our general model, we start with an arbitrary finite set of individuals. Each individual is characterized by a pretest infection probability and two costs associated with being falsely treated as infectious or not infectious. Such "error costs" subsume economic costs (e.g., the loss in productivity of quarantining a non-infectious individual) and health costs (e.g., the increase in infection rate, and so an increase in both hospitalization and fatality rates, due to non isolating an infectious individual). We take these costs as primitives of the model possibly stemming from a larger intertemporal problem of epidemic control. The authority assigns at most one test to each individual from a finite portfolio of tests that differ in their sensitivities and specificities.[3] The objective is to minimize the sum of the expected losses over all individuals.

Proposition 1 provides a simple algebraic condition on test allocation under which no pairwise test permutation is payoff-improving. We then apply the result to three particular scenarios that are relevant in the context of COVID-19. In each of these scenarios, Proposition 1 implies a particular monotonicity property of the optimal test allocation with respect to individuals' characteristics.

First, we consider individuals homogenous in both losses but heterogeneous in their pretest probabilities.We define the slope of a test to be the loss-weighted difference between its sensitivity and specificity. Then, the slope of the test applied to an individual is non-decreasing in her pretest probability. Second, suppose individuals are homogenous in their pretest probabilities and losses stemming from a false-positive error, but they differ in their false-negative losses. Then, the sensitivity of a test assigned to an individual is nondecreasing in her loss from a false-negative error. Finally, if pretest probabilities and false-negative losses are homogenous but false-positive losses are heterogeneous, then the specificity of the

[3]We discuss the question of sequential testing of the same individual to verify test results in section 4.2 below.

test applied to an individual is nondecreasing in her false-positive loss.

We build on the monotonicity of the optimal allocation in Section 4 where we analyze the marginal benefit of an additional test that expands the authority's test portfolio, and we provide an algorithmic solution to the test-allocation problem. We also show how to determine when it is beneficial to carry out confirmatory testing (i.e. using a follow-up test sequentially on the same individual) instead of using at most one test on each person. The monotonicity of the optimal test allocation simplifies the analysis of these problems since it reduces the set of allocations one needs to consider.

Medical diagnostic tests differ in employed screening mechanisms. By considering only their sensitivities and specificities, we implicitly assume that the analyzed tests are identical along other relevant dimensions; for instance, they may all test directly for the virus presence or all test for antibodies.[4] Tests for the virus presence, relevant for the diagnosis of the early-stage of COVID-19, include precise but slow PCR tests and faster but relatively imprecise LAMP tests and antigen tests. Larremore et al. (2020) argue that precision of testing is secondary to its frequency for COVID-19 surveillance. To sustain high volume of testing for the virus presence, the gold standard PCR testing may need to be massively accompanied with less precise methods such as antigen tests, leading to substantial heterogeneity of test precision. The family of antibody tests includes relatively precise but slow ELISA tests and imprecise but cheap serological rapid-tests. Although serological tests are ineffective in the diagnosis of the early stage of the disease, countries with limited budgets have been and may continue to be dependent on these tests with notable heterogeneous precisions; see Figure 1.

We rely on the standard economic framework that measures the value of information to the extent that it guides choice under uncertainty; see Marschak (1959), Arrow (1998) and Radner and Stiglitz (1984) for early contributions. In the context of testing for an infection, this approach has been applied in Boozer and Philipson (2000); Kasy and Teytelboym (2020) adapt this approach to sequential disease testing. Galeotti et al. (2020) explain the economic concept of information value on COVID-19 testing examples. The test-allocation problem is akin to the rational-inattention problem of Sims (2003) of constrained optimization over information structures, but our decision-maker faces additional constraints implied by the discrete nature of the medical tests.

Recent contributions in epidemiology, computer science and economics investigate the value of COVID-19 testing, typically within variants of the SIR diffusion model, e.g., Acemoglu et al. (2020), Grassly et al. (2020), Berger et al. (2020), Cleevely et al. (2020), Piguillem and Shi (2020), Gollier and Gossner (2020), Jonnerby et al. (2020) and Brotherhood

[4]Toxvaerd (2020) considers the heterogeneity of tests' ability to screen the stage of the disease but abstracts from the heterogeneity in precision.

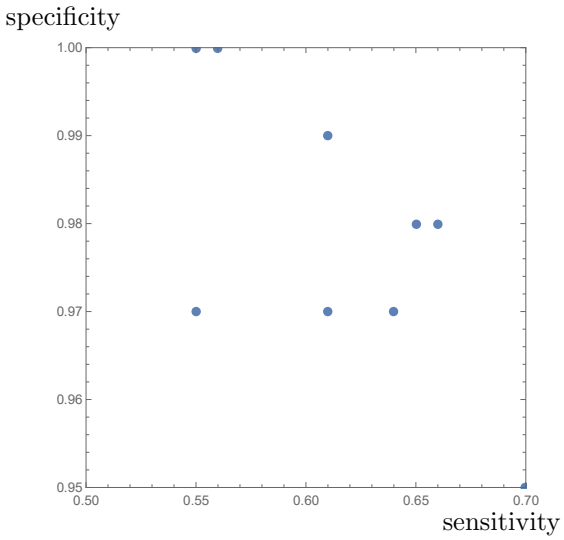

Figure 1: Precision parameters of eight serological tests as validated in Adams et al. (2020).

et al. (2020). Broadly speaking, those papers analyse to what extent testing, with one homogenous test type, allows to relax social distancing for a given infection-flow target; often testing is random across the population. Relative to this literature, our main contribution is the formulation of the test-allocation problem for the heterogeneous tests portfolio. This allows us to derive policy prescriptions on allocation of tests that differ in their specificities and sensitivities.

2 Test-allocation Problem

A public-health authority, referred to here as decision-maker (DM), chooses an action $a_i \in \{0,1\}$ for each individual $i \in \mathcal{I} = \{1,\dots,I\}$ and receives payoff $\sum_{i\in\mathcal{I}} u_i(a_i,\theta_i)$, where each $\theta_i \in \{0,1\}$ is a private health state of the individual i unknown to the DM. The DM assigns prior – in medical terminology pre-test – probability $p_i \in [0,1]$ to $\theta_i = 1$ for each i. The states are independent across the individuals. To avoid trivialities, we assume that neither action is dominant and label the actions so that the optimal choice in state θ is $a = \theta$. Let $\ell_i^\theta = u_i(\theta,\theta) - u_i(1-\theta,\theta) > 0$ be the loss from the decision error in state $\theta_i = \theta$ for the individual i and let $\ell_i = (\ell_i^0, \ell_i^1)$.

The DM can employ tests t from a finite set $\mathcal{T}$. Each test t is a Blackwell experiment that delivers a signal $x \in \{0,1\}$ with interior probability $t(x \mid \theta)$ when applied to an individual in health state θ. We assume that $\mathcal{T}$ includes a trivial test, denoted $\emptyset$, that generates a signal independent of θ; applying test $\emptyset$ to an individual is equivalent to not testing her. It is feasible

for the authority to test nobody, i.e., the number of trivial tests is at least I. Without loss of generality, we label the signals generated by each test t so that $t(1 \mid \theta)/t(0 \mid \theta)$ increases in θ, and refer to $t(1 \mid 1)$ as the *sensitivity* and to $t(0 \mid 0)$ as the *specificity* of the test t. The results of the tests are conditionally independent across the individuals. The DM assigns to each individual i a test, updates her belief about the individual based on the test applied and the test result, and chooses an action $a_i \in \{0, 1\}$.

We define the value of a test in the standard manner. Let

$$v_i(q) = \max_{a \in \{0,1\}} \{q u_i(a, 1) + (1 - q) u_i(a, 0)\}$$

be the value of the DM with belief q with respect to the choice of the action a_i. The value of the test t applied to individual i with pretest probability p is

$$V_i(p, t) = \mathrm{E}\left[v_i\left(q_{t,p}(x)\right)\right] - v_i(p),$$

where $q_{t,p}(x) = \frac{pt(x|1)}{pt(x|1)+(1-p)t(x|0)}$ is the posterior probaility – in medical terminology post-test probability – formed after the test t returns result x for an individual with the pretest probability p. The expectation is with respect to the signal x.

The *test-allocation problem* consists of finding a one-to-one test-allocation rule $\tau : \mathcal{I} \longrightarrow \mathcal{T}$ that solves

$$\max_{\tau} \sum_{i \in \mathcal{I}} V_i\left(p_i, \tau(i)\right). \tag{1}$$

Distinct costs of the two types of errors allow the authority to accommodate the trade-off between the socioeconomic cost of isolation ℓ_i^0 and epidemiological cost of contagion ℓ_i^1. Individual heterogeneity across these costs may reflect distinct degrees of socioeconomic disruptions and varying connectedness and hence various spreading potential across individuals. Although these costs are exogenous, our model can be embedded into a broader optimization. The authority's broader dynamic problem may consist of minimizing the flow of economic destruction under a constraint of keeping the effective reproduction factor of the disease under a chosen threshold. Since the constraint requires the prevention of a sufficient number of infections at any given time, the damage $-u_i(0, 1)$ from failing to isolate an infectious individual i would then be proportional to the product of the expected number of infections caused by i and the shadow price of the epidemiological constraint; see Appendix A.

We assume that if a test $t \neq \emptyset$ has no value for individual i, $V_i(p_i, t) = 0$, then it is not assigned to i and, instead, the trivial test $\emptyset$ is assigned to i. Note that the test-allocation problem, combined with the a posteriori optimal action choice, is equivalent to the maximization of the DM's payoff.

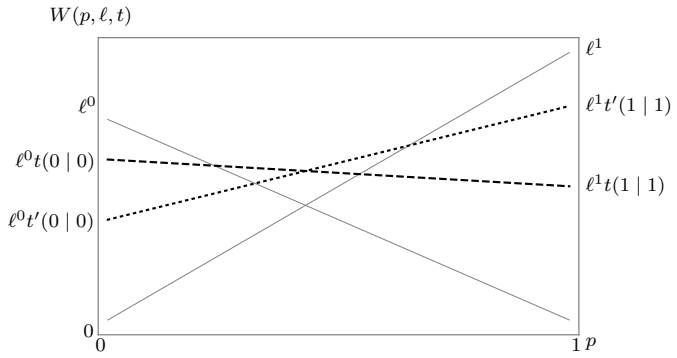

Figure 2: Modified values $W(p, \ell, t)$ for a population with individual-independent losses $\ell_i = \ell = (\ell^0, \ell^1)$. The two full lines correspond to the trivial tests $\emptyset^0$ and $\emptyset^1$. Dashed and dotted lines correspond to non-trivial tests.

We proceed with a useful transformation of the test-allocation problem. Instead of the trivial test $\emptyset$, we introduce two trivial tests $\emptyset^0$ and $\emptyset^1$. The trivial test $\emptyset^x$ always returns the signal x, for $x = 0, 1$. There is a non-binding supply of both $\emptyset^x$, and all the other tests are in the same supply as before. That is, we let $\mathcal{T}'$ to contain all non trivial tests tests $t \neq \emptyset$ from $\mathcal{T}$, I copies of $\emptyset^0$ and $\emptyset^1$, and we remove all trivial tests $\emptyset$. This is useful because when the DM allocate tests in $\mathcal{T}'$, we can assume, without loss of generality, that the DM chooses for each individual i an action a_i equal to the result of the test applied to i. That is, each test t is equivalent to the stochastic choice rule $t(a \mid \theta)$.

We define the *modified value of the test* $t \in \mathcal{T}'$ applied to an individual i with pretest probability p to be sum of the test's sensitivity and specificity weighted by the loss values and pretest probabilities of both health states. Formally,

$$W(p, \ell, t) = p\ell^1 t(1 \mid 1) + (1 - p)\ell^0 t(0 \mid 0).$$

Figure 2 plots the modified values of tests for a special case in which losses ℓ_i are the same for all individuals i. The *modified test-allocation problem* consists of finding a one-to-one test-allocation rule $\tau : \mathcal{I} \longrightarrow \mathcal{T}'$ that solves

$$\max_{\tau} \sum_{i \in \mathcal{I}} W\left(p_i, \ell_i, \tau(i)\right). \tag{2}$$

The allocation problem (1) and the modified allocation problem (2) are equivalent. Let p_i^* given by $p_i^* \ell_i^1 = (1 - p_i^*)\ell_i^0$ denote the belief at which the DM is indifferent between actions $a_i = 0$ and $a_i = 1$.

Lemma 1. *Suppose that τ solves problem (1) and τ' solves problem (2).*

1. *For a nontrivial test $t \neq \emptyset$, $\tau(i) = t$ if, and only if, $\tau'(i) = t$.*
2. *When $p_i < p_i^*$, then $\tau(i) = \emptyset$ if, and only if, $\tau'(i) = \emptyset^0$. When $p_i > p_i^*$, then $\tau(i) = \emptyset$ if and only if $\tau'(i) = \emptyset^1$.*

Proof of Lemma 1. For any p and t such that $V_i(p,t) > 0$, the DM who applies t to an individual i with pretest probability p chooses a_i according to the stochastic choice function $t(a_i \mid \theta_i)$. Hence,

$$\begin{aligned} V_i(p,t) &= p(t(1 \mid 1)u_i(1,1) + t(0 \mid 1)u_i(0,1)) + (1-p)(t(0 \mid 0)u_i(0,0) + t(1 \mid 0)u_i(1,0)) - v_i(p) \\ &= W(p, \ell_i, t) + pu_i(0,1) + (1-p)u_i(1,0) - v_i(p). \end{aligned}$$

Similarly, when $p \leq p_i^*$ then $V_i(p,\emptyset) = W(p, \ell_i, \emptyset^0) + pu_i(0,1) + (1-p)u_i(1,0) - v_i(p)$ and when $p \geq p_i^*$ then $V_i(p,\emptyset) = W(p, \ell_i, \emptyset^1) + pu_i(0,1) + (1-p)u_i(1,0) - v_i(p)$. Thus, when $\tau : \mathcal{I} \to \mathcal{T}$ and $\tau' : \mathcal{I} \to \mathcal{T}'$ are such that $\tau(i) = t = \tau'(i)$ for all $t \neq \emptyset, \emptyset^0, \emptyset^1$, and if $\tau(i) = \emptyset$ and $\tau'(i)$ optimally allocates $\emptyset^0$ or $\emptyset^1$, then the objectives achieved by τ and τ' in problems (1) and (2), respectively, differ only by a term independent of the tests' allocations. □

From now on, we always refer to the modified test-allocation problem, cease to refer to the modification and write $\mathcal{T}$ instead of $\mathcal{T}'$.

The next result provides a simple necessary condition for optimality of allocation. Let $\mathbf{w} = (w^0, w^1) = ((1-p)\ell^0, p\ell^1)$ and let $\mathbf{t}$ stand for the vector $(t(0 \mid 0), t(1 \mid 1))$. Let "$\cdot$" stand for the scalar product.

Proposition 1. *If τ solves the test-allocation problem, then*

$$0 \leq (\mathbf{w}_i - \mathbf{w}_j) \cdot (\boldsymbol{\tau}(i) - \boldsymbol{\tau}(j)) \text{ for all } i, j \in \mathcal{I}.$$

Proof of Proposition 1. Optimality of τ implies that for all $i, j \in \mathcal{I}$,

$$\begin{aligned} 0 &\leq W(p_i, \ell_i, \tau(i)) + W(p_j, \ell_j, \tau(j)) - W(p_i, \ell_i, \tau(j)) - W(p_j, \ell_j, \tau(i)) \\ &= \mathbf{w}_i \cdot \boldsymbol{\tau}(i) + \mathbf{w}_j \cdot \boldsymbol{\tau}(j) - \mathbf{w}_i \cdot \boldsymbol{\tau}(j) - \mathbf{w}_j \cdot \boldsymbol{\tau}(i) \\ &= (\mathbf{w}_i - \mathbf{w}_j) \cdot (\boldsymbol{\tau}(i) - \boldsymbol{\tau}(j)). \end{aligned}$$

□

3 Applications

In what follows, we apply Proposition 1 to three particular populations, each heterogeneous only along one dimension. In these scenarios, Proposition 1 implies simple monotonicity properties of the optimal allocations. To illustrate the economic content of the results, we refer to individual i with $\theta_i = 1$ as infectious, to action $a_i = 1$ as to quarantining i and to action 0 as not quarantining the individual.

We start with the DM who has no individual-specific information on either of the two losses but possesses individual-level information on individuals' health statuses. For instance, individuals may have or lack symptoms or may have reported different contact histories, and the DM maps these pieces of information to heterogeneous pretest probabilities p_i. Assuming homogenous losses ℓ^0 and ℓ^1, let the *slope* of the test t be defined as

$$\sigma_t = t(1 \mid 1)\ell^1 - t(0 \mid 0)\ell^0.$$

That is, the slope of test t is the loss-weighted difference between its sensitivity and specificity.[5]

Corollary 1. *Suppose $\ell_i^0 = \ell_j^0$ and $\ell_i^1 = \ell_j^1$ for all $i, j \in \mathcal{I}$.*

1. *Slopes of the optimally allocated tests are nondecreasing in the individuals' pretest probabilities. That is, if $p_i > p_j$, then $\sigma_{\tau(i)} \geq \sigma_{\tau(j)}$.*

2. *Individuals with sufficiently low or high pretest probabilities are not tested. That is, there exists $\underline{p} < \overline{p}$ such that, if $p_i < \underline{p}$, then the DM chooses $a_i = 0$ without testing individual i. If $p_i > \overline{p}$, then the DM chooses $a_i = 1$ without testing i. If $\underline{p} < p_i < \overline{p}$, then the DM applies a non-trivial test to i and chooses a_i equal to the test's result.*

The latter statement follows from the fact that the two trivial tests, $\emptyset^0$ and $\emptyset^1$, have the extreme slopes $-\ell^0$ and ℓ^1, respectively, across all the tests in $\mathcal{T}$. Intuitively, opportunity cost of a test exceeds the value of information for individuals with pretest probabilities close to 0 or 1, and thus near certain types are left untested.

Next, we consider a population for which the DM does not have individual-specific information on the health statuses, and thus she attaches a same pretest infection probability

[5] Bergemann et al. (2018) introduce the slope of the test (in their terminology "differential informativeness") in a related context. They consider a seller who offers a menu of tests to a buyer with private information. They show in a binary action - binary state setting that incentive-compatibility requires the slope of the test to be nondecreasing in the pretest probability of the type whom the test is allocated to. The intuition for their result can be gleaned from our Figure 2: Incentive compatibility implies that the test allocated to each type p_i must be on the upper envelope of the modified test values (net of prices).

$p_i = p$ to all individuals. The considered group of individuals is also homogenous in their meeting rates (e.g., they all work in the same location/plants and live in the same city). Hence, the loss from leaving infectious individuals unquarantined is homogenous within the population, i.e., $\ell_i^1 = \ell_j^1$ for all $i, j \in \mathcal{I}$. However, quarantine costs are heterogeneous across individuals. For example, for those individuals who can work from home the cost of being quarantined is lower than for those who cannot work from home. The DM is aware of this heterogeneity, i.e, the DM knows individual-specific losses ℓ_i^0 stemming from the false-positive errors.

Corollary 2. *Suppose $p_i = p_j$ and $\ell_i^1 = \ell_j^1$ for all $i, j \in \mathcal{I}$.*

1. *Specificities of the optimally allocated tests are nondecreasing in the individuals' false-positive losses. That is, if $\ell_i^0 > \ell_j^0$, then $\tau(i)(0 \mid 0) \geq \tau(j)(0 \mid 0)$.*

2. *Individuals with sufficiently low or high false-positive losses are not tested. That is, there exists $\underline{\ell} < \overline{\ell}$ such that, if $\ell_i^0 < \underline{\ell}$, then the DM chooses $a_i = 1$ without testing individual i. If $\ell_i^0 > \overline{\ell}$, then the DM chooses $a_i = 0$ without testing i. If $\underline{\ell} < \ell_i^0 < \overline{\ell}$, then the DM applies a non-trivial test to i and chooses a_i equal to the test's result.*

The second part of the corollary follows from the fact that the two trivial tests, $\emptyset^0$ and $\emptyset^1$, have the extreme specificities 1 and 0, respectively, across all the tests in $\mathcal{T}$.

Finally, we assume that the DM has no individual-specific information on the health statuses nor on the quarantine costs. However, the DM has information on the social connectivity/meeting rates of the individuals. Those with high meeting rates spread the virus to many others when they are infectious and not quarantined, hence they generate large losses. In this case, p and ℓ^0 are homogenous across the population and ℓ_i^1 differ across i.[6]

Corollary 3. *Suppose $p_i = p_j$ and $\ell_i^0 = \ell_j^0$ for all $i, j \in \mathcal{I}$.*

1. *Sensitivities of the optimally allocated tests are nondecreasing in the individuals' false-negative losses. That is, if $\ell_i^1 > \ell_j^1$, then $\tau(i)(1 \mid 1) \geq \tau(j)(1 \mid 1)$.*

2. *Individuals with sufficiently low or high false-negative losses are not tested. That is, there exists $\underline{\ell} < \overline{\ell}$ such that, if $\ell_i^1 < \underline{\ell}$, then the DM chooses $a_i = 0$ without testing individual i. If $\ell_i^1 > \overline{\ell}$, then the DM chooses $a_i = 1$ without testing i. If $\underline{\ell} < \ell_i^1 < \overline{\ell}$, then the DM applies a non-trivial test to i and chooses a_i equal to the test's result.*

[6]The UK government has recently announced a pilot for family members to get regular testing for safe care home visits. Relatives of those living in care homes will be tested before reuniting with their loved ones in care homes (see https://www.gov.uk/government/news/pilot-for-family-members-to-get-regular-testing-for-safer-care-home-visits for the official press release). Visitors will be offered either a PCR or Ag test. Since a false-positive error is relatively costly in the context of the care home visit, our result suggests that the visitors should be prioritized in the access to the PCR testing.

The three corollaries are not exhaustive of practical circumstances in which Proposition 1 implies monotonicity of the optimal test allocation. Suppose for instance that all the tests in the portfolio have a same specificity and differ only in their sensitivities.[7] Let the population differ in all three parameters l_i^0, l_i^1, and p_i. Then, sensitivity of the test assigned to individual i is non-decreasing in $l_i^1 p_i$.

4 Marginal Benefit of a Test

We now characterize the marginal benefit of a test that becomes newly available relative to the current optimized test allocation. This benefit has two parts: First, the direct improvement achieved by the new test relative to the test it replaces. Second, the indirect benefit from second-best use of the replaced test, i.e. the opportunity cost of the replaced test. We provide a recursive characterization of the opportunity cost accounting for the further replacements induced by the new application of the first replaced test. The characterizations of the marginal benefit and of the opportunity cost are greatly simplified by the monotonicity of the optimal allocation since the monotonicity reduces the number of all possible reoptimizations we need to consider. To that end, we proceed here with one of the three applications from the previous section that possess the monotonicity structure of their solutions. Though the method below applies to all these three settings, we formulate it for the first setting in which both loss values ℓ^0 and ℓ^1 are homogenous and the pretest probabilities p_i differ across individuals. We then use the characterization of the opportunity cost in an algorithmic solution for the optimal test allocation. We conclude the section with two illustrative examples. Subsection 4.1 computes marginal benefits of particular tests for a specific portfolio of serological tests. Subsection 4.2 analyzes whether it is worthwhile to verify the results of relatively imprecise antigen tests by relatively precise PCR tests, accounting for the opportunity cost of the verification test.

We fix the sequence of the pretest probabilities, $p_1, \ldots, p_I$ and assume without loss of generality that it is nondecreasing. Let $\mathcal{J}$ be a subset of the set of individuals $\mathcal{I} = \{1, \ldots, I\}$ and $\mathcal{T}$ be the set of available tests. We write $(\mathcal{J}, \mathcal{T})$ for the test-allocation problem that assigns tests from $\mathcal{T}$ to individuals in $\mathcal{J}$, let $\tau^{\mathcal{J},\mathcal{T}}$ denote the optimal test allocation and let $\mathcal{V}(\mathcal{J}, \mathcal{T})$ be the value in this problem induced by $\tau^{\mathcal{J},\mathcal{T}}$. In what follows we omit the upper index and write τ whenever we refer to $\tau^{\mathcal{I},\mathcal{T}}$.

Suppose that the set of the available tests $\mathcal{T}$ is expanded by a test t^*. We define the

[7] This is a realistic approximation for PCR and Ag tests for SARS-COV-2 that have approximately 100% specificity and differ in sensitivities. See `https://www.fda.gov/news-events/press-announcements/coronavirus-covid-19-update-fda-authorizes-first-antigen-test-help-rapid-detection-virus-causes`.

marginal benefit of test t^* in the test allocation problem $(\mathcal{J}, \mathcal{T})$ to be

$$B\left(t^*, \mathcal{J}, \mathcal{T}\right) = \mathcal{V}\left(\mathcal{J}, \mathcal{T} \cup \{t^*\}\right) - \mathcal{V}(\mathcal{J}, \mathcal{T}).$$

To characterise the marginal benefit $B\left(t^*, \mathcal{I}, \mathcal{T}\right)$ of test t^* in the problem $(\mathcal{I}, \mathcal{T})$, we let

$$\pi_i = B\left(\tau(i), \mathcal{I} \setminus \{i\}, \mathcal{T} \setminus \{\tau(i)\}\right), \tag{3}$$

and refer to it as to the opportunity cost of the test $\tau(i)$ allocated to the individual i in the solution to the problem $(\mathcal{I}, \mathcal{T})$. It captures the marginal benefit that the test allocated to i could provide if it is removed from i and it is made available to the residual individuals in $\mathcal{I} \setminus \{i\}$.

The additivity of the payoffs implies that the marginal benefit of the test t^* is

$$B\left(t^*, \mathcal{I}, \mathcal{T}\right) = \max\left\{\max_{i \in \mathcal{I}}\{W\left(p_i, t^*\right) - W\left(p_i, \tau(i)\right) + \pi_i\}, 0\right\}, \tag{4}$$

where we omit ℓ from the argument of $W(p, \ell, t)$. That is, the welfare effect of the replacement of the test $\tau(i)$ by t^* is the sum of (a) the direct increase of value obtained for the individual i, i.e., $W\left(p_i, t^*\right) - W\left(p_i, \tau(i)\right)$ and (b) the marginal benefit of the test $\tau(i)$ in the residual allocation problem over individuals in $\mathcal{I} \setminus \{i\}$, i.e., the opportunity cost π_i. The marginal benefit of the test t^* is obtained by finding the individual for whom this welfare effect is largest (if non-negative, otherwise the test is disposed of).

We provide a recursive characterisation of π_i that relies on the monotonicity result from Corollary 1. Finding the best alternative use of the test $\tau(i)$ among the individuals in $\mathcal{I} \setminus \{i\}$ is, without additional structure, a hard problem.[8] We show though that monotonicity greatly reduces the set of possible re-optimizations of the allocation over the residual agents $\mathcal{I} \setminus \{i\}$, when test $\tau(i)$ is removed from the individual i. In particular, this re-optimization consists either of a sequence of adjacent individuals on the right of i each passing their originally allocated test to their right-adjacent neighbour or of a sequence of adjacent individuals on the left of i each passing their originally allocated test to their left-adjacent neighbour. Accordingly, we recursively define the *left-hand* and *right-hand* opportunity costs λ_i and ρ_i of the test $\tau(i)$ allocated to i to be the the values of the test $\tau(i)$ when the test reallocation

[8]The opportunity cost is similar to a Vickrey-Clarke-Groves tax. In a general problem, calculating all of the VCG taxes is an intractable combinatorial problem. Here, the monotonicity of optimal allocation allows for tractable characterization.

is restricted to a sequence of tests' shifts all in one direction: $\lambda_1 = 0$ and $\rho_I = 0$, and

$$\begin{aligned} \lambda_i &= \max\{0, W(p_{i-1}, \tau(i)) - W(p_{i-1}, \tau(i-1)) + \lambda_{i-1}\} \text{ for } i > 1, \\ \rho_i &= \max\{0, W(p_{i+1}, \tau(i)) - W(p_{i+1}, \tau(i+1)) + \rho_{i+1}\} \text{ for } i < I. \end{aligned} \tag{5}$$

Thus, λ_i is the benefit of replacing the test $\tau(i-1)$ allocated to $i-1$ by the test $\tau(i)$ accounting for the fact that the test $\tau(i-1)$ becomes available for $i-2$, etc; ρ_i is the analogous maximal benefit attainable in a series of one-step rightward tests' replacements.

Lemma 2. *The opportunity cost of the test $\tau(i)$ allocated to the individual i is*

$$\pi_i = \max\{\lambda_i, \rho_i\}. \tag{6}$$

If two individuals i and j are allocated a test with a same specificity and sensitivity, $\tau(i) = \tau(j)$, then $\pi_i = \pi_j$.

Proof of Lemma 2. Let us start with observing the following *no-recall* property of the optimal allocation: If a test $t \in \mathcal{T}$ has not been allocated to an individual in the optimal allocation of $(\mathcal{J}, \mathcal{T})$, then this test t will not be allocated in the problem $(\mathcal{J}, \mathcal{T} \cup \{t^*\})$. That is, we denote (with some abuse of notation) the set of tests employed in the optimal allocation of $(\mathcal{J}, \mathcal{T})$ to individuals in $\mathcal{J}' \subseteq \mathcal{J}$ by $\tau^{\mathcal{J},\mathcal{T}}(\mathcal{J}') = \left\{t \in \mathcal{T} : \exists i \in \mathcal{J}' \text{ such that } \tau^{\mathcal{J},\mathcal{T}}(i) = t\right\}$. We claim that the following no-recall property holds

$$\mathcal{V}(\mathcal{J}, \mathcal{T} \cup \{t^*\}) = \mathcal{V}\left(\mathcal{J}, \tau^{\mathcal{J},\mathcal{T}}(\mathcal{J}) \cup \{t^*\}\right).$$

The property holds if $|\mathcal{J}| = 1$. Suppose that the no-recall property holds for all sets of individuals with size $|\mathcal{J}| - 1$. If t^* is not allocated in the problem $(\mathcal{J}, \mathcal{T} \cup \{t^*\})$, then the property holds. Now suppose that t^* is allocated to an individual $i \in \mathcal{J}$ in the problem $(\mathcal{J}, \mathcal{T} \cup \{t^*\})$. Then the allocation to $\mathcal{J} \setminus \{i\}$ solves problem $(\mathcal{J} \setminus \{i\}, \mathcal{T})$ and only tests from $\tau^{\mathcal{J},\mathcal{T}}(\mathcal{J})$ are allocated in this problem by the induction hypothesis, as needed.

The no-recall property allows to rewrite the expression in (3) for the opportunity cost of test $\tau(i)$ allocated to individual i as follows

$$\pi_i = \mathcal{V}(\mathcal{I} \setminus \{i\}, \tau(\mathcal{I})) - \mathcal{V}(\mathcal{I} \setminus \{i\}, \tau(\mathcal{I} \setminus \{i\})).$$

The monotonicity of the optimal allocation in the problem $(\mathcal{I} \setminus \{i\}, \tau(\mathcal{I}))$ implies that $\tau(i)$ will be either disposed of with or allocated only to $i-1$ or $i+1$ (when several individuals in the left or right neighborhoods have a same pretest probability then restriction to this one-step reallocation is without loss).

Assume that $\tau(i)$ is allocated to $i-1$ in the problem $(\mathcal{I}\setminus\{i\},\tau(\mathcal{I}))$. Then, by the monotonicity again, $\tau(i-1)$ can be disposed of with or allocated only to $i-2$. In the latter case, $i-2$ can be disposed of with or allocated only to $i-3$, etc. The chain of replacements terminates when the last replaced test is disposed of. Simple optimization over the length of this leftward chain of replacements secures payoff $\mathcal{V}(\mathcal{I}\setminus\{i\},\tau(\mathcal{I}\setminus\{i\}))+\lambda_i$ in the problem $(\mathcal{I}\setminus\{i\},\tau(\mathcal{I}))$. An analogous argument applies if $\tau(i)$ is allocated to $i+1$ in the problem $(\mathcal{I}\setminus\{i\},\tau(\mathcal{I}))$. Then, optimization over the length of the rightward chain of replacements secures payoff $\mathcal{V}(\mathcal{I}\setminus\{i\},\tau(\mathcal{I}\setminus\{i\}))+\rho_i$. Optimal choice over the leftward and rightward chains of replacements or disposal of the test $\tau(i)$ implies (6).

For the other statement in the proposition, note that the monotonicity of allocation and the definitions of λ_i and ρ_i imply that if $\tau(i)=\tau(j)$, then $\lambda_i=\lambda_j$ and $\rho_i=\rho_j$. □

The next result summarizes.

Proposition 2. *The marginal benefit $B(t^*,\mathcal{I},\mathcal{T})$ of the test t^* is given by* (4)*, where the opportunity costs π_i are given by* (5) *and* (6)*.*

Let us now turn to the characterization of the optimal test allocation. The test-allocation problem is a special case of a complete-information allocation problem solved in Koopmans and Beckmann (1957) by linear-programming techniques. Here, we exploit additional simplicity stemming from the monotonicity of solution that is not assumed in the general setting of Koopmans and Beckmann. Our solution highlights the role of the opportunity cost of a test and features simple replacement chains induced by gradual expansion of the test portfolio.

The monotonicity result and the characterization of the opportunity costs allow for a simple algorithmic solution of the test-allocation problem. We label the tests in the set $\mathcal{T}=\{t_1,\ldots,t_T\}$ monotonically so that slope σ_k of the test t_k is nondecreasing in k. The trivial non-informative tests that have extreme slopes are at the beginning and the end of the sequence while the informative tests are in its middle. The algorithm starts with allocating the first I (trivial) tests from the sequence to the individuals, and then it expands the test portfolio by adding one additional test from the sequence in each further step and reoptimizes the allocation. In each step, if the newly available test is not disposed of, then it must be applied to the individual $i=I$ with the highest pretest probability since this test has a higher slope than all the previously allocated tests. The newly added test should be applied to the individual I if it improves upon the previously allocated test to I accounting for its opportunity cost.

That is, in step 1 of the algorithm, we allocate the first I tests from the sequence monotonically according to $\tau^1(i)=t_i$. In step $l=2,\ldots,T-I+1$, we compute the left-hand costs

λ_i^{l-1} of the test allocated to individuals $i \in \mathcal{I}$ under the test allocation τ^{l-1} from the step $l-1$. Let us consider the test t_{I+l-1}; this test has not been considered in the $l-1$ previous steps and hence it has a weakly higher slope than all tests assigned in the allocation τ^{l-1}. If $W\left(p_I, \tau^{l-1}(I)\right) - \lambda_I^{l-1} > W\left(p_I, t_{I+l-1}\right)$, then we dispose of the test t_{I+l-1}, set τ^l to τ^{l-1} and terminate the step l. Otherwise, if $W\left(p_I, \tau^{l-1}(I)\right) - \lambda_I^{l-1} \leq W\left(p_I, t_{I+l-1}\right)$, then we (a) set $\tau^l(I) = t_{I+l-1}$, (b) find maximal i^* such that $\lambda_{i^*}^{l-1} = 0$, (c) dispose of the test $\tau^{l-1}(i^*)$, set $\tau^l(i) = \tau^{l-1}(i+1)$ for all $i \geq i^*$, (d) set $\tau^l(i) = \tau^{l-1}(i)$ for all $i < i^*$, and terminate the step l.

4.1 Example

This section illustrates our marginal-benefit characterization in a simple example that considers four rapid serology tests from Table 1. Two of them, manufactured by Guangzhou Wondfo Biotech and Zhuhai Livzon Diagnostics, were purchased by the Indian government in April 2020.[9] The other two tests are particular test brands validated in Adams et al. (2020) (who do not reveal the brands) that we dub as the "sensitive" and "specific" tests highlighting their comparative advantages.

We suppose that the DM faces two subpopulations of 1000 individuals each with pretest probabilities 5% and 10%, respectively, and is endowed with 750 Wondfo tests and 750 Livzon tests. Such antibody tests can be used for several purposes such as establishing immunity (cf. Grassly et al., 2020) or contact tracing and linking clusters retrospectively (Winter and Hegde, 2020). We assume here that they are used for the latter and that hence false negatives are more costly than false positives; in particular we assume that $\ell^0 = 1$ and $\ell^1 = 4$. The optimal test allocation assigns all Livzon tests to the high-probability subpopulation, tests the residual 250 high-probability individuals with the Wondfo test, assigns the remaining Wondfo tests to 500 low-probability individuals, and leaves the last 500 low-probability individuals untested; see Figure 3.

Let us verify that this indeed is the optimal test allocation. Let t_w and t_l denote the Wondfo and Livzon test types and $p^1 = 0.05$ and $p^2 = 0.1$ be the two pretest probabilities. All tests must be allocated at optimum since the test value $V(p,t) > 0$ for all four combinations of $(p,t) \in \{p^1, p^2\} \times \{t_w, t_l\}$. The monotonicity result of Corollary 1 implies that we only need to optimize over the number of the untested individuals in, say, the high-probability group. That is, we need to verify unprofitability of only one particular deviation that removes the

[9] See New York Times report on the purchase of the two tests; `https://www.nytimes.com/reuters/2020/04/27/world/asia/27reuters-health-coronavirus-india-kits.html`. We retrieved the parameters values for the Wondfo test from a validation study at `https://www.finddx.org/covid-19/dx-data/` and the parameters for the Livzon test from the validation study at `https://pellecome.com/wp-content/uploads/2020/04/4-Evaluation-Report-Livzon-Dx-rapid-test.pdf` on May 11, 2020.

	Sensitivity	Specificity	Marginal benefit
Wondfo	69%	99.1%	0.13
Livzon	78.7%	99.7%	0.174
"sensitive" test	70%	95%	0.097
"specific" test	61%	99%	0.11

Table 1: Four serological tests. See footnote 9 for the validation studies for the first two tests and Adams et al. (2020) for validation of the last two tests. Since we are unable to verify the details of the validation studies, these tests' precision parameters are illustrative.

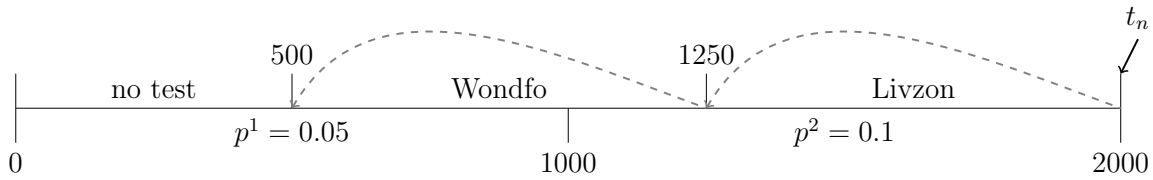

Figure 3: Optimal allocation and optimal replacement chain caused by replacement of one Livzon test with the "sensitive" test t_n. Individuals are placed on the line according to pretest probability p, which weakly increases from left to right.

Livzon test from a high-probability individual and leaves this individual untested. The gain from this deviation, $-V(p^2, t_l) + \pi_l = -0.64$, is negative, as needed, where π_l is (with some abuse of notation) the opportunity cost of the Livzon test computed according to Lemma 2.

Let us expand the test portfolio with one piece of the "sensitive" test from Table 1, denoted by t_n. To compute its marginal benefit, we need to derive first the opportunity costs of both Wondfo and Livzon tests and, then, compute the replacement benefit $W(p_i, t_n) - W(p_i, \tau(i)) + \pi_i$ for each of the four groups of individuals.[10]

We illustrate the case in which the new test t_n replaces the Livzon test assigned to an individual with high pretest probability. The direct payoff effect is $W(p^2, t_n) - W(p^2, t_l) = -0.077$. Additionally, one copy of the Livzon test becomes available and the resulting optimal reallocation of the tests within the set of individuals $\mathcal{I} \setminus \{i\}$ increases the DM's payoff by an amount equal to the opportunity cost of the Livzon test, $\pi_l = 0.174$. This reallocation involves two replacement steps: first, the newly available Livzon test is applied to a high-probability individual who was previously assigned a Wondfo test, and, second, her Wondfo test is reallocated to a low-probability individual who has not been previously tested; see

[10]The optimal allocation partitions the population into (i) those with pretest probability 0.05 who are not being tested, (ii) those with a pretest probability 0.05 tested with Wondfo test, (iii) those with pretest probability of 0.1 tested with Wondfo, and (iv) those with a pretest probability of 0.1 tested with Livzon.

Figure 3 for this chain of reallocations.[11] The total benefit of these two replacement steps is summarized by π_l and recursively defined by (5) and (6). The marginal benefit, $0.097 = -0.077 + 0.174$, of the test t_n is the maximum of the net replacement values across the four groups.

Similarly, the marginal benefit of the "specific" test from Table 1, denoted $\hat{t}_n$, is 0.11. Hence, in this example, the "specific" test $\hat{t}_n$ is a more valuable addition to the DM's test portfolio than the "sensitive" test t_n. Intuitively, since the prevalence among tested individuals is rather low, the DM prefers to expand her portfolio with a test that accurately screens healthy individuals. The DM should purchase $\hat{t}_n$ as long as its cost does not exceed (roughly) 11% of the cost of falsely "detecting" a previous infection. The comparison of the marginal benefits of the two tests, t_n and $\hat{t}_n$, reverses in favour of the "sensitive" test when the prevalence rates are high. In the same setting but with the prevalence rates of the two subpopulations being 30% and 40%, the marginal benefits of the tests t_n and $\hat{t}_n$ are 0.12 and 0.04, respectively.

4.2 Confirmatory testing

If a test result is inconclusive, a health authority can verify it with a follow-up test. While such a confirmatory test decreases uncertainty for a given patient, it also has an opportunity cost. For COVID-19, such confirmatory testing is advised by many national and international health agencies. Our characterization of the marginal benefit of a test can be used to establish whether such testing is fruitful enough to justify its opportunity cost.

We can think of the whole procedure of first-round testing with possible confirmatory testing as a *compound test* with sensitivity and specificity equal to the correct results of the compound testing procedure conditional on the infected and healthy status, respectively. Use of such a compound test consumes the constituent tests in a proportion that depends on the pretest probability of the tested individual and the sensitivity and specificity of the first used test. For example, confirmatory PCR testing after a positive antigen test result consumes as many PCR tests as we expect (true and false) positive antigen test results. Once the sensitivity and specificity of the compound test are derived, we can then use the results of the preceding sections to compare the marginal benefit of the compound test with the sum of the marginal benefits of the constituent tests.

Consider a health authority that faces individuals that are either symptomatic but have no known contact (pretest probability 40%) or who are contacts of confirmed cases but asymptomatic (pretest probability 5%). To avoid uncertainty over the number of the con-

[11] These two reallocation steps can be implemented as a chain of many one-step replacements.

firmatory tests, we assume here that there is a continuum of individuals, half of whom are in each group. As in the introduction, the health authority has PCR tests (97% sensitivity and 100% specificity) and Ag tests (90% sensitivity and 99% specificity). However, there are only PCR tests for 25% of the overall population, and antigen tests for another 50%. We continue to assume, illustratively, that a false negative decision is four times as costly as a false positive.

The optimal test allocation in absence of confirmatory testing is that all PCR tests are used on the symptomatic group, the remainder of that group gets tested with the Ag tests, and the remaining Ag tests are randomly distributed among the asymptomatic contacts. Can it be beneficial to use some of the available PCR tests to confirm results of the antigen testing instead? Two such methods are discussed for the case of COVID-19:

Confirming positive Ag tests from low-pretest-probability individuals with PCR Positive Ag test results of the low-pretest-probability individuals are often incorrect. Confirming such positive test results with highly specific PCR tests is hence recommended, among others, by the WHO, the US CDC and Health Canada.[12] In our example, such confirmatory testing consumes an average of $0.05t^{Ag}(1 \mid 1) + 0.95t^{Ag}(1 \mid 0) = 0.0545$ PCR tests for every antigen test.[13] Furthermore, the compound sensitivity is $t^{cp}(1 \mid 1) = t^{Ag}(1 \mid 1)t^{PCR}(1 \mid 1) = 0.873$ and the compound specificity is $t^{cp}(0 \mid 0) = t^{Ag}(0 \mid 0) + t^{Ag}(1 \mid 0)t^{PCR}(0 \mid 0) = 1$. Since the compound test has a low slope, it is applied to the low-pretest types and one additional such compound test brings a marginal benefit of 0.1746 . Instead of this compound test, the health authority could use an additional Ag test on an asymptomatic individual as well as an additional 0.0545 PCR tests on symptomatic individuals, where the latter frees up another 0.0545 Ag tests for use on asymptomatic individuals. Reallocating optimally these additional tests lead to a marginal benefit of $0.1862 = 0.1705 + 0.0545 \times 0.2885$, where 0.1705 and 0.2885 are marginal benefits of the Ag and PCR test, respectively. Confirming positive Ag tests is hence not beneficial in this situation.

This conclusion changes, however, if prevalence in the asymptomatic group is lower. For example, if the prevalence is only 1%, as could be the case with screening the general population, then the compound test provides marginal benefit 0.035, which is higher than the total marginal benefit 0.029 obtained when the constituent tests are deployed separately. Confirming positive antigen tests with PCR becomes more attractive for two reasons: First, the highly specific compound test becomes relatively more attractive compared to a simple antigen test, whereas its lower sensitivity matters less. Second, it consumes fewer PCR tests,

[12] E.g., CDC (2020): "When the pretest probability is low, those persons who receive a positive antigen test should isolate until they can be confirmed by RT-PCR."

[13] This is under the assumption that tests' results are conditionally independent.

since there are fewer positive Ag results to confirm.

Confirming negative antigen tests from high-pretest-probability individuals with PCR This is recommended by, among others, the CDC and WHO.[14] In our example, we can calculate the compound test parameters analogously to above to find that such a compound test provides a marginal benefit of 0.3257, whereas the constituent tests, if optimally allocated, have a combined marginal benefit of 0.3534. Again, confirmatory testing does not justify its opportunity costs.

This changes for a high enough pretest probability of the symptomatic group. If it is at 60%, then the compound test has a marginal benefit of 0.4033, compared to the total of 0.3267 for the constituent tests. Similarly to above, there are two effects at work that make confirmatory testing more attractive: As the pretest probability among symptomatic people increases, the highly sensitive compound test becomes more valuable relative to the other tests, and it also consumes fewer PCR tests.

Considering the marginal benefits does not in itself tell us what the optimal allocation allowing for confirmatory testing is. But for a given allocation, checking the marginal benefits of creating or disassembling compound tests will reveal whether the given allocation is indeed optimal or can be improved by confirmatory testing. As we have seen, this depends crucially on the pretest probabilities of the tested individuals, but in a more subtle way than existing guidelines suggest.

5 Discussion

Our analysis can be extended in various directions. First, we assume that individuals' health statuses are independent. Our results continue to hold when the correlations are weak so that the optimal decision for each one individual is not affected by the tests' results of the others. Furthermore, our results hold in the presence of correlations if the DM is constrained to individual action choices that do not depend on others' tests results. This latter condition may be relevant in practice due to logistical constraints. In general, however, the presence of correlations in health statuses may revert our results. As an example, consider a case in which all individuals have the same losses and differ in pretest probabilities. Two individuals with nearly median pretest probabilities get tested in the optimal allocation when health statuses are independent by Corollary 1, but if their health statuses are perfectly correlated, then

[14] E.g., WHO (2020): "A negative Ag-RDT result cannot completely exclude an active COVID-19 infection, and, therefore, repeat testing or preferably confirmatory [PCR] testing should be performed whenever possible ..., particularly in symptomatic patients."

testing both is suboptimal.

Second, we abstract from individuals' incentives to get tested and to reveal private information about their infection probabilities. Since individual and social benefits of testing and isolation may differ, incentive compatibility may be a substantial part of the practical test-allocation problem. Bergemann et al. (2018) study incentive compatibility in a market for information. Atkeson et al. (2020) point out in the context of COVID-19 testing that the precision of the post-test posterior information enhances compliance with quarantine measures.

A Appendix

We sketch here a problem of the public health authority (the decision maker, DM) who wishes to minimize economic costs of quarantines keeping the number of new infections sufficiently low. The first-order conditions of this problem coincide with our main model.

There is a continuum of individuals $i \in [0,1]$, where each individual has a pre-test probability $p_i \in (0,1)$ of being infected.[15] The DM decides whether to quarantine each individual. Quarantining i, infectious or not, has an economic cost c_i; falsely releasing i if infected means i will infect s_i others. The DM has a continuous portfolio of tests $t_j(x \mid \theta)$, indexed by $j \in [0,T]$, which includes the two trivial tests. The DM's test allocation is a one-to-one mapping $j(i) : [0,1] \to [0,T]$. The test result is always implemented (this is without loss of generality, since the two trivial tests are available).

The total number of infections in an allocation j is $\int_0^1 p_i t_{j(i)}(0 \mid 1) s_i di$. We assume that the DM wants to keep this spreading below some maximum level $\bar{L}$ (where the value $\bar{L}$ may be time dependent and an outcome of intertemporal optimization that we do not model). The DM's optimization problem is hence:

$$\min_j \int_0^1 \left(p_i t_{j(i)}(1 \mid 1) + (1-p_i) t_{j(i)}(1 \mid 0)\right) c_i di \tag{7}$$

$$\text{s.t.:} \int_0^1 p_i t_{j(i)}(0 \mid 1) s_i di \leq \bar{L}. \tag{8}$$

That is, unlike in our main model, the DM's problem exhibits interdependency across individuals since feasibility of allocation of the tests to a subpopulation depends on the test allocation for the residual population.

Let us assume that the epidemiological constraint (8) is tight enough to be binding, and

[15] The assumption of the continuous population instead of the discrete one eliminates uncertainty over the infection rate caused by imprecise testing.

let μ be its positive shadow price. Then, the first order condition of the problem (7) subject to (8) is the same as that of our main model with the payoff functions $u_i(a_i, \theta_i)$ equal to $u_i(0,0) = 0$, $u_i(1,0) = -c_i$, $u_i(0,1) = -\mu s_i$ and $u_i(1,1) = -c_i$.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. *NBER Working Paper*.

Adams, E. R., M. Ainsworth, R. Anand, M. I. Andersson, K. Auckland, J. K. Baillie, E. Barnes, S. Beer, J. Bell, T. Berry, et al. (2020). Antibody testing for covid-19: A report from the national covid scientific advisory panel. *medRxiv*.

Arrow, K. (1998). J. 1971. essays in the theory of risk bearing. *Chicago: Markham*.

Atkeson, A., M. Droste, M. Mina, and J. H. Stock (2020). Economic benefits of covid-19 screening tests. *NBER Working Paper* (w28031).

Bergemann, D., A. Bonatti, and A. Smolin (2018). The design and price of information. *American Economic Review 108*(1), 1–48.

Berger, D., K. Herkenhoff, and S. Mongey (2020). An seir infectious disease model with testing and conditional quarantine. *Becke Friedman Institute Working Paper*.

Boozer, M. A. and T. J. Philipson (2000). The impact of public testing for human immunodeficiency virus. *Journal of Human Resources*, 419–446.

Brotherhood, L., P. Kircher, C. Santos, and M. Tertilt (2020). An economic model of the covid-19 epidemic: The importance of testing and age-specific policies. *CEPR Working Paper*.

CDC (2020). Interim guidance for rapid antigen testing for sars-cov-2, 4 september 2020.

Cleevely, M., D. Susskind, D. Vines, L. Vines, and S. Wills (2020). A workable strategy for covid-19 testing: stratified periodic testing rather than universal random testing. *COVID Economics, CEPR Press*, 44–70.

Dinnes, J., J. J. Deeks, A. Adriano, S. Berhane, C. Davenport, S. Dittrich, D. Emperador, Y. Takwoingi, J. Cunningham, S. Beese, et al. (2020). Rapid, point-of-care antigen and molecular-based tests for diagnosis of sars-cov-2 infection. *Cochrane Database of Systematic Reviews* (8).

Galeotti, A., J. Steiner, and P. Surico (2020). An economic approach to testing. Mimeo.

Gollier, C. and O. Gossner (2020). Group testing against covid-19. *Covid Economics*.

Grassly, N., M. Pons Salort, E. Parker, P. White, K. Ainslie, M. Baguelin, S. Bhatt, A. Boonyasiri, O. Boyd, N. Brazeau, et al. (2020). Report 16: Role of testing in covid-19 control.

Jonnerby, J., P. Lazos, E. Lock, F. Marmolejo-Cossio, C. B. Ramsey, M. Shukla, and D. Sridhar (2020). Maximising the benefits of an acutely limited number ofcovid-19 tests. *ResearchGate*.

Kasy, M. and A. Teytelboym (2020). Adaptive Targeted Infectious Disease Testing. *Oxford Review of Economic Policy*.

Koopmans, T. C. and M. Beckmann (1957). Assignment problems and the location of economic activities. *Econometrica: journal of the Econometric Society*, 53–76.

Larremore, D. B., B. Wilder, E. Lester, S. Shehata, J. M. Burke, J. A. Hay, M. Tambe, M. J. Mina, and R. Parker (2020). Test sensitivity is secondary to frequency and turnaround time for covid-19 surveillance. *medRxiv*.

Marschak (1959). Remarks on the economics of information. Technical report, Cowles Foundation for Research in Economics, Yale University.

Maxim, L. D., R. Niebo, and M. J. Utell (2014). Screening tests: a review with examples. *Inhalation toxicology 26*(13), 811–828.

Piguillem, F. and L. Shi (2020). The optimal covid-19 quarantine and testing policies. *EIEF Working Paper 20/04*.

Radner, R. and J. Stiglitz (1984). A nonconcavity in the value of information. *Bayesian models in economic theory 5*, 33–52.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Toxvaerd, F. (2020). Social distancing with asymptomatic infection: beliefs, fatalism and testing. Mimeo.

WHO (2020). Antigen-detection in the diagnosis of sars-cov-2 infection using rapid immunoassays: interim guidance, 11 september 2020. Technical report, World Health Organization.

Winter, A. K. and S. T. Hegde (2020). The important role of serology for covid-19 control. *The Lancet Infectious Diseases*.