

Optimal Test Allocation*

Jeffrey Ely
Northwestern U

Andrea Galeotti
LBS

Jakub Steiner
Cerge-Ei and Zurich U

July 10, 2020

Abstract

A health authority chooses a binary action for each of several individuals that differ in their pre-test probabilities of being infected 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 and provide an algorithmic solution for the test-allocation problem.

1 Introduction

Public-health authorities and organizations from around the world have to make decisions about many individuals under uncertainty over their Covid-19 infection statuses with limited test portfolios of heterogeneous qualities. We argue in this note that matching of tests types to individual types improves efficiency of testing relative to the current practice of uniform test allocation. The insights of this paper can guide allocation of tests on the national level, within the workplace, universities and schools. Next, we illustrate this in a simple example.

Optimization over test allocation may have a large impact on the error rates of the screening for disease. For illustration, let a health authority screen 2000 individuals. Using freely observable information about their residences, symptoms, etc., the authority categorizes the individuals into two equally sized low- and high-risk groups with respective Covid-19 prevalences of 10% and 70%. The authority assigns status “healthy” or “Covid-19 case” to each

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

individual with each error being equally costly. In the absence of tests, only the high-risk individuals are categorized as Covid-19 cases and $1000 \times 0.1 + 1000 \times (1 - 0.7) = 400$ errors occur in expectation.

Let us endow the authority with two test types that differ in their sensitivities and specificities, 750 pieces of each type.¹ The first test type, to which we refer as (relatively) *specific* has specificity of 99% and sensitivity of 61%. The second test type is (relatively) *sensitive* with specificity of 95% and sensitivity of 70%. Recall that a test’s sensitivity and specificity are the likelihoods that it correctly screens for the presence and absence of the antibodies, respectively. We assume that each individual is tested at most once. A naive strategy is to distribute the tests randomly uniformly across all the individuals, to choose the action implied by her test result for each individual (which is ex-post optimal) and to choose the ex-ante optimal action for the untested individual. This allocation reduces the number of expected errors by 66 to 334. However, the authority can do better by allocating the 750 sensitive tests to 750 high-risk individuals and the 750 specific tests to the low-risk ones. This change in the allocation further reduces the expected number of errors by about half of the previous error reduction – by 29.25 to 304.75.

The last test allocation turns out to minimize the number of the expected errors. It prioritizes high-risk individuals in access to the sensitive test since, for them, the test’s ability to correctly screen an infected individual is more relevant than for the low-risk individuals. Vice versa, low-risk individuals are assigned with the specific test because, for them, the test’s ability to correctly screen healthy individuals is relevant. In contrast with the common practice, the optimal allocation does not test all high-risk individuals (the allocation that randomly uniformly distributes all the tests among 1500 riskiest individuals leads to 331.25 expected errors).

In our general model, we allow for an arbitrary finite set of individuals that differ in their pre-test infection probabilities and in the losses associated with the two types of decision errors. The authority can assign at most one test to each individual from a finite portfolio of tests that differ in their sensitivities and specificities. The objective is to minimize the sum of the expected losses. Proposition 1 provides a simple algebraic condition on test allocation under which no pairwise test permutation is payoff-improving. We then apply the result

¹ These are, in fact, two serological tests validated in Adams et al. (2020). Serological tests inform various types of public-health decisions. Presence or absence of antibodies indicates individual’s past infectious status and thus may navigate whose contacts to trace and test for infection, Winter and Hegde (2020). Presence of IgM antibodies indicates acute infection and thus may call for isolation of the individual, Jacofsky et al. (2020). If a country periodically tests health and other key workers with PCR tests, then positive serological test result indicates that further PCR screening of the individual is not necessary. Finally, presence of IgG antibodies may indicate some level of immunity of the individual. See Grassly et al. (2020) on both health-care workers’ screening and on immunity passports.

to three particular scenarios that are relevant to the current Covid-19 debates. 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 pre-test probabilities, such as in the introductory example. 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 nondecreasing in her pre-test probability. Second, suppose individuals are homogenous in their pre-test probabilities and losses stemming from the 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 the false-negative error. Finally, if pre-test probabilities and false-negative losses are homogenous but false-positive losses are heterogeneous, then the specificity of the 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. The monotonicity of the optimal test allocation simplifies the analysis of both 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.² 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 current golden-standard of PCR testing may need to be massively accompanied with less precise methods, leading to substantial heterogeneity of test precision. The family of the antibody tests include relatively precise but slow ELISA tests and imprecise but cheap serological rapid-tests. Although the 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)

²Toxvaerd (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).

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), and 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), Cleevly et al. (2020), Piguillem and Shi (2020), Gollier and Gossner (2020), Jonnerby et al. (2020) and Brotherhood 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 different 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 $t(x | \theta)$ 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 $t(1 | \theta)/t(0 | \theta)$ increases in θ , and refer to $t(1 | 1)$ as to *sensitivity* and to $t(0 | 0)$ as to *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

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

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 pre-test probability p is

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

where $q_{t,p}(x) = \frac{pt(x|1)}{pt(x|1) + (1-p)t(x|0)}$ is the posterior – in medical terminology, post-test probability – formed after the test t returns result x for an individual with the pre-test 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 $\tau : \mathcal{I} \rightarrow \mathcal{T}$ that solves

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

We assume that if a test t 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.

We proceed with a useful transformation of the test-allocation problem. Instead of the




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 t and t' , respectively.

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 $t(a | \theta)$.

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

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

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 $\tau : \mathcal{I} \rightarrow \mathcal{T}'$ that solves

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

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 pre-test probability p chooses a_i according to the stochastic choice function $t(a_i | \theta_i)$. Hence,

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

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 $\tau : \mathcal{I} \rightarrow \mathcal{T}$ and $\tau' : \mathcal{I} \rightarrow \mathcal{T}'$ are such that $\tau(i) = t = \tau'(i)$ for all $t \neq \emptyset, \emptyset^0, \emptyset^1$, and if $\tau(i) = \emptyset$ and then $\tau'(i)$ 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. \square

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 $(t(0 | 0), t(1 | 1))$. 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 $i, j \in \mathcal{I}$,

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

\square

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 infected, 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 pre-test probabilities p_i . Assuming homogenous losses ℓ^0 and ℓ^1 , let the *slope* of the test t be defined as

$$\sigma_t = t(1 | 1)\ell_1 - t(0 | 0)\ell_0.$$

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

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' pre-test probabilities. That is, if $p_i > p_j$, then $\sigma_{\tau(i)} \geq \sigma_{\tau(j)}$.*
2. *Individuals with sufficiently low or high pre-test probabilities are not tested. That is, there exists $\underline{p} < \bar{p}$ such that, if $p_i < \underline{p}$, then the DM chooses $a_i = 0$ without testing individual i . If $p_i > \bar{p}$, then the DM chooses $a_i = 1$ without testing i . If $\underline{p} < p_i < \bar{p}$, 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} .

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 pre-test infection probability $p_i = p$ to all individuals. The considered group of individuals is also homogenous in their meeting rates (e.g., they all work in a same location/plants and live in a same city). Hence, the loss from leaving infected 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 $\tau(i)(0 | 0) \geq \tau(j)(0 | 0)$.*
2. *Individuals with sufficiently low or high false-positive losses are not tested. That is, there exists $\underline{\ell} < \bar{\ell}$ such that, if $\ell_i^0 < \underline{\ell}$, then the DM chooses $a_i = 1$ without testing individual i . If $\ell_i^0 > \bar{\ell}$, then the DM chooses $a_i = 0$ without testing i . If $\underline{\ell} < \ell_i^0 < \bar{\ell}$, 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 infected 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 .

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

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.³ 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 $\ell_i^1 p_i$.

4 Marginal Benefit of a Test

The test-allocation problem is an allocation problem in which the DM has complete information about the surplus from each matched pair of the individual and the test; see

³This is a realistic approximation of antigen 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>.

Koopmans and Beckmann (1957) for the linear-programming solution for a general class of these problems. Here, we proceed with simple observations for any of the three applications from the previous section, exploiting 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 pre-test probabilities p_i differ across individuals.

We aim to characterize the marginal benefit of a test that becomes newly available relative to the current DM's test portfolio. We will then describe a simple algorithmic solution to the test-allocation problem that builds on the monotonicity of the optimal allocation. We conclude the section with an illustrative example.

We fix the sequence of the pre-test probabilities, p_1, \dots, p_I and assume without loss of generality that it is nondecreasing. Let \mathcal{J} be a subset of $\mathcal{I} = \{1, \dots, I\}$ and \mathcal{T} be a set of the available tests. We write $(\mathcal{J}, \mathcal{T})$ for the test-allocation problem defined by these two sets, 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 marginal benefit of test t^* to be

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

We now characterise $B(t^*, \mathcal{I}, \mathcal{T})$. To this end, we define

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

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 $(\mathcal{I}, \mathcal{T})$.

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

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

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., $W(p_i, t^*) - W(p_i, \tau(i))$ 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. This monotonicity greatly reduces the set of possible re-optimizations of the allocation over $\mathcal{I} \setminus \{i\}$, when $\tau(i)$ is removed from the individual i and becomes available for $\mathcal{I} \setminus \{i\}$. In particular, the 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 there is a sequence of adjacent individuals on the left of i each passing the originally allocated test to their left-adjacent neighbour. We denote λ_i and ρ_i the *left-hand* and *right-hand* opportunity costs of the test $\tau(i)$ allocated to i , respectively. They are defined in simple recursions as follows: $\lambda_1 = 0$ and $\rho_I = 0$, and

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

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.

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 has not been allocated to an individual in the optimal allocation of $(\mathcal{J}, \mathcal{T})$, then this test 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 $\tau^{\mathcal{J}, \mathcal{T}}(\mathcal{J}') = \{t \in \mathcal{T} : \exists i \in \mathcal{J}' \text{ such that } \tau^{\mathcal{J}, \mathcal{T}}(i) = t\}$. We claim that the following no-recall property holds

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

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 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 with or allocated only to $i - 2$. In latter case, $i - 2$ can be disposed with or allocated only to $i - 3$, etc. The chain of replacements terminates when the last replaced test is disposed with. 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 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 $\rho_i = \rho_j$. \square

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).*

The monotonicity result and the characterization of the opportunity costs allow for a simple algorithmic solution of the test-allocation problem. Let us label the tests in the set $\mathcal{T} = \{t_1, \dots, t_T\}$ 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.

In step 1 of the algorithm, allocate the first I tests from the sequence monotonically according to $\tau^1(i) = t_i$. In step $l = 2, \dots, T - I + 1$, compute the left-hand costs $\lambda_i(\tau^{l-1})$ according to (5) for all $i \in \mathcal{I}$ for the allocation τ^{l-1} . 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 $W(p_I, \tau^{l-1}(I)) - \lambda_I(\tau^{l-1}) > W(p_I, t_{I+l-1})$ then dispose of the test t_{I+l-1} , set τ^l to τ^{l-1} and terminate the step l . Otherwise, if

	Sensitivity	Specificity	Marginal benefit
Wondfo	69%	99.1%	0.13
Livzon	78.7%	99.7%	0.17
“sensitive” test	70%	95%	0.097
“specific” test	61%	99%	0.11

Table 1: Four serological tests. See footnote 4 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.

$W(p_I, \tau^{l-1}(I)) - \lambda_I(\tau^{l-1}) \leq W(p_I, t_{I+l-1})$, then (a) set $\tau^l(I) = t_{I+l-1}$, (b) find maximal i^* such that $\lambda_{i^*}(\tau^{l-1}) = 0$, (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-tests tests from Table 1. Two of them, manufactured by Guangzhou Wondfo Biotech and Zhuhai Livzon Diagnostics, were purchased by the Indian government in April 2020.⁴ The other two tests are the “sensitive” and “specific” tests from the example in the introduction; recall that those tests have been validated in Adams et al. (2020) (who do not reveal the two tests’ brands).

The DM faces two subpopulations of 1000 individuals each with the pre-test probabilities 5% and 10%, respectively, and is endowed with 750 Wondfo test and 750 Livzon tests. The false-negative error is four times costlier than the false-positive error; $\ell_0 = 1$ and $\ell_1 = 4$ where the cost unit corresponds to the economic cost of quarantining an individual. The optimal test allocation assigns all Livzon tests to the high-risk subpopulation, tests the residual 250 high-risk individuals with Wondfo test, assigns the remaining Wondfo tests to 500 low-risk individuals, and leaves the last 500 low-risk 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 pre-test probabilities. All tests must be allocated at optimum since $V(p, t) > 0$ 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, high-risk group. That is, we

⁴See 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.




Figure 3: Optimal allocation and optimal replacement chain caused by replacement of one Livzon test with the “sensitive” test t_n .

need to verify unprofitability of only one particular deviation that removes the Livzon test from a high-risk 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 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 $W(p_i, t_n) - W(p_i, \tau(i)) + \pi_{\tau(i)}$ for each of the four groups of individuals.⁵

We illustrate the case in which the test t_n replaces the Livzon test assigned to an individual i with the high pre-test 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, $\pi_l = 0.17$. This reallocation involves two replacement steps: first, the newly available Livzon test is applied to a high-risk individual who was previously assigned Wondfo test, and, second, her Wondfo test is reallocated to a low-risk individual who has not been previously tested; see Figure 3 for this chain of reallocations.⁶ 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, 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

⁵The optimal allocation partitions the population into (i) those with pre-test probability 0.05 who have not being tested, (ii) those with a pre-test probability 0.05 tested with Wondfo test, (iii) those with pre-test probability of 0.1 tested with Wondfo, and (iv) those with a pre-test probability of 0.1 tested with Livzon.

⁶This two reallocation steps can be implemented as a chain of many one-step replacements.

(roughly) 11% of the quarantine cost. The comparison of the marginal benefits of the two tests, t_n and \hat{t}_n , reverses 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.

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 pre-test probabilities. Two individuals with nearly median pre-test probabilities get tested in the optimal allocation when health statuses are independent by Corollary 1, but if their health statuses are perfectly correlated, then 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 related problem.

Finally, we allow the authority to test each individual at most once and do not consider sequential testing. Sequences of diagnostic test results are difficult to interpret because correlations of the tests' errors are less understood than the marginal error distributions. Additionally, sequential testing is logistically impractical under some conditions. Yet, it has been used in Covid-19 context; Morris and Strack (2017) and Hébert and Woodford (2017) may be useful for the first steps in this direction.

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*.
- 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*.
- Cleevly, 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.
- 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.
- Hébert, B. and M. Woodford (2017). Rational inattention and sequential information sampling. Technical report, National Bureau of Economic Research.
- Jacofsky, D., E. M. Jacofsky, and M. Jacofsky (2020). Understanding antibody testing for covid-19. *The Journal of Arthroplasty*.
- 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 of covid-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.
- Morris, S. and P. Strack (2017). The wald problem and the equivalence of sequential sampling and static information costs. *Unpublished manuscript, June*.
- 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.
- Winter, A. K. and S. T. Hegde (2020). The important role of serology for covid-19 control. *The Lancet Infectious Diseases*.