

WORKING PAPER · NO. 2020-85

Pooled Testing for Quarantine Decisions

Elliot Lipnowski and Doron Ravid

NOVEMBER 2020

Pooled Testing for Quarantine Decisions

Elliot Lipnowski^{*,†}
Columbia University

Doron Ravid
University of Chicago

November 3, 2020

Abstract

We study optimal testing to inform quarantine decisions for a population exhibiting a heterogeneous probability of carrying a pathogen. Because test supply is limited, the planner may choose to test a pooled sample, which contains the specimens of multiple individuals (Dorfman, 1943). We characterize the unique optimal allocation of tests. This allocation features *assortative batching*, whereby agents of differing infection risk are never jointly tested. Moreover, the planner tests only individuals whose prior quarantine decision is the most uncertain. Finally, individuals with higher infection risk are tested in smaller batches, because such tests minimize the informational externality of group testing.

JEL Classification: D04, D61, I18, D62

^{*}E-mail: e.lipnowski@columbia.edu, dravid@uchicago.edu.

[†]We thank Amnon Amir, Ned Augenblick, Sylvain Chassang, Georgy Egorov, Navin Kartik, Pietro Ortolova, and Leeat Yariv for their comments and suggestions. Advait Ganapathy and Alexander Haberman provided excellent research assistance.

1 Introduction

A key policy tool for addressing an emerging pandemic is the creation of social distance. Although one can require the entire population to shelter in place, such a coarse policy bears enormous economic costs. An alternative solution is to use selective quarantines. Ideally, one would regularly test everyone, releasing individuals who test negative and quarantining those who test positive. Testing an entire country’s population, however, requires an unrealistically large testing capacity. One path toward relieving testing shortages is to use the same test to assess multiple individuals. Such group tests pool together samples collected from different people, returning a positive outcome whenever one of the group’s members is infected. The idea of testing a pooled sample goes back at least as far as Dorfman (1943), who suggested it as an efficient way to screen soldiers for syphilis. Since its introduction, pooled testing has been used in various domains, including the detection of HIV (e.g., Cahoon-Young et al., 1989; Kline et al., 1989; Behets et al., 1990; Archbold et al., 1991), avian flu (e.g., Arnold et al., 2013), and rare genes (e.g., Shental, Amir, and Zuk, 2010). Some recent studies have shown pooled testing can be effective in detecting a person infected with COVID-19 in groups of 32 (Yelin et al., 2020), 48 (Shental et al., 2020), or 57 (Theagarajan, 2020) individuals. Thus, group testing holds the promise of significantly increasing a state’s testing capacity. Our goal is to understand the optimal way to use group testing to aid a country’s quarantine policy.

When pooling samples from multiple individuals, one must choose not only whom to test, but also how to group the tested individuals. The reason is that the information produced by pooled tests depends jointly on our prior knowledge about all of the group’s members. For intuition, suppose we had a test that could perfectly detect whether the tested sample contains the pathogen. Applied to a single person, the test simply tells us whether the person is infected. Things are different when samples from two people, say, Ann and Bob, are pooled together. Whereas a negative result still indicates neither is infected, much less is known following a positive result: In this case, all we know is that *at least one* person is infected. We do not know whether both individuals are infected, or only one of them is—and, if only one, who. Which state is most likely depends on our prior knowledge about the tested individuals. Clearly, we are more likely to believe Bob is infected if he exhibits symptoms than if he does not. Moreover, knowing Bob is symptomatic also influences our inference about Ann, because the pooled test was likely to come back positive irrespective of whether she is infected. Hence, our prior information about one member of a pooled sample affects what we learn about all of the group’s members. As such, correctly dividing individuals into groups is crucial for our ability to benefit from pooled testing.

To illustrate our analysis and results, consider a planner who wants to maximize the expected number of individuals she releases back to the workforce, but will only release an individual who is known to be healthy. The planner has four individuals in her jurisdiction: two of high risk, and two of low risk. A high-risk individual is infected with probability 0.6, and a low-risk individual is infected with probability 0.1 (infection status being independent across people). The planner has two test kits, each of which can detect the presence of the virus without error. Without pooling samples, the best the planner can do is test the two

low-risk individuals, in which case she expects to release $2 \times 0.9 = 1.8$ individuals.

Suppose now that the planner can use each kit to test a pooled sample from up to two individuals, and each individual may be tested at most once. How should she allocate her two test kits across the four agents? Observe that not all pooling is beneficial for the planner. In particular, pooling together the low- and high-risk individuals is worse than single-sample testing: Because the probability that a low-risk and a high-risk person are simultaneously healthy is only $0.9 \times 0.4 = 0.36$, pooling each high-risk individual with a low-risk person releases only $4 \times 0.36 = 1.44$ people in expectation. It turns out that combining a low-risk individual's sample with the sample of a high-risk person is never optimal. In other words, we show the optimal test allocation involves *assortative batching*; that is, only individuals with the same infection risk are pooled. Observe that such assortative pooling does better than single-sample testing in the current example: By using one kit to test the two low-risk individuals, and the other kit on the two high-risk individuals, one releases $2 \times 0.9^2 + 2 \times 0.4^2 = 1.94$ people on average. Our analysis also shows the optimal allocation batches lower-risk individuals into larger groups than individuals with a higher probability of being infected. Consistent with this result, the optimal static test allocation in this example uses one kit to test a single high-risk individual, and one kit to test a pooled sample of the two low-risk people. In expectation, the planner releases $2 \times 0.9^2 + 0.4 = 2.02$ individuals.

We solve for the optimal allocation of potentially pooled tests in a simple static environment with a continuum of agents, each of whom is either infected or healthy. Health is independent across agents, with different agents being infected with different probabilities given their observable characteristics. Whereas each agent's infection probability is known, each individual's realized health state is not. To learn more about each agent's state, a planner has a small mass of tests at her disposal. Each test can be applied to a sample taken from up to K individuals, and each individual can be included in at most one test. A test returns a negative outcome if all of its samples come from healthy individuals. If at least one of the tested individuals is infected, the test returns a positive outcome. After tests are applied, the planner decides which individuals to quarantine and which to release. The planner gains a fixed benefit from releasing people to work, but suffers a larger per-person loss from releasing infected individuals. Quarantined individuals generate a normalized payoff of zero to the planner, whether infected or healthy.¹

Our analysis differs from the existing medical and statistical literature on group testing by focusing on quarantine decisions, which can be meaningfully informed by imperfect information. By contrast, most medical studies on group testing focus on the possibility of reducing the number of tests required for obtaining an individual diagnosis regarding each tested agent's infection status.² Whereas an individualized diagnosis is important for

¹Our paper does not address the problem of optimal test use for immediate diagnostic purposes. Therefore, the test capacity in our model might best be interpreted as the stock of tests remaining for the general population once hospitals' diagnostic needs are met.

²See, for example, Dorfman (1943), Hwang (1972), Cahoon-Young et al. (1989), Kline et al. (1989), Litvak, Tu, and Pagano (1994), Phatarfod and Sudbury (1994), Arnold et al. (2013), Bilder and Tebbis (2012), Liu et al. (2012), Aprahamian, Bish, and Bish (2019), Shental et al. (2020), and Sinnott-Armstrong, Klein, and Hickey (2020).

treatment, the added benefit of testing more subjects may be worth the added uncertainty regarding an individual’s infection status when deciding whether they should be placed in quarantine. The difference in focus justifies a difference in assumptions: In contrast to the existing medical literature, which requires multiple tests to be run on the same person, our analysis assumes each person can be tested at most once.³

As mentioned above, we show batching should be *assortative*; namely, only the samples of equally risky individuals should be batched together. For intuition, observe that batching creates a negative inference externality in the form of false positives: A person may get a positive result because she was batched with an infected individual despite being healthy. As such, the cost of batching Ann with Bob increases with the probability that she is healthy and Bob is infected. By matching individuals who cause large (resp. small) negative externalities with those who are less (more) sensitive to them, assortative batching minimizes the loss of information imposed by batched testing.

Given assortative batching, the planner’s problem reduces to making two choices for each level of infection risk: what fraction of agents of said infection risk should be tested, and what batch size to use for them. This reduction enables a complete solution of the optimal test allocation for a given population. We present one such solution in Figure 1, which illustrates the optimal testing policy for a specific example. The figure depicts both which agents are tested and how the optimal size of group tests varies with the tested agents’ respective infection probabilities.

We show that the planner tests only individuals whose optimal quarantine decisions are most uncertain. More precisely, testing only individuals whose infection risk lies in an interval around the planner’s quarantine indifference threshold is optimal. Observe this feature holds for the optimal policy depicted in Figure 1. Roughly speaking, the value of a test increases with the probability that the tested individual’s outcome differs from the prior, that is, the probability that an a priori quarantined (resp. released) individual is released (quarantined). Because this probability is higher the closer an individual’s infection risk is to the quarantine threshold, testing around this threshold maximizes the planner’s benefits from each test.

Figure 1 also shows that among tested individuals, group size declines with an agent’s probability of infection. We prove that this feature always holds under optimal testing. The reason is that the negative externality generated by batching increases with the probability of infection. Moreover, we show that for sufficiently high infection probabilities, the negative externalities from batching are so large that batching is never optimal. For example, the planner should never batch test individuals with a prior infection probability above 0.5. To see why, consider a planner who needs to allocate a single test among two quarantined persons, each of whom is infected with probability p . Whereas using the test on a single sample would release $1 - p$ people in expectation, the average number of people the planner releases by pooling the two samples is $2 * (1 - p)^2$. Thus, pooling the two samples together is beneficial if and only if $1 - p < 2 * (1 - p)^2$, that is, $p < 0.5$. Pursuing this logic to its nat-

³Consistent with this assumption, United States federal law distinguishes between screening and diagnostic tests, designating pooled tests for screening so long as they do not involve retesting (Mandavilli, 2020).




Figure 1: Optimal testing policy when pool tests of up to ten agents can be conducted; agents’ infection risks are uniformly distributed from 0% to 30%; the supply of tests is exactly enough to individually test 1 in 10 agents (hence, barely enough to test all agents with maximal grouping); and the planner’s preferences exhibit indifference to the quarantine decision for an agent of infection risk $\frac{1}{15} \approx 6.67\%$.

ural conclusion, we show that, in general, pooling k individuals whose individual infection probability is above $1/k$ is never optimal.

Within the economics literature, the closest work to our own is a recent collection of papers studying some form of optimal test allocation to inform quarantine decisions. In its focus on group testing, for instance, our work is related to that of Gollier (2020), Gollier and Gossner (2020), and Augenblick et al. (2020). Gollier (2020) shows that batching every tested agent into small group tests can outperform individualized testing for quarantine decisions and quantifies the benefits of such an improvement. Gollier and Gossner (2020) explore the benefits of group testing both for assessing population-level prevalence for a disease agent and for increasing the number of individuals who can be reliably assessed with scarce tests through adaptive testing. Taking a numerical approach, Augenblick et al. (2020) highlight a complementarity between frequent screening and group testing, and argue that machine learning can aid optimal group composition.

A few other recent studies look at optimal non-batched testing for informing quarantine decisions (Deb et al., 2020; Ely, Galeotti, and Steiner, 2020; Kasy and Teytelboym, 2020). Deb et al. (2020) study the joint design of testing and financial incentives for a binary self-quarantine decision that induces (through agents’ self-selection based on private information) coarse targeted testing. Kasy and Teytelboym (2020) study how costly perfect tests should be dynamically administered to both inform individuals’ quarantine decisions and learn about how disease prevalence covaries with observable characteristics. Ely, Galeotti,

and Steiner (2020) study how a finite budget of heterogeneously imperfect tests should be allocated to heterogeneous agents. Because different group sizes and compositions are differentially informative of a given individual’s health, the test heterogeneity in Ely, Galeotti, and Steiner (2020) is conceptually related to our model. Compared to their paper, our focus on group testing requires us to address two novel issues: (i) a lack of separability across jointly tested agents due to inference externalities; and (ii) a quantity-quality trade-off due to the information degradation that arises from testing larger groups.

Finally, our work is thematically related to the growing economics literature, building on the classic SIR model and its cousins, that incorporates testing into models of aggregate quarantine decisions (e.g., Rowthorn and Toxvaerd, 2012; Acemoglu et al., 2020; Alvarez, Argente, and Lippi, 2020; Berger, Herkenhoff, and Mongey, 2020; Brotherhood et al., 2020; Eichenbaum, Rebelo, and Trabandt, 2020; Piguillem and Shi, 2020).

2 Model

A planner must allocate scarce tests to a continuum population of heterogeneous agents who may be infected by some disease pathogen. Agents have risk types p distributed via μ , a measure over $(0, 1)$.⁴ An agent of type p has a probability p of being infected. For convenience, we assume μ is atomless. Our set of agents is “doubly infinite”: Every type p in the support of μ is held by a continuum mass of agents. Each agent’s type is common knowledge, and neither the planner nor the agents know which agents are infected. We can interpret an agent’s type as reflecting demographic characteristics, symptoms, recently visited countries, and so on. We assume no aggregate uncertainty exists, with the health of any finite set of agents being independent.

The planner has access to a mass $T \geq 0$ of tests. As described in our introduction, a given test can be administered to a batch of k distinct agents, for any $k \in K := \{1, \dots, K\}$, where $K \in \mathbb{N}$ is the test capacity (a technological constraint). We assume no agent can be tested more than once. Our baseline model also assumes away any imperfection in testing not derived from batching: If k people are batch tested, the test outcome is positive if *at least* one of them is infected, and is negative otherwise. Each test’s outcome is publicly observed. Thus, if k agents of types $p_1, \dots, p_k$ are batch tested, the posterior probability that a tested agent of type p_i is infected is 0 if the test is negative, and

$$q_{i,k}^+(p_1, \dots, p_k) := \frac{p_i}{1 - \prod_{j=1}^k (1 - p_j)} \quad (1)$$

if the test is positive. We call the probability in (1) the agent’s **positive posterior**. An untested agent has posterior probability equal to his type.

After allocating her tests, the planner observes the tests’ outcomes, and decides which agents (be they tested or untested) to quarantine and which to release. We normalize the planner’s payoff from quarantining an agent to zero. An unquarantined agent creates a benefit, for instance, an economic or social benefit, of $b > 0$. However, if that agent turns out

⁴Throughout, the term “measure” refers to a finite, positive, countably additive Borel measure.

to be infected, the planner additionally suffers a loss of $c > b$.⁵ Therefore, if an agent's posterior probability of illness is $q \in [0, 1]$, the planner derives an expected payoff of zero from quarantining him, and $b - cq$ otherwise. Consequently, the planner's optimal quarantine decision is simple: Quarantine an agent if and only if his posterior illness probability is above $\bar{q} := \frac{b}{c}$. Thus, an agent whose post-test infection probability is q gives the planner an expected payoff of

$$v(q) := \max\{b - cq, 0\} = (b - cq)\mathbf{1}_{q < \bar{q}}.$$

Consider now the planner's payoff from testing k individuals whose prior infection probabilities are $p_1, \dots, p_k$. With probability $\prod_{i=1}^k (1 - p_i)$, the test indicates all agents are healthy, and so the posterior probability any agent is infected is zero. In this case, the planner gets a payoff of $v(0)$ per agent. With complementary probability, the test is positive, in which case the planner's posterior for each tested agent is given by (1). Hence, the planner's total expected payoff from the test is

$$\sum_{i=1}^k \left\{ \left[\prod_{j=1}^k (1 - p_j) \right] v(0) + \left[1 - \prod_{j=1}^k (1 - p_j) \right] v \circ q_{i,k}^+(p_1, \dots, p_k) \right\}.$$

Subtracting the planner's payoff from these individuals had she conducted no test at all, $\sum_{i=1}^k v(p_i)$, gives the test's net benefit,

$$\beta_k(p_1, \dots, p_k) := \sum_{i=1}^k \left\{ \left[\prod_{j=1}^k (1 - p_j) \right] v(0) + \left[1 - \prod_{j=1}^k (1 - p_j) \right] v \circ q_{i,k}^+(p_1, \dots, p_k) - v(p_i) \right\}.$$

We now formulate our planner's problem. Let $\mathcal{T}$ be the set of all K -tuples $\vec{\tau} = (\tau_k)_{k=1}^K$, where each τ_k is a measure over $(0, 1)^k$. We interpret $\vec{\tau}$ as a testing strategy, with each τ_k describing the set of tests that are used to batch test a group of k distinct agents. Our planner's problem is

$$\begin{aligned} \max_{(\tau_k)_{k \in \mathcal{T}}} & \sum_{k=1}^K \int_{(0,1)^k} \beta_k(p_1, \dots, p_k) d\tau_k(p_1, \dots, p_k) \\ \text{s.t.} & \sum_{k=1}^K \tau_k((0, 1)^k) \leq T \end{aligned} \tag{S}$$

$$\sum_{k=1}^K \sum_{i=1}^k \text{marg}_i \tau_k \leq \mu. \tag{D}$$

Thus, the planner wishes to maximize the total expected net benefit from her tests, subject to two constraints. First, she must satisfy the *supply constraint*, (S), which says the number of tests used must be below the planner's testing capacity. Second, she must respect her

⁵Note our model is equivalent to one in which the benefit accrues only from unquarantined healthy agents, with the cost of an unquarantined infected agent being $c - b$.

demand constraint, (D), a measure inequality saying that the number of agents tested of any given type cannot surpass the total number of agents of that type in the population. In what follows, we describe the solution to the above problem.

3 The Optimal Testing Policy

In this section, we characterize the optimal testing policy. We begin by showing the planner tests only agents whose quarantine status depends on the test's outcome. We then show the optimal testing policy is assortative, meaning each test is applied to a homogeneous group of agents. Using this result, we replace the supply constraint with its implied shadow cost, which allows us to solve the planner's problem via pointwise maximization. The resulting solution exhibits three properties. First, the planner only tests agents whose type lies in an interval around the quarantine threshold. Second, no test contains more than $1/p$ agents of type p . And third, agents with higher infection risk are tested in smaller groups.

We begin by simplifying the planner's objective. Consider a batched test that tests k agents with risk types $p_1, \dots, p_k$. Say the test is **pivotal** for the i^{th} agent if the decision of whether to quarantine him is different for a positive versus a negative test. Because a negative test perfectly clears him, this condition is equivalent to the positive posterior $q_{i,k}^+(p_1, \dots, p_k)$ being at least $\bar{q}$. We now use pivotality to rewrite the direct benefit associated with the i^{th} agent's allocative decision,

$$\left[\prod_{j=1}^k (1 - p_j) \right] v(0) + \left[1 - \prod_{j=1}^k (1 - p_j) \right] v\left(\frac{p_i}{1 - \prod_{j=1}^k (1 - p_j)}\right) - v(p_i).$$

First, if the test is not pivotal for agent i , the test's result has no effect on behavior and hence no payoff consequence.⁶ Next, observe that whenever the test is pivotal for i , v takes value zero when evaluated at i 's positive posterior, and value b after a negative test. We can therefore rewrite the planner's problem as

$$\begin{aligned} \max_{(\tau_k)_{k \in \mathcal{T}}} & \sum_{k=1}^K \int_{(0,1)^k} \sum_{i=1}^k \mathbf{1}_{q_{i,k}^+(p_1, \dots, p_k) \geq \bar{q}} \left\{ b \left[\prod_{j=1}^k (1 - p_j) \right] - v(p_i) \right\} d\tau_k(p_1, \dots, p_k) \quad (\text{O}) \\ \text{s.t. } & (\text{S}), (\text{D}). \end{aligned}$$

We now state our first result regarding optimal testing: An agent should be tested only if his quarantine status depends on the test's outcome. The key to the proposition's proof is showing that adding agents for whom a test is not pivotal does not increase the test's value to the planner.

Proposition 1 (PIVOTALITY).

The planner optimally tests only agents for whom the test is pivotal.

⁶Mathematically, because the value function v is affine on an interval containing the agent's prior type, zero (which arises from a negative test), and the agent's positive posterior, this payoff term is exactly zero.

It is apparent from (O) that no direct benefit arises from non-pivotally testing an agent, and so any test that is not pivotal for any agent may as well not be run. To show one should remove agents from a test that is not pivotal for them, we further verify group testing only generates negative informational externalities. To do so, suppose a test is pivotal for agent i . Then, the associated value from testing him is $\left[\prod_{j=1}^k (1 - p_j)\right] b - v(p_i)$. Because both this value and agent i 's positive posterior increase when other agents are removed, reducing the batch size only increases the direct value from testing agent i .

Next, we establish the planner does not benefit from batching heterogeneous agents into one test.

Theorem 1 (ASSORTATIVE BATCHING).

The planner optimally tests agents of different types using different tests.

Intuitively, the risk of a positive test realization from a high-risk agent becomes less costly if said agent is batched with other high-risk agents.⁷ To understand the theorem's proof, imagine the planner intends to run k different tests, each of which batches an identical set of k heterogeneous agents.⁸ That is, each individual test is for a batch of non-identical agents, but the k -tuple of types $(p_1, \dots, p_k)$ is the same for each of the k tests. Without loss, say $p_1 \leq \dots \leq p_k$. Given Proposition 1, we may as well assume each of the k tests in question is pivotal for each of its k tested agents.

Let us now observe that the planner would benefit from instead batching agents assortatively, that is, from running k different tests of k agents in which no two different-type agents are tested together. We begin by explaining that the assortative test preserves pivotality. Consider first the agents of type p_1 . In the new test, each such agent is paired with $k - 1$ agents of a lower type than in his original test. Thus, the probability that agent p_1 receives a positive test result despite being healthy decreases. It follows that the assortative test increases p_1 's positive posterior, and so preserves pivotality.

Next, to show re-batching preserves test pivotality for agents of higher types, we define the homogeneous-updating rule q^+ taking each batch size, $k \in K$, and agent type $p \in (0, 1)$, to the agent's positive posterior from a test pooling k agents of type p together,

$$q^+(p, k) := q_{i,k}^+(p, \dots, p) = \frac{p}{1 - (1-p)^k}.$$

In the appendix, we show q^+ is increasing in its first argument. Therefore, knowing the assortative test is pivotal for agent p_1 implies the test is pivotal for all other agents under consideration.

⁷A similar force arises in problems of optimal team composition with complementarities (Prat, 2002). As our appendix shows, applying insights from that literature helps establish that assortative batching remains optimal even when testing is imperfect. See section 4

⁸Our double-continuum population assumption guarantees that if a given test of k agents is run, the planner without loss runs continuum many such identical tests. The appropriate assortativity notion in a finite-agent setting, in which each agent could potentially have a distinct infection probability, would require that only agents of *consecutive* types are jointly tested. Studying Dorfman testing (see section 4), Aprahamian, Bish, and Bish (2019) show that such consecutive batching is optimal with finitely many agents.

We now argue the above re-batching raises the principal's objective. By adding up the direct net payoff effects from testing each of the k^2 agents in (O), one can see this payoff ranking follows if re-batching the agents raises the average probability of testing negative, that is, if

$$\sum_{i=1}^k \frac{1}{k} (1 - p_i)^k > \prod_{i=1}^k (1 - p_i).$$

Simple algebraic manipulation shows the above inequality is equivalent to the arithmetic-mean-geometric-mean inequality (see the Appendix).

Next, we apply Theorem 1 to further simplify the planner's objective. To that end, define $\mathcal{T}^A$ to be the set of all K -tuples $\vec{\tau} = (\tau_k)_{k=1}^K$, where each τ_k is a measure over $(0, 1)$. In a mild abuse of notation, we can view each $\vec{\tau} \in \mathcal{T}^A$ as an element of $\mathcal{T}$ by identifying each τ_k with the unique measure over $\{\vec{p} \in (0, 1)^k : p_1 = \dots = p_k\}$ with marginals equal to τ_k . Proposition 1 tells us the planner optimally restricts attention to testing policies under which each test is pivotal for every tested agent, whereas Theorem 1 says these tests optimally batch only homogeneous agents. Hence, adding a pivotality constraint for each tested agent and specializing the objective (O), supply constraint (S), and demand constraint (D) to testing strategies that respect those two results, the planner problem can be simplified to

$$\max_{(\tau_k)_{k \in K} \in \mathcal{T}^A} \sum_{k=1}^K \int_{(0,1)} k \left[(1-p)^k b - v(p) \right] d\tau_k(p) \quad (\text{O}^A)$$

$$\text{s.t. } \tau_k\{q^+(\cdot, k) < \bar{q}\} = 0 \quad \forall k \in K \quad (\text{P}^A)$$

$$\sum_{k=1}^K \tau_k(0, 1) \leq T \quad (\text{S}^A)$$

$$\sum_{k=1}^K k \tau_k \leq \mu. \quad (\text{D}^A)$$

We now replace the global supply constraint with the implied shadow cost of exhausting testing resources. In principle, even with tests having a fixed marginal cost rather than an aggregate constraint, our planner's problem would remain non-separable across agents due the informational externality agents exert on each other in a group test. However, building on Theorem 1, our next result shows that, nevertheless, our planner's problem can be solved individually for each *type* of agent. To state the result, define for each agent type $p \in (0, 1)$ and shadow cost $\gamma \geq 0$ the associated **marginal testing value**,

$$\varphi(p, k, \gamma) := (1 - p)^k b - v(p) - \frac{\gamma}{k} \text{ for each } k \in K,$$

the **optimal batch sizes**,

$$k^*(p, \gamma) := \operatorname{argmax}_{k \in K} \varphi(p, k, \gamma) \text{ s.t. } q^+(p, k) \geq \bar{q},$$

and the **testing index**,

$$\varphi^*(p, \gamma) := \max_{k \in K} \varphi(p, k, \gamma) \text{ s.t. } q^+(p, k) \geq \bar{q}.$$

The following theorem shows how, up to calibrating one global cost parameter, the planner can make efficient decisions—both about whom to test and about how much to engage in batched testing—separably across agent types.

Theorem 2 (INDEXABILITY).

Feasible testing policy $\vec{\tau} \in \mathcal{T}^A$ is optimal if and only if it admits a shadow cost $\gamma \geq 0$ such that

1. *No agent of type p with $\varphi^*(p, \gamma) < 0$ is tested,*
2. *Every agent of type p with $\varphi^*(p, \gamma) > 0$ is tested in a homogeneous k -batch for some $k \in k^*(p, \gamma)$, and*
3. *Either (S^A) holds with equality or $\gamma = 0$.*

Moreover, a unique optimal testing policy exists.

The preceding results make proving Theorem 2 straightforward. First, given the demand constraint (D^A) , one can view the planner as deciding, for each type p and each batch size k , the fraction of p -type agents who will be tested in a batched test of k (identical) agents. Given such a formulation, observe that the objective (O^A) , demand constraint (D^A) , and pivotality constraint (P^A) are all separable across agents' types—even though the benefits accrued from a batched test entail externalities. Hence, replacing the constraint (S^A) with an augmented objective, a standard Lagrangian result (specifically, the convex multiplier rule; see Pourciau, 1980, 1983) implies the planner may optimize the modified objective pointwise across agent types.

Next, we turn to proving some qualitative features of the form of optimal testing—concerning both who is tested and the degree to which they are optimally batched. First, we show the set of tested agents is an interval around the quarantine threshold.

Proposition 2 (THRESHOLD TESTING).

In the optimal testing policy, the tested agents are those in some interval including $\bar{q}$.

Because the shadow cost of a test is the same for all agents, intuitively, tests should be used on those for whom they provide the most instrumental value. Under individual testing (i.e., without batching), this property is established by both Ely, Galeotti, and Steiner (2020) and Kasy and Teytelboym (2020). As Kasy and Teytelboym (2020) observe, the value of testing is maximized at the quarantine cutoff and is concave in an individually tested agent's type. This concavity, however, does not extend to batches of sizes larger than 1, because $(1 - p)^k$ is strictly convex in p whenever $k > 1$. To prove Proposition 2, we show that the marginal testing value is nonetheless strictly quasiconcave in the range where it is positive, and so the set of types whose testing value is above the shadow cost forms an interval.

Next, we derive an upper bound on a tested agent type's group size, which holds irrespective of the capacity of tests.

Proposition 3 (LIMITED BATCHING).

In the optimal testing policy, no tested agent of type p is tested in a k -batch with $k \geq 1/p$.

We prove this proposition by considering the planner’s benefits from decreasing a test’s batch size. These benefits are most transparent for individuals whose type lies above the planner’s optimal quarantine threshold, $p > b/c$. Reducing the batch size for such agents from k to $k - 1$ changes the expected number of released individuals by $(1 - p)^{k-1}(kp - 1)$, and so benefits the planner whenever $k > 1/p$. In the appendix, we show reducing the batch size helps the planner even more for agents whose type lies below the threshold, because such agents generate a release benefit that outweighs its cost even if left untested.

We conclude this section by showing the degree of batching is always decreasing in an agent’s type. That is, agents of lower infection risk are tested in larger batches.

Proposition 4 (DECREASING BATCH SIZES).

In the optimal testing policy, the test batch size is decreasing in the type of tested agents.

The proof of the above proposition shows that, for any given shadow cost, the optimal batch size decreases in an agent’s type. Such a structure would follow immediately from Theorem 2 if the marginal testing value φ were submodular in agent type and batch size, and the unconstrained-optimal batch size always happened to satisfy the pivotality constraint. However, neither of these two properties holds. Still, we show they hold in the relevant range, that is, whenever the marginal testing value is positive and the batching limits of Proposition 3 are respected.

Contrasting Proposition 4 with the intuition behind Proposition 2 is worthwhile. Proposition 2’s intuition is that costly information resources should be deployed for the types of agents whose best allocative decision is most uncertain, that is, those around the threshold. However, Proposition 4 advocates decreased batching—in particular, enhanced information—as one moves to the right of the quarantine threshold. This apparent contradiction arises because the two propositions focus on different marginal trade-offs. At the margin, Proposition 2 asks which agents we should test at a given batch size if an additional test is made available. The answer is the currently-untested agent whose quarantine decision is most likely to change due to the test. By contrast, Proposition 4 asks which types should be tested in larger batches if we must do so to free up an additional test. The answer is the type for which additional batching would cause the smallest information externality—that is, those who are most likely to be healthy.

4 Discussion of Modeling Assumptions

In this section, we discuss some of our modeling assumptions and briefly describe some possible extensions. Formal supporting arguments for any nontrivial claims are in Appendix C.

Imperfect tests. Our analysis assumes tests are perfect; that is, they always reveal whether the pooled sample contains an infected individual. In practice, tests for infection are often imperfect, exhibiting both false positives, indicating an infection when there is none, and false negatives, failing to detect an infection even when it is present.

We extend our model to incorporate imperfect tests in Appendix section C. There, we allow for errors to occur both at the analysis stage, where the pooled sample may be wrongly labeled as positive or negative, and at the sampling stage, with each person’s sample failing to represent the individual’s actual infection state. We show such errors do not affect the optimality of assortative batching. The economic intuition is the same as in the baseline model: Pooling samples from individuals with a low likelihood of infection minimizes the informational externality due to group testing. Formally, this intuition corresponds to the planner’s total benefit from a batched test being a symmetric and supermodular function of the tested individuals’ infection probabilities (see Prat, 2002). Knowing that batching is assortative, one can again replace the test-supply constraint with a shadow cost to obtain an indexability result similar to Theorem 2. However, the resulting index when tests are imperfect is substantially more complicated than the one introduced in Theorem 2. Still, one can show Propositions 2 and 4 continue to hold whenever tests exhibit only false positives.

Richer heterogeneity. In our model, the sole source of heterogeneity between agents is their probability of infection. Another way in which agents may differ from each other is in the benefits and costs associated with being released from quarantine. Specifically, suppose each agent i has his own idiosyncratic infection probability p_i , generates a benefit of b_i if released, and costs the planner c_i if released while infected. If the benefit varies inversely with infection risk, that is, $(p_i - p_j)(b_i - b_j) \leq 0$ for any two agents i and j , assortatively re-batching group tests continues to yield a payoff improvement if it preserves pivotality. Intuitively, such a relationship can only increase the benefits of assortativity, which pools agents whose release benefit is highest with agents who are the least likely to produce a release-impeding positive test result.⁹ As such, Theorem 1 remains intact as long as agents’ optimal threshold $\bar{q} = \frac{b}{c}$ is also decreasing in agents’ infection risk, that is, $(p_i - p_j)\left(\frac{b_i}{c_i} - \frac{b_j}{c_j}\right) \leq 0$ for any two agents i and j , which, for instance, holds if costs are homogeneous or proportional to release benefits. In this case, a set of assortatively batched tests of a given size is pivotal for all agents if and only if the test administered to the lowest-risk agents is. It follows that one can extend Theorem 1 to this more general setting whenever optimal quarantine thresholds and release benefits are appropriately monotone with infection probability.

Even supposing the relationships between sources of heterogeneity described above, the exact structure of an optimal testing policy can change substantially. For example, Proposition 4 does not extend: A higher infection probability may no longer be associated with a smaller batch size. For example, in the extreme special case in which b is heterogeneous and p is not, it is nearly immediate that the marginal testing value φ is submodular in these benefits and the batch size. As such, with enough heterogeneity in the release benefit and little enough heterogeneity in infection risk, the order described in Proposition 4 reverses.

⁹Consistent with this intuition, assortative batching is typically not optimal when high-risk types are the ones who generate a high benefit when released.

General quarantine thresholds. We assumed the planner decides whom to quarantine and whom to release. But noncompliance of individual agents (in either direction), political frictions in policymaking, or other factors may leave agents’ quarantine decision out of the hands of our planner. A simple way to capture such a friction in the model is to make $\bar{q} \in (0, 1)$ an exogenous parameter, which may differ from the planner’s ideal threshold of $\frac{b}{c}$. Many of our results remain intact in this more general environment. Specifically, Proposition 1 and Theorem 1 go through without change, because their proofs do not depend on the quarantine threshold. These results, in turn, imply Theorem 2’s proof remains valid, too. Threshold testing (Proposition 2) also extends—even though our planner does not always prefer higher-quality information when the threshold is suboptimal.

In contrast to our other results, whether Propositions 3 and 4 remain true depends on the quarantine threshold. Whenever the threshold lies below the planner’s optimum (i.e., the case of excessive quarantining), both propositions go through without change, with their proofs adapting nearly immediately. If, however, the quarantine threshold is too lax, the two propositions may not hold on the interval of types just below the threshold—those whom the planner would like to quarantine but, for exogenous reasons, must release unless they receive a pivotal test. Because the change in these individuals’ quarantine status when they test positive in a pivotal test generates a discrete benefit to the planner, she may prefer to saturate the pivotality constraint. For instance, consider the following case of an extremely lenient threshold: Every type of agent is below the fixed quarantine threshold, and the cost of releasing infected individuals is high enough that the optimal threshold is nearly zero. When the cost is sufficiently high and tests are sufficiently scarce, each tested type of agent is tested according to the maximum batch size feasible under pivotality, reversing the batch-size order suggested by Proposition 4.

Comparison to Dorfman testing. In our analysis, we made a many-to-one assumption: Each agent’s sample can be included in at most one test. This assumption differentiates us from existing pooled-testing methods, which rely on running multiple tests on the same individual. The most prominent¹⁰ among the existing methodologies is the testing policy suggested by Dorfman (1943). According to Dorfman’s (1943) method, each agent in a given group test is retested if the original test is positive. We now compare our many-to-one solution to the optimal implementation of Dorfman’s (1943) method within our setting.¹¹

We begin by noting that, unlike our many-to-one testing policy, the pivotality considerations expressed in our Proposition 1 are not a concern under Dorfman’s (1943) regime. The reason is that under Dorfman’s (1943) regime, every tested agent eventually receives an error-free result regarding their infection status: Either the agent’s initial test came back negative, or the agent is retested using an individualized test. Since we assume tests are error-free, it follows that Dorfman’s (1943) testing method reveals each tested agent’s true

¹⁰This method has been applied in a variety of environments, see for example Cahoon-Young et al. (1989), Strarner, Caglioti, and Strong (2000), Dodd, Notari IV, and Stramer (2002), Gaydos (2005), and U.S. Food and Drug Administration (2020).

¹¹Aprahamian, Bish, and Bish (2019) analyze the optimal way to implement Dorfman’s (1943) testing policy within a finite population.

infection status, regardless of the prior probability that agent is infected.

It turns out the optimal implementation of Dorfman testing shares the assortative batching feature of our optimal many-to-one test allocation. The reason is that the same procedure that minimizes the negative inference externality under many-to-one testing also increases the efficiency of pooling tests according to Dorfman’s (1943) method. While we show this result formally in the appendix, we now provide an intuitive explanation. Recall the modification behind Theorem 1, which resorts agents from non-assortative tests of the same batch size into assortative tests of equal size in a way that lowers the average probability of obtaining a positive group test result, while holding fixed the population whose sample is tested in a given batch size, as well as the number of initial tests conducted for each batch size. Under Dorfman’s (1943) regime, this modification reduces the expected number of pooled tests of a given size that result in retesting. Moreover, the outcome distribution for each individual remains unchanged, because Dorfman’s (1943) policy precisely determines each tested agent’s infection status. It follows this modification improves upon any non-assortative test allocation, and so the optimal policy under Dorfman’s (1943) regime must be assortative.

Once we know that assortative batching is optimal, we can apply a multiplier argument similar to Theorem 2 to obtain a problem that can be solved separably across agents of different risk types. The resulting problem turns out to share the same single-peaked structure that leads to our interval testing result (Proposition 2). One can also show that the optimal implementation of Dorfman’s (1943) method exhibits batch sizes that decrease in the tested-agent’s infection risk, similar to our Proposition 4, though the argument is different. We refer the reader to the appendix for details.

Unlike optimal many-to-one policy, the optimal implementation of Dorfman’s (1943) method yields a single optimal test size for each risk category. Specifically, conditional on being tested, each agent’s optimal batch size depends only on his infection-risk. For an explanation, recall that under Dorfman’s (1943) algorithm, the composition of an agent’s pool only impacts the expected number of tests required to precisely diagnose the agent’s infection status. Since batching is assortative, this number is a function of only two variables: the probability the agent is infected, and the size of the pool. Therefore, under Dorfman’s (1943) method, it is always optimal to use the pool size that minimizes the expected number of tests required for diagnosing each risk type whenever that type is tested. By contrast, in many-to-one testing, the optimal pool size for each agent resolves a quantity-quality trade-off, which depends on all of the model’s parameters: the supply of tests (T), the benefit from released agents (b), the loss from releasing sick agents (c), and the virus’ spread throughout the population (μ). Thus, whereas the optimal implementation of Dorfman’s (1943) regime is qualitatively similar to optimal many-to-one testing, the two deliver different policy prescriptions, which differ in the way they respond to the underlying environment.

References

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston. 2020. “A multi-risk sir model with optimally targeted lockdown.” Tech. rep., National Bureau of

Economic Research.

- Alvarez, Fernando E, David Argente, and Francesco Lippi. 2020. “A simple planning problem for covid-19 lockdown and smart tracing.” Tech. rep., National Bureau of Economic Research.
- Aprahamian, Hrayr, Douglas R Bish, and Ebru K Bish. 2019. “Optimal risk-based group testing.” *Management Science* 65 (9):4365–4384.
- Archbold, E, S Mitchel, D Hanson, B Hull, LM Galo, and O Monzon. 1991. “Serum-Pooling Strategies for HIV Screening: Experiences in Trinidad, Ecuador, and the Phillippines.” In *VII International Conference on AIDS Abstract Book*, vol. 2. 331.
- Arnold, ME, MJ Slomka, VJ Coward, S Mahmood, PJ Raleigh, and IH Brown. 2013. “Evaluation of the Pooling of Swabs for Real-Time PCR Detection of Low Titre Shedding of Low Pathogenicity Avian Influenza in Turkeys.” *Epidemiology & Infection* 141 (6):1286–1297.
- Augenblick, Ned, Jonathan Kolstad, Ziad Obermeyer, and Ao Wang. 2020. “Group Testing in a Pandemic: The Role of Frequent Testing, Correlated Risk, and Machine Learning.” *Working paper* .
- Behets, Frieda, Stefano Bertozzi, Mwamba Kasali, Mwandagalirwa Kashamuka, L Atikala, Christopher Brown, Robert W Ryder, and Thomas C Quinn. 1990. “Successful Use of Pooled Sera to Determine HIV-1 Seroprevalence in Zaire with Development of Cost-Efficiency Models.” *AIDS (London, England)* 4 (8):737–741.
- Berger, David W, Kyle F Herkenhoff, and Simon Mongey. 2020. “An seir infectious disease model with testing and conditional quarantine.” Tech. rep., National Bureau of Economic Research.
- Bilder, Christopher R and Joshua M Tebbs. 2012. “Pooled-Testing Procedures for Screening High Volume Clinical Specimens in Heterogeneous Populations.” *Statistics in medicine* 31 (27):3261–3268.
- Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2020. “An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies.” *Working paper* .
- Cahoon-Young, B, A Chandler, T Livermore, J Gaudino, and R Benjamin. 1989. “Sensitivity and specificity of pooled versus individual sera in a human immunodeficiency virus antibody prevalence study.” *Journal of Clinical Microbiology* 27 (8):1893–1895.
- Deb, Rahul, Mallesh Pai, Akhil Vohra, and Rakesh Vohra. 2020. “Testing Alone Is Insufficient.” *Available at SSRN 3593974* .

- Dodd, RY, EP Notari IV, and SL Stramer. 2002. “Current Prevalence and Incidence of Infectious Disease Markers and Estimated Window-Period Risk in the American Red Cross Blood Donor Population.” *Transfusion* 42 (8):975–979.
- Dorfman, Robert. 1943. “The detection of defective members of large populations.” *The Annals of Mathematical Statistics* 14 (4):436–440.
- Eichenbaum, Martin S, Sérgio Rebelo, and Mathias Trabandt. 2020. “The Macroeconomics of Testing and Quarantining.” Tech. rep., National Bureau of Economic Research.
- Ely, Jeffrey, Andrea Galeotti, and Jakub Steiner. 2020. “Optimal Test Allocation.” *Working paper*.
- Galichon, Alfred. 2018. *Optimal transport methods in economics*. Princeton University Press.
- Gaydos, Charlotte A. 2005. “Nucleic Acid Amplification Tests for Gonorrhea and Chlamydia: Practice and Applications.” *Infectious Disease Clinics* 19 (2):367–386.
- Gollier, Christian. 2020. “Optimal group testing to exit the covid confinement.” Tech. rep., Technical report, Toulouse School of Economics.
- Gollier, Christian and Olivier Gossner. 2020. “Group Testing Against COVID-19.” *Covid Economics* 1 (2):32–42.
- Hwang, FK. 1972. “A Method for Detecting All Defective Members in a Population by Group Testing.” *Journal of the American Statistical Association* 67 (339):605–608.
- Kasy, Maximilian and Alexander Teytelboym. 2020. “Adaptive Targeted Infectious Disease Testing.” *Oxford Review of Economic Policy* URL <https://doi.org/10.1093/oxrep/graa018>. Graa018.
- Kline, R L, T A Brothers, R Brookmeyer, S Zeger, and T C Quinn. 1989. “Evaluation of human immunodeficiency virus seroprevalence in population surveys using pooled sera.” *Journal of Clinical Microbiology* 27 (7):1449–1452.
- Litvak, Eugene, Xin M Tu, and Marcello Pagano. 1994. “Screening for the Presence of a Disease by Pooling Sera Samples.” *Journal of the American Statistical Association* 89 (426):424–434.
- Liu, Aiyi, Chunling Liu, Zhiwei Zhang, and Paul S Albert. 2012. “Optimality of Group Testing in the Presence of Misclassification.” *Biometrika* 99 (1):245–251.
- Mandavilli, Apoorva. 2020. “Federal Officials Turn to a New Testing Strategy as Infections Surge.” *New York Times*.
- Phatarfod, RM and Aidan Sudbury. 1994. “The Use of a Square Array Scheme in Blood Testing.” *Statistics in Medicine* 13 (22):2337–2343.

- Piguillem, Facundo and Liyan Shi. 2020. “Optimal COVID-19 quarantine and testing policies.” *Working paper* .
- Pourciau, BH. 1980. “Modern multiplier rules.” *The American Mathematical Monthly* 87 (6):433–452.
- . 1983. “Multiplier rules and the separation of convex sets.” *Journal of Optimization Theory and Applications* 40 (3):321–331.
- Prat, Andrea. 2002. “Should a team be homogeneous?” *European Economic Review* 46 (7):1187–1207.
- Rowthorn, Bob RE and Flavio Toxvaerd. 2012. “The optimal control of infectious diseases via prevention and treatment.” *Working paper* .
- Shental, Noam, Amnon Amir, and Or Zuk. 2010. “Identification of rare alleles and their carriers using compressed sequencing.” *Nucleic acids research* 38 (19):e179–e179.
- Shental, Noam, Shlomia Levy, Shosh Skorniakov, Vered Wuvshet, Yonat Shemer-Avni, Angel Porgador, and Tomer Hertz. 2020. “Efficient high throughput SARS-CoV-2 testing to detect asymptomatic carriers.” *medRxiv* .
- Sinnott-Armstrong, Nasa, Daniel Klein, and Brendan Hickey. 2020. “Evaluation of Group Testing for SARS-CoV-2 RNA.” *medRxiv* .
- Strarner, Susan L, Sally Caglioti, and DM Strong. 2000. “NAT of the United States and Canadian Blood Supply.” *Transfusion* 20040:1165–1168.
- Theagarajan, Lakshmi N. 2020. “Group Testing for COVID-19: How to Stop Worrying and Test More.” *arXiv preprint arXiv:2004.06306* .
- U.S. Food and Drug Administration. 2020. “Coronavirus (COVID-19) Update: FDA Issues First Emergency Authorization for Sample Pooling in Diagnostic Testing 2020.” URL <https://www.fda.gov/news-events/press-announcements/coronavirus-covid-19-update-fda-issues-first-emergency-authorization-sa> Accessed on October 28th, 2020.
- Yelin, Idan, Noga Aharoni, Einat Shaer-Tamar, Amir Argoetti, Esther Messer, Dina Berenbaum, Einat Shafran, Areen Kuzli, Nagam Gandali, Tamar Hashimshony et al. 2020. “Evaluation of COVID-19 RT-qPCR test in multi-sample pools.” *medRxiv* .

A Appendix: Proofs of Main Results

A.1 Proof of Proposition 1

We begin by formalizing the sense in which removing agents from a batched test improves the quality of information for the remaining agents.

Lemma 1 (INFORMATIVE TESTING).

Removing some agents from a group test strictly increases

1. *Each remaining agent's positive posterior from the test, and*
2. *The probability of a negative test result.*

Proof. Say the agents' types in a k -batched test are $p_1, \dots, p_k \in (0, 1)$. The probability of a negative test is $N := \prod_{j=1}^k (1 - p_j)$, which strictly increases when some agents are removed from the test. The positive posterior of the i^{th} agent is $\frac{p_i}{1-N}$, which strictly increases when N increases. $\square$

Although formalizing it entails some notational cost, the proof of the proposition is straightforward given the above lemma.

Proof of Proposition 1. Given a testing policy $\vec{\tau} \in \mathcal{T}$, we construct an alternative testing policy that yields a higher objective (O) and only pivotally tests agents. For each $k \in K$ and $J \subseteq \{1, \dots, k\}$, let $\pi_{k,J} : \mathbb{R}^k \rightarrow \mathbb{R}^{|J|}$ be the projection onto the J -coordinates in order, and let $\Theta_{k,J}$ be the set of all $\vec{p} = (p_i)_{i=1}^k \in (0, 1)^k$ such that the agents who find a batched test with type profile $\vec{p}$ pivotal are those in J , that is, such that

$$\left\{ i \in \{1, \dots, k\} : \frac{p_i}{1 - \prod_{j=1}^k (1 - p_j)} \geq \bar{q} \right\} = J.$$

We now define the new testing policy that arises from taking τ and removing all agents whose respective tests are not pivotal for them. Formally, let

$$\vec{\tau}^* := \left(\sum_{k=\hat{k}}^K \sum_{J \subseteq \{1, \dots, k\}: |J|=\hat{k}} \tau_k [(\cdot) \cap \Theta_{k,J}] \circ \pi_{k,J}^{-1} \right)_{\hat{k}=1}^K$$

be the alternative testing policy. It is immediate from the construction that $\vec{\tau}^* \in \mathcal{T}$ and that $\vec{\tau}^*$ satisfies the supply (S) and demand (D) constraints because τ does. Moreover, given the first part of Lemma 1, for every $\hat{k} \in \{1, \dots, K\}$, the measure $\tau_{\hat{k}}^*$ over $\hat{k}$ -batched tests comprises only tests that are pivotal for every tested agent. Finally, the second part of Lemma 1 therefore implies $\vec{\tau}^*$ yields a strictly higher objective (O) than $\vec{\tau}$ if any positive measure of tests under τ was pivotal for a proper nonempty subset of the tested agents. $\square$

A.2 Proof of Theorem 1

We begin by making some observations about the comparative statics of belief updating from positive results of batched tests.

Lemma 2 (COMPARATIVE STATICS OF UPDATING).

An agent's positive posterior from a test is:

1. *Strictly decreasing in the types of the other agents batched into the same test.*
2. *Strictly decreasing in k and strictly increasing in p , if the test is a homogeneous test of $k > 1$ different agents of type p .*

Proof. The first result is immediate from the formula for a positive posterior. It is also immediate that $q^+(p, k)$ is strictly decreasing in k . All that remains, then, is to show $q^+(p, k)$ strictly increases with p when $k > 1$. But this feature follows readily from

$$q^+(p, k) = \frac{1-(1-p)}{1-(1-p)^k} = \frac{1}{\sum_{\ell=0}^{k-1} (1-p)^\ell}.$$

□

Let us briefly comment on the one nontrivial piece of the above lemma. Consider two different hypothetical batch tests, each with the same number of batched agents, each homogeneous, one testing high-risk agents and the other testing low-risk agents. Ranking the associated positive posterior from the two tests involves a trade-off. On one hand, the higher-risk agents have a higher prior probability of being infected, and thus will have a higher posterior probability of illness given any outcome in any fixed experiment about one's own illness. On the other hand, being batched with higher-risk agents leads to a different experiment about one's own illness, with a higher incidence of false positives, so that a positive test should be less suggestive of illness. The lemma establishes that the former effect always dominates.

Next, we record a useful computation concerning the probability of a negative outcome of a batched test. It says that, given k groups of k identical agents who can be sorted into k different batch tests of k agents, the expected number of negative tests is maximized by making each batched test homogeneous, rather than making the tests identically heterogeneous.

Lemma 3 (COMPARATIVE STATICS OF BATCHING).

Given $k \in K$ and a list of k types $p_1, \dots, p_k \in (0, 1)$ that are not all identical,

$$k^2 \prod_{i=1}^k (1 - p_i) < k \sum_{i=1}^k (1 - p_i)^k.$$

Proof. This inequality is the AM–GM inequality, applied to $(1 - p_1)^k, \dots, (1 - p_k)^k$. □

Assortative batching is conceptually straightforward to establish given its supporting lemmas.

Proof of Theorem 1. Given a testing policy $\vec{\tau} \in \mathcal{T}$, which we without loss take to be pivotal for every tested agent by Proposition 1, we construct an alternative testing policy with assortative batching that yields a higher objective (O). For each $k \in K$, let $\iota_k : \mathbb{R} \rightarrow \mathbb{R}^k$ be the diagonal inclusion, that is, the map taking $p \mapsto (p, \dots, p)$. Define now the new testing policy that arises from taking τ and splitting every k copies of the same k -batched test into k different homogeneous-batch tests that together test the same k^2 agents. Formally, let

$$\vec{\tau}^* := \left(\left[\frac{1}{k} \sum_{i=1}^k \text{marg}_i \tau_k \right] \circ \iota_k^{-1} \right)_{\hat{k}=1}^K$$

be the alternative testing policy. It is immediate from the construction that $\vec{\tau}^* \in \mathcal{T}$ and that $\vec{\tau}^*$ satisfies the supply (S) and demand (D) constraints because τ does.

Now, let us observe that the new testing policy still satisfies pivotality. To do so, it suffices to show that, for any $p_1, \dots, p_k \in (0, 1)$, we have $\frac{p_i}{1-(1-p_i)^k} \geq \bar{q}$ for every $i \in \{1, \dots, k\}$ whenever $\frac{p_i}{1-\prod_{j=1}^k (1-p_j)} \geq \bar{q}$ for every $i \in \{1, \dots, k\}$. But, letting $i^* \in \text{argmin}_{i \in \{1, \dots, k\}} p_i$, this feature follows from

$$\frac{p_{i^*}}{1-\prod_{j=1}^k (1-p_j)} \leq \frac{p_{i^*}}{1-(1-p_{i^*})^k} \leq \frac{p_i}{1-(1-p_i)^k} \quad \forall i \in \{1, \dots, k\},$$

where the two inequalities come directly from the two parts of Lemma 2.

Finally, having shown that every administered test under $\vec{\tau}^*$ is pivotal for every tested agent, Lemma 3 implies $\vec{\tau}^*$ yields a strictly higher objective (O) than $\vec{\tau}$ if $\vec{\tau}$ entails a positive measure of non-assortative tests. $\square$

A.3 Proof of Theorem 2

That optimality of a given testing rule can be expressed in terms of an endogenous shadow cost follows readily from a standard multiplier rule for convex-optimization problems. Existence and uniqueness then follow from direct arguments using this characterization.

Proof of Theorem 2. To see this result, recasting the planner's problem in slightly different language is convenient. For a given $\vec{\tau} \in \mathcal{T}^A$, the constraint (D^A) says that each of $\{\tau_k\}_{k=1}^K$ is absolutely continuous with respect to μ and that the scaled Radon–Nikodym derivatives $f_k := k \frac{d\tau_k}{d\mu}$ satisfy $\sum_{k=1}^K f_k \leq_{\text{a.e.}} 1$. The constraint (S^A) then says that $\sum_{k=1}^K \int \frac{f_k}{k} d\mu \leq T$, and the constraint (P^A) says that $[q^+(\cdot, k) - \bar{q}] f_k \geq_{\text{a.e.}} 0$ for every $k \in K$. Moreover, the objective (O^A) can be rewritten as $\sum_{k=1}^K \int \varphi(p, k, 0) f_k(p) d\mu(p)$. Therefore, defining the convex set

$$\mathcal{F} := \left\{ f \in [L^1(0, 1)]^K : f_k \geq 0 \text{ and } [q^+(\cdot, k) - \bar{q}] f_k \geq 0 \quad \forall k \in K, \text{ and } \sum_{k=1}^K f_k \leq 1 \right\},$$

we can equivalently express the planner's problem as

$$\begin{aligned} \max_{f \in \mathcal{F}} \quad & \sum_{k=1}^K \int \varphi(\cdot, k, 0) f_k(\cdot) \, d\mu \\ \text{s.t.} \quad & \sum_{k=1}^K \int \frac{f_k}{k} \, d\mu \leq T. \end{aligned}$$

Now, applying the convex multiplier rule (and noticing $0 \in \mathcal{F}$ satisfies $\sum_{k=1}^K \int \frac{0}{k} \, d\mu < T$), a given feasible $f^* \in \mathcal{F}$ is optimal if and only if some $\gamma \geq 0$ exists such that

- $f^* \in \operatorname{argmax}_{f \in \mathcal{F}} \left\{ \gamma T + \sum_{k=1}^K \int \varphi(\cdot, k, \gamma) f_k(\cdot) \, d\mu \right\};$
- $\gamma = 0$ if $\sum_{k=1}^K \int \frac{f_k^*}{k} \, d\mu < T$.

The optimality characterization then follows from noticing that, for $p \in (0, 1)$ and $k \in K$, the vector $w^* \in \mathbb{R}^K$ maximizes $\sum_{k=1}^K w_k \varphi(p, k, \gamma)$ over all $w \in \mathbb{R}_+^K$ with $([q^+(\cdot, k) - \bar{q}] w_k)_{k=1}^K \geq 0$ if and only if

- $w_k^* = 0$ for every $k \in K \setminus k^*(p, \gamma);$
- $w^* \geq 0$, with equality if $\varphi^*(p, \gamma) < 0;$
- $\sum_{k=1}^K w_k^* = 1$ if $\varphi^*(p, \gamma) > 0.$

Finally, we verify existence and uniqueness. If $T \geq \mu(0, 1)$, existence and uniqueness are trivial because testing every agent without batching uniquely (up to measure zero) achieves the first-best outcome, so suppose $T < \mu(0, 1)$.

As a starting observation, for any $\gamma \geq 0$, the nonzero piecewise-polynomial function $\varphi(\cdot, k, \gamma) - \varphi(\cdot, k', \gamma)$ can have at most finitely many roots for any two distinct $k, k' \in K$. Hence, at most finitely many types can have more than one optimal batch size for any given multiplier γ . In particular, for any given $\gamma \geq 0$ and Borel $f : (0, 1) \times K \rightarrow \mathbb{R}$, the number $\int f(p, k^*(p, \gamma)) \, d\mu(p)$ is well defined.

Now, for each type $p \in (0, 1)$, the function $\varphi(p, k, \cdot)$ is continuous and strictly decreasing for each $k \in K$ with $q^+(p, k) \geq \bar{q}$, and so the finite maximum $\varphi^*(p, \cdot)$ is continuous and strictly decreasing too. Therefore, the functions $\gamma \mapsto \int_{\{\varphi^*(\cdot, \gamma) \geq 0\}} \frac{1}{k^*(\cdot, \gamma)} \, d\mu$ and $\gamma \mapsto \int_{\{\varphi^*(\cdot, \gamma) > 0\}} \frac{1}{k^*(\cdot, \gamma)} \, d\mu$ are upper semicontinuous and lower semicontinuous, respectively, and both are nonincreasing. But non-atomicity of μ tells us the (continuous and piecewise strictly monotone) transformation $\varphi(\cdot, k, \gamma)$ has non-atomic μ -distribution too for each $k \in K$. It follows that $\varphi^*(\cdot, \gamma)$ is μ -non-atomically distributed. Therefore, the function

$$\begin{aligned} m : \mathbb{R}_+ &\rightarrow \mathbb{R} \\ \gamma &\mapsto \int_{\{\varphi^*(\cdot, \gamma) \geq 0\}} \frac{1}{k^*(\cdot, \gamma)} \, d\mu = \int_{\{\varphi^*(\cdot, \gamma) > 0\}} \frac{1}{k^*(\cdot, \gamma)} \, d\mu \end{aligned}$$

is continuous and nonincreasing. As $m(0) > T > 0 = \lim_{\gamma \rightarrow \infty} m(\gamma)$, the intermediate value theorem delivers some γ^* for which $m(\gamma) = T$. An optimal testing policy then exists in which exactly the types in $\{\varphi^*(\cdot, \gamma) > 0\}$ are tested, and each such type p is tested in a $k^*(p, \gamma^*)$ -batch.

Hence, all that remains is to verify uniqueness. To that end, fix some witnessing multiplier $\gamma^* \geq 0$ for which some testing satisfies the optimality conditions in the statement of the theorem. Observe that, in fact, *every* optimal testing rule satisfies said optimality conditions for γ^* . Indeed, satisfying the optimality conditions corresponds exactly to forming a Nash equilibrium of the two-player zero-sum game in which Maximizer chooses a testing policy in $\mathcal{F}$, Minimizer chooses a multiplier in $\mathbb{R}_+$, and the former's objective is $(f, \gamma) \mapsto \gamma T + \sum_{k=1}^K \int \varphi(\cdot, k, \gamma) f_k(\cdot) d\mu$. But because the set of Nash equilibria of any two-player zero-sum game is a product set of strategy profiles, every optimal testing rule is witnessed as such by γ^* . Uniqueness will then follow if we can show a unique optimal $f^* \in \operatorname{argmax}_{f \in \mathcal{F}} \left\{ \gamma^* T + \sum_{k=1}^K \int \varphi(\cdot, k, \gamma^*) f_k(\cdot) d\mu \right\}$ exists. But because $k^*(\cdot, \gamma^*)$ is almost everywhere single-valued (as noted above), the optimal batch size is unique for all but a null set of tested agents. Moreover, $\varphi(\cdot, k, \gamma^*)$ is a nonzero polynomial on either side of $\bar{q}$, so that the finite maximum $\varphi^*(\cdot, \gamma^*)$ has finitely many zeroes. The optimal set of agents tested is then also unique up to a null set. $\square$

A.4 Proof of Proposition 2

We now show that the set of types the planner tests is an interval around the threshold $\bar{q}$. The proof shows that, among types of at least as high risk as the lowest-risk tested types, the marginal testing value is single-peaked in one's type, with its peak occurring at the threshold. Consequently, the result follows from Theorem 2.

The following lemma shows such single-peakedness applies separately to each batch size for the range of types where that batch size is worth employing.

Lemma 4 (SINGLE-PEAKED INDEX AS TYPE VARIES).

For each $\gamma \geq 0$ and $k \in K$, the set $I_k := \{\varphi(\cdot, k, \gamma) \geq 0\}$ is an interval that has $\bar{q}$ in its closure if it is nonempty, and $\varphi(\cdot, k, \gamma)$ is strictly increasing on $I_k \cap (0, \bar{q}]$ and strictly decreasing on $I_k \cap (\bar{q}, 1)$.

Proof. First, for $p \in [\bar{q}, 1)$, we have $\varphi(p, k, \gamma) = b(1 - p)^k - \frac{\gamma}{k}$, which strictly decreases with p . Therefore, $I_k \cap [\bar{q}, 1)$ is convex, it contains $\bar{q}$ if it is nonempty, and $\varphi(\cdot, k, \gamma)$ is strictly decreasing there.

Next, consider any $p \in (0, \bar{q}]$ such that $\varphi(p, k, \gamma) \geq 0$. Let us show that $\frac{\partial}{\partial p} \varphi(p, k, \gamma) > 0$, which will deliver the lemma. To that end, observe that every $\tilde{p} \in (0, \bar{q}]$ has

$$\begin{aligned} \varphi(\tilde{p}, k, \gamma) &= b(1 - \tilde{p})^k - (b - c\tilde{p}) - \frac{\gamma}{k} \\ &= b \left[1 - \frac{\tilde{p}}{q^+(p, k)} \right] - (b - c\tilde{p}) - \frac{\gamma}{k} \\ &= \tilde{p} \left[c - \frac{b}{q^+(\tilde{p}, k)} \right] - \frac{\gamma}{k}. \end{aligned}$$

In particular,

$$\varphi(p, k, \gamma) \geq 0 \implies c - \frac{b}{q^+(p, k)} \geq \frac{\gamma}{kp} \geq 0,$$

and the second part of Lemma 2 implies $c - \frac{b}{q^+(\cdot, k)}$ is strictly increasing too. Therefore,

$$\frac{\partial}{\partial p} \varphi(p, k, \gamma) = 1 \left[c - \frac{b}{q^+(p, k)} \right] + p \frac{\partial}{\partial p} \left[c - \frac{b}{q^+(p, k)} \right] > 0,$$

as required. $\square$

Now, we show the above lemma readily delivers interval testing.

Proof of Proposition 2. Let us consider an optimal testing policy, and a witnessing multiplier γ as delivered by Theorem 2. Consider any $p \in (0, 1)$ such that $\varphi^*(p, \gamma) \geq 0$. Then some $k \in K$ exists such that $q^+(p, k) \geq \bar{q}$ and $\varphi(p, k, \gamma) = \varphi^*(p, \gamma)$. Consider any p' such that $p < p' \leq \bar{q}$ or $\bar{q} < p' < p$. First, observe that $q^+(p', k) \geq \bar{q}$: This inequality is immediate from $q^+(p', k) \geq p'$ if $p \geq \bar{q}$, and from the second part of Lemma 2 if $p' > p$. Therefore, $\varphi^*(p', \gamma) \geq \varphi(p', k, \gamma)$. But Lemma 4 then says $\varphi(p', k, \gamma) > \varphi(p, k, \gamma)$, delivering the proposition. $\square$

A.5 Proof of Proposition 3

We first record a simple efficiency property of optimal testing: The planner should not be able to remove some tested agents from a batched test and derive the same value. Intuitively, if she could, she could then profitably test the removed agents for a strict improvement.

Lemma 5 (TEST EFFICIENCY).

Suppose $\gamma \geq 0$ and $p \in (0, 1)$ have $\varphi^*(p, \gamma) > 0$. Then, every $k \in k^*(p, \gamma)$ has

$$\operatorname{argmax}_{k' \in \{1, \dots, k\}} [k'(1-p)^{k'}] = \{k\}.$$

Proof. Take an arbitrary $k' \in K$ with $k' < k$. First observe that

$$b(1-p)^k - \frac{\gamma}{k} = \varphi(p, k, \gamma) + v(p) \geq \varphi(p, k, \gamma) = \varphi^*(p, \gamma) > 0.$$

It follows that

$$\begin{aligned} k' [\varphi(p, k', \gamma) - \varphi(p, k, \gamma)] &= k' \left[b(1-p)^{k'} - \frac{\gamma}{k'} \right] - k' \left[b(1-p)^k - \frac{\gamma}{k} \right] \\ &> k' \left[b(1-p)^{k'} - \frac{\gamma}{k'} \right] - k \left[b(1-p)^k - \frac{\gamma}{k} \right] \\ &= b [k'(1-p)^{k'} - k(1-p)^k]. \end{aligned}$$

Therefore, because Lemma 2 says $q^+(p, k') \geq \bar{q}$ if $q^+(p, k) \geq \bar{q}$, having $k'(1-p)^{k'} \geq k(1-p)^k$ would contradict $k \in k^*(p, \gamma)$. $\square$

We now show, as a direct consequence of the above lemma, a limit on the batch size of a given agent.

Proof of Proposition 3. Given $\gamma \geq 0$ and an agent type $p \in (0, 1)$ with $\varphi^*(p, \gamma) > 0$, our goal is to show $k^*(p, \gamma) \subseteq \{1, \dots, \kappa(p)\}$, where $\kappa(p) := \max\left(K \cap [1, \frac{1}{p}]\right)$.

First, any $k \in K$ with $k > \kappa(p)$ has

$$k(1-p)^k - (k-1)(1-p)^{k-1} = (1-p)^{k-1}(1-kp) \leq 0,$$

so that $k \mapsto k(1-p)^k$ is weakly decreasing on $\{k \in K : k \geq \kappa(p)\}$.

Consider now an arbitrary $k \in K$ with $k > \kappa(p)$, so that the above calculation shows $k(1-p)^k \leq \kappa(p)(1-p)^{\kappa(p)}$. Our goal is to show $k \notin k^*(p, \gamma)$. By definition of k^* , we have nothing to show if $q^+(p, k) < \bar{q}$ or $\varphi(p, k, \gamma) \leq 0 < \varphi^*(p, \gamma)$; so assume otherwise. But then, Lemma 5 tells us $\varphi(p, \kappa(p), \gamma) > \varphi(p, k, \gamma)$. But Lemma 2 implies $q^+(p, \kappa(p)) \geq \bar{q}$ because $q^+(p, k) \geq \bar{q}$, so that $k \notin k^*(p, \gamma)$. $\square$

A.6 Proof of Proposition 4

In this subsection, we derive some intermediate results concerning the comparative statics of optimal batch sizes.

First, although our index is not generally single-peaked in the size of a batch, the following lemma establishes that it is so on the relevant range, that is, for the set of batch sizes that outperform not testing an agent at all (given the shadow cost of testing).

Lemma 6 (SINGLE-PEAKED INDEX AS BATCH SIZE VARIES).

For any $\gamma \geq 0$ and $p \in (0, 1)$, the set $\{k \in K : \varphi(p, k, \gamma) > 0\}$ is an interval in K , and the function $\varphi(p, \cdot, \gamma)$ is strictly quasiconcave on said interval.

Proof. Let $\zeta := \frac{1}{2} \log \frac{1}{1-p} > 0$, and define $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ by letting $f(0) := -v(p)$ and $f(z) := be^{-\frac{2\zeta}{z}} - v(p) - \gamma z$ for each $z > 0$. The function f is clearly smooth on $(0, \infty)$, and is therefore, applying L'Hôpital's rule at $z = 0$, continuous on its whole domain. Moreover, because the function $k \mapsto \frac{1}{k}$ is strictly monotone on K and each $k \in K$ has $\varphi(p, k, \gamma) = f\left(\frac{1}{k}\right)$, it suffices to show $\{f > 0\}$ is an interval and f is strictly quasiconcave on said interval. To that end, note each $z > 0$ has

$$f'(z) = \frac{2b\zeta}{z^2} e^{-\frac{2\zeta}{z}} - \gamma \implies f''(z) = \frac{4b\zeta}{z^4} e^{-\frac{2\zeta}{z}} (\zeta - z),$$

so that f is strictly convex on $[0, \zeta]$ and strictly concave on $[\zeta, \infty)$. Now, let z_* be the minimizer of $f|_{[0, \zeta]}$. Because $f|_{[z_*, \zeta]}$ is strictly increasing and $f|_{[\zeta, \infty)}$ is strictly concave (hence strictly quasiconcave), it follows directly that $f|_{[z_*, \infty)}$ is strictly quasiconcave. The result then follows if $f|_{[0, z_*]} \leq 0$. But the latter property indeed holds because $f|_{[0, z_*]}$ is decreasing and $f(0) \leq 0$. $\square$

We now demonstrate that where the constraint (P^A) is not relevant, the optimal batch size is decreasing in an agent's type. This feature would obtain immediately if the marginal testing value exhibited decreasing differences in the batch size and an agent's type. The proof shows that decreasing differences indeed obtains in the relevant range of batch sizes, where the relevant range is that suggested by Proposition 3.

Lemma 7 (DECREASING BATCHING WITHOUT PIVOTALITY CONCERNS).

Suppose $\gamma \geq 0$; $0 < p < p' < 1$; $\varphi^*(p, \gamma), \varphi^*(p', \gamma) > 0$; $k \in k^*(p, \gamma)$; $k' \in k^*(p', \gamma)$; and $q^+(p, k') \geq \bar{q}$. Then, $k \geq k'$.

Proof. As in the proof of Proposition 3, let $\kappa(p') := \max\left(K \cap [1, \frac{1}{p'}]\right)$. First note that if $k \geq \kappa(p')$, the result follows directly from Proposition 3. Let us now focus on the case in which $k < \kappa(p')$, and assume for a contradiction that $k < k'$. Proposition 3 then tells us $k' \in \{k+1, \dots, \kappa(p')\}$. Moreover, Lemma 2 then says $q^+(p', k) \geq q^+(p', k') \geq \bar{q}$, implying $\varphi(p', k', \gamma) \geq \varphi(p', k, \gamma)$ by definition of $k^*(p', \gamma) \ni k'$.

We now proceed to show $\varphi(p, k', \gamma) > \varphi(p, k, \gamma)$, which will (because $q^+(p, k') \geq \bar{q}$ by hypothesis) contradict $k \in k^*(p, \gamma)$, delivering the lemma.

To that end, observe any $\tilde{p} \in [p, p']$ and $\tilde{k} \in \{k+1, \dots, k'\}$ have

$$\begin{aligned} \frac{\partial}{\partial \tilde{p}} \left[(1 - \tilde{p})^{\tilde{k}} - (1 - \tilde{p})^{\tilde{k}-1} \right] &= -\frac{\partial}{\partial \tilde{p}} \left[\tilde{p}(1 - \tilde{p})^{\tilde{k}-1} \right] \\ &= -(1 - \tilde{p})^{\tilde{k}-2} \left[(1 - \tilde{p}) - \tilde{p}(\tilde{k} - 1) \right] \\ &= (1 - \tilde{p})^{\tilde{k}-2} \left[\tilde{p}\tilde{k} - 1 \right] \\ &\leq (1 - \tilde{p})^{\tilde{k}-2} \left[p'\kappa(p') - 1 \right] \\ &< 0. \end{aligned}$$

Summing over $\tilde{k}$ and integrating over $\tilde{p}$ then yields

$$(1 - p')^{k'} - (1 - p')^k < (1 - p)^{k'} - (1 - p)^k.$$

Finally, we see that

$$\begin{aligned} 0 &\leq \varphi(p', k', \gamma) - \varphi(p', k, \gamma) \\ &= \left[b(1 - p')^{k'} - \frac{\gamma}{k'} \right] - \left[b(1 - p')^k - \frac{\gamma}{k} \right] \\ &< \left[b(1 - p)^{k'} - \frac{\gamma}{k'} \right] - \left[b(1 - p)^k - \frac{\gamma}{k} \right] \\ &= \varphi(p, k', \gamma) - \varphi(p, k, \gamma), \end{aligned}$$

as desired. $\square$

The next lemma extends the previous one to account for the pivotality constraint on batched testing. Roughly, Lemma 7 determines how optimal batch sizes vary with type wherever the pivotality constraint is not binding, and Lemma 6 determines how they vary where pivotality is binding.

Lemma 8 (DECREASING BATCHING WHEN PIVOTALITY IS SLACK).

Suppose $\gamma \geq 0$; $0 < p < p' < 1$; $\varphi^*(p, \gamma), \varphi^*(p', \gamma) > 0$; $k \in k^*(p, \gamma)$; $k' \in k^*(p', \gamma)$; and $q^+(p, k+1) \geq \bar{q}$. Then, $k \geq k'$.

Proof. Define $I := [p, p'] \cap \{\varphi^*(\cdot, \gamma) > 0\}$, which is an interval by Proposition 2. For each $\tilde{k} \in K \setminus \{1\}$, Lemma 2 and the intermediate value theorem deliver a unique $p_{\tilde{k}}$ such that $q^+(p_{\tilde{k}}, \tilde{k}) = \bar{q}$. Letting $p_1 := 0$, Lemma 2 tells us $p_1 < \dots < p_K$.

Throughout the proof, we fix an arbitrary selector $k^{**} : I \rightarrow K$ of $k^*(\cdot, \gamma)|_I$. It suffices to show that k^{**} is weakly decreasing.

To that end, let us establish by induction on $\tilde{k} \in K$ that k^{**} is weakly decreasing on $I \cap (0, p_{\tilde{k}}]$. Such monotonicity is vacuous whenever $I \cap (0, p_{\tilde{k}}) = \emptyset$, and in particular for the base case of $\tilde{k} = 1$. Toward the inductive step, suppose $\tilde{k} \in K \setminus \{1\}$ is such that k^{**} is weakly decreasing on $I \cap (0, p_{\tilde{k}-1}]$, and focus on the nontrivial case in which $I \cap (0, p_{\tilde{k}}) \neq \emptyset$. By Lemma 7, k^{**} is weakly decreasing on $I \cap [p_{\tilde{k}-1}, p_{\tilde{k}})$, and so on $I \cap (0, p_{\tilde{k}})$.

Hence, all that remains for the inductive step is to show $k^{**}(p_{\tilde{k}}) \leq \lim_{\tilde{p} \nearrow p_{\tilde{k}}} k^{**}(\tilde{p}) =: k_-$. Assume otherwise for a contradiction. Then, $k^{**}(p_{\tilde{k}}) > k_-$ has $\varphi(p_{\tilde{k}}, k^{**}(p_{\tilde{k}}), \gamma) > \varphi(p_{\tilde{k}}, k_-, \gamma), 0$. By continuity of φ , therefore, some $p_- \in [p_{\tilde{k}-1}, p_{\tilde{k}})$ exists such that

$$\varphi(p_-, k^{**}(p_{\tilde{k}}), \gamma) > \varphi(p_-, k_-, \gamma), 0$$

and $k_- = k^{**}(p_-)$. But then, Lemma 6 implies $\varphi(p_-, k_- + 1, \gamma) > \varphi(p_-, k_-, \gamma)$ too. Because Lemma 2 implies $q^+(p_-, k_- + 1, \gamma) \geq q^+(p, k + 1, \gamma) \geq \bar{q}$, the payoff ranking contradicts $k_- \in k^*(p_-, \gamma)$.

We have thus established that k^{**} is weakly decreasing on $I \cap (0, p_K]$. Finally, Lemma 7 says k^{**} is weakly decreasing on $I \cap [p_K, 1)$, and so on all of I . $\square$

Finally, we observe that the previous lemma delivers a simple qualitative structure on the optimal batch size as one varies an agent's type: Conditional on testing, higher types receive a higher-quality (i.e., less batched) test.

Proof of Proposition 4. Consider the optimal testing rule, and let γ be a witnessing shadow cost as delivered by Theorem 2. It suffices to show $k \geq k'$ for any $0 < p < p' < 1$ with $\varphi^*(p, \gamma) > 0$; $\varphi^*(p', \gamma) > 0$; $k \in k^*(p, \gamma)$; and $k' \in k^*(p', \gamma)$.

Assume otherwise for a contradiction. By Lemma 8, it must be that $q^+(p, k + 1) < \bar{q}$, and therefore (because $k + 1 \leq k'$ and by Lemma 2) that $q^+(p, k') < \bar{q}$. The intermediate value theorem delivers some $p'' \in (p, p']$ with $q^+(p'', k') = \bar{q}$, and Lemma 7 says $k^*(p'', \gamma) \subseteq \{k', \dots, K\}$. But Lemma 2 implies $k^*(p'', \gamma) \subseteq \{1, \dots, k'\}$, so that $k^*(p'', \gamma) = \{k'\}$. Therefore,

$$\begin{aligned} \varphi^*(p'', \gamma) &= \varphi(p'', k', \gamma) \\ &= b(1 - p)^{k'} - (b - cp'') - \frac{\gamma}{k'} \\ &\leq cp'' - b[1 - (1 - p)^{k'}] \\ &= cp'' - b \frac{p''}{q^+(p'', k')} \\ &= -\frac{p''}{\bar{q}}[b - c\bar{q}] \\ &= 0, \end{aligned}$$

contradicting Lemma 4. $\square$

B Appendix: Explicit Form of Optimal Testing

Although we have focused on the qualitative features of optimal batch sizes (in Propositions 2, 3, and 4), our arguments in fact solve the planner’s problem rather completely. We briefly detail this solution in the present section.

First, observe the marginal value of testing decreases with the shadow cost γ for every batch size and agent type, so that the excess supply of tests will always increase with γ . Therefore, if we can compute the optimal testing policy associated with any given γ , correctly calibrating γ (as delivered by Theorem 2) to any desired degree of accuracy is straightforward. Thus, the planner’s problem essentially reduces to solving Theorem 2’s auxiliary program for a given γ . By Proposition 2, we need only consider γ such that $\varphi^*(\bar{q}, \gamma) > 0$.

Toward analyzing this problem, it suffices to consider test policies with weakly decreasing batch size by Proposition 4. Note that continuity of the marginal testing value implies that any type at which the batch size switches is one for which the two batch sizes generate the same marginal testing value. Then, because the marginal testing value is strictly single-peaked (by Lemma 6) in the batch size wherever positive, these two batch sizes must be consecutive; otherwise, batch sizes between them would be strictly better. Then, for any given $k \in \{1, \dots, K-1\}$, let $p_k(\gamma) \in [0, \frac{1}{k+1}]$ be uniquely defined by $\varphi(p_k, k, \gamma) = \varphi(p_k, k+1, \gamma)$; here, the bound on p_k comes from Proposition 3 and uniqueness comes from strict monotonicity of $\varphi(\cdot, k, \gamma) - \varphi(\cdot, k+1, \gamma)$ over this interval. Then, the lowest batch size $k_-(\gamma)$ used is the lowest $k \in \{1, \dots, K-1\}$ such that $p_k(\gamma) \geq \bar{q}$ and $\varphi(p_k(\gamma), k, \gamma) > 0$ if such a k exists, and is otherwise the lowest $k \in \{1, \dots, K-1\}$ such that $p_k(\gamma)$ is below the quarantine threshold. Similarly, the highest batch size $k_+(\gamma)$ used is the highest $k \in \{2, \dots, K\}$ such that $p_k(\gamma) \leq \bar{q}$ and $\varphi(p_{k-1}(\gamma), k, \gamma) > 0$ if such a k exists, and is otherwise the highest $k \in \{2, \dots, K\}$ such that $p_{k-1}(\gamma)$ is above the quarantine threshold. Then, the optimal batch sizes $\{k_-(\gamma), \dots, k_+(\gamma)\}$ of tested types are decreasing with cutoffs given by $p_{k_+(\gamma)-1}, \dots, p_{k_-(\gamma)}$. Finally, the set of tested types is the interval $(\underline{p}(\gamma), \bar{p}(\gamma))$, where $\underline{p}(\gamma) \in [0, p_{k_+(\gamma)}]$ and $\bar{p}(\gamma) \in [p_{k_-(\gamma)}, \frac{1}{p_{k_-(\gamma)}}]$ are uniquely determined by equations $\varphi(\underline{p}(\gamma), k_+(\gamma), \gamma) = \varphi(\bar{p}(\gamma), k_-(\gamma), \gamma) = 0$.

C Appendix: Proofs for Section 4

This section contains supporting analysis for the some extensions, discussed in section 4, of our model.

C.1 Imperfect Testing

Consider a model of imperfect testing in which up to two sources of imperfection exist. A test subject’s swab generates a binary interim test outcome, which is independent across all agents; and then, the laboratory generates a binary composite test outcome, which is informative of whether at least one of the subjects has a positive interim test outcome. In what follows, we fix a batch size $k \in K \setminus \{1\}$, which our notation will suppress. Parametrize the testing technology by four error probabilities $\lambda_-, \lambda_+, \sigma_-, \sigma_+ \geq 0$ such that $\lambda_- + \lambda_+, \sigma_- +$

$\sigma_+ < 1$. Here, σ_+ and σ_- are the false-positive rate and false-negative rate in the swab, and λ_+ and λ_- are the false-positive rate and false-negative rate in the laboratory. So, if a group test pools k subjects, of whom $\ell \in \{0, \dots, k\}$ are infected, the probability of a negative test outcome is¹²

$$\lambda_- + (1 - \lambda_- - \lambda_+)(1 - \sigma_+)^{k-\ell} \sigma_-^\ell.$$

It is nearly immediate that the pivotality property extends to this more general technological setting, and an identical argument to that of the benchmark setting establishes that sorting k^2 agents from k identical tests into k different homogeneous tests will preserve pivotality. Therefore, letting $\pi(\vec{p}|\lambda_-, \lambda_+, \sigma_-, \sigma_+)$ denote the expected net benefit from running a test of k individuals with heterogeneous type vector $\vec{p} = (p_1, \dots, p_k)$, assortative batching will follow if we establish that $\sum_{i=1}^k \pi(p_i \mathbf{1}_k | \lambda_-, \lambda_+, \sigma_-, \sigma_+) > k\pi(\vec{p} | \lambda_-, \lambda_+, \sigma_-, \sigma_+)$.

We begin by establishing the inequality for the special case of perfect aggregation in the laboratory and no false positives from the swab. Toward proving the payoff ranking in this case, observe that, with swabs having a false-negative rate $\alpha \in [0, 1)$, the benefit generated by a test is

$$\pi(\vec{p}|0, 0, \alpha, 0) = kb \prod_{j=1}^k [1 - (1 - \alpha)p_j] - \sum_{i=1}^k \alpha p_i c \prod_{j \in \{1, \dots, k\} \setminus \{i\}} [1 - (1 - \alpha)p_j],$$

because αp_i is the probability that agent i is infected and is nevertheless undetected by his swab, whereas $1 - (1 - \alpha)p_j$ is the probability that agent j is not detected to be positive by his swab. Now, let $\pi_b(\vec{p}|\alpha) := kb \prod_{j=1}^k [1 - (1 - \alpha)p_j]$ and $\pi_c(\vec{p}|\alpha) := \pi(\vec{p}|0, 0, \alpha, 0) - \pi_b(\vec{p}|\alpha)$. Applying the AM-GM inequality to $\left([1 - (1 - \alpha)^k]\right)_{i=1}^k$ yields $\sum_{i=1}^k \pi_b(p_i \mathbf{1}_k | \alpha) > k\pi_b(\vec{p}|\alpha)$. Moreover, $\pi_c(\cdot|\alpha)$ is obviously symmetric in its arguments, and is supermodular because every cross-partial derivative is nonnegative. Therefore (see Prat, 2002, Corollary 1), $\sum_{i=1}^k \pi_c(p_i \mathbf{1}_k | \alpha) \geq k\pi_c(\vec{p}|\alpha)$, so that summing the two inequalities yields the desired payoff ranking.¹³

Now, considering again all four kinds of error, we verify the required payoff ranking to establish assortative batching. Given a type profile $\vec{p} \in (0, 1)^k$ and a set $J \subseteq \{1, \dots, k\}$ of agents, let $P(J|\vec{p}) := \prod_{i=1}^k [p_i \mathbf{1}_{i \in J} + (1 - p_i) \mathbf{1}_{i \notin J}]$, the probability that the set of infected agents is exactly J . But then, the net benefit created by a batched test for k individuals with

¹²To unify notation, we adopt the convention that $0^0 = 1$ wherever it appears.

¹³One could also, rather than appealing to the AM-GM inequality, apply Prat (2002) directly to $\pi(\vec{p}|0, 0, \alpha, 0)$ for a weak inequality, establishing that assortative batching is (potentially not uniquely) optimal. Alternatively, a slight modification of his proof shows that a strict payoff ranking is generated by strict supermodularity.

types $\vec{p}$ is

$$\begin{aligned}
\pi(\vec{p}|\lambda_-, \lambda_+, \sigma_-, \sigma_+) &= (1 - \lambda_- - \lambda_+) \sum_{J \subseteq \{1, \dots, k\}} P(J|\vec{p}) (kb - |J|c) (1 - \sigma_+)^{k-|J|} \sigma_-^{|J|} \\
&\quad + \lambda_- \left(kb - c \sum_{i=1}^k p_i \right) \\
&= (1 - \lambda_- - \lambda_+)(1 - \sigma_+)^k \sum_{J \subseteq \{1, \dots, k\}} P(J|\vec{p}) (kb - |J|c) \left(\frac{\sigma_-}{1 - \sigma_+} \right)^{|J|} \\
&\quad + \lambda_- \left(kb - c \sum_{i=1}^k p_i \right) \\
&= (1 - \lambda_- - \lambda_+)(1 - \sigma_+)^k \pi(\vec{p}|0, 0, \frac{\sigma_-}{1 - \sigma_+}, 0) + \lambda_- \left(kb - c \sum_{i=1}^k p_i \right).
\end{aligned}$$

Therefore, the difference between $\sum_{i=1}^k \pi(p_i \mathbf{1}_k | \lambda_-, \lambda_+, \sigma_-, \sigma_+)$ and $k\pi(\vec{p} | \lambda_-, \lambda_+, \sigma_-, \sigma_+)$ is proportional to that between $\sum_{i=1}^k \pi(p_i \mathbf{1}_k | 0, 0, \frac{\sigma_-}{1 - \sigma_+}, 0)$ and $k\pi(\vec{p} | 0, 0, \frac{\sigma_-}{1 - \sigma_+}, 0)$, which our above work demonstrates to be strictly positive. Assortative batching follows.

Now, given assortative batching, indexability extends easily to the setting with testing errors, with the modified marginal testing value given by

$$\begin{aligned}
\varphi(p, k, \gamma) &= b(1 - \lambda_- - \lambda_+)(1 - \sigma_+)^k \left[1 - \left(1 - \frac{\sigma_-}{1 - \sigma_+} \right) p \right]^{k-1} \left\{ 1 - \left[1 + \left(\frac{c}{b} - 1 \right) \frac{\sigma_-}{1 - \sigma_+} \right] p \right\} \\
&\quad - [v(p) - \lambda_-(b - cp)] - \frac{\gamma}{k}.
\end{aligned}$$

If testing exhibits no false negatives, this marginal value can be expressed as

$$\hat{b}(1 - \hat{p})^k - v\left(\frac{\hat{p} - \sigma_+}{1 - \sigma_+}\right) - \frac{\gamma}{k},$$

where $\hat{b} := b(1 - \lambda_+) > 0$ and $\hat{p} := 1 - (1 - \sigma_+)(1 - p) \in (\sigma_+, 1)$. Therefore, analysis identical to that of the benchmark model shows this extended model exhibits threshold testing, a swab-error-adjusted modification of limited batching, and ordered batching.

C.2 Richer Heterogeneity

We verify here that a suitable generalization of Lemma 3 obtains (from which assortative batching follows) if the benefit of releasing an agent is negatively associated with his probability of infection. To that end, for a batch size $k \in K$, consider a vector of infection risks $\vec{p} \in (0, 1)^k$ and a vector of benefits $\vec{b} \in (0, \infty)^k$ such that $p_1 \leq \dots \leq p_k$ and $b_1 \geq \dots \geq b_k$, with $p_1 < p_k$. Our goal is to show that

$$k \sum_{i=1}^k b_i \prod_{j=1}^k (1 - p_j) < \sum_{i=1}^k kb_i (1 - p_i)^k.$$

Assume without loss that $\sum_{i=1}^k b_i = 1$. In this case, modifying a textbook proof of the AM-GM inequality shows

$$\begin{aligned}
\log \left[\sum_{i=1}^k b_i (1 - p_i)^k \right] &> \sum_{i=1}^k b_i \log \left[(1 - p_i)^k \right] \text{ (by Jensen's inequality)} \\
&= \sum_{i,j=1}^k b_i \delta_{ij} \log \left[(1 - p_j)^k \right] \\
&\geq \sum_{i,j=1}^k b_i \frac{1}{k} \log \left[(1 - p_j)^k \right] \text{ (by Galichon, 2018, Theorem 4.3(i))} \\
&= \log \prod_{j=1}^k (1 - p_j) \\
&= \log \left[\sum_{i=1}^k b_i \prod_{j=1}^k (1 - p_j) \right],
\end{aligned}$$

as required.

C.3 General Quarantine Thresholds

Suppose a fixed threshold $\bar{q} \in (0, 1)$ is such that $\mu[\bar{q}, 1) = 0$. We verify below that if the optimal quarantine threshold $\frac{b}{c}$ is close enough to zero, and tests are sufficiently scarce, optimal batch size is weakly increasing in the order of tested agents in Proposition 4. To that end, suppose $KT < \mu(0, 1)$. Because the benefit b affects the optimal test policy only via the ratio $\frac{c}{b}$, we without loss normalize $b = 1$ and focus on the case in which c is sufficiently high. Moreover, we index any objects that depend on the (one moving) parameter c by said parameter.

Define the function $\bar{k} : (0, 1) \rightarrow \mathbb{R}$ by taking each type $p \in (0, 1)$ to the highest $k \in K$ such that $q^+(p, k) \geq \bar{q}$, which is weakly increasing by Lemma 2. Because $\int_{\underline{p}}^{\bar{q}} \frac{1}{k} d\mu \geq \int_{\underline{p}}^{\bar{q}} \frac{1}{K} d\mu > T$, the intermediate value theorem delivers some cutoff $\underline{p} \in (0, \bar{q})$ with $\int_{\underline{p}}^{\bar{q}} \frac{1}{k} d\mu = T$.

We show that if c is sufficiently high, it is optimal to test exactly the types in $(\underline{p}, \bar{q})$ and put each such type p in a batch of size $\bar{k}(p)$. In what follows, assume $c > \frac{1}{\underline{p}} \left[1 - (1 - \underline{p})^{\bar{k}(\underline{p})} \right]$, and let

$$\gamma_c := \bar{k}(\underline{p}) \left\{ c \underline{p} - \left[1 - (1 - \underline{p})^{\bar{k}(\underline{p})} \right] \right\},$$

so that $\varphi_c(\underline{p}, \bar{k}(\underline{p}), \gamma_c) = 0$. Now, observe

$$\begin{aligned}
& \frac{1}{c} \inf_{p \in (\underline{p}, \bar{q}), k \in \{1, \dots, K-1\}} [\varphi_c(p, k+1, \gamma_c) - \varphi_c(p, k, \gamma_c)] \\
&= \inf_{p \in (\underline{p}, \bar{q}), k \in \{1, \dots, K-1\}} \frac{1}{c} \left\{ [(1-p)^{k+1} - (1-p)^k] + \gamma_c \left(\frac{1}{k+1} - \frac{1}{k} \right) \right\} \\
&= \inf_{p \in (\underline{p}, \bar{q}), k \in \{1, \dots, K-1\}} \left[\frac{\gamma_c}{ck(k+1)} - \frac{1}{c} p(1-p)^k \right] \\
&\geq \frac{\gamma_c}{c(K-1)K} - \frac{1}{c} \\
&= \frac{\bar{k}(\underline{p})}{(K-1)K} \left\{ \underline{p} - \frac{1}{c} [1 - (1-\underline{p})^{\bar{k}(\underline{p})}] \right\} - \frac{1}{c} \\
&\geq \frac{\bar{k}(\underline{p})}{(K-1)K} \left(\underline{p} - \frac{1}{c} \right) - \frac{1}{c} \\
&\xrightarrow{c \rightarrow \infty} \frac{\bar{k}(\underline{p})}{(K-1)K} \underline{p} \\
&> 0.
\end{aligned}$$

Therefore, for sufficiently large c , each $p \in (\underline{p}, \bar{q})$ has $\varphi_c(p, \cdot, \gamma_c)$ strictly increasing and so $k_c^*(p, \gamma_c) = \{\bar{k}(p)\}$. In particular, these features imply $\varphi_c^*(\underline{p}, \gamma_c) = \varphi_c(\underline{p}, \bar{k}(\underline{p}), \gamma_c) = 0$. Lemma 4 then tells us $\varphi_c^*(\cdot, \gamma_c)$ is negative below $\underline{p}$ and positive between $\underline{p}$ and $\bar{q}$. By Theorem 2, therefore, the given testing policy is optimal.

C.4 Comparison to Dorfman Testing

We first cast the problem of optimal Dorfman testing in similar language to our main analysis. The planner chooses, for each batch size $k \in K$, measure τ_k over profiles of agents to be tested in batches of size of k . A tested agent of type $p \in (0, 1)$ is sure to receive a test that fully resolves his infection status—either a negative test in the initial group test, or an individual test if the former is positive. Therefore, the planner will extract net value $pv(1) + (1-p)v(0) - v(p) = b(1-p) - v(p)$ from testing said agent, irrespective of who is jointly tested. The demand constraint, that each agent is tested in at most one group test, is exactly as it was for the case of many-to-one testing. The supply constraint is now more complex. The planner's expected number of tests from a single test of k individuals is no longer 1, but must additionally be augmented by k times the probability of a positive test outcome if $k > 1$ —that is, if the test is not an individual test. The supply constraint says that this augmented number of tests does not exceed the testing capacity.¹⁴ Hence, the Dorfman

¹⁴Because of our continuum assumption, the total number of tests used will be deterministic.

planner's problem is

$$\begin{aligned}
& \max_{(\tau_k)_{k \in \mathcal{T}}} \sum_{k=1}^K \int_{(0,1)^k} \sum_{i=1}^k [b(1-p_i) - v(p_i)] d\tau_k(p_1, \dots, p_k) \\
& \text{s.t.} \quad \sum_{k=1}^K \int_{(0,1)^k} \left\{ 1 + \mathbf{1}_{k>1} k \left[1 - \prod_{j=1}^k (1-p_j) \right] \right\} d\tau_k(p_1, \dots, p_k) \leq T \quad (\text{S}_{\text{Dorfman}}) \\
& \quad \sum_{k=1}^K \sum_{i=1}^k \text{marg}_i \tau_k \leq \mu. \quad (\text{D})
\end{aligned}$$

Let us now observe now that assortative batching is optimal under Dorfman testing as well. Just as in the case of many-to-one testing, we consider a modification that “symmetrizes” a given non-assortative testing policy. Under the resorted policy, the same agents are tested in a batch of size k for any k , but each collection of k identical tests of k heterogeneous agents is replaced with k different tests each of k homogeneous agents. Because the objective and (D) depend only on the marginal distributions of each τ_k , neither is affected. However, again applying the arithmetic-mean-geometric-mean inequality, such re-sorting strictly relaxes (S_{Dorfman}). Thus, assortative batching is optimal, strictly so if first-best is not attainable.

Given assortative batching, the planners problem can be recast as

$$\begin{aligned}
& \max_{(\tau_k)_{k \in \mathcal{T}^A}} \sum_{k=1}^K \int_{(0,1)} k [b(1-p) - v(p)] d\tau_k(p) \quad (\text{O}_{\text{Dorfman}}^A) \\
& \text{s.t.} \quad \sum_{k=1}^K \int_{(0,1)} \left\{ 1 + \mathbf{1}_{k>1} k [1 - (1-p)^k] \right\} d\tau_k(p) \leq T \quad (\text{S}_{\text{Dorfman}}^A) \\
& \quad \sum_{k=1}^K k \tau_k \leq \mu. \quad (\text{D}^A)
\end{aligned}$$

Finally, just as with many-to-one testing, we can replace the supply constraint with the shadow cost of a test. For any $p \in (0, 1)$ and shadow cost $\gamma \geq 0$ define the marginal testing value,

$$\begin{aligned}
\varphi_{\text{Dorfman}}(p, k, \gamma) &:= b(1-p) - v(p) - \frac{\gamma}{k} \left\{ 1 + \mathbf{1}_{k>1} k [1 - (1-p)^k] \right\} \\
&= [b(1-p) - v(p)] + \gamma \left\{ \mathbf{1}_{k>1} \left[(1-p)^k - 1 \right] - \frac{1}{k} \right\}
\end{aligned}$$

for each $k \in K$, the optimal batch sizes,

$$k_{\text{Dorfman}}^*(p, \gamma) := \arg\max_{k \in K} \varphi_{\text{Dorfman}}(p, k, \gamma),$$

and the testing index,

$$\varphi_{\text{Dorfman}}^*(p, \gamma) := \max_{k \in K} \varphi_{\text{Dorfman}}(p, k, \gamma).$$

Then, following nearly identically the argument for many-to-one testing, a feasible Dorfman testing policy $\vec{\tau} \in \mathcal{T}^A$ is optimal if and only if it admits a shadow cost $\gamma \geq 0$ such that¹⁵

1. No agent of type p with $\varphi_{\text{Dorfman}}^*(p, \gamma) < 0$ is tested,
2. Every agent of type p with $\varphi_{\text{Dorfman}}^*(p, \gamma) > 0$ is tested in a homogeneous k -batch for some $k \in k^*(p, \gamma)$, and
3. Either (S_{Dorfman}^A) holds with equality or $\gamma = 0$.

Now, just as for many-to-one testing, we can study the marginal testing index to establish further properties of the optimal testing policy. In particular, as a virtually identical analysis shows that this testing index strictly increases (resp. decreases) to the left (right) of the threshold $\bar{q}$, it follows that the set of tested agents is again an interval containing this threshold.

Next, observe that whenever first-best is unattainable (and hence the shadow cost satisfies $\gamma > 0$ at the optimum), any batch size $k \in \{2, \dots, K\}$ has $\varphi_{\text{Dorfman}}(p, 1, \gamma) - \varphi_{\text{Dorfman}}(p, k, \gamma) = \gamma \left[\frac{1}{k} - (1-p)^k \right]$, which has the same sign as $p - p_k$, where $p_k := 1 - \left(\frac{1}{k}\right)^{\frac{1}{k}}$. Thus, an optimal testing policy never uses batch size of k for agents of type greater than p_k , as individually testing each such agent uses fewer tests. Moreover, any tested agent (barring a measure-zero set of agents) is individually tested if and only if his type is greater than $\max_{k \in \{2, \dots, K\}} p_k$.

Now, we establish that batch sizes decrease with infection risk whenever first-best is unattainable. Assume otherwise for a contradiction. Just as in the case of many-to-one matching, the difference $\varphi_{\text{Dorfman}}(\cdot, k', \gamma) - \varphi_{\text{Dorfman}}(\cdot, k, \gamma)$ is a nonzero polynomial for distinct $k', k \in K$, and so switches signs at only finitely many points. Hence, some type $p^* \in (0, 1)$ and batch sizes $k \in \{1, \dots, K-1\}$ and $k' \in \{k+1, \dots, K\}$ exist such that a batch size of k (resp. k') is optimal for agents of type slightly below (resp. above) p^* . Given the previous paragraph's characterization of when individual testing is optimal, it must be that $k > 1$ and $p^* < p_k$. Defining $\ell := k' - k$, and letting $f := \frac{1}{\gamma} [\varphi_{\text{Dorfman}}(\cdot, k', \gamma) - \varphi_{\text{Dorfman}}(\cdot, k, \gamma)] : (0, 1) \rightarrow \mathbb{R}$, direct computation shows $f'(p) = (1-p)^{k-1} \left[k - (k+\ell)(1-p)^\ell \right]$. Therefore, $f'(p)$ has the same sign as $p - p_*$, where $p_* := 1 - \left(\frac{k}{k+\ell}\right)^{\frac{1}{\ell}}$. Thus, f is strictly quasiconvex with minimizer p_* . As f switches from negative to positive at p^* , then, it follows that $p_* < p^* < p_k$. Observe now that

$$\begin{aligned} f(p_k) &= \frac{\ell}{k(k+\ell)} - (1-p_k)^k \left[1 - (1-p_k)^\ell \right] \\ &= \frac{\ell}{k(k+\ell)} - \frac{1}{k} \left[1 - (1-p_k)^\ell \right] \\ &= \frac{1}{k} \left[(1-p_k)^\ell - \frac{k}{k+\ell} \right], \end{aligned}$$

which is strictly negative because $(1-p_k)^1 - \left(\frac{k}{k+\ell}\right)^{\frac{1}{\ell}} = p_* - p_k$ is. But then, as p^* lies between the minimizer of f and another point at which f is strictly negative, we have $f(p^*) < 0$, contradicting the definition of p^* .

¹⁵Note that, when the supply constraint binds, each tested type's optimal batch size under Dorfman testing is independent of the test supply. This feature does not hold for the case of many-to-one testing.