An Economic Framework for Vaccine Prioritization[*]

Mohammad Akbarpour[†] [ⓡ] Eric Budish[‡] [ⓡ] Piotr Dworczak[§]
[ⓡ] Scott Duke Kominers[¶]

May 15, 2021

Abstract

We propose an economic framework for determining priority rules for allocating a scarce supply of vaccines that gradually become available during a public health crisis, such as the Covid-19 pandemic. Agents differ in observable and unobservable characteristics, and the designer maximizes a social objective function over all feasible mechanisms accounting for those characteristics, as well as agents' endogenous behavior in the face of the pandemic. The framework emphasizes the role of externalities and incorporates equity as well as efficiency concerns. Our results provide an economic justification for providing the vaccines immediately and for free to some groups of agents, while at the same time showing that a carefully constructed pricing mechanism can improve outcomes by screening for individuals with the highest private and social benefits of receiving the vaccine.

Keywords: Covid-19, vaccination, mechanism design, inequality

JEL codes: C78, D47, D61, D63, D82

1 Introduction

In response to the Covid-19 pandemic, multiple effective vaccines have been developed with unprecedented speed. Nevertheless, vaccine supply chains are constrained due to production and logistical challenges, making vaccines a scarce resource in the short run. The scarcity creates a clear market design problem: How should we stage and sequence the delivery of vaccines? And to what extent should we allocate vaccines for free with rationing, as opposed to using prices to match demand with supply? Solving this problem effectively requires incorporating externalities, ethical considerations, and equity concerns—making it impossible to apply off-the-shelf economic tools. For example, to what extent do we want to vaccinate certain groups of agents (say, doctors) before anyone else? What if there are multiple groups that deserve special consideration (front-line workers, teachers, elderly); should we vaccinate these groups sequentially or in parallel?

The goal of this paper is to clarify these issues by proposing a framework that can inform decisions about prioritization across different groups. The framework is "economic" in the sense that it poses an allocation problem with constraints and an objective function, and provides a characterization of the optimal mechanism that can be implemented using a combination of priorities, prices, and lotteries. More specifically, the paper makes three broad contributions. First, it provides a way to evaluate the vaccination priority rules that are already being used during the current Covid-19 pandemic; we show that certain features of these protocols can indeed be justified as optimal, while others could be improved upon. Second, our framework can be used to inform policy: either in the context of the current pandemic in countries that have not yet finalized vaccination schemes, or in the context of future public health crises that many experts view as inevitable. Third, our analysis provides a blueprint for analyzing related problems in which a scarce resource of differentiated quality must be allocated optimally while taking into account externalities, endogenous behavior of agents, and equity considerations.

Before discussing our findings, we briefly sketch the model: Prior to receiving a vaccine, each agent chooses an action reflecting her behavior during the pandemic. For simplicity, we model the behavior as a binary choice: The agent may choose to take precautions that minimize the probability of contracting the virus (the "safe" action) at the cost of significantly limiting her in-person interactions (understood broadly as any activities, related to work or leisure, that create meaningful risk of infection); or she may choose to continue engaging in these interactions, therefore incurring a greater risk of infection (the "risky" action). The specific interactions the agent engages in—as well as the implicit cost of taking precautions—depend on the agent's type. In particular, the agent's choice is determined by the ranking of

the private *health benefit* of not contracting the virus and the private *socio-economic benefit* of in-person interactions. These benefits incorporate any constraints the agent might face: For example, if someone must work in-person in a high risk environment to not lose their job, this is modeled as a high socio-economic benefit of in-person interactions.

The agent's private choice generates externalities—the safe action generates social health benefits by slowing down the spread of the virus, while the risky action generates benefits associated with the agent's economic and social activity. Receiving the vaccine is modeled as unleashing both types of benefits (associated with the safe and risky actions) at the same time. All these benefits are expressed in dollar values. To address inequality and ethical concerns, we assume that each agent is associated with a welfare weight capturing the societal value of giving that agent a dollar; for example, a higher weight may be attached to agents who are less wealthy, those who are disproportionately harmed by the pandemic, or those perceived as playing key roles in fighting the pandemic.

The mechanism designer has access to some observable information about agents that we refer to as *labels*, which might include, for example, the agent's profession, age, income, or neighborhood. Additionally, the designer may truthfully elicit the agents' private willingness to pay for a vaccine by setting up differentiated price schedules for different speeds of receiving the vaccine. (While in our model we express prices in terms of monetary payments, a special case of our model allows interpreting the payments as exerting costly efforts such as queueing; see Section 8 for discussion of this point.) Overall, the designer may choose any vaccination schedule (i.e., a potentially random time of receiving the vaccine for every agent) as long as it is feasible (given the exogenous dynamic supply of vaccines), as well as incentive-compatible and individually-rational under the associated payment rules. We assume that the designer maximizes a utilitarian objective—consisting of the sum of all agents' utilities, including the externalities, weighted by their welfare weights. The utilitarian objective also allows for a positive weight on revenue which, may be especially desirable if—as in some developing-world contexts—funds raised in the priced market can be used to purchase more vaccines for public delivery or to provide other public services. (In an extension, we also consider an alternative health objective that includes only the total private and social health benefits.)

The classical economic solution to allocation problems is an auction or a market system—such a method ensures that goods are allocated to the agents with the highest *willingness to pay* (*WTP*).[1] Yet in the context of vaccines, and indeed in our framework, the standard argument in favor of pure market allocation (or efficient auctions) is insufficient because of the presence of both externalities and heterogeneous social welfare weights. For example,

---

[1]See Pancs (2020) for how an auction mechanism could be augmented to account for allocative externalities in the context of vaccine allocation.

the WTP of a young healthy doctor may be lower in comparison to the social value of vaccinating that doctor since this also protects the patients that she is in contact with. Moreover, a standard market allocation is oblivious to the ethical considerations reflected in welfare weights—such as the idea that doctors should be protected specially because their job both exposes them to risk and is central to addressing the consequences of the pandemic.

For these and other reasons, many countries have opted for a pure priority system in place of a market allocation, identifying groups of agents with similar observable characteristics (such as job or age) and offering vaccines sequentially to these groups free of charge, starting from groups with the highest health benefits, positive externalities, and moral desert. (In our model, a *group* is defined as the set of agents with the same label.) We show that such a sequential priority system is indeed optimal *if the designer cannot use prices.* In particular, when vaccines are distributed for free, it is suboptimal to have overlaps in vaccination schedules for different groups; instead, the optimal schedule induces a strict ordering over groups, based on a one-dimensional sufficient statistic derived from the objective function for each group. The sufficient statistic used to determine the optimal priority ordering is a compact measure aggregating both private health and socio-economic benefits, as well as the associated externalities; it also reflects the social welfare weights that could capture ethical and equity considerations.

A potential concern with this method, however, is that observable information may not fully reveal the relevant variables. For example, two individuals similar on observables may have different family or social circumstances that lead to large differences in their value for being able to take the risky action (e.g., acute social isolation versus close family or friends in one's pod). Or, a ride-share driver with high WTP to receive a vaccine might be choosing the safe action (not driving) but really need the money from driving; alternatively, their high WTP might indicate that they are taking the risky action (driving) and thus have a high private health benefit, as well as a high health externality because they are interacting with many different people. More generally, the benefits of vaccinating a given person depend on unobservable choices (e.g., the level of precaution that agent is taking prior to getting a vaccine) and preferences (e.g., the value for being able to engage in in-person interactions). This points to potential benefits of prices to supplement priority rules for some groups—a carefully designed price system has the advantage of screening for unobserved characteristics that could reveal relevant information about the private and social values of vaccination. Of course, prices can at most reveal WTP, which means they are far from perfect as a screening device in our framework. In particular, WTP is not only affected by the true "value" of getting vaccinated but also by other factors such as the opportunity cost of money. As a result, the optimal pricing system has to optimally trade-off screening benefits against the

resulting redistributive consequences, given the designer's objective. The optimal mechanism with prices takes a hybrid form, combining priorities based on observable information with price-induced screening of the privately observed WTP.

This is where our framework has the most power. By employing the methods developed in Akbarpour ⓡ Dworczak ⓡ Kominers (2020) (henceforth ADK), we can derive the exact form of the optimal mechanism with prices. The optimal mechanism specifies how the vaccines are allocated across groups—sets of agents with the same observable characteristics (labels)—and then how the vaccines are allocated within each group. Within a group, agents are partitioned according to WTP into blocks, with each block either matched to available vaccination times assortatively at increasing prices (a "market allocation") or vaccinated in random order without a payment (a "free allocation"). Overall, this solution describes a complete vaccination schedule, along with the supporting payments.

One of our key insights is that the social benefits of vaccinating a given agent depend crucially on that agent's endogenous action choice prior to receiving the vaccine that can be indirectly inferred through the label and willingness to pay. For example, most doctors are effectively required to undertake in-person interactions by nature of their profession; thus, a label associated with being a doctor indicates that vaccination will have a *health benefit* for that doctor, as well as a positive *health externality* by helping protect the people that doctor interacts with. In contrast, many college professors have been able to teach from home during the pandemic. Thus, vaccinating such individuals will have a *socio-economic benefit* for them, as well as a positive *socio-economic externality*, because it enables these agents to take actions (such as advising and teaching in-person) that would otherwise be avoided due to health risk.

Using our framework, we can identify sufficient conditions under which it is optimal for a group of agents to receive the vaccine immediately and for free even when prices could be used (what we call "priority allocation"). These conditions require that the label that defines the group is associated with a sufficiently high positive externality or a sufficiently high welfare weight; additionally, the weight on revenue should be sufficiently small. Thus, our result justifies priority allocation to front-line health workers; they have a high health externality (which is the relevant externality since, by definition, these workers are forced to choose the risky action) as well as a high welfare weight due to the social awareness of their key role in fighting the pandemic.

In contrast, in developing countries whose ability to purchase vaccines in the international market depends on their internal revenue generation (so that the weight on revenue could be substantial), it may be desirable to vaccinate multiple groups simultaneously, in what effectively is a combination of a subsidized public allocation program and a private market.

4

For example, it may be optimal to allocate vaccines to front-line health workers at low or zero prices, while at the same time offering vaccines at high prices to the general population. Only once groups with high externalities and welfare weights are vaccinated will the prices for the general population be reduced.

Moreover, even if a pure market allocation is suboptimal, judicious use of prices for parts of the allocation can still enhance efficiency substantially. As in any economic problem, using prices gives priority to people with the highest willingness to pay. In our context, high private benefits to vaccination may stem from both health benefits (for those who engage in in-person interactions prior to getting a vaccine) and socio-economic benefits (for those who choose the safe action). For example, prices may help allocate vaccines early on to people with privately known co-morbidities or other idiosyncratic health concerns. Indeed, if an agent is concerned about their health, then either that agent is taking a significant health risk by engaging in in-person interactions, or she is sacrificing her (potentially large) socio-economic benefits by avoiding them. Either way, the agent is likely to have a significant willingness to pay. In contrast, healthy people who are not concerned about becoming infected with Covid-19 will have a low willingness to pay for a vaccine *regardless* of their socio-economic benefit.[2] When private health concerns are not directly observable, using prices to allocate vaccines is the only way to ensure that the former set of agents receives vaccines ahead of the latter.

Another insight revealed by our framework is that willingness to pay can be used beyond its traditional role of identifying agents with a high private benefit. In our context, WTP could also be informative about the externalities associated with vaccinating a given agent, which may work both in favor or against using prices to allocate vaccines, depending on the group. For a simple illustration of the first possibility, note that many non-essential workers might have private information about their company's plans for returning to in-person work. Because an unvaccinated agent returning to in-person work induces a negative health externality, vaccinating this agent early on is socially valuable; at the same time, being forced to return to in-person work raises the individual's private benefit and hence willingness to pay. These two forces combined create a positive correlation between the social and the private values which a price system can effectively exploit. While in our framework we only model individual choices, the same logic may be applied to justify why it may be desirable to offer vaccines at carefully chosen prices to corporations and other organizations. Indeed, this would allow the institutions that have a particularly high value for returning to in-person interactions to secure earlier access to vaccines for their members,

---

[2]Indeed, even if the socio-economic benefit is large, they will simply choose the risky action, and vaccination will only give them a small health benefit.

and hence avoid the potential adverse health consequences of reopening.

However, pricing becomes suboptimal when the correlation between willingness to pay and externalities is negative. For an illustration, let us revisit the gig-economy workers example. As argued earlier, their health externality may be large in case they choose to perform jobs requiring substantial in-person interactions. Under reasonable assumptions, workers that are less wealthy are more likely, on the margin, to continue performing risky tasks; for example, ride-share drivers may afford to stop driving for some time if ride-sharing is only a supplementary source of income for them, but are unlikely to do so if their livelihood depends on it. At the same time, factors such as low income or a challenging financial situation may imply a low WTP for a vaccine, especially if the *private* health benefit of a worker is small. As a result, using prices could lead to allocating vaccines to gig-economy workers with lower than average health externalities. In such cases, free allocation with rationing would perform better by reaching the high-health-externality workers with higher probability.

Our framework predicts that vaccinating groups of agents sequentially (e.g., health workers, then teachers) is optimal when allocation within these groups is free (and relies on rationing). However, when prices are used to provide earlier access to agents with highest WTP, this is no longer the case. It is in general optimal to have overlaps in the schedule (e.g., some teachers receive the vaccines before some health workers). The intuition is simple: Under a free allocation, vaccinating each agent within a group provides the same expected social benefit since a free allocation leads to *random* order of vaccination within each group; in contrast, under a price allocation, *agents with relatively higher value are vaccinated first—* the marginal value of vaccinating the highest-WTP teacher may easily exceed the marginal value of vaccinating the lowest-WTP doctor.

Finally, our framework makes it possible to compare how the optimal mechanism depends on the designer's objective function. For illustration, consider the case in which the designer is only concerned about health outcomes, that is, she ignores the socio-economic benefits in her calculations. Under this pure-health objective, it is only socially beneficial to vaccinate agents who would otherwise choose the risky action—if someone already chooses the safe action prior to vaccination, then vaccinating that agent generates the socio-economic benefits (the agent is now able to engage in in-person interactions) but not the health benefits. As a result, agents with labels that reveal a high probability of choosing the safe action will have low priority in the optimal vaccination schedule. For example, such groups may include the elderly and people with potential comorbidities. In contrast, healthy college students are much more likely to choose the risky action, and thus vaccinating them may have a substantial positive health externality (even if their private health benefit is small). Thus,

paradoxically, precisely when health is the objective, the optimal schedule may prioritize those individuals who are not at risk themselves, and are thus likely to spread the virus prior to receiving the vaccine. Under the utilitarian objective, the priority could reverse: A utilitarian designer takes into account that by vaccinating individuals with the highest risk she allows them to engage in in-person interactions that generate both private socio-economic gains and the socio-economic externality. These gains may be particularly large if such individuals receive high welfare weights—which is plausible because they are among those most adversely affected by the pandemic.

1.1 Related literature

Within the economics literature, there are a few recent proposals for the problem of vaccine allocation. Most prominently, Pathak et al. (2021) developed a model of reserve design and associated multiple-category priority system for use in the allocation of vaccines, ventilators, and other scarce health resources. They show how using multiple group-specific reserves makes it possible to balance distinct ethical goals, such as a desire to target vaccines in ways that both reach the most vulnerable and reach those with high health externalities.[3] This work has been quite influential in shaping vaccine allocation in practice, being included in the National Academies of Sciences, Engineering, and Medicine (2020) Framework for Equitable Allocation of Covid-19 Vaccine, and, in part as a result, being incorporated in numerous states' allocation policies.[4] They have also spurred a medical ethics literature looking at different ways to implement multiple reserve systems (e.g., Pathak et al. (2020); Sönmez et al. (2020); Makhoul and Drolet (2021)), and associated innovation in market design theory (e.g., Delacrétaz (2020)).

Our approach is complementary to that of Pathak et al. (2021): Pathak et al. (2021) take social priorities as an exogenous input reflecting underlying but unmodeled ethical considerations; in contrast, we take a more classical economic approach of maximizing a welfare function, and thus our priorities are an endogenous feature of the optimal mechanism. This means that our priorities emerge from the economic dynamics involving agents' characteristics and choices of behavior in the face of the pandemic, as well as the resulting externalities. Our framework can still incorporate ethical considerations indirectly through the social wel-

---

[3]Grigoryan (2021) has built on the ideas of Pathak et al. (2021) by introducing a match quality term to the allocation model, leading to a mechanism for allocating vaccines to different groups of agents in ways that maximize group-specific efficacy while respecting prioritization goals.

[4]Pathak et al. (2021) also had significant impact on medical resource allocation earlier in the pandemic—for example, the University of Pittsburgh Medical Center (UPMC) explicitly developed a weighted lottery system for ventilators based on their work (University of Pittsburgh Department of Critical Care Medicine, May 28, 2020).

fare weights in the objective function. In this sense, we can think of our model as being informative about the priorities that could be used as an input to the types of reserve systems Pathak et al. (2021) propose. However, the structure of the priorities arising in our framework are slightly different from those considered by Pathak et al. (2021), depending on the tools available to the designer. If the mechanism cannot use prices, then our framework (generically) predicts that groups should be strictly ordered, with no overlap in the vaccination schedule. This arises because of linearities in our model: Under a welfare function with constant welfare weights, there is no force preventing the designer from wanting to give absolute priority to the group with the highest social value. In contrast, the reserve-design of Pathak et al. (2021) allows for overlaps in vaccination schedules for different groups, which corresponds to the idea of nonlinearities in social preferences arising from the need to balance multiple ethical frameworks. If the mechanism can use prices, then we obtain a similar conclusion to Pathak et al. (2021) for a different reason: Rather than reflecting ethical considerations, the overlaps we obtain in the presence of prices are a consequence of decreasing marginal value of vaccinating individuals within groups when prices allow for efficient sorting.

Our work also complements and in many cases directly reinforces existing analyses of medical and ethical reasons for prioritizing certain groups for early Covid-19 vaccine allocation. Persad et al. (2020) developed an ethical framework for vaccine allocation based on the principles of "benefiting people and limiting harm"; "prioritizing disadvantaged populations"; and "preclud[ing] consideration of differences [. . .] when doing so would not help prevent harm or prioritize disadvantaged groups." They find that these principles particularly support prioritizing "health care workers; other essential workers and people in high-transmission settings; and people with medical vulnerabilities associated with poorer COVID-19 outcomes[. . .]." At least at a loose level, this concords with the predictions of our model: Healthcare workers naturally have high value for vaccination because of high exposure risk, and also have high health externalities because their job explicitly requires interacting with others, including some who may themselves be at risk. Essential workers and others in high-transmission settings are likewise at high risk themselves and have high health externalities because they must engage in in-person interactions. And then those with medical vulnerabilities have high individual value of vaccination, but may not have health externalities because being at risk leads them to take actions that minimize in-person interaction. Thus our model is a rationale for prioritizing all three of these groups, but for prioritizing healthcare and essential workers the most.

Bubar et al. (2021) used an age-stratified SEIR model to investigate the impact of prioritizing different groups for Covid-19 vaccination. Their results highlighted the value of

prioritizing younger, higher-contact individuals in order to reduce incidence of the disease, but found prioritizing older adults to be more effective at reducing mortality and (often) overall years of life lost (see also Rahmandad (2021)). Our framework can rationalize both of these arguments: the former in terms of health externalities, and the latter in terms of high individual health value of vaccination.[5] Schmidt et al. (2020b) and Bibbins-Domingo et al. (2021) demonstrated how targeting Covid-19 vaccines according to observables such as neighborhood characteristics can help prioritize socially vulnerable populations—in the language of our framework, this corresponds to prioritizing populations with high welfare weight as revealed by their label (see also Schmidt et al. (2021)). Schmidt (2020) and Schmidt et al. (2020a), meanwhile, presented evidence for assigning high welfare weights to disadvantaged populations both because they are generally under resourced and because they face especially high Covid-19 incidence.[6]

Conceptually, the prior work just described focused only on priority mechanisms and observable types, whereas we focus on agents' privately observed health and socio-economic benefits of being vaccinated and ask the question of whether and when rationing can be superior to more market-based mechanism. To our knowledge, Brito et al. (1991) were the first to study a version of this trade-off in the context of vaccination; like us, they highlight the role of market mechanisms in identifying who has highest demand for vaccination, but point out that given health externalities from vaccination, competitive equilibrium allocation may be suboptimal. They then show how to design a tax and subsidy scheme that makes use of the revelation implications of who chooses to be vaccinated to find an especially efficient allocation. Pancs (2020), meanwhile, analyzed a fully market-based solution, modeling the problem of vaccine allocation as a "position auction" (cf. Varian (2007); Edelman et al. (2007)), in which agents can bid for positions in the vaccine queue. Crucially, in the auction proposed by Pancs (2020), agents can also bid on behalf of others; under the VCG payment rule, such an auction is truthful and leads to an efficient allocation of vaccines even in the presence of externalities. However, this mechanism requires each agent to correctly estimate—and then communicate—the value that she places on vaccinating any other agent. Our approach to externalities is different: we assume that the designer can estimate, given a set of observables about an agent, the total externality that vaccinating that agent has on the rest of the society. As a result, our mechanism does not require agents to compute and communicate a multi-dimensional vector of values. Additionally, we allow for heterogeneous

---

[5]Goldstein et al. (2021) argued that vaccinating the oldest individuals also maximizes the years of future life saved—although their model does not include health externalities of vaccination, which significantly increase the importance of vaccinating younger individuals.

[6]Emanuel et al. (2020a), meanwhile, presented a prioritization scheme for cross-country vaccine distribution which, like our framework, considers both health and economic harms.

welfare weights that let us capture redistributive and moral concerns.

In terms of methods, we rely on the tools developed by ADK. Because our model differs from the setting of ADK along a few important dimensions, Appendix A explains how to apply these tools in our context. In their model of allocation under redistributive concerns, ADK built upon and extended a number of important prior contributions; most notably, Weitzman (1977) was first to argue that a market mechanism is not optimal when agents' needs are not well expressed by willingness to pay—an idea fundamental to the trade-offs considered in this paper. Moreover, Condorelli (2012, 2013) provided a comprehensive mechanism design framework that allows for a rich set of objective functions for the designer. Nevertheless, the vaccine prioritization problem requires a set of modeling assumptions whose combination is unique to the setting of ADK: heterogeneous quality (understood here as some vaccines being available earlier than others), flexible preferences of the designer over revenue, and groups of agents with the same observable characteristics.

The inclusion of externalities connects our work to mechanism design with allocative externalities (see, e.g., Jehiel et al. (1996); Jehiel and Moldovanu (2001)). However, we model the externality simply as constant benefit that the designer receives by vaccinating a given agent, and hence this does not lead to the interesting complications associated with strategic interactions between the agents studied in previous work.

And finally, our work also ties in with economic frameworks for allocating scarce Covid-19 tests accounting for transmission externalities (e.g., Deb et al. (2020); Kasy and Teytelboym (2020); Ely et al. (2020a)), as well as research that has sought to give guidance on how to trade off health and economic impacts in pandemic response (e.g., Akbarpour et al. (2020); Budish (2020); Ely et al. (2020b)).

2 Framework

A designer controls the allocation of vaccines to a unit mass of agents. The vaccines become gradually available over time: Let the function $A : [0, \infty) \rightarrow [0, 1]$ describe their availability, where $A(t)$ is interpreted as the cumulative mass of vaccines available at t.

Before receiving a vaccine, each agent privately decides how to react to the pandemic. In the model, each agent takes a binary decision $a \in \{\text{Safe, Risky}\}$. We interpret the choice of $a = \text{Safe}$ as the agent taking precautions that significantly impact the agent's in-person activities in order to minimize the risk of infection (e.g., staying at and working from home, avoiding public transit, and social distancing). The choice of $a = \text{Risky}$, meanwhile, represents the agent choosing to engage in in-person interactions; this can incorporate both work and leisure activities, and the specific activities depend on the agent's type. For example,

for a medical professional, choosing $a =$ Risky might simply represent seeing patients as normal, while $a =$ Safe could mean seeing patients online instead.[7] For a retiree, the decision might be between self-isolating at home (Safe) or seeing their family and friends. For a student, both Safe and Risky may entail some amount of in-person interaction (such as going to class), but Safe represents minimizing that interaction as much as possible (e.g., by choosing not to attend large social gatherings).[8] In practice, the level of precaution is more naturally thought of as a continuous variable but we model it as a binary choice to simplify the analysis; our qualitative conclusions continue to hold if a is chosen from a larger set, as discussed in Section 8. We assume that the decision a is not directly observable but—as we describe soon—we allow for observable information that can perfectly predict the decision for some agents.

Intuitively, the decision between Safe and Risky depends on the comparison between the Covid-related risk that the agent would incur by engaging in in-person interactions and the private disutility that the agent suffers from taking the precautions. The agent's decision may be socially inefficient because she ignores the externalities that the decision creates: When choosing $a =$ Risky, the agent increases the probability of infection for all other agents; when choosing $a =$ Safe, the agent deprives other agents of the benefits of interacting with her in-person (e.g., patients of a doctor working from home experience a decrease in the quality of the service). To capture all these considerations and their interaction with optimal vaccine policy, we decompose the agent's description by separating Covid-related consequences from all other payoff consequences, and by separating private gains from externalities. Specifically, each agent is described by her *characteristics* that we express in dollar values (to ensure that they can be compared to one another):

- v: the private *socio-economic benefit* of choosing $a =$ Risky relative to $a =$ Safe, not including Covid-related risk. That is, under the (hypothetical) assumption that she is not going to contract the virus either way, v is the maximal amount of dollars the agent is willing to pay to engage in in-person interactions relative to taking all precautions. For example, v measures the utility the agent derives from working in-person, going to the gym, seeing friends and family, eating out, etc.

- v_{ex}: the positive *socio-economic externality* generated by the agent choosing $a =$ Risky

---

[7]Note that the example of doctors also highlights how in referring to the Risky action as "risky" we are just referencing inherent risk in the activity; there is no value judgment intended. Moreover, of course, agents in practice may choose different activity patterns in different parts of their lives—for example, many front-line workers in healthcare and other industries (e.g., grocery workers and teachers) are doing jobs that are "risky" while taking maximal precaution in their lives outside of work.

[8]The meaning of Safe and Risky could also depend on the mandatory precautions imposed by authorities. For simplicity, we do not model the interactions between vaccine allocation and other policy tools.

relative to $a =$ Safe, not including Covid-related risk. That is, v_{ex} is the value to society of the agent engaging in in-person interactions under the (hypothetical) assumption that this has no influence on the infection risk for other agents. For example, if the agent is a kindergarten teacher, v_{ex} captures the benefits that children and their parents receive when the teacher chooses to work.

- h: the private *health benefit* of not being at risk of infection. That is, h is the maximal amount of dollars the agent is willing to pay in order to avoid the Covid-related risk to the extent possible. For example, h may depend on the agent-specific risk of infection, presence of potential comorbidities, quality-adjusted life-years (QALYs), etc.

- h_{ex}: the positive *health externality* generated if the agent minimizes the risk of her own infection. That is, h_{ex} is the value to society of the agent not being a potential spreader of the virus. For example, by choosing to work from home, the agent decreases the Covid-related risk for her co-workers.[9]

- λ: a *social welfare weight* measuring how much a dollar of value given to the agent contributes to a social welfare function to be described. All previous values are expressed in dollars, and hence v and h are affected by the agent's opportunity cost of money (that could depend, for example, on income and wealth). The parameter λ converts these dollar values into social values that can be compared *across* individuals. For example, λ allows for redistributive preferences of the designer based on factors such as income, socioeconomic status, belonging to a minority etc.

We consider a few examples next to clarify the meaning of our concepts. **Front-line health workers** have a relatively high h because they are at a high risk of infection if they choose $a =$ Risky (of course, h will vary by age and health status); still, their v is typically even higher because their job, by definition, cannot be done remotely (that is, v captures the fact that they would lose their job if they chose $a =$ Safe). Their externalities h_{ex} and v_{ex} are both large because front-line health workers provide a tremendous value to their patients by seeing them in person. Additionally, the social perception of their sacrifice and heroism may be reflected in high λ. For **ride-share drivers**, the ranking of v and h may depend on whether they have other sources of income; if driving is their main job or if their savings are low, v may be high. As a result, we may expect that poor drivers are more likely to choose $a =$ Risky. Because ride-share drivers are in close contact with many people, their h_{ex} is high; their v_{ex} may be relatively low due to existence of alternative means of transportation,

---

[9]Note that this health externality is measured only with respect to spreading the virus directly. Health impacts through other channels driven by the agent's activity (e.g., if the agent is a doctor treating patients) are incorporated into v_{ex}.

decreased mobility during a pandemic, and higher elasticity of labor supply (than in case of health workers whose short-run supply is almost completely inelastic). A **healthy college student** is an example of someone with low h but high h_{ex}; young healthy people are unlikely to suffer serious consequences of an infection but they may play a role in transmitting the virus to more vulnerable populations, especially if their social network is wide. Finally, a **software engineer with potential comorbidities** will typically have h higher than v and a relatively low v_{ex} because their job can be performed effectively out of home.

In practice, agent characteristics are partially observable. For example, an individual's job may be observed, and it will reveal some information about the characteristics (as argued above); yet, factors such as attitudes, beliefs, and lifestyle may be agents' private information. To obtain a compact description of observability, we assume that the designer observes a label i for each agent, and knows the joint distribution of characteristics, that is, she can form a belief about $(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda)$ conditional on observing i (in particular, it is possible that i perfectly reveals some characteristics). The label i belongs to a finite set I that captures all observable features of agents that the designer can condition her allocation on. We refer to all agents with the same label i as *group i*. An example of a label i could be "a doctor, below 60 years old, with no underlying health conditions."

We make two additional assumptions. First, the agent can at least observe v and h—her private benefits. (It does not matter for our analysis whether the agent observes anything else.) Second, the externalities v_{ex} and h_{ex} are independent of v and h conditional on i. The latter condition states that, conditional on observable information, each of the two externalities generated by the agent have no systematic relationship to her private benefits. This assumption is most likely violated to some extent in practice; but it underscores the point that decisions taken by privately-optimizing agents will not in general be aligned with the social objective (that will include the externalities).[10]

Next, we proceed to specifying the payoffs. We make a strong assumption that the health consequences of receiving a vaccine are the same as those of choosing $a = \text{Safe}$. As a result, every vaccinated individual enjoys utility $v + h$. This assumption could be false for many reasons;[11] we nevertheless make it for simplicity and because it seems to capture the gist of the problem. In the absence of a vaccine, the agent compares v and h to determine her action: She chooses $a = \text{Safe}$ if $h > v$; otherwise, she chooses $a = \text{Risky}$. The agent ignores her externalities when making that decision. For now, we assume that both h and v are non-

---

[10]The assumption could be false, for example, when agents care directly about the health or welfare of others; in this case, we would expect positive correlation of h with h_{ex} and v with v_{ex}. Our methods could easily handle this more general case but the interpretation of the model would be less transparent.

[11]Depending on an individual, either a vaccine or taking maximal precautions could provide a higher level of protection agaists the virus.

negative; we relax that assumption in Section 8, where we explain how we could handle the case in which the designer might choose to pay the agents with negative h to get vaccinated. To reduce the number of cases to consider, we assume that in each group i there is at least one agent with $h = 0$.[12] This can be justified if some agents either do not believe that the vaccine is effective or recently had Covid (previous infection is believed to provide some level of immunity). Each agent's utility is quasi-linear in a monetary payment p. We assume that when an agent receives a vaccine at time t, she enjoys its benefits for a fraction $\delta(t)$ of the total duration of the pandemic, where δ is a strictly decreasing function with $\delta(0) = 1$.

$$\delta(t) \underbrace{[v + h]}_{\text{post-vaccination utility}} + (1 - \delta(t)) \underbrace{[\max\{v,\, h\}]}_{\text{pre-vaccination utility}} - p \qquad (2.1)$$

Let $V(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda, t, p)$ be the benefit to the designer of vaccinating an agent with characteristics $(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda)$ at time t and at a price p. The designer then maximizes the expectation of this function with respect to the population distribution of types, with t and p specified by the chosen mechanism. Let $\mathbf{1}_{\text{Risky}}$ and $\mathbf{1}_{\text{Safe}}$ denote the event that an agent chooses $a = \text{Risky}$ and $a = \text{Safe}$, respectively, prior to receiving a vaccine. In our baseline scenario, we focus on a utilitarian objective function (in Section 8, we discuss an alternative objective function and how our results would change):

$$V(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda, t, p) := \delta(t) \left(\mathbf{1}_{\text{Safe}}(\lambda(v - p) + v_{\text{ex}}) + \mathbf{1}_{\text{Risky}}(\lambda(h - p) + h_{\text{ex}}) \right) + \alpha p. \quad (2.2)$$

If a given agent chose $a = \text{Safe}$ prior to receiving the vaccine, then giving the vaccine to that agent unleashes the private benefit v and a positive externality v_{ex}; if, however, the agent chose $a = \text{Risky}$ prior to receiving the vaccine, then giving the vaccine to that agents unleashes the private health benefit h and a positive health externality h_{ex}.[13]

Additionally, the designer places a weight $\alpha \geq 0$ on revenue generated by the mechanism. In practice, α is determined by how the designer uses the monetary surplus. If revenue subsidizes the federal budget or is given back to agents as a lump-sum transfer, then the most natural specification is for α to be equal to the average social welfare weight. If revenue is used to finance free vaccines to poorer communities (in the presence of an implicit budget constraint), then α could be above the average welfare weight. The weight α could be 0

---

[12]Formally, conditional on any $i \in I$, the lower bound of the support of the distribution of h is 0.

[13]Note that we implicitly assumed that the average social welfare weight in the population is 1, and that the externality of the agent is distributed uniformly across all other agents—this justifies why the coefficients on v_{ex} and h_{ex} in the objective function (2.2) are 1. In practice, some agents might exert a stronger externality on a subset of the population (e.g., doctors' externality on their patients)—we could capture this by introducing another dimension of the agents' type at the cost of further complicating our model.

under an alternative interpretation of our model in which agents "pay" for the vaccine by burning utility, e.g., by queueing (in that case, p is interpreted as the time spent in line required to obtain the vaccine).

3 Allocation mechanisms

In many countries, Covid-19 vaccines have been so far allocated using a simple priority schedule. The population is divided into several groups based on observable and verifiable criteria (what we refer to as "labels" in our model). These groups are ordered from most to least critical, and vaccines are allocated for free to agents within each group, with more critical groups receiving vaccines earlier. Apart from the composition and ordering of the groups, important policy debates pertain to whether there should be overlap in the vaccination schedule of various groups. It has also been proposed that using prices could lead to a more efficient allocation.

In our framework, we allow the designer to optimize over all feasible allocation mechanisms.[14] By the Revelation Principle, for the sake of finding the optimal mechanism, we may imagine that the designer asks agents to report their characteristics $(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda)$ subject to incentive-compatibility constraints. Because the designer observes $i \in I$ for each agent, these incentive constraints are imposed only on the support of $(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda)$ conditional on i. Each agent must receive a non-negative utility from participating. As a function of the report, the agent is promised a (potentially random) time of vaccination, and is charged a payment. The mechanism must respect physical feasibility constraints, in that it cannot allocate more vaccines before time t than the availability $A(t)$, for any t. We also assume that all vaccines must be allocated as soon as they become available,[15] and that prices set by the mechanism are non-negative.[16]

Using the fact that the time of receiving a vaccine only matters for payoffs via $\delta(t)$, we will rephrase our model with $q = \delta(t)$ referred to as the *quality* $q \in [0, 1]$ of the vaccine. That is, the highest-quality vaccine $q = 1$ is available immediately, while the lowest-quality vaccine $q = 0$ becomes available when it no longer has any value (alternatively, $q = 0$ can be interpreted as not getting a vaccine at all). Given the availability schedule A, we can define the corresponding distribution F of quality q that the designer distributes among agents. (The parameter q could also capture additional dimensions of quality, e.g., the effectiveness

---

[14]We omit the formal definitions; see ADK for the technical details.

[15]This seems to be the right assumption given the application; moreover, it is without loss of generality as long as the value of vaccinating every agent is non-negative.

[16]Relaxing this assumption might be reasonable when agents can have negative h—see Section 8 for a discussion.

of the vaccine in preventing an infection, or its side effects.) From now on, we will treat F and q as primitives of our model but we will interchangeably use the "time interpretation."

The fundamental difficulty facing the designer choosing an allocation mechanism is that the parameters entering the objective function (2.2) are not directly observable. Therefore, the designer must rely on information that is observable—the labels i—as well as on information that can be elicited through the mechanism itself. The first observation is that an incentive-compatible mechanism with transfers can only elicit information about agents' willingness to pay derived from their primitive private types.[17]

Lemma 1. *It is optimal for the designer to condition the allocation of vaccines only on labels* i *and willingness to pay* r, *where*

$$r = \min\{v, h\}.$$

Moreover, if the designer is constrained not to use prices, then it is optimal to condition the allocation of vaccines only on labels i.

Lemma 1 is intuitive: Equation (2.1) reveals that the agent's private value for getting vaccinated is

$$(v + h) - \max\{v, h\} = \min\{v, h\} = r.$$

In other words, r is the maximal price that the agent is willing to pay for receiving a vaccine immediately. Two agents with the same label and willingness to pay are behaviorallly indistinguishable with regards to any mechanism with prices. If the mechanism attempted to condition the allocation on additional dimensions of the type (e.g., on the unobserved h_{ex} or v_{ex}), each agent would simply report characteristics associated with the most preferential treatment by the mechanism, and the effective allocation would not vary with these dimensions (conditional on i and r). Hence, the designer might as well focus on mechanisms in which the allocation only depends on i and r.

The key consequence of Lemma 1 is that what matters for determining the optimal vaccine allocation is the expected benefit that the designer gets by vaccinating an agent with label i and WTP r. Given the linearity of payoffs in vaccine quality $q = \delta(t)$, the mechanism must only specify the expected quality allocated to an agent with WTP r in group i that we will denote $Q_i(r)$. Under our assumption that prices are non-negative, and that the lower bound on r is 0 in each group, the price p paid by type r of label i in an incentive-compatible mechanism is uniquely pinned down, given any allocation Q_i.[18] Overall, a sufficient statistic to evaluate the objective (2.2) under an incentive-compatible mechanism is $Q_i(r)V_i(r)$, where

---

[17]See Jehiel and Moldovanu (2001), Che et al. (2013), and Dworczak ⓣ Kominers ⓣ Akbarpour (2020) for proofs of closely related results.

[18]This follows from the payoff equivalence theorem, see for example Milgrom (2001).

$V_i(r)$ is the expected per-unit-of-quality social benefit from allocating a vaccine to an agent with WTP r in group i. Under regularity conditions, we can compute $V_i(r)$ explicitly.

Suppose that WTP has a continuous distribution conditional on i, fully supported on $[0, \bar{r}_i]$; let G_i be its CDF, and let γ_i be its inverse hazard rate.[19]

Lemma 2. *The expected per-unit-of-quality social benefit from allocating a vaccine to an agent with WTP r in group i in an incentive-compatible mechanism is given by*

$$V_i(r) = \underbrace{\Lambda_i(r) \cdot \gamma_i(r)}_{private\ utility} + \underbrace{\alpha(r - \gamma_i(r))}_{revenue} + \underbrace{v_{ex}^i \cdot \mathbb{P}(a = Safe|i,\ r) + h_{ex}^i \cdot \mathbb{P}(a = Risky|i,\ r)}_{externality}. \quad (3.1)$$

In the above, $\Lambda_i(\tau) = \mathbb{E}[\lambda|i, r \geq \tau]$ is the expected welfare weight on all agents in group i with WTP above τ, and $v_{ex}^i := \mathbb{E}[v_{ex}|i]$ and $h_{ex}^i = \mathbb{E}[h_{ex}|i]$ are the expected externalities conditional on the label i.

The first component of $V_i(r)$ is the private-utility term that consists of the inverse hazard rate of WTP (which measures information rents) multiplied by $\Lambda_i(r)$ which is the best estimate—given the designer's information—of the welfare weight placed on agents with WTP above r. Intuitively, in an incentive-compatible mechanism, changing the utility of type r has consequences for the utility of all higher types, and hence these payoff consequences must be properly weighted. Since the true weights λ are not observable, the designer can only infer them based on i and r. The second component in the objective function is the usual virtual surplus term that captures revenue maximization. The last component is the externality term, where $v_{ex}^i := \mathbb{E}[v_{ex}|i]$ and $h_{ex}^i := \mathbb{E}[h_{ex}|i]$ are the best estimates of externalities conditional on the label i. By our earlier assumption, v_{ex} and h_{ex} are independent of v and h, and thus also independent of r (hence, we do not need to condition on r to find the best estimates of the externalities). The key part of the objective is the estimation $\mathbb{P}(a = \text{Risky}|i, r)$ of the probability of the unobserved event that the agent chooses $a = \text{Risky}$ prior to receiving the vaccine. This is intuitive: the higher the probability that the agent chooses $a = \text{Risky}$, the higher the relative weight on the health externality h_{ex}^i unleashed by vaccinating this agent, and the lower the weight on the socio-economic externality v_{ex}^i.

By using monetary transfers, the designer can elicit information about willingness to pay when allocating vaccines. If, however, prices are set to zero, the allocation may no longer depend on WTP, and only the label i can be used (Lemma 1). In that case, the relevant statistic that determines the optimal allocation is the expected per-unit-of-quality social benefit from allocating a vaccine to a *random* agent in group i.

---

[19]We impose an additional regularity condition that is not needed for this result but will be used later that the inverse hazard rate is continuous and equal to 0 at the upper bound $\bar{r}_i$ for each i.

Lemma 3. *The expected per-unit-of-quality social benefit from allocating a vaccine to a random agent in group i is given by*

$$\overline{V}_i := \mathbb{E}[V_i(r)|\, i] = \mathbb{E}[\lambda \cdot r |\, i] + h_{ex}^i \cdot \mathbb{P}(a = Risky|\, i) + v_{ex}^i \cdot \mathbb{P}(a = Safe|\, i). \qquad (3.2)$$

Note that the private-utility term reduces to the expectation of λr since agents receive the vaccine without a payment. For the same reason, the revenue term drops out. The relevant externality benefit depends on the size of the two externalities h_{ex}^i and v_{ex}^i in group i, and which behavior (Safe versus Risky) is more likely given the label i.

4   Optimal Priority without Prices

We first solve the problem assuming that the designer does not charge monetary transfers for the vaccines. This has been the dominant practice in most countries. In the next section, we ask how introducing prices alters the optimal mechanism. We refer to the zero-price allocation within each group as *free allocation*. If F_i is the cdf of vaccine quality allocated to group i, free allocation means that $Q_i(r) = \int_0^1 q dF_i(q)$, so that the expected time of receiving the vaccine is the same for each agent with label i. In particular, free allocation involves *rationing* if there are not sufficiently many vaccines for everyone within a group, and *randomization* if the there is dispersion in quality (timing) of vaccines available for a group.

Our first result is that when the designer does not use prices, it is optimal to vaccinate groups one by one with no overlaps, with the order determined by the sufficient statistic from Lemma 3.

Result 1. *Suppose that the allocation within each group is free. Then, it is optimal to vaccinate groups sequentially in the order of decreasing $\overline{V}_i$. That is, if $\overline{V}_j > \overline{V}_k$, then all agents in group j are vaccinated before all agents in group k.*

The intuition for Proposition 1 is straightforward. Under free allocation, every vaccinated agent with the same label has the same expected contribution to the social objective function because the order of vaccination within a group is random. Thus, there is no reason to alternate between two groups: Instead, the designer always obtains a higher marginal value from vaccinating an agent from group i with higher $\overline{V}_i$.

The form of $\overline{V}_i$ predicted by Lemma 3 reveals the determinants of priority under free allocation. First, priority is given to groups for which the label reveals high welfare weights λ and high willingness to pay. High welfare weights could be attached, for example, to agents

who are poor, particularly adversely affected by the pandemic, or those playing a key role in fighting the pandemic. Second, priority depends on the expected externality revealed by the label. Crucially, which externality benefit (h_{ex}^i or v_{ex}^i) is relevant depends on what the label reveals about the expected behavior a of agents in the group.

For illustration, consider group i to consist of front-line doctors and nurses. Because members of this group are at risk precisely because they are providing front-line care, it is natural to assume that society attaches a high weight λ to agents in that group (see, for example, Emanuel et al., 2020b). This label is also associated with a high health externality h_{ex}^i—which is the relevant externality because these individuals are engaging with Covid-19 patients directly ($\mathbb{P}(a = \text{Risky}|i) \approx 1$). Thus, Result 1 predicts that front-line health workers should receive the vaccines early on. If there are no groups j with higher $\overline{V}_j$, then all front-line health workers should be vaccinated before vaccines are made available to any other group.

For a different application of Result 1, consider the problem of whether priority should be given to group j consisting of people who are at high risk in case of infection (e.g., the elderly) or to group k of people who are most likely to spread the virus (e.g., students living in dormitories). In group j, the benefit h is high by definition, and so r is relatively high as well (at least on average). In contrast, since $r = \min\{v, h\}$ and h is low for most young, healthy individuals, r is typically low in group k. To simplify, let us approximate

$$\mathbb{E}[\lambda \cdot r \,|\, k] \approx 0;$$
$$\mathbb{P}(a = \text{Risky}|j) \approx 0; \text{ and}$$
$$\mathbb{P}(a = \text{Risky}|k) \approx 1.$$

Then, group j has priority over k if and only if $\mathbb{E}[\lambda \cdot r \,|\, j] + v_{\text{ex}}^j > h_{\text{ex}}^k$. Thus, group j should receive the vaccines earlier if their average welfare-weighted WTP plus the socio-economic externality exceeds the health externality of group k. For instance, for elderly people living in nursing homes, v_{ex}^j captures the value of family members being able to visit them. At the same time, the health externality h_{ex}^k could be relatively low, if students living in dormitories interact mostly with other young healthy individuals. Thus, the utilitarian objective may naturally support prioritizing the elderly (and others at high risk) over students (and others who interact primarily with people with low risk of serious illness). In contrast, if k is the group of public transit drivers (or drivers of ride-sharing platforms), then k may be associated with a larger health externality h_{ex} because those drivers interact with many riders of all ages; this could potentially lead them to have a higher priority than some high-risk individuals.

19

5   Optimal Priority with Prices

In this section, we describe the optimal priority when the designer can use prices. The main difference to the case of free allocation is that the designer can now screen based on WTP, and hence the marginal social benefit of vaccinating an agent from group i may vary with r (see Lemma 2). To screen, the designer charges higher prices for higher-quality vaccines (that are available earlier) ensuring assortative matching between WTP and quality within a group. We will refer to this method as a *market allocation* since it coincides with what a competitive market would achieve under the assumption of group-specific market clearing. Formally, if F_i is the cdf of vaccine quality allocated to group i, a market allocation means that $Q_i(r) = F_i^{-1}(G_i(r))$, where F_i^{-1} is the generalized inverse of the cdf F_i.

The first question we ask is whether it might still be optimal to vaccinate some groups of agents immediately and for free (which we call *priority allocation*), even when monetary transfers are feasible. Let μ_i be the mass of agents in group i, and recall that $A(0)$ is the mass of vaccines available immediately.

Result 2. *Suppose that $A(0) \geq \sum_{j \in J} \mu_j$. Then, it is optimal for groups $J \subset I$ to receive priority allocation (all agents with $i \in J$ receive a vaccine immediately and for free) if*

$$\min_{j \in J,\, x} \left\{ \mathbb{E}[V_j(r) \,|\, r \leq x] \right\} \geq \max_{i \notin J,\, x} \left\{ \mathbb{E}[V_i(r) \,|\, r \geq x] \right\}.$$

Moreover, this condition is necessary when $A(0) = \sum_{j \in J} \mu_j$, that is, when there are exactly enough vaccines for groups J at time 0.

Result 2 states that all the agents in groups J get priority if (1) the designer has enough vaccines to vaccinate all agents in groups J immediately, and (2) the minimal marginal value of vaccinating an agent belonging to J is higher than the highest marginal value the designer could obtain from any agent outside of J. To understand the exact form of the second condition, imagine a situation in which all agents in J are vaccinated, and none of the agents in $I \setminus J$ are vaccinated. Then, under the binding capacity constraint, optimality requires that the designer cannot benefit from taking away one vaccine from groups J and allocating it in the best possible way to groups $I \setminus J$. The "best possible way" of allocating the vaccine takes into account incentive constraints; for example, when group i has no vaccines, and the designer wants to allocate a single vaccine to that group, she can allocate it to the highest-WTP type $\bar{r}_i$ by simply setting a price equal to $\bar{r}_i$; however, if she wants to allocate it to some type $r < \bar{r}_i$, the best she can do is to set a price r and ration uniformly at random (this maximizes the probability that r gets that single vaccine among all incentive-compatible mechanisms). Thus, the maximal marginal value from allocating a single vaccine

to groups $I \setminus J$ is equal to the maximum over $i \notin J$ and all incentive-compatible lotteries that the the designer could use to allocate that vaccine. Similarly, the marginal cost of taking away one vaccine from groups J can be found as the minimum over $i \in J$ and all lotteries such that "subtracting" that lottery from the optimal mechanism still results in an incentive-compatible mechanism.

To identify interpretable conditions for some groups to receive priority allocation, let

$$T_{\text{ex}}^i(r) := v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe}|i, r) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky}|i, r)$$

denote the total externality in group i as a function of r. A simple calculation shows that

$$\mathbb{E}[V_j(r)|\, r \leq x] = \mathbb{E}[\lambda r |\, j, r \leq x] + [\Lambda_j(x) - \alpha]x\frac{1 - G_j(x)}{G_j(x)} + \mathbb{E}(T_{\text{ex}}^j(r)|j, r \leq x), \qquad (5.1)$$

and

$$\mathbb{E}[V_i(r)|\, r \geq x] = \mathbb{E}[\lambda r |\, i, r \geq x] + [\alpha - \Lambda_i(x)]x + \mathbb{E}(T_{\text{ex}}^i(r)|i, r \geq x). \qquad (5.2)$$

By Result 2, for groups J to receive priority allocation it must be that the value of (5.1) is uniformly higher (over $j \in J$ and x) than the value of (5.2) (over $i \notin J$ and x). Both (5.1) and (5.2) consist of three terms capturing the welfare effects of taking one vaccine from group j (by decreasing the allocation probability uniformly for types $r \leq x$) and allocating it to group i (using a uniform lottery over types $r \geq x$). The first term quantifies the social value of the resulting change in the private utility, excluding payments. The second term quantifies the social value of the change in payments: the direction of this effect depends on the ranking of the average welfare weights $\Lambda_i(x)$ and the weight on revenue α (note that $x(1 - G_j(x))$ in (5.1) is the increase in revenue gathered from types above x when the allocation probability of types below x decreases; x in (5.2) is the price charged to implement the lottery in which types above x receive the vaccine). The third term quantifies the social value of the change in the expected externality for a group.

Based on the above discussion and Result 2, providing priority allocation to groups J is more likely to be optimal when (1) these groups are associated with high welfare weights λ, (2) the designer is not too concerned about revenue (α is relatively low), and (3) groups J have high externality. This has a few implications. First, although a high welfare weight λ raises the value of (5.1), it is never a sufficient force on its own: This is because $\mathbb{E}[\lambda r |\, j, r \leq x]$ is 0 when $x = 0$, reflecting our assumption that in each group there are some individuals with low WTP. This is intuitive: The welfare weight has bite only when an agent gets a strictly positive utility from vaccination. Second, a low weight on revenue is needed because the designer has the option to sell vaccines to high-WTP agents in non-prioritized groups.

21

Indeed, (5.2) is lower bounded by $\alpha \max_{i \notin J} \bar{r}_i$, where $\bar{r}_i$ could be at the order of thousands or even millions of dollars if there are very wealthy individuals. Third, the externality term is likely the most significant potential contribution to (5.1) being high *uniformly* over x: It suffices that the label j is highly predictive of $a =$ Safe and v_{ex}^j is high, or that the label j is highly predictive of $a =$ Risky and h_{ex}^j is high.

Consider j to be the group of front-line doctors and nurses. As we already argued in Section 4, this group is likely to be associated with high welfare weights λ, and a high health externality h_{ex}^j. Moreover, because this label reveals that $a =$ Risky with high probability (by definition, these agents work directly with Covid-19 patients), we can think of $\mathbb{P}(a =$ Risky$|j, r)$ as being approximately 1 (in particular, almost constant in r). Therefore, the assumptions of Result 2 are likely to hold—indicating that the entire group j should be prioritized—unless the designer places a high weight α on revenue.

If the designer does place a high weight α on revenue (which could be the case in a developing country that can buy more vaccines overall if it raises more revenue), the conclusion must be modified. The designer could benefit from selling early access to vaccines to wealthy people with high WTP. We formalize this in the following result.

Result 3. *Suppose that it is optimal to use a market allocation within group i and a free allocation within group j. If $V_i(\bar{r}_i) > \overline{V}_j > V_i(0)$, then it is optimal to start vaccinating agents in group i first, then to vaccinate all agents in group j, and then to vaccinate the remaining agents in group i.*

The intuition for Result 3 is straightforward. Under free (random) allocation, every vaccinated agent has the same expected contribution to the social objective function. In contrast, when a market allocation is optimal, the most "valuable" agents within a group are vaccinated first. Thus, for any group with free allocation, once it is optimal to start vaccinating that group, all agents in the group should receive the vaccine before proceeding to any other group. In contrast, for any group with a market allocation, the schedule could be more spread out, with the possibility of simultaneous vaccination with another market-allocation group as well as a "pause" during which some free-allocation group receives the vaccines. Of course, Result 3 is incomplete in that it does not specify *when* it is optimal to use a market versus free allocation within groups—we return to this issue in the next section.

When the weight on revenue α is high, Result 3 could apply: The designer first offers vaccines at high prices to the general population. Then, high-externality groups (e.g., doctors and nurses) are vaccinated free of charge. Finally, the vaccines are again allocated using a market mechanism, with prices gradually decreasing over time. We can even determine the threshold price at which the first stage should stop: If J denotes the high-externality groups,

then that price $p^\star$ is determined by

$$V_{I\setminus J}(p^\star) = \mathbb{E}[V_J(r)], \tag{5.3}$$

assuming that $V_{I\setminus J}(r) = \sum_{i\notin J} V_i(r)$ is non-decreasing. The left hand-side of (5.3) can be approximated by

$$V_{I\setminus J}(p^\star) \approx \mathbb{E}[\lambda|r = p^\star,\, i \notin J]p^\star + (\alpha - \mathbb{E}[\lambda|r = p^\star,\, i \notin J])\, p^\star = \alpha p^\star$$

if $p^\star$ is close to the maximal WTP, since the externality term is not too high by definition (the private-utility term cancels out because type $p^\star$ can be charged approximately its WTP if it is close to the maximal WTP). The right-hand side of (5.3) is equal to

$$\mathbb{E}[V_J(r)] = \sum_{j\in J} \left\{ \mathbb{E}[\lambda r|\, j] + \mathbb{E}[T_{\text{ex}}^j(r)|j] \right\},$$

which is a special case of (5.1) with x set to the maximal WTP (the revenue term drops out because the price for these groups is zero). If we set α to be equal to the average welfare weight, then we conclude that the price in the early-access stage should be

$$p^\star \approx \sum_{j\in J} \left\{ \mathbb{E}[\lambda r|\, j] + \mathbb{E}[T_{\text{ex}}^j(r)|j] \right\}.$$

Since both the welfare weight λ and the expected externality in groups including doctors and nurses are high, $p^\star$ should be much higher than the average WTP of doctors and nurses, and higher than the social value of the expected externality, probably placing the estimate in the range of hundreds of thousands of dollars, if not more.

The policy of selling vaccines early to "millionaires" does not seem to be popular in practice. A potential explanation is that the welfare weight on millionaires is zero. However, our framework shows that this is not enough: The derivation is unchanged even if $\mathbb{E}[\lambda|r = p^\star] = 0$. Another explanation is that the weight on revenue α is low, at least in developed countries. However, unless the weight is 0 (which seems unlikely), $\alpha \bar{r}_i$ could still be large. The most likely explanation, in our view, is that such a policy would have some degree of "repugnance."[20]

In our discussion above, we described the last stage as using a market mechanism to allocate remaining vaccines to the "rest" of the population. However, even when using pricing, the designer could still rely on labels to guide the allocation. The next result casts

---

[20]See Roth (2007) for a discussion of repugnance as a constraint on markets.

light on optimal priority across groups when a market allocation is used.

Result 4. *Suppose that it is optimal to use a market allocation within groups i and j. Then, it is optimal to vaccinate group i before group j if and only if*

$$V_i(0) \geq V_j(\bar{r}_j).$$

If instead we have $V_i(\bar{r}_i) > V_j(\bar{r}_j) > V_i(0) > V_j(0)$, then it is optimal to start vaccinating agents in group i first, then to vaccinate agents in both groups for some time, and then to vaccinate the remaining agents in group j.

Result 4 stays in sharp contrast to Result 1. Under market allocation, there is in general significant overlap in the vaccination times for various groups. This is because careful pricing selects the agents with the highest value to be vaccinated earlier, and thus the marginal social value of allocating a vaccine to a certain group varies with how many agents in that group have been vaccinated already. For example, if $V_i(0) = V_j(0)$ and $V_i(\bar{r}_i) = V_j(\bar{r}_j)$, then groups i and j are vaccinated simultaneously. Nevertheless, this priority schedule requires prices to vary with the group identity: For example, if agents in group i have on average a higher WTP than agents in group j , then simultaneous vaccination can only be achieved if agents in group i face higher prices for the vaccines.

6 Market versus Free Allocation

In the previous section, we took as given the optimality of the allocation method within groups, and we focused on the implications for determining optimal priority. We now return to the issue of the optimal allocation within groups. Taking as given the pool of vaccines allocated to a given group i, the designer could allocate all these vaccines at a price of 0 (free allocation) thus ensuring that no one has to pay for it but not being able to screen based on WTP. Alternatively, the designer could charge higher prices for higher-quality vaccines ensuring assortative matching between WTP and quality. There is also a host of hybrid mechanisms that combine randomized allocation with assortative matching for various intervals of WTP.

Result 5. *If $V_i(r)$ is non-decreasing, it is optimal to use a market allocation within group i. If $V_i(r)$ is non-increasing, it is optimal to use a free allocation within group i. In all other cases, a hybrid mechanism is optimal.[21]*

---

[21] A closely related result, albeit in a different setting and under a different objective function, is established by Condorelli (2013).

The intuition for the result is simple: A market allocation achieves an assortative matching between WTP and vaccine quality. Thus, such an allocation is optimal when higher-WTP agents contribute more to the social objective function. When it is the lower-WTP agents that contribute more, the first-best allocation would induce an anti-assortative matching; that, however, is not possible given the screening devices that the designer possesses. The best she can do in that case is to induce zero correlation between WTP and vaccine quality which is achieved by having a free allocation (with uniform rationing).

We will not pursue optimality of various hybrid mechanisms here.[22] Instead, we are interested in identifying distinct economic forces that work in favor of free (random) versus market (assortative) allocation. Monotonicity of $V_i(r)$ is determined by two distinct forces (see (3.1) in Lemma 2):

1. **Private utilities+revenue.** The term $\Lambda_i(r)\gamma_i(r) + \alpha(r - \gamma_i(r))$ is the basic ingredient of the welfare function analyzed in ADK. We summarize the main findings. Suppose first that the designer does not have redistributive concerns or WTP does not reveal any inequality in welfare weights ($\Lambda_i(r)$ is constant in r). Then, we can easily recognize some familiar cases. In the transferable-utility case, it would be customary to set α to be equal to the average welfare weight within group i (revenue is internally redistributed as a lump-sum payment), and then the private-utility+revenue term reduces to r and is hence trivially increasing regardless of distributional assumptions. This scenario corresponds to the core economic intuition that markets are "efficient"—they maximize total WTP. If instead α is 0, then we are in the scenario of "costly screening" or "money burning" that has also been extensively studied in the literature[23]—the optimal allocation depends on the monotonicity of the inverse hazard rate γ_i. Since the inverse hazard rate is decreasing for many commonly used distributions, this case typically results in a free allocation. This is intuitive: If raising revenue has no value for the designer, then any positive price charged to an agent is a pure social loss. Economically, this case is relevant when implicit prices are non-monetary, e.g., agents have to stay in line to obtain the vaccine—we revisit this scenario in Section 8. A more general conclusion is that—under regular distributions—a market allocation is more likely to be optimal when α is higher (relative to the average welfare weight for a given group). Finally, let us consider the effect of redistributive preferences. If the designer prefers to redistribute towards poorer agents, $\Lambda_i(r)$ could be decreasing. This is because we might expect a positive correlation between wealth and willingness to pay, everything else

---

[22]Computing the optimal hybrid mechanism is not mathematically difficult (see ADK for details) but it is not clear whether such complicated mechanisms could be adopted in practice for vaccine allocation.

[23]See, for example, Hartline and Roughgarden (2008), Condorelli (2012), and Chakravarty and Kaplan (2013).

being equal. If this effect is strong enough, it could make the private-utility+revenue term decreasing, leading to a free allocation. Apart from the strength of the primitive redistributive preferences of the designer (as expressed by the dispersion in weights λ), the key determinant of the steepness of $\Lambda_i(r)$ is the informativeness of the label i. To see that, note that $\Lambda_i(r)$ captures the *residual* correlation between WTP and welfare weights, conditional on i. If the label i defines a relatively narrow homogeneous group (e.g., doctors of a certain specialty and certain age), then $\Lambda_i(r)$ is unlikely to vary a lot with r. In contrast, if i describes a highly heterogeneous group (e.g., all people below 65 years of age, excluding front-line health workers), then WTP can pick up a large part of the variability in welfare weights. Summarizing, if the designer has strong redistributive preferences, she could opt for free allocation if labels fail to accurately identify those with the highest welfare weights.

2. **Externalities.** The externality term $v_{\text{ex}}^i + (h_{\text{ex}}^i - v_{\text{ex}}^i) \cdot \mathbb{P}(a = \text{Risky}|i, r)$ is novel to this paper. For groups i with high externalities (doctors, nurses, teachers etc.), this will likely be the dominating term. The monotonicity of this term depends on two factors: (1) which externality effect, h_{ex}^i or v_{ex}^i, is stronger for group i, and (2) whether $\mathbb{P}(a = \text{Risky}|i, r)$ is increasing or decreasing in r. Higher r reflects a higher need for a vaccine. However, this higher need could be both associated with the desire to engage in in-person interactions for agents who chose $a = \text{Safe}$, as well as the desire to protect one's health for agents who chose $a = \text{Risky}$. Thus, ex-ante, it is unclear in which direction the effect should go. Higher r could also reflect higher wealth, everything else fixed. Thus, it seems reasonable to assume that $\mathbb{P}(a = \text{Risky}|i, r)$ is decreasing in r, reflecting the fact that wealthier agents can either "afford" to stay at home due to savings, or have jobs that are easier to perform remotely. Then, if the health externality h_{ex}^i is larger than the socio-economic externality v_{ex}^i for group i, the externality term is decreasing, and thus free allocation is optimal for group i. But if the socio-economic externality is higher, market allocation is optimal.

We consider three examples to illustrate the above discussion of optimal within-group allocation. First, let i be the group of ride-share drivers, cashiers, or other front-line workers. As argued before, such groups have a particularly high health externality h_{ex}^i that could easily dominate private-utility and revenue considerations (at least for relatively young and healthy individuals) and could also dominate the socio-economic externality v_{ex}^i. Thus, the designer would like to target the vaccines toward workers who are most likely to choose $a = \text{Risky}$ since vaccinating these agents yields the highest expected externality gain. If poorer workers are more likely to continue working, this could justify providing the vaccines

for free (and rationing if necessary) since a market price would tilt the allocation towards richer workers—the ones who are more likely to choose Safe.

For the second example, let i be the group of small and medium enterprise owners. Their decision a may be whether they temporarily close down their business, which directly affects their employees—thus, their v_{ex}^i may be high. At the same time, for many businesses, h_{ex}^i may be relatively low if the employees are relatively low risk and mostly interact with one another. In such cases, it may be socially efficient to keep the business open even if it is privately optimal for the owner to suspend operations. Thus, the designer wants to target the vaccines towards business owners (and their employees) who are most likely to choose $a = $ Safe. If business owners who have larger savings are more likely to stay at home, then it becomes optimal to use a market allocation (under the same assumption that there is a positive correlation between wealth and WTP).

Finally, imagine that i describes the group of "all remaining agents" once all high-priority groups have been vaccinated. Since externalities may play a smaller role, the key distinction will now be whether the designer is concerned about revenue or not. If the weight α is relatively high, a market allocation will typically be optimal due to its revenue-maximization and efficient-allocation properties. However, if the designer can identify a subgroup j for which the average welfare weight is far above the weight on revenue (e.g., individuals living in a poor neighborhood), then a free allocation may be preferred for j.

As was already argued in Sections 4 and 5, the optimal allocation within each group influences the optimal priority across the groups. Suppose that $V_i(r)$ is non-decreasing for each $i \in I$. Then, the optimal mechanism is a "tiered market allocation." That is, the designer allocates a pool of vaccines to each group, and then a market price guides the allocation within each group by clearing the group-specific market (independently of all the other groups). Of course, in line with Result 4, groups with higher contributions to the social objective function receive more vaccines early on. Thus, market-clearing prices within such groups are generally lower than in other groups. The condition in Result 2 for priority allocation to groups J simplifies in this special case to $\min_{i \in J} V_i(0) \geq \max_{i \notin J} V_i(\bar{r}_i)$. However, in general, there is significant overlap across groups: Agents with high WTP from groups with low externality may receive a vaccine before agents with low WTP from groups with high externality. Suppose instead that $V_i(r)$ is non-increasing for each $i \in I$. Then, it is optimal not to use prices, and the optimal mechanism reduces to the sequential free allocation from Section 4: Groups i are ordered from highest to lowest $\overline{V}_i$ and vaccinated sequentially. (This is consistent with Result 2 whose second condition reduces to $\min_{i \in J}\{\mathbb{E}[V_i(r)]\} \geq \max_{i \notin J}\{\mathbb{E}[V_i(r)]\}$ in that case.) The meaning of free allocation depends on the ranking of the group in the order: "Early" groups receive priority allocation (are vaccinated immediately)

while "later" groups may be heavily rationed.

Illustrative example. Suppose there are three labels and the distribution of marginal values of vaccination in groups is such that $\overline{V}_1 > \overline{V}_2 > \overline{V}_3$. Suppose for now that prices cannot be used. Then, as vaccines become available, the optimal mechanism first randomly allocates vaccinates to all members of group 1, then to all members of group 2, and then to all members of group 3. The marginal value for society of this allocation is depicted in the left panel of Figure 6.1. Note that these marginal values are constant within each group.


Figure 6.1: An example of optimal mechanism when prices can (right panel) or cannot (left panel) be used with three groups.

Now suppose prices can be used. Let us assume that $V_1(r)$ and $V_2(r)$ are non-increasing and $V_3(r)$ is non-decreasing, and that $\overline{V}_1 > V_3(\bar{r}) > \overline{V}_2$, meaning that the marginal value of vaccinating a member of group 3 with the highest WTP is more than the average value of vaccinating a member of group 2, but less than average marginal value for group 1. In this case, the optimal vaccination schedule proceeds as follows (see the right panel of Figure 6.1): We first vaccinate all members of group 1 for free through rationing. Then, we use a price schedule for group 3 that decreases over time at a rate that ensures an assortative matching between the highest-WTP agents in group 3 and vaccines becoming available between t_1 and t_2. When the current price is such that the marginal value of an agent in group 3 that would purchase at that price is equal to $\overline{V}_2$, we pause the allocation in group 3 (by freezing the price), and vaccinate all members of group 2 for free via rationing. Once they are all vaccinated, we resume the declining price schedule for group 3, thus vaccinating the remaining members of that group in an assortative fashion.

The principle benefit of the price mechanism in this example is that it allows to vaccinate the high-marginal-value individuals from group 3 earlier than the low-marginal-value individuals, resulting in a modified priority in which some agents from group 3 are vaccinated before group 2. When free allocation is used within group 3, all agents in group 3 are vaccinated after group 2 because there is no way to identify these high-marginal-value individuals. This analysis also illustrates how the use of prices can lead to strictly higher welfare overall, since more social value is unlocked earlier on in the distribution process.

7 Alternative "pure-health" objective function

In preceding sections, we have been focusing on a typical (at least to economists) welfare function that aggregates all agents' utilities. However, in popular discourse, other objectives are being considered. Here, we apply our methods to a "pure-health" objective function that only takes into account the private health benefits and the health externality. That is, we set

$$V(v,\, v_{\mathrm{ex}},\, h,\, h_{\mathrm{ex}},\, \lambda,\, t,\, p) = \delta(t)\mathbf{1}_{\mathrm{Risky}}(\lambda h + h_{\mathrm{ex}}). \tag{7.1}$$

Note that if the designer only cares about health outcomes, then (under our assumptions) she only benefits from vaccinating agents who chose $a = \mathrm{Risky}$ prior to being vaccinated. Moreover, the private benefit h is multiplied by the social welfare weight λ because h is expressed in dollar value. As such, h could be influenced by, for example, the agent's wealth. Thus, the designer converts these private values to "social" values before aggregation.

While the health objective is of special interest, we emphasize that it imposes strong ethical consequences. To see that sharply, imagine an individual who will die for sure if they choose $a = \mathrm{Risky}$ (h is extremely high) but who nevertheless suffers from being forced to stay at home (v is high). Under the health objective, the value of vaccinating such an individual is 0. The reason is that this individual receives the health benefit h and generates the externality h_{ex} regardless of whether they are vaccinated or not. In contrast, vaccinating such an individual would be highly desirable under the utilitarian objective.

All of our formal results continue to hold with $V_i(r)$ defined as

$$V_i(r) = r \cdot \mathbb{E}\left[\lambda \,|i,\, r,\, \mathrm{Risky}\right] \cdot \mathbb{P}(\mathrm{Risky}|i,\, r) + h_{\mathrm{ex}}^i \cdot \mathbb{P}(\mathrm{Risky}|i,\, r).$$

The first component of the health objective corresponds to the private health benefit and is a product of three terms. To understand them, note that the health benefit of vaccination is only obtained by agents who chose $a = \mathrm{Risky}$ prior to receiving the vaccine. The last term $\mathbb{P}(\mathrm{Risky}|i,\, r)$ is the best estimate of the probability of this event conditional on information

available to the designer. The second term $\mathbb{E}\left[\lambda\,|i,\,r,\,\text{Risky}\right]$ is the best estimate of the social welfare weight which now additionally conditions on the fact that the agent chose $a = \text{Risky}$. Finally, the first term r is the agent's willingness to pay which is *equal to* her health benefit conditional on choosing $a = \text{Risky}$. Thus, a somewhat surprising conclusion is that when the designer cares only about health outcomes, WTP is closely aligned with her objective. As a consequence, if the designer does not care about inequality within group i (λ does not vary with r conditional on i) and the assessed probability $\mathbb{P}(\text{Risky}|i,\,r)$ does not depend on r (for example, because i already reveals that $\mathbb{P}(\text{Risky}|i,\,r) = 1$), then a market allocation is optimal.

A free allocation may be preferred under the health objective when the probability of choosing $a = \text{Risky}$ is strongly decreasing in WTP r, especially if this is reinforced by a decrease in the expected welfare weight λ with r, and when the health externality h_{ex}^i is large. The previously considered group of ride-share drivers may serve as an illustration. Since it is the poorest drivers that are most likely to be forced to continue driving (choosing $a = \text{Risky}$), it is natural to expect that $\mathbb{P}(\text{Risky}|i,\,r)$ will be decreasing in r in that group. The health externality is large because ride-share drivers come in close contact with many people. Hence, it is optimal to use a free allocation. The idea is that a free allocation avoids tilting the allocation towards richer drivers; this would be suboptimal because these drivers are more likely to choose $a = \text{Safe}$ prior to being vaccinated, and hence generate a smaller positive health externality after vaccination.

To determine the optimal priority of groups, we compute the analogs of (5.1) and (5.2) as
$$\mathbb{E}[V_j(r)\,|\,r \le x] = \mathbb{E}\left[\lambda\,r\,\mathbf{1}_{\text{Risky}}\,|i,\,r \le x\right] + h_{\text{ex}}^i \cdot \mathbb{P}(\text{Risky}|i,\,r \le x). \tag{7.2}$$

and
$$\mathbb{E}[V_i(r)\,|\,r \ge x] = \mathbb{E}\left[\lambda\,r\,\mathbf{1}_{\text{Risky}}\,|i,\,r \ge x\right] + h_{\text{ex}}^i \cdot \mathbb{P}(\text{Risky}|i,\,r \ge x). \tag{7.3}$$

Thus, a key determinant of prioritized groups under the health objective is the label-revealed probability of choosing $a = \text{Risky}$. The reason is that—under our simplifying assumptions— vaccinating individuals who choose $a = \text{Safe}$ has no health benefit.

The health objective supports even more strongly the idea that doctors and nurses should receive priority allocation. These groups have a high probability of choosing Risky, and a high health externality. However, the health objective function is less likely than the utilitarian objective to support priority allocation to groups associated with high socio-economic externalities, e.g., teachers—especially if they teach remotely ($a = \text{Safe}$ with high probability). More generally, Result 2 likely applies to groups with a high private and social health value whose observables reveal the action $a = \text{Risky}$ with high probability—for

example, front-line workers. A high priority would be given to groups like first-responders, cashiers, and delivery workers whose jobs cannot be done remotely. Lowest priority would be given to people who are likely to choose $a = $ Safe, perhaps such as college professors.

Under the health objective, the answer to the question of whether we should sell vaccines to millionaires is different than before. Indeed, consider some individual with WTP $r = 1000000. Under the health objective, the question of whether that individual should receive a vaccine depends crucially on

$$\mathbb{E}[\lambda|i,\ r = \$1000000,\ \text{Risky}] \text{ and } \mathbb{P}(\text{Risky}|i,\ r = \$1000000);$$

if either one of these terms is close to 0, the answer is "no." The first term can be 0 if the designer has an explicit redistributive motive in that she is not concerned with the welfare of the very rich. It can also be low if λ is thought of as correcting for the "confounding" effect of individual wealth on the "private value for health." If one individual has a WTP $100 for the vaccine, and another has a WTP $100,000, we are unlikely to regard vaccinating that second individual to be a thousand times more socially valuable. The second term can be 0 if very rich people are almost sure to choose $a = $ Safe. If any one of these two terms is low enough, the optimal mechanism does not include an initial stage in which vaccines are sold at high prices.

Finally, we note that the health objective yields an intuitive, but somewhat paradoxical insight about vaccinating the elderly versus college students. Because people who are at high risk when infected are much more likely to choose $a = $ Safe, their contribution to the health objective function is low. In contrast, healthy and young people are likely to choose $a = $ Risky, which means that their contribution is large. Thus, under the health objective, potential spreaders of the virus should be vaccinated before people with highest risks. This is despite the fact that we reached the opposite conclusion under the utilitarian objective (as discussed in the preceding section). *The paradox is that if the goal is to maximize overall population health, then it is especially important to vaccinate the individuals who are at relatively low private health risk because those are the agents who are likely to take risky actions.* That conclusion could not be found in a model that does not endogenize the action choices of the agents.

8   Concluding remarks

Our baseline framework focuses on the main trade-offs associated with the choice of vaccine prioritization. We made simplifying assumptions to emphasize the novel insights: the idea

that screening based on willingness to pay may reveal important information about externalities, the role of redistributive preferences and revenue, and the importance of accounting for endogeneity of individual responses to the pandemic. Below, we discuss additional points and extensions.

Continuous choice of precaution level a. Naturally, there are more than two ways in which individuals can react to the pandemic, corresponding to a larger set of actions a to choose from. Our model could easily accommodate such cases at the cost of complicating notation and interpretation. The main difference is that instead of making an inference about the probability of choosing $a =$ Risky, $\mathbb{P}(\text{Risky}|i, r)$, the designer would try to estimate the distribution of the action a conditional on i and r. If $a \in [0, 1]$ with higher a corresponding to more risky behavior, the conclusions from a linear model would depend on $\mathbb{E}[a|i, r]$, with similar intuitions as in the current specification.

Paying agents for getting vaccinated. In the baseline model, willingness to pay was assumed to be non-negative. When we relax that assumption, because of externalities, it may makes sense to pay agents with negative health benefit h in exchange for them agreeing to getting vaccinated.

There are several cases to be considered. If the objective function $V_i(r)$ is non-decreasing in some group i, then the optimal allocation would be to (eventually) vaccinate all agents with $r \geq r_i^\star$, where $r_i^\star < 0$ is the threshold WTP at which $V_i(r_i^\star) = 0$, that is, at which the private disutility and revenue loss become equal to the positive externality from vaccination. If $V_i(r)$ is decreasing (and non-negative in expectation), on the other hand, then all agents in group i should be vaccinated, and it becomes necessary to compensate everyone in that group (including agents with positive WTP) by an amount required to convince the most-negative-WTP agents to get the vaccine. A decreasing $V_i(r)$ may indeed arise when low (negative) r is related to skepticism about the pandemic resulting in ignoring the recommended safety measures; for example, agents who believe that vaccines are harmful may also be more likely to believe that wearing masks is unhealthy. Note, however, that a globally decreasing $V_i(r)$ is unlikely to arise when the lowest WTP is very negative—in that case, the cost of vaccinating everyone becomes prohibitively high (that is, we should start thinking of α as being relatively large) and the threshold type is chosen to optimally trade-off the revenue loss against the externalities.

Elastic supply of vaccines. While our model assumes a fixed supply of vaccines, we can indirectly model the supply effects via the weight on revenue α. Especially for developing

countries, monetary costs may be the bottleneck in expanding the available supply. Consequently, such countries may want to set a high α in their objective function to capture the positive effects of revenue on total supply. This favors a market allocation. The most likely outcome is the co-existence of public market where vaccines are allocated at low or zero prices to groups with highest externalities, and a private market with relatively high market-clearing prices that generate substantial revenue.

Queueing. While we interpreted our model as featuring monetary payments and prices, an alternative interpretation is that agents "pay" by engaging in a costly activity, such as queueing. In that case, we have $\alpha \leq 0$ since the designer does not benefit from this (inherently wasteful) activity. The case $\alpha < 0$ captures the negative impact that physical queueing could have on the spread of the virus. All mathematical results continue to hold in such a model. However, our results must be interpreted accordingly. For example, a "market allocation" now means that people who spend the most time in the queue get the vaccine first. A free allocation means that there is no queue and a lottery decides about priority. Some of our assumption may naturally be reversed. For example, we argued that poorer agent may have a lower willingness to pay r, everything else being equal. When values are measured in terms of disutility from waiting in a line, it may well be the case that poorer agents are associated with higher r, which now becomes "willingness to queue."

Decentralized implementation. In our approach so far, we have focused on what the optimal allocation (and potential payments) are; we have not discussed how they can be implemented in practice. A pure priority system, like the one described in Section 4, requires centralized control over the implementation mechanism. Because the allocation is based on labels, it must be ensured that individuals receive the vaccines only at their prescribed time. Revenue-maximizing entities (like private pharmacies) may lack the incentives or ability to verify eligibility. In contrast, a pure market allocation could be—at least in principle—achieved in a decentralized fashion by a competitive market. The reason is that, under sufficiently fierce competition, the homogeneity of vaccines would ensure that revenue-maximizing firms would sell them at prices implementing the efficient allocation. In intermediate cases, when prices depend on labels, achieving a decentralized implementation would be far more challenging. In certain cases, it could be possible to do that by issuing label-specific "coupons"—vaccines are sold at pharmacies and other outlets at a list price, but agents in eligible categories receive a coupon that entitles them to a discount.[24]

---

[24]A similar implementation has been used in the context of registration for vaccine appointments, with specialized codes that entitled individuals in certain groups to move up in the queue.

References

AKBARPOUR, M., C. COOK, A. MARZUOLI, S. MONGEY, A. NAGARAJ, M. SACCAROLA, P. TEBALDI, S. VASSERMAN, AND H. YANG (2020): "Socioeconomic network heterogeneity and pandemic policy response," NBER Working Paper No. 27374.

ASHLAGI, I., F. MONACHOU, AND A. NIKZAD (2020): "Optimal Dynamic Allocation: Simplicity through Information Design," Available at SSRN.

BIBBINS-DOMINGO, K., M. PETERSEN, AND D. HAVLIR (2021): "Taking Vaccine to Where the Virus Is—Equity and Effectiveness in Coronavirus Vaccinations," in JAMA Health Forum, American Medical Association, vol. 2, e210213–e210213.

BRITO, D. L., E. SHESHINSKI, AND M. D. INTRILIGATOR (1991): "Externalities and compulsary vaccinations," Journal of Public Economics, 45, 69–90.

BUBAR, K. M., K. REINHOLT, S. M. KISSLER, M. LIPSITCH, S. COBEY, Y. H. GRAD, AND D. B. LARREMORE (2021): "Model-informed COVID-19 vaccine prioritization strategies by age and serostatus," Science, 371, 916–921.

BUDISH, E. (2020): "Maximize utility subject to $R \leq 1$: A simple price-theory approach to Covid-19 lockdown and reopening policy," Tech. rep., nBER Working Paper No. 28093.

BULOW, J. AND J. ROBERTS (1989): "The Simple Economics of Optimal Auctions," Journal of Political Economy, 97, 1060–1090.

CHAKRAVARTY, S. AND T. R. KAPLAN (2013): "Optimal allocation without transfer payments," Games and Economic Behavior, 77, 1 – 20.

CHE, Y.-K., W. DESSEIN, AND N. KARTIK (2013): "Pandering to Persuade," American Economic Review, 103, 47–79.

CONDORELLI, D. (2012): "What money can't buy: Efficient mechanism design with costly signals," Games and Economic Behavior, 75, 613 – 624.

——— (2013): "Market and non-market mechanisms for the optimal allocation of scarce resources," Games and Economic Behavior, 82, 582–591.

DEB, R., M. PAI, A. VOHRA, AND R. VOHRA (2020): "Testing alone is insufficient," .

DELACRÉTAZ, D. (2020): "Processing Reserves Simultaneously," Working Paper.

DWORCZAK, P. ⓡ S. D. KOMINERS ⓡ M. AKBARPOUR (2020): "Redistribution through Markets," Econometrica (forthcoming).

EDELMAN, B., M. OSTROVSKY, AND M. SCHWARZ (2007): "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," American economic review, 97, 242–259.

ELY, J., J. STEINER, AND A. GALEOTTI (2020a): "Optimal Test Allocation," Working Paper.

——— (2020b): "Rotation as Contagion Mitigation," Working Paper.

EMANUEL, E. J., G. PERSAD, A. KERN, A. BUCHANAN, C. FABRE, D. HALLIDAY, J. HEATH, L. HERZOG, R. LELAND, E. T. LEMANGO, ET AL. (2020a): "An ethical framework for global vaccine allocation," Science, 369, 1309–1312.

EMANUEL, E. J., G. PERSAD, R. UPSHUR, B. THOME, M. PARKER, A. GLICKMAN, C. ZHANG, C. BOYLE, M. SMITH, AND J. P. PHILLIPS (2020b): "Fair Allocation of Scarce Medical Resources in the Time of Covid-19," New England Journal of Medicine, 382, 2049–2055.

GOLDSTEIN, J. R., T. CASSIDY, AND K. W. WACHTER (2021): "Vaccinating the oldest against COVID-19 saves both the most lives and most years of life," Proceedings of the National Academy of Sciences, 118.

GRIGORYAN, A. (2021): "Effective, Fair and Equitable Pandemic Rationing," Working Paper.

HARTLINE, J. D. AND T. ROUGHGARDEN (2008): "Optimal mechanism design and money burning," in Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, 75–84.

JEHIEL, P. AND B. MOLDOVANU (2001): "Efficient Design with Interdependent Valuations," Econometrica, 69, 1237–1259.

JEHIEL, P., B. MOLDOVANU, AND E. STACCHETTI (1996): "How (Not) to Sell Nuclear Weapons," The American Economic Review, 86, 814–829.

KASY, M. AND A. TEYTELBOYM (2020): "Adaptive targeted infectious disease testing," Oxford Review of Economic Policy, 36, S77–S93.

Kleiner, A., B. Moldovanu, and P. Strack (2020): "Extreme Points and Majorization: Economic Applications," Working Paper.

M. Akbarpour ⓡ Dworczak, P. ⓡ S. D. Kominers (2020): "Redistributive Allocation Mechanisms," Working Paper.

Makhoul, A. T. and B. C. Drolet (2021): "A Reserve System for the Equitable Allocation of a Severe Acute Respiratory Syndrome Coronavirus 2 Vaccine," CHEST, 159, 1292–1293.

Milgrom, P. (2001): Putting Auction Theory to Work, Cambridge University Press.

Muir, E. V. and S. Loertscher (2020): "Monopoly pricing, optimal randomization, and resale," Working Paper.

Myerson, R. B. (1981): "Optimal auction design," Mathematics of Operations Research, 6, 58–73.

National Academies of Sciences, Engineering, and Medicine (2020): "Framework for Equitable Allocation of COVID-19 Vaccine," .

Pancs, R. (2020): "A Vaccine Auction," Working Paper.

Pathak, P. A., T. Sönmez, and M. U. Ünver (2020): "Improving ventilator rationing through collaboration with experts on resource allocation," JAMA Open, 3, e2012838–e2012838.

Pathak, P. A., T. Sönmez, M. U. Ünver, and M. B. Yenmez (2021): "Fair allocation of vaccines, ventilators and antiviral treatments: Leaving no ethical value behind in health care rationing," Working Paper.

Persad, G., M. E. Peek, and E. J. Emanuel (2020): "Fairly prioritizing groups for access to COVID-19 vaccines," JAMA, 324, 1601–1602.

Rahmandad, H. (2021): "Behavioral Responses to Risk Promote Vaccinating High-contact Individuals First," .

Ray, D. ⓡ A. Robson (2018): "Certified Random: A New Order for Coauthorship," American Economic Review, 108, 489–520.

Roth, A. E. (2007): "Repugnance as a Constraint on Markets," Journal of Economic perspectives, 21, 37–58.

SCHMIDT, H. (2020): "Vaccine Rationing and the Urgency of Social Justice in the Covid-19 Response," Hastings Center Report, 50, 46–49.

SCHMIDT, H., P. PATHAK, T. SÖNMEZ, AND M. U. ÜNVER (2020a): "Covid-19: How to prioritize worse-off populations in allocating safe and effective vaccines," British Medical Journal, 371, m3795.

SCHMIDT, H., P. A. PATHAK, M. A. WILLIAMS, T. SONMEZ, M. U. ÜNVER, AND L. O. GOSTIN (2021): "Rationing Safe and Effective COVID-19 Vaccines: Allocating to States Proportionate to Population May Undermine Commitments to Mitigating Health Disparities," Working Paper.

SCHMIDT, H., U. ÜNVER, M. WILLIAMS, P. PATHAK, T. SÖNMEZ, AND L. GOSTIN (2020b): "What Prioritizing Worse-Off Minority Groups for COVID-19 Vaccines Means Quantitatively: Practical, Legal and Ethical Implications," Working Paper.

SÖNMEZ, T., P. A. PATHAK, M. U. ÜNVER, G. PERSAD, R. D. TRUOG, AND D. B. WHITE (2020): "Categorized Priority Systems: A New Tool for Fairly Allocating Scarce Medical Resources in the Face of Profound Social Inequities," Chest.

UNIVERSITY OF PITTSBURGH DEPARTMENT OF CRITICAL CARE MEDICINE (May 28, 2020): "Model Hospital Policy for Fair Allocation of Scarce Medications to Treat COVID-19," https://www.ccm.pitt.edu/?q=content/model-hospital-policy-fair-allocation-medications-treat-covid-19.

VARIAN, H. R. (2007): "Position auctions," international Journal of industrial Organization, 25, 1163–1178.

WEITZMAN, M. L. (1977): "Is the price system or rationing more effective in getting a commodity to those who need it most?" Bell Journal of Economics, 8, 517–524.

A   Proofs

Our model can be solved using techniques developed in Akbarpour ⓡ Dworczak ⓡ Kominers (2020).[25] Even though ADK use a different objective function, their method applies for any

---

[25]These techniques be seen as a generalization of the ironing technique developed by Myerson (1981). Following the intuitive approach to ironing developed by Bulow and Roberts (1989), Hartline and Roughgarden (2008) applied it to a problem with multiple goods, and Condorelli (2012) to multiple goods with heterogeneous quality. Muir and Loertscher (2020) rely on similar techniques to solve a problem of a revenue-

objective function as long as it is linear in the quality of the good allocated to every agent. As explained in Section 3, the timing of the vaccination in our framework is mathematically equivalent (under the transformation $q = \delta(t)$) to the quality of the good in ADK.

A.1 Proof of Lemmas 1 - 3

To prove Lemma 1, note that we can write the utility of an agent who receives a vaccine with quality q at price p as (see equation 2.1)

$$q \min\{v,\, h\} - p + \text{const.}$$

Thus, by defining $r = min\{v,\, h\}$, we obtain that r is the willingness to pay for quality in the model of ADK. The first part of Lemma 1 then follows immediately from Claim 1 of ADK (analogous results are proven in Jehiel and Moldovanu (2001), Che et al. (2013), and Dworczak Ⓡ Kominers Ⓡ Akbarpour (2020)). The second part of Lemma 1 follows from the observation that if two agents with the same label i and types $(v,\, v_{\text{ex}},\, h,\, h_{\text{ex}},\, \lambda)$ and $(v',\, v'_{\text{ex}},\, h',\, h'_{\text{ex}},\, \lambda')$, respectively, receive different outcomes in a mechanism without prices, then it must be (by incentive-compatibility) that they receive the same *expected* quality. Because the properties of mechanisms in our model depend only on the expected-quality schedules, it is without loss of optimality to assume that the optimal mechanism only conditions the allocation of vaccines on the labels.[26]

To prove Lemma 2, note that the designer's payoff from allocating a vaccine with quality q at price p to an agent with type $(v,\, v_{\text{ex}},\, h,\, h_{\text{ex}},\, \lambda)$ is given by (see equation 2.2)

$$\lambda(qr - p) + \alpha p + q\left(\mathbf{1}_{\text{Safe}} v_{\text{ex}} + \mathbf{1}_{\text{Risky}} h_{\text{ex}}\right).$$

By the revelation principle and Lemma 1, in the problem with prices, the designer can restrict attention to direct mechanisms of the form $(Q_i(r),\, t_i(r))_{i \in I, r \in [0, \bar{r}_i]}$, where $t_i(r)$ is the payment charged to agent with label i and WTP r.[27] The expected payoff for the designer

---

maximizing seller in the presence of resale; Ashlagi, Monachou, and Nikzad (2020) show that these methods can be also used in designing the optimal dynamic allocation in a multi-good environment by optimizing over how much information is disclosed about different types of objects; finally, Kleiner, Moldovanu, and Strack (2020) demonstrate that all these procedures can be obtained as a special case of a general property of extreme points that arise in optimization problems involving majorization constraints.

[26]One might think that the designer could expand the set of *feasible* mechanisms by implementing the same expected quality to two agents using two different lotteries; but ADK show that this is not the case in a model with a continuum population of agents.

[27]Of course, the optimization problem has a feasibility constraint stating that the expected-quality schedules $Q_i(r)$ are jointly feasible given the primitive distribution of quality F; see ADK for details.

from using such a mechanism is

$$\sum_{i \in I} \mu_i \int_0^{\bar{r}_i} \left\{ \underbrace{\lambda_i(r) U_i(r)}_{\text{agents' weighted utility}} + \alpha \underbrace{t_i(r)}_{\text{revenue}} + \underbrace{Q_i(r) \mathbb{E}[\mathbf{1}_{\text{Safe}} v_{\text{ex}} + \mathbf{1}_{\text{Risky}} h_{\text{ex}} | i, r])}_{\text{externalities}} \right\} dG_i(r),$$

where $U_i(r) = Q_i(r) r - t_i(r)$, and μ_i is the mass of agents with label i. It follows that our objective function differs from that analyzed in ADK only by the additive term $Q_i(r) \mathbb{E}[\mathbf{1}_{\text{Safe}} v_{\text{ex}} + \mathbf{1}_{\text{Risky}} h_{\text{ex}} | i, r]$. Moreover, it follows from our assumption that the externalities v_{ex} and h_{ex} are independent of v and h conditional on i that

$$\mathbb{E}[\mathbf{1}_{\text{Safe}} v_{\text{ex}} + \mathbf{1}_{\text{Risky}} h_{\text{ex}} | i, r] = v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} | i, r) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} | i, r).$$

ADK show that in an inventive-compatible mechanism with non-negative transfers (as is assumed here)

$$\sum_{i \in I} \mu_i \int_0^{\bar{r}_i} \left\{ \lambda_i(r) U_i(r) + \alpha t_i(r) \right\} dG_i(r) = \sum_{i \in I} \mu_i \int_0^{\bar{r}_i} \tilde{V}_i(r) Q_i(r) dG_i(r),$$

where $\tilde{V}_i(r) = \Lambda_i(r) \gamma_i(r) + \alpha(r - \gamma_i(r))$. It follows immediately that in our setting the designer's objective is

$$\sum_{i \in I} \mu_i \int_0^{\bar{r}_i} V_i(r) Q_i(r) dG_i(r),$$

where $V_i(r)$ is defined by (3.1). Moreover, all the results of ADK apply to our setting by substituting $\tilde{V}_i(r)$ for $V_i(r)$ defined by (3.1).

Finally, to prove Lemma 3, it suffices to observe that giving a vaccine with quality q to a *random* agent in group i has a social benefit $\mathbb{E}[V_i(r)|i]$ which is

$$\int_0^{\bar{r}_i} \left(\Lambda_i(r) \gamma_i(r) + \alpha(r - \gamma_i(r)) + v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} | i, r) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} | i, r) \right) dG_i(r).$$

A simple calculation (using integration by parts) shows that the first term in the integrand integrates out to $\mathbb{E}[\lambda \cdot r | i]$, the second term disappears (revenue in a mechanism without prices is 0), while the last two terms are simply $v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} | i) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} | i)$.

A.2 Proof of Results 1 - 5

To derive Results 1 - 5, we first restate the results of ADK in our context. To identify an optimal mechanism with prices, we proceed in two steps:

1. First, vaccines are allocated optimally "across" groups: F is split into $|I|$ CDFs F_i.

2. Then, vaccines are allocated optimally "within" groups: For each label i, the vaccines in F_i are allocated according to the expected-quality schedule $Q_i(r)$.

We will refer to the two steps above as the "within problem" and the "across problem," respectively (see ADK for formal definitions of these optimization problems).

For the setting without prices (Section 4), the within problem becomes trivial—vaccines are allocated uniformly at random, so that $Q_i(r) = \int q \, dF_i(q)$ for all r and any i. This observation allows us to prove Result 1.

Proof of Result 1. When the allocation is free within each group (the designer cannot use prices), the value from allocating a unit of quality to group i is simply $\overline{V}_i$, as defined in Lemma 3. Therefore, the across problem can be formally written as

$$\max_{(F_i)_{i \in I}} \sum_{i \in I} \mu_i \overline{V}_i \int_0^1 q \, dF_i(q), \tag{A.1}$$

$$\text{s.t.} \quad \sum_{i \in I} \mu_i F_i(q) = F(q), \, \forall q \in Q. \tag{A.2}$$

It follows immediately that in any optimal solution, $\max(\text{supp}(F_i)) \leq \min(\text{supp}(F_j))$ whenever $\overline{V}_i < \overline{V}_j$ which corresponds to the statement that all agents in group j receive a weakly higher quality vaccine (are vaccinated earlier) than any agent in group i. (If $\overline{V}_i = \overline{V}_j$, then the order of vaccination of groups i and j does not matter for the designer's expected payoff, so vaccinating the two groups sequentially, in any order, is optimal.) This finishes the proof of Result 1.

Next, we prove Results 2 - 5. To this end, we restate two theorems from ADK: The first one describes the solution to the within problem (with prices), while the second one describes the solution to the across problem. The statements differ slightly from ADK due to two differences in the settings. First, we do not allow for free disposal; second, the results can be simplified because we assume that the lower bound of the distribution of r is zero in each group (while ADK allow for an arbitrary non-negative lower bound $\underline{r}_i$).

Theorem 1 (ADK). *Define*

$$\Psi_i(t) := \int_t^1 V_i(G_i^{-1}(x)) dx.$$

The value of the within problem for group i (for a fixed F_i) is given by

$$\int_0^1 co(\Psi_i)(F_i(q))dq,$$

where $co(\Psi_i)$ denotes the concave closure of Ψ_i. An optimal solution is given by an expected-quality schedule $Q_i^\star(r) = \Phi_i^\star(G_i(r))$, where $\Phi_i^\star$ is non-decreasing and satisfies

$$\Phi_i^\star(x) = \begin{cases} \frac{\int_a^b F_i^{-1}(y)dy}{b-a} & \text{if } x \in (a, b) \text{ and } (a, b) \text{ is a maximal interval on which } co(\Psi_i) > \Psi_i, \\ F_i^{-1}(x) & \text{otherwise,} \end{cases}$$

for almost all x.

Theorem 2 (ADK). *Let $s_i(x) \equiv co(\Psi_i)'(x)$. There exists a non-increasing function $S(q)$ such that for all i and q, the optimal solution $(F_i^\star)_{i \in I}$ to the across problem satisfies*

$$\begin{cases} F_i^\star(q) = 0 & \text{if } s_i(0) < S(q), \\ F_i^\star(q) = 1 & \text{if } s_i(1) > S(q), \\ F_i^\star(q) \text{ solves } s_i(F_i^\star(q)) = S(q) & \text{otherwise.} \end{cases}$$

Moreover, $S(q) = \max_{i: F_i^\star(q)<1} s_i(F_i^\star(q))$.

Proof of Result 2. The first condition for priority allocation to groups J, $A(0) \geq \sum_{j \in J} \mu_j$ is clearly necessary, since if it does not hold, it is not feasible to vaccinate all agents in groups J immediately (at time 0). If that condition holds, then a sufficient condition for priority allocation to J can be deduced directly from Theorem 2. Indeed, all agents in groups J receive the vaccines before any other group if, for all $j \in J$, $i \notin J$, the slope of $co(\Psi_j)$ at 1 is lower than the slope of $co(\Psi_i)$ at 0. Moreover, this condition is necessary when $A(0) = \sum_{j \in J} \mu_j$. The slope of $co(\Psi_j)$ at 1 is equal to $\min_x \mathbb{E}[V_j(r)|r \leq x]$, while the slope of $co(\Psi_i)$ at 0 is equal to $\max_x \mathbb{E}[V_i(r)|r \geq x]$, by direct calculation, proving the result.

Proof of Result 5. We prove this result first because the supporting arguments are also needed in the proofs of Results 3 and 4. The proof follows directly from Theorem 1. When $V_i(r)$ is non-decreasing, Ψ_i is concave, and hence $co(\Psi_i) = \Psi_i$ everywhere. Thus, $Q_i^\star(r) = F_i^{-1}(G_i(r))$, corresponding to a market allocation. When $V_i(r)$ is non-increasing, Ψ_i is convex, and hence $co(\Psi_i) > \Psi_i$ on the interior of the domain.[28] Thus, $Q_i^\star(r) = \int_0^1 q dF_i(q)$, corresponding to a free allocation.

---

[28]Except for the knife-edge case in which $V_i(r)$ is constant; but then any allocation method is optimal.

Proof of Result 3. This result follows directly from Theorem 1 and Theorem 2. For a group j with free allocation, $\mathrm{co}(\Psi_j)$ is linear, so that the slope $s_j(x)$ is constant in x, equal to $-\overline{V}_j$. And for a group i with market allocation, we have $\mathrm{co}(\Psi_i) = \Psi_i$, and so $s_i(0) = -V_i(0)$, while $s_i(1) = -V_i(\bar{r}_i)$.

Proof of Result 4. This result follows directly from Theorem 1 and Theorem 2, given that for a group i with market allocation, we have $s_i(0) = -V_i(0)$ and $s_i(1) = -V_i(\bar{r}_i)$.