

AN ECONOMIC FRAMEWORK FOR VACCINE PRIORITIZATION*

MOHAMMAD AKBARPOUR[®]
ERIC BUDISH[®]
PIOTR DWORCZAK[®]
SCOTT DUKE KOMINERS

We propose an economic framework for determining the optimal allocation of a scarce supply of vaccines that become gradually 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 welfare function over all feasible mechanisms—accounting for agents’ characteristics, as well as their 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 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. The solution casts light on the classic question of whether prices or priorities should be used to allocate scarce public resources under externalities and equity concerns. *JEL Codes:* C78, D47, D61, D63, D82.

*The[®] symbol indicates that the authors’ names are in certified random order, as described by Ray[®] Robson (2018). The authors thank Susan Athey, Péter Biró, Matthew Cortland, Glenn Ellison, Ezekiel Emanuel, Simon Finster, Drew Fudenberg, Navid Ghaffarzadegan, Gregg Gonsalves, Anup Malani, Paul Milgrom, Romans Pancs, Parag Pathak, Canice Prendergast, Hazhir Rahmandad, Alvin Roth, Tayfun Sönmez, Alex Tabarrok, Nikhil Vellodi, Robert Wilson, the editor (Robert Barro), several referees, and numerous seminar and conference audiences for insightful conversations and comments, and thank Xiaoyun Qiu for excellent research assistance. Akbarpour, Dworzak, and Kominers gratefully acknowledge support from the Washington Center for Equitable Growth. Akbarpour gratefully acknowledges the support of an Alfred P. Sloan Fellowship. Dworzak gratefully acknowledges the support received from the European Union under ERC Starting Grant IMD-101040122. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. An abstract of this work appeared in the *Proceedings of the 23rd ACM Conference on Economics and Computation*.

© The Author(s) 2023. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

The Quarterly Journal of Economics (2024), 359–417. <https://doi.org/10.1093/qje/qjad022>. Advance Access publication on June 12, 2023.

I. INTRODUCTION

Economists typically prescribe prices to guide the allocation of scarce resources, arguing that the implicit selectivity of the price system helps allocate resources to those who value them the most. However, in many contexts, considerations such as fairness, equity, or consumption externalities provide arguments against using prices—and indeed, many ethicists and policy makers opt for schemes that allocate resources free of charge to certain selected groups. The price system, they argue, directs resources to those who are able to pay the most, which may not match up with true needs or moral desert.¹

The question of whether to use prices or priorities played out in the context of allocating vaccines during the COVID-19 pandemic: multiple effective vaccines were developed with unprecedented speed; nevertheless, vaccine supply chains were (and to some degree remain) constrained due to production and logistical challenges—making vaccines a scarce resource in the short run.² Although prices could help identify individuals with the highest private values for vaccines, most countries opted for a priority system with rationing.

This article contributes to the debate on allocating vaccines by employing mechanism-design tools to characterize the socially optimal allocation scheme. We derive the optimal scheme from economic primitives—a specification of social and individual preferences over vaccination and a strategic environment in which agents have private information and make endogenous choices about which actions to take absent receiving a vaccine. Crucially, our framework incorporates ethical and equity concerns, as well as externalities, but allows the designer to use prices. The key insight is that although social considerations may indeed limit the role that prices play in the optimal mechanism, they are typically not sufficient to rule out prices completely. As a result, a priority system with rationing may coexist with a pricing scheme; such a hybrid mechanism allows the designer to leverage observable information while screening for unobservable characteristics. A secondary contribution is thus to cast light on the classical question of prices versus priorities; we show that posing the question

1. See, for example, [Satz \(2010\)](#) and [Sandel \(2012\)](#), and, in the context of vaccine allocation, [Persad, Peek, and Emanuel \(2020\)](#).

2. For a discussion of the issues around vaccine production capacity, see [Castillo et al. \(2021\)](#), [Athey et al. \(2022\)](#), [Bown \(2022\)](#), [Budish et al. \(2022\)](#), and [Kominers and Tabarrok \(2022\)](#).

as a dichotomy is misleading and provide a framework that can be used to determine the optimal combination of prices and priorities in allocation problems with equity concerns and externalities.

Before discussing our detailed findings, we briefly sketch the framework. Before 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 decrease 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).³ The specific interactions the agent engages in—and the implicit cost of taking precautions—depend on the agent’s type. In particular, the agent’s choice is determined by the comparison of the private health benefit of not contracting the virus and the private socioeconomic benefit of in-person interactions.

The agent’s private choice generates externalities—the safe action leads to public health benefits by slowing down the spread of the virus, and the risky action leads to benefits associated with the agent’s economic and social activity. Receiving the vaccine is modeled as providing both types of benefits at the same time.

Both the health and the socioeconomic benefits are measured in units of money (dollars). To address inequality and ethical concerns, we assume that each agent is associated with a welfare weight capturing the social value of giving that agent a unit of money; 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.

A designer chooses a mechanism that allocates available vaccines over time. The designer does not observe agents’ characteristics or action choices directly, but she can use three types of tools to guide the allocation process. First, she may condition the allocation on observable information about agents that we refer to as labels. Our framework allows for any set of labels, which in

3. What the “risky action” represents may be context dependent. For example, a doctor performing her tasks in person is thought of as taking the risky action even if she follows all required safety protocols (the “safe action” would be to see her patients remotely)—whereas for an ordinary person, taking precautions such as wearing a mask and social distancing might be interpreted as the “safe action.”

practice might include profession, age, income, or neighborhood. Second, we allow the designer to charge prices. Prices may depend on labels and can vary over time. Third, the designer can rely on randomization. Overall, viewing the allocation process as a direct revelation mechanism, the designer may choose any vaccination schedule (i.e., a potentially random vaccination time for every agent) as long as it is incentive compatible and individually rational. 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.⁴ 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.⁵

If the labels available to the designer were perfectly revealing of agents' characteristics, the solution would be conceptually straightforward. First, each agent would be assigned a numerical score—a social value of being vaccinated—capturing that agent's contribution along the three dimensions that the social objective function aggregates: private utility (weighted by welfare weights), revenue, and externalities. Second, agents would be vaccinated in the decreasing order of their social values.

In practice, however, labels are unlikely to reveal all information relevant to determining the social values. For example, two people with similar observable characteristics may have different family or social circumstances that lead to differences in their value for being able to take the risky action—such as acute social isolation versus having close family or friends in one's "COVID pod." This introduces a potential role for prices; when a positive price is charged for a vaccine, agents with the same label may self-select into being a buyer or a nonbuyer based on their private willingness to pay. As long as willingness to pay is correlated with welfare-relevant characteristics, prices allow the designer to elicit additional relevant information. In this example, a social planner

4. Our analysis is purely normative—the objective functions of real-life policy makers may be more complicated and partly shaped by political factors that our model does not capture. For more on this point, see [Boettke and Powell \(2021\)](#).

5. That said, organizations providing vaccines to developing countries (such as COVAX, or other international donors) may not be willing to do so if the vaccines are being used for general revenue-raising, since their explicit objective may be to effect an in-kind transfer. This was true for the COVID-19 pandemic, where COVAX even to some degree blocked countries from exchanging their COVAX-allocated vaccines for other vaccines ([Budish et al. 2022](#)).

may prefer to give the vaccine first to agents facing acute social isolation, and only then to agents whose friends and relatives are part of their COVID pod. As long as the former group tends to have higher willingness to pay for the vaccine than the latter, prices will allow the designer to achieve that allocation. However, self-selection based on willingness to pay need not always be socially desirable. Willingness to pay is also strongly linked to wealth, and thus the price allocation will in general skew the allocation toward wealthier agents. Resolving this trade-off underlies the characterization of the optimal mechanism.

To understand how the optimal mechanism is determined, it is helpful to decompose the allocation process into two steps. First, available vaccines are split across groups—sets of agents sharing the same observable characteristics (i.e., the same label); second, vaccines are allocated to individual agents in the groups. The second step—which we refer to as the “within-group” problem—specifies, for each group separately, whether vaccines are allocated randomly or using prices. To determine the optimal allocation method, the designer forecasts the unobserved social values of vaccination by conditioning on the group label and agents’ willingness to pay. Prices are used when agents with higher willingness to pay have higher expected social values, conditional on the label. The solution to the within-group problem then shapes the first step of the mechanism, which we refer to as the “across-group” problem. Here, the designer forecasts social values by conditioning on the labels, taking into account the optimal within-label allocation method. Groups whose labels reveal higher expected social values receive vaccines earlier. However, the optimal schedule need not be a total ordering of the groups; overlaps in the schedule are possible when prices are used in some groups. Overall, the two steps describe a complete vaccination schedule, along with the supporting payments. In the remainder of this section, we provide an overview of the key implications of this characterization.

First, as we just explained, the optimal mechanism relies on prices when willingness to pay—conditional on some label—correlates positively with the social objective. Typically, the private-utility and revenue components of the social objective will be a source of positive correlation. Assuming that the designer has redistributive preferences, welfare weights may naturally be negatively correlated with willingness to pay, at least in groups with substantial wealth heterogeneity. However, the correlation between willingness to pay and externalities may be positive or negative, and thus the presence of externalities may work in

favor of or against using prices to allocate vaccines, depending on the group.

For a simple illustration of the former possibility, note that many nonessential 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 such an 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 thus 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.⁶

For an illustration of the latter possibility, consider the group of gig economy workers. A gig economy worker's health externality may be large when they choose to perform jobs requiring substantial in-person interactions. Under reasonable assumptions, workers who are less wealthy are more likely, on the margin, to continue performing risky tasks; for example, ride-share drivers may stop driving for some time if ride-sharing is only a supplementary source of income for them, but are unlikely to be able to afford 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 willingness to pay 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.

Second, uncovering the social values requires inferring agents' endogenous action choices prior to receiving the vaccine. This is because the agent's action determines which types of private benefits and externalities are achieved through vaccination. Labels can play an important role here. For example, most doctors are effectively forced to undertake in-person interactions by the nature of their profession; thus, a label associated with being a

6. 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, and hence avoid the potential adverse health consequences of reopening.

doctor indicates that vaccination will have a health benefit for that doctor and a positive health externality by helping protect the people whom that doctor interacts with. As a result, vaccinating doctors will tend to have a high health value regardless of their willingness to pay. In contrast, many college professors have been able to teach from home during the pandemic. Thus, vaccinating professors will have a socioeconomic benefit for them, as well as a positive socioeconomic externality, because it enables them to take actions (such as advising and teaching in-person) that would otherwise be avoided due to health risk.

Following this line of reasoning, we identify sufficient conditions under which it is optimal for a group of agents to receive vaccines immediately and free of charge—what we call the absolute priority allocation. These conditions require that the label that defines the group be associated with a sufficiently high positive externality and/or a sufficiently high welfare weight; in addition, the weight on revenue should be sufficiently small. Our result may justify absolute priority allocation for front-line health workers: they have a high health externality (which is the relevant externality since these workers are in effect forced to choose the risky action by nature of their jobs) and a high welfare weight due to their key role in fighting the pandemic.

In developing economies—whose ability to purchase vaccines in the international market may depend on revenue generation—the weight on revenue could be substantial. In such cases, it may be desirable to vaccinate multiple groups simultaneously, in what is effectively a combination of a subsidized public allocation program and a private market. 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. Then, once groups with high externalities and welfare weights are vaccinated, prices for the general population are reduced.

Finally, we find that the mode of within-group allocation (price-based versus free allocation) affects the allocation across groups of agents. Vaccinating groups of agents sequentially (e.g., health workers first and then teachers) is optimal when allocation in these groups is free (and relies on rationing). However, when prices are used to provide earlier access to agents with the highest willingness to pay, this is no longer the case. It is then generally optimal to have overlaps in the schedule—for example, teachers with high willingness to pay might receive the vaccines

before health workers with low willingness to pay. The intuition is simple: under a free allocation, vaccinating each agent in a group provides the same expected social benefit because a free allocation leads to random order of vaccination within each group. By contrast, under a price-based allocation, agents with higher social values are vaccinated first—and the marginal social value of vaccinating the teacher with the highest willingness to pay may easily exceed the marginal value of vaccinating the doctor with the lowest willingness to pay.

In summary, the consideration of welfare weights and externalities can, under some conditions, justify the use of a pure priority system even if prices could be used. However, as standard economic intuition suggests, prices play an important role in screening for agents with the highest private and—sometimes—social benefits of vaccination. Thus, rather than treating prices and priorities as two alternatives, we should think of them as complementary tools in distributing scarce vaccine resources in a socially optimal way.

Although our framework focuses on the role of price in ensuring optimal sorting, there are of course other advantages of the price system that may be relevant in the context of vaccine allocation. First, because we treat supply as exogenous in our analysis, we ignore the potential positive influence of prices on the availability of vaccines; for example, pharmaceutical firms' incentives for vaccine discovery and production depend crucially on how they are compensated.⁷ Second, a price-based allocation may be significantly simpler to implement than a priority-based system, because the latter may require the government to process and verify vast amounts of personal data. Finally, prices allow individuals to make choices without revealing sensitive information about themselves; in contrast, an allocation based on labels forces agents to compromise their privacy to secure a better position in the line. Of course, our key finding is that some degree of pricing may emerge as an optimal allocation method solely for the screening benefits, that is, even when the additional benefits of prices are omitted. At the same time, using prices to allocate a potentially life-saving treatment may raise moral concerns. We return to these additional considerations regarding prices in [Section VII](#).

7. See, for example, [Castillo et al. \(2021\)](#), [Athey et al. \(2022\)](#), and [Kominers and Tabarrok \(2022\)](#).

I.A. *Relationship to Alternate Approaches*

Our work connects mechanism design to the existing analyses of medical and ethical reasons for prioritizing certain groups for early vaccine allocation. In the context of COVID-19, these reasons are discussed in detail in the [National Academies of Sciences, Engineering, and Medicine \(2020\)](#) (NASEM) framework, which has been quite influential in shaping vaccine allocation in practice. The NASEM framework is based on three foundational principles—maximum benefit, equal concern, and mitigation of health inequities (see [Persad, Peek, and Emanuel 2020](#) for a related discussion). To operationalize these principles, the study developed four risk-based criteria that determine priorities among population groups: (i) risk of acquiring infection, (ii) risk of severe morbidity and mortality, (iii) risk of negative societal impact, and (iv) risk of transmitting infection to others. The final recommendation was a four-phased approach to COVID-19 vaccine allocation, starting from front-line health workers and people with severe risk of morbidity and mortality.⁸ Our approach is related to this ethical framework in multiple ways. First and foremost, our model explains how the NASEM operational principles can be mapped into an otherwise standard economic framework of maximizing a welfare function in an economy populated by strategic and privately informed agents. Our model captures all the risk factors just described, and thus the mechanism we identify characterizes the optimal trade-off among them. In particular, in our framework, risks (i) and (ii) are captured by the private health benefit parameter; risk (iii) is modeled as the private socioeconomic benefit and the socioeconomic externality; and risk (iv) is the health externality. Because we account for private information of agents, our framework paves the way for employing the ethical principles codified by NASEM under the realistic assumption that the underlying risk-based criteria are not perfectly observable.

8. The NASEM framework also introduced a vaccine reserve, along the lines proposed by [Pathak et al. \(2022b\)](#), under which a share of vaccines allocated at any given time would be reserved “for deployment by [the Centers for Disease Control and Prevention] for use in areas of special need (identified through a vulnerability index, such as the SVI or the CCVI) or epidemiological ‘hot spots’” ([National Academies of Sciences, Engineering, and Medicine 2020](#), 105). This is broadly consistent with our findings, as it corresponds to prioritizing individuals with high risk (perhaps because they are unable to choose the safe action) and/or high welfare weight.

Moreover, at least at an informal level, the four-phase NASEM allocation scheme accords with the findings of our model in the case in which prices are ruled out. Indeed, our framework suggests that the optimal pure-priority vaccination schedule is characterized by “phases” in which only certain groups are eligible to be vaccinated. Moreover, the groups of agents that our framework identifies as having the highest priority are similar to those highlighted by NASEM. One of our results suggests that it might be optimal to vaccinate front-line health workers first because of their high health externalities and high social welfare weights.

The key distinction between our framework and the standard ethical, purely priority-based approach is that we do not rule out prices in principle. In this sense, our framework brings the classic economic idea that prices can identify those who value an object the most into the vaccine allocation problem. More importantly, we argue that prices can actually help better satisfy the ethical goals of vaccine allocation—at least in the probabilistic sense—because individuals’ willingness to pay can be informative not just about their private health and socioeconomic benefits but also about their unobserved externalities and welfare weights. We elaborate on this point in [Section VII](#). At the same time, we show that the priority-based system—as if prices were banned to begin with—can sometimes emerge as the optimal mechanism (especially when the “screening benefit” of prices is not that large).

In the matching theory literature, [Pathak et al. \(2022b\)](#) developed a model of reserve design and associated multiple-category priority system for use in the allocation of vaccines and other scarce health resources (see also [Pathak et al. 2022a](#)).⁹ Our approach is distinct from that of [Pathak et al. \(2022b\)](#) in two ways. First, we take a different approach to balancing various ethical goals that might arise in the context of allocating vaccines. [Pathak et al. \(2022b\)](#) treat ethical goals as incommensurate and show how to satisfy them in parallel with a mechanism in which vaccines are matched to agents subject to a constraint that a sufficient number of vaccines must be allocated to multiple “categories” (with each category reflecting a different ethical

9. The practical impact of this work is discussed at <https://www.covid19reservesystem.org/>. This work also spurred a medical ethics literature looking at different ways to implement multiple-reserve systems (e.g., [Pathak, Sönmez, and Ünver 2020](#); [Makhoul and Drolet 2021](#); [Sönmez et al. 2021](#)), and associated innovation in market design theory (e.g., [Delacrétaz 2020](#); [Grigoryan 2021](#)).

goal). By contrast, we model ethical goals via welfare weights, which allows us to take a more classical approach of maximizing a welfare function. Second, our mechanism is designed to deal with the problem of private information. As such, we allow prices to potentially guide the allocation scheme, which leads to material differences between our mechanism and the reserve-based priority mechanisms. Some of our conclusions are similar but have different underlying intuitions. For example, the reserve design of Pathak et al. (2022b) creates overlaps in vaccination schedules for different categories as a consequence of incommensurability of ethical considerations. We also obtain that an overlapping schedule may be optimal; in our case, this is a consequence of decreasing marginal values of vaccinating individuals in a given group when prices allow for efficient sorting.¹⁰

The question we ask is closely related to the problem of optimal targeting of vaccines (see Gans 2022 for a general discussion). 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 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 2022; and again Gans 2022). Our framework can rationalize both arguments—the former in terms of health externalities, and the latter in terms of high individual health value of vaccination.¹¹ Similarly, Vellodi and Weiss (2021a, 2021b) analyzed optimal targeting of a policy intervention (including but not limited to targeting vaccines) in a model where agents are heterogeneous, choose an endogenous response to the pandemic, and exert externalities. Schmidt et al. (2020b) and Bibbins-Domingo, Petersen, and Havlir (2021) demonstrated how targeting COVID-19 vaccines according to observables such as neighborhood characteristics can help prioritize socially

10. When we shut down the price channel, our framework predicts that groups should be strictly ordered, with no overlap in the vaccination schedule. This occurs because of linearities in our model; in essence, we assume that allocating a vaccine dose via rationing to two agents with similar observables generates the same expected social value, regardless of who gets the vaccine first.

11. Goldstein, Cassidy, and Wachter (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.

vulnerable populations—in the language of our framework, this corresponds to prioritizing populations with high welfare weights as revealed by their label (see also [Schmidt et al. 2021](#)). [Schmidt \(2020\)](#) and [Schmidt et al. \(2020a\)](#), meanwhile, argued in favor of assigning high welfare weights to disadvantaged populations both because they are generally underresourced and because they face especially high COVID-19 incidence.¹²

[Kutasi et al. \(2021\)](#) studied the decision faced by a designer who has access to various types of vaccines differing in quality. Among other questions, [Kutasi et al. \(2021\)](#) asked whether and how to allocate lower-quality vaccines that become available early on in the pandemic. They emphasized how the differentiated quality and timing may be used to screen for agent characteristics such as unobserved comorbidities or ability to work remotely; this is similar to how adding prices in our framework with a single type of vaccine allows the designer to improve screening.

To our knowledge, [Brito, Sheshinski, and Intriligator \(1991\)](#) were the first to study the trade-off between private and social values 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 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 efficient allocation. [Pancs \(2020\)](#) analyzed a fully market-based solution, modeling the problem of vaccine allocation as a “position auction” (see [Edelman, Ostrovsky, and Schwarz 2007](#); [Varian 2007](#)), in which agents can bid for positions in the vaccine queue. Crucially, in the auction proposed by [Pancs \(2020\)](#), agents also bid on behalf of others; thus, for the auction to achieve efficiency, each agent must correctly estimate and communicate the value she places on vaccinating all other agents. Our approach to externalities is different. We assume that the designer estimates—given her information on the agent—the total health and socioeconomic externalities that vaccinating that agent has on the rest of society, and the only private information our mechanism elicits from agents is their own willingness to pay. Finally, [Kang and Zheng \(2023\)](#) suggested a different type of market-like

12. [Emanuel et al. \(2020a\)](#) proposed a prioritization scheme for cross-country vaccine distribution which, like our framework, considers both health and economic harms (see also [Budish et al. 2022](#); [Dogan and Raghavan 2022](#)).

solution to the vaccine allocation problem—one in which agents can trade priority endowments.

In terms of methods, we build on the framework developed by Akbarpour[®] Dworzak[®] Kominers (forthcoming) (ADK); Appendix A explains how to adapt these methods to our model. In their model of allocation under redistributive concerns, ADK 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 article. Condorelli (2013) showed how the ironing technique of Myerson (1981) can be used to determine whether a market or a nonmarket mechanism should be used.

Finally, the inclusion of externalities connects our work to mechanism design with allocative externalities (see Jehiel, Moldovanu, and Stacchetti 1996; as well as Jehiel and Moldovanu 2001, who focused on strategic interactions between a small number of agents). Most closely related are contemporaneous papers that model externalities in large populations. Ostrizek and Sartori (2023) proposed a model that—similarly to the current article—incorporates externalities into a screening framework. Kang (2020) and Pai and Strack (2022) analyzed a mechanism-design setting where an agent's consumption of a good creates a negative externality (e.g., gasoline consumption) and examined the optimality of a Pigouvian tax scheme, rationing, and quantity ceilings.

II. 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 time t .

Before receiving a vaccine, each agent privately decides how to react to the pandemic. In the model, we assume that each agent takes a binary decision $a \in \{\text{Safe}, \text{Risky}\}$.¹³ We interpret the choice of $a = \text{Safe}$ as the agent taking precautions that significantly adjust in-person activities to minimize the risk of infection (e.g., staying at and working from home, avoiding public transit, and social distancing). The choice of $a = \text{Risky}$

13. For simplification, we conduct our analysis in a static framework in which the agent cannot condition her action on the current state of the pandemic.

represents the agent choosing to continue engaging in in-person interactions normally; 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 $a = \text{Risky}$ might simply represent seeing patients as normal, whereas $a = \text{Safe}$ could mean seeing patients online instead.¹⁴ For a retiree, the decision might be between self-isolating at home or seeing their family and friends. For a student, $a = \text{Safe}$ and $a = \text{Risky}$ may entail some amount of in-person interaction (e.g., going to class), but $a = \text{Safe}$ represents minimizing that interaction as much as possible (e.g., by choosing not to attend large social gatherings). 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. We assume that the decision a is not directly observable but—as we describe soon—we allow for observable information that may reveal information about it.

Intuitively, the decision between Safe and Risky depends on the comparison between the COVID-related risk the agent would incur by engaging in in-person interactions and the private disutility the agent suffers from taking precautions.¹⁵ The agent's decision may be socially inefficient because she ignores the externalities that her decision creates: when choosing $a = \text{Risky}$, the agent increases the probability of infection for all other agents; when choosing $a = \text{Safe}$, the agent deprives other agents of the benefits of interacting with her in-person (e.g., patients of a doctor working from home may experience a decrease in the quality of care). 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,

14. The example of doctors highlights how in referring to the $a = \text{Risky}$ action as “risky” we are just referencing inherent infection risk associated with the activity; there is no value judgment intended. Moreover, 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) may be doing jobs that are “risky” while taking maximal precaution in their lives outside of work.

15. Throughout, we use the language of COVID specifically because our illustrative examples are framed in reference to the COVID-19 pandemic; however, our framework can be used in the context of any infectious disease.

each agent is described by five characteristics, which we express in dollar values (to ensure that they can be compared to one another):

- v : the private *socioeconomic benefit* of choosing $a = \text{Risky}$ relative to $a = \text{Safe}$, 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, and so forth.
- v_{ex} : the positive *socioeconomic externality* generated by the agent choosing $a = \text{Risky}$ relative to $a = \text{Safe}$, 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 the reduction in infection risk associated with choosing $a = \text{Safe}$ relative to $a = \text{Risky}$, ignoring all other aspects of the agent's utility. For example, h may depend on the agent-specific risk of infection, presence of potential comorbidities, and expected quality-adjusted life-years.
- h_{ex} : the positive *health externality* generated if the agent minimizes the risk of her own infection by choosing $a = \text{Safe}$ instead of $a = \text{Risky}$. That is, h_{ex} is the value to society of the agent reducing her risk of spreading the virus. For example, by choosing to work from home, the agent decreases the COVID-related risk for her co-workers.¹⁶
- 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.¹⁷ All previous values are expressed in dollars, and hence v and h are affected by

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

17. This way of modeling social preferences has been used extensively in public finance; see [Saez and Stantcheva \(2016\)](#) for a general treatment.

the agent's opportunity cost of money (that could depend, for example, on the agent's 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, and so forth.

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 $a = \text{Risky}$ (of course, h will still vary by age and health status); yet, 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 $a = \text{Safe}$). Their externalities v_{ex} and h_{ex} are both large because front-line health workers provide a tremendous value to their patients by seeing them in person, yet in doing so, they interact with many people who face significant health risk. In addition the social perception of their moral desert 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 poorer drivers are more likely to choose $a = \text{Risky}$. Because ride-share drivers who do drive come into close contact with many people, their h_{ex} is high; their v_{ex} may be relatively low due to the existence of alternative means of transportation, decreased mobility during a pandemic, and higher elasticity of labor supply. A healthy college student may be an example of someone with low h but high h_{ex} ; young healthy people are generally less likely to suffer serious consequences from infection but may still play a role in transmitting the virus to more vulnerable populations, especially if their social networks are broad. A software engineer with potential comorbidities might have h higher than v and a relatively low v_{ex} because their job can be performed effectively from home. Finally, a CEO of a large company may have a large v_{ex} , as well as a high v and h due to low opportunity cost of money, but a relatively small λ if the designer has strong redistributive preferences.

In practice, agent characteristics are partially observable. For example, an individual's job may be observed, and that reveals some information about that agent's characteristics; yet factors such as attitudes, beliefs, and lifestyle may still be 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, i could be arbitrarily informative about some characteristics).¹⁸ The label i belongs to a finite set I that captures all observable features of agents on which the designer can condition her allocation; 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, we assume that agents can observe their private benefits v and h . (It does not matter for our analysis whether agents can observe anything else.) Second, we assume that 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, the externalities generated by an agent have no systematic relationship to that agent’s private benefits. This assumption could be violated to some extent in practice, but it underscores the point that decisions made by privately optimizing agents will not in general be aligned with the social objective (which will include the externalities).¹⁹

Next we proceed to specifying the payoffs. We assume that the health consequences of receiving a vaccine are the same as those of choosing $a = \text{Safe}$. As a result, each vaccinated individual enjoys utility $v + h$. This strong assumption simplifies our arguments while capturing the gist of the problem (we will flag some results whose interpretation could change if the assumption were relaxed). 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 nonnegative; we relax that assumption in [Section VII](#), where we explain how we could handle the case in which the designer might choose to pay the agents with negative h to become vaccinated. To reduce the number of cases to consider, we assume that in

18. For technical reasons, we assume that the distribution of characteristics conditional on each i is continuous.

19. 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 can easily handle this more general case, but the interpretation of the model would be less transparent.

each group i there is at least one agent with $h = 0$.²⁰ 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 $\delta(0) = 1$. Thus, the agent’s utility is given by²¹

$$(1) \quad U(v, h, t, p) := \underbrace{\delta(t) [v + h]}_{\text{post-vaccination utility}} + (1 - \delta(t)) \underbrace{[\max\{v, h\}]}_{\text{pre-vaccination utility}} - p.$$

We focus on a utilitarian objective function for the designer (in Section VI, we discuss an alternate “pure health” objective function and how our results would change). Specifically, we let $V(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda, t, p)$ be the social value of vaccinating an agent with characteristics $(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda)$ at time t and at a price p . If $\mathbf{1}_{\text{Risky}}$ and $\mathbf{1}_{\text{Safe}}$ respectively denote the event that an agent chooses $a = \text{Risky}$ and $a = \text{Safe}$ prior to receiving a vaccine, then

$$\begin{aligned} V(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda, t, p) & \\ & := \lambda U(v, h, t, p) + \delta(t) (\mathbf{1}_{\text{Safe}} v_{\text{ex}} + \mathbf{1}_{\text{Risky}} h_{\text{ex}}) + \alpha p \\ (2) \quad & \stackrel{\text{up to a constant}}{=} \delta(t) (\mathbf{1}_{\text{Safe}}(v + v_{\text{ex}}) + \mathbf{1}_{\text{Risky}}(h + h_{\text{ex}})) + (\alpha - \lambda)p. \end{aligned}$$

The designer maximizes the expectation of this function with respect to the population distribution of types, with t and p specified by the mechanism. The social value (2) encodes an important observation: if an agent chooses $a = \text{Safe}$ before receiving the vaccine, then giving the vaccine to that agent unleashes the socioeconomic private benefit v and the socioeconomic externality v_{ex} ; by contrast, if an agent chooses $a = \text{Risky}$ prior to receiving the vaccine, then giving the vaccine to that agent unleashes the private health benefit h and the health externality h_{ex} .

20. Formally, conditional on any $i \in I$, the lower bound of the support of the distribution of h is 0. This can be justified if some agents either do not believe that the vaccine is effective or have recently had COVID (previous infection is believed to provide some level of immunity).

21. Since the time frame for the vaccine allocation problem is relatively short, we ignore discounting of the monetary payment.

In addition, 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. Revenue could be used to purchase more vaccines,²² especially in the context of developing economies. From the perspective of a small country, if an international price of a vaccine is P_{vac} , and an extra vaccine is allocated for free to a poorer community generating a total value of V_{vac} , then α should be set to $\frac{V_{\text{vac}}}{P_{\text{vac}}}$ and could easily exceed the average welfare weight. Finally, the weight α could be 0 under an alternative interpretation of our model in which agents “pay” for the vaccine by “burning” utility, for example, by queueing (in that case, p is interpreted as the time spent in line required to obtain the vaccine)—a possibility that we revisit in [Section VII](#).

III. ALLOCATION MECHANISMS

In many countries, COVID-19 vaccines were allocated using a simple priority schedule. In such an allocation scheme, the population is divided into several groups based on observable and verifiable criteria (what we referred to as “labels” in our model). These groups are ordered from most to least critical, and vaccines are allocated for free to agents in 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, and the scope for using prices.

We allow the designer to optimize over all feasible allocation mechanisms, possibly using prices. By the revelation principle, for the sake of finding the optimal mechanism, we may imagine that the designer asks agents to report their characteristics $(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda)$; then, as a function of the report, the agent is promised a (potentially random) time of vaccination, and is charged a monetary transfer.²³ The mechanism must satisfy

22. This can be seen as a reduced-form way of relaxing our assumption of fixed supply of vaccines; see [Section VII](#) for an extended discussion.

23. As will become clear, agents’ reports will only depend on h and v , and hence this formulation does not require agents to know their own vector of characteristics beyond private benefits.

incentive compatibility constraints, that is, it must be optimal for each agent to report truthfully. 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 nonnegative utility from participating. The mechanism must respect physical feasibility constraints, in that, for any t , it cannot allocate more vaccines before time t than the availability $A(t)$. We also assume that all vaccines must be allocated as soon as they become available, and that prices set by the mechanism are nonnegative.²⁴

Using the fact that the time of receiving a vaccine only matters for payoffs via $\delta(t)$, we 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 cumulative distribution function F of quality q .²⁵ From now on, we treat F and q as primitives of our model, but we interchangeably use the “time” interpretation.²⁶

The fundamental difficulty facing the designer is that the parameters entering the [objective \(2\)](#) 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 (WTP) derived from their primitive private types.²⁷

24. The first of these assumptions means that the designer should not purposefully delay vaccination to increase prices and raise more revenue. The second assumption is with loss of generality when agents can have negative h ; we discuss this case in [Section VII](#).

25. The assumption that F is a proper CDF is without loss of generality because we can always assume that the designer has an unlimited number of zero-quality vaccines.

26. This transformation allows us to rely on the formalization of the mechanism design problem proposed by ADK; we omit the formal statement of the optimization problem for brevity.

27. See [Jehiel and Moldovanu \(2001\)](#), [Che, Dessein, and Kartik \(2013\)](#), and [Dworzak, Kominers, and Akbarpour \(2021\)](#) for proofs of closely related results.

LEMMA 1. It is optimal for the designer to condition the allocation of vaccines only on agents' labels i and WTP 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 the labels i .

Lemma 1 is intuitive: [expression \(1\)](#) 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 an agent is willing to pay for receiving a vaccine immediately. Two agents with the same label and WTP are behaviorally indistinguishable with regard to any mechanism with prices. If the mechanism attempted to condition the allocation on additional dimensions of the type (e.g., on the unobserved externalities 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 , which we denote $Q_i(r)$. Under our assumption that prices are nonnegative, and that the lower bound on h (and hence r) is zero 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 .²⁸ Overall, a sufficient statistic to evaluate the [objective \(2\)](#) under an incentive-compatible mechanism is $Q_i(r)V_i(r)$, where $V_i(r)$ 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.

28. This follows from the payoff equivalence theorem; see, for example, [Milgrom \(2004\)](#).

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.

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

$$\begin{aligned}
 V_i(r) &= \underbrace{\Lambda_i(r) \cdot \gamma_i(r)}_{\text{private utility}} + \underbrace{\alpha(r - \gamma_i(r))}_{\text{revenue}} \\
 (3) \quad &+ \underbrace{v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} \mid i, r) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} \mid i, r)}_{\text{externality}},
 \end{aligned}$$

where $\Lambda_i(r) := \mathbb{E}_{\tilde{r} \sim G_i} [\lambda \mid i, \tilde{r} \geq r]$ is the expected welfare weight on all agents in group i with WTP above r , and $v_{\text{ex}}^i := \mathbb{E}[v_{\text{ex}} \mid i]$ and $h_{\text{ex}}^i = \mathbb{E}[h_{\text{ex}} \mid i]$ are the expected socioeconomic and health externalities, respectively, 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 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 $v_{\text{ex}}^i = \mathbb{E}[v_{\text{ex}} \mid i]$ and $h_{\text{ex}}^i = \mathbb{E}[h_{\text{ex}} \mid i]$ 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 $\mathbb{P}(a = \text{Risky} \mid i, r)$ 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 socioeconomic externality v_{ex}^i .

For technical reasons, we assume that $V_i(r)$ is continuous in r for every i , and that the inverse hazard rate $\gamma_i(r)$ is continuous and equal to zero at the upper bound \bar{r}_i for each i .

By using monetary transfers, the designer can elicit information about WTP 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 .

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

$$\begin{aligned} \bar{V}_i := \mathbb{E}[V_i(r)] &= \mathbb{E}[\lambda \cdot r \mid i] + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} \mid i) \\ (4) \qquad \qquad \qquad &+ v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} \mid i).^{29} \end{aligned}$$

Note that the private-utility term reduces to the expectation of λr because 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 .

IV. OPTIMAL ALLOCATION WITHOUT PRICES

We first solve the problem assuming that the designer does not charge monetary transfers for the vaccines, which 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 $Q_i(r) = \int_0^1 q dF_i(q)$, for all r , 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 sufficient vaccines for everyone in group i (that is, when F_i has an

29. We adopt a convention that whenever the expectation operator is applied to a random variable indexed by a label i , the distribution is computed conditional on i ; for example, $\mathbb{E}[V_i(r)]$ denotes the expectation of $V_i(r)$ over r distributed according to G_i .

atom at zero), and randomization if there is dispersion in quality (timing) of vaccines available for group i .

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 in each group is free. Then it is optimal to vaccinate groups sequentially in the order of decreasing \bar{V}_i . That is, if $\bar{V}_j > \bar{V}_k$, under an optimal mechanism, all agents in group j are vaccinated before all agents in group k .

The intuition for [Result 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 in 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 the group i with higher \bar{V}_i .

The form of \bar{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 WTP. High welfare weights could be attached, for example, to agents who are poor, are particularly adversely affected by the pandemic, or are 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, we suppose that group i comprises front-line health workers. 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 people are engaging with COVID-19 patients directly ($\mathbb{P}(a = \text{Risky} | i) \approx 1$). Thus, [Result 1](#) suggests that front-line health workers should receive the vaccines early on. If there are no groups j with higher \bar{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 a group j consisting of people who are at high risk in case of infection (e.g., the elderly) or to a group k of people who are most likely to spread the virus (e.g., students living in dormitories). In group j , the health benefit h is high by definition, and so r is relatively high as well (at least on average, given that most people have a non-negligible socioeconomic benefit v). By contrast, since $r = \min\{v, h\}$ and h is low for most young, healthy people, r is typically low in group k . To simplify, let us approximate

$$\mathbb{E}[\lambda \cdot r \mid k] \approx 0;$$

$$\mathbb{P}(a = \text{Risky} \mid j) \approx 0; \text{ and}$$

$$\mathbb{P}(a = \text{Risky} \mid k) \approx 1.$$

Then, group j has priority over k if and only if $\mathbb{E}[\lambda \cdot r \mid j] + v_{\text{ex}}^j > h_{\text{ex}}^k$. Thus, group j should receive the vaccines earlier if their average welfare-weighted WTP plus their socioeconomic externality exceeds the health externality of group k . For instance, for elderly people living alone, 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 for students if they live in dormitories and 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 at 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 people.

V. OPTIMAL ALLOCATION WITH PRICES

In this section, we describe the optimal allocation 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](#)). Screening is achieved by charging higher prices for higher-quality vaccines (i.e., vaccines that are available earlier)—this ensures an increasing relationship between WTP and quality within a group.

An important special case is that of assortative matching between WTP and quality. Formally, if F_i is the CDF of vaccine quality allocated to group i , assortative matching means that $Q_i(r) = F_i^{-1}(G_i(r))$, where F_i^{-1} is the generalized inverse of the CDF F_i . We refer to this allocation method as a *market allocation* because it coincides with what a competitive market would achieve under the assumption of group-specific market clearing.

V.A. When to Use Prices

Before analyzing how allowing for prices affects the optimal priority across groups, we focus on a simpler question: fixing the pool of vaccines allocated to a given group i , when should the designer opt for free versus market allocation?

RESULT 2. If $V_i(r)$ is nondecreasing, then it is optimal to use a market allocation within group i . If $V_i(r)$ is nonincreasing, then it is optimal to use a free allocation within group i . In all other cases, a hybrid mechanism (in which agents are partitioned into intervals according to WTP, and allocation is either random or assortative in each interval) is optimal.³⁰

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 who contribute more, the first-best allocation would induce an anti-assortative matching; however, that is not possible due to incentive compatibility constraints. The best the designer 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). When $V_i(r)$ is nonmonotone, the optimal mechanism combines regions of random and assortative matching. [Appendix C](#) describes the exact form of the optimal mechanism in this case; here, we focus on identifying distinct economic forces that work in favor of free (random) versus

30. A random allocation is a generalization of free allocation: Allocation $Q_i(r)$ is random on an interval of WTP if $Q_i(r)$ is constant over that interval (the associated price is also constant but may be strictly positive). The economic intuition behind [Result 2](#) is similar to that underlying [Corollaries 2, 4, and 6 of Condorelli \(2013\)](#).

market (assortative) allocation. Recall from [Lemma 2](#) that

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

The monotonicity of $V_i(r)$ is thus determined by two forces that we discuss next.

1. *Private Utility + Revenue.* The term $\Lambda_i(r) \cdot \gamma_i(r) + \alpha(r - \gamma_i(r))$ coincides with the welfare function analyzed by ADK; we summarize the intuitions relevant for the context of the vaccine allocation problem.

Suppose that the designer does not have redistributive concerns (i.e., $\Lambda_i(r)$ is constant in r). In the canonical transferable utility case, α is set to the average welfare weight in group i (revenue is internally redistributed as a lump-sum payment). Then, the term $\Lambda_i(r) \cdot \gamma_i(r) + \alpha(r - \gamma_i(r))$ reduces to r and is thus always increasing. This scenario corresponds to the core economic intuition that markets are “efficient”—they maximize total WTP. Instead, setting α to zero corresponds to the scenario of “costly screening” that has also been extensively studied in the literature,³¹ where the optimal allocation depends on the monotonicity of the inverse hazard rate γ_i . Because the inverse hazard rate is decreasing for many commonly used distributions, setting $\alpha = 0$ often leads to the optimality of free allocation; economically, this case is relevant when implicit prices are nonmonetary, for example, agents have to wait in line to obtain the vaccine (we revisit this scenario in [Section VII](#)). Beyond the two extreme cases, a useful observation is that, for regular distributions (in the sense of [Myerson 1981](#)), a market allocation becomes more socially preferred when α is higher.

Now, suppose that the designer does have redistributive preferences. If the designer prefers to redistribute toward poorer agents, then $\Lambda_i(r)$ might naturally be decreasing because of a positive correlation between wealth and WTP, everything else being equal. If this effect is strong enough, it could make the

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

(private utility + revenue) term decreasing, supporting 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 this, 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 substantially with r . By 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. In summary, if the designer has strong redistributive preferences, she could opt for free allocation when labels fail to accurately identify those with the highest welfare weights.

2. *Externalities.* For groups i with high externalities (such as doctors, nurses, or teachers), the externality term may be large relative to the other two terms and hence likely to determine the monotonicity of $V_i(r)$. To understand the impact of the externality term, it is convenient to rewrite it as $v_{\text{ex}}^i + (h_{\text{ex}}^i - v_{\text{ex}}^i) \cdot \mathbb{P}(a = \text{Risky} | i, r)$. Monotonicity thus depends on (i) which externality effect, h_{ex}^i or v_{ex}^i , is stronger for group i ; and (ii) whether $\mathbb{P}(a = \text{Risky} | i, r)$ is increasing or decreasing in r . Let us consider the second factor first. Higher WTP r reflects a higher need for a vaccine, which could both be associated with the desire to engage in in-person interactions (for agents who chose $a = \text{Safe}$), as well as with the desire to protect one's health (for agents who chose $a = \text{Risky}$). Thus, monotonicity of $\mathbb{P}(a = \text{Risky} | i, r)$ in r depends on which of the two possibilities becomes more likely as r increases. Since higher r is also associated with higher wealth, it is perhaps more natural to assume that $\mathbb{P}(a = \text{Risky} | i, r)$ is decreasing in r ; if wealthier agents can either "afford" to stay at home due to savings or have jobs that are easier to perform remotely, then higher WTP is indicative of a higher probability of choosing $a = \text{Safe}$.

If $\mathbb{P}(a = \text{Risky} | i, r)$ is indeed decreasing in r , then the monotonicity of the externality term depends on the sign of $h_{\text{ex}}^i - v_{\text{ex}}^i$. If the health externality h_{ex}^i is larger than the socioeconomic externality v_{ex}^i for group i , considerations related to externalities motivate using free allocation of vaccines; but if

the socioeconomic externality is higher, the market allocation is preferred.

We consider three examples to illustrate the preceding discussion of optimal within-group allocation. First, let i be the group of ride-share drivers, cashiers, or other front-line workers. As argued already, 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 people) and could also dominate the socioeconomic externality v_{ex}^i . Thus, the designer would like to target the vaccines toward workers who are most likely to choose $a = \text{Risky}$ because vaccinating those 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 market pricing would tilt the allocation toward richer workers—the ones who are more likely to choose *Safe*.

For the second example, let i be the group of owners of small- and medium-sized enterprises. Their decision a may be whether to 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 each other. 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 may want to target the vaccines toward business owners (and their employees) who are most likely to choose $a = \text{Safe}$. 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. 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., people living in a poor neighborhood), then a free allocation may be preferred for j .

V.B. How Prices Affect Group Priorities

We now return to the question of optimal priorities across groups. Unlike in Section IV, we now allow the designer to *optimally* allocate vaccines within groups, potentially relying on prices.

Our first observation is that it might still be optimal to vaccinate some groups of agents immediately and for free—what we call *absolute priority allocation*—regardless of the optimal within-group allocation. We formalize this observation by assuming that the mass of vaccines available at time $t = 0$, i.e., $A(0)$, is large enough that some groups could receive the vaccine immediately. If it is optimal to vaccinate a group immediately, then the allocation must be free in the group since those agents all receive the same allocation.³² Let μ_i denote the mass of agents in group i .

RESULT 3. Suppose that $A(0) \geq \sum_{j \in J} \mu_j$. Then, it is optimal for groups $J \subset I$ to receive absolute priority allocation (i.e., for agents with $i \in J$ to receive vaccines immediately and for free) if

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

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

Result 3 states that—assuming it is feasible to do so—all the agents in groups in J receive absolute priority if the minimal marginal value of vaccinating an agent belonging to a group in J is higher than the maximal marginal value the designer could obtain from any agent outside of J , where the marginal values are computed subject to incentive compatibility constraints. To understand the exact form of this condition, imagine a situation in which all agents in groups in J are vaccinated, and none of the agents in groups 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 in J

32. Here, we rely on our assumption that the lower bound of the support of h in each group is zero; if instead WTP is bounded away from zero, it may be optimal—depending on the value of α —to charge a constant price equal to the joint lowest WTP in the group.

and allocating it in the best possible way to groups in $I \setminus J$. 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 in $I \setminus J$ is equal to the joint maximum over $i \notin J$ and all incentive-compatible lotteries that the designer could use to allocate that vaccine. Similarly, the marginal cost of taking away one vaccine from groups in 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 absolute priority allocation, let

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

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

$$\begin{aligned} \mathbb{E}[V_j(r) \mid r \leq x] &= \mathbb{E}[\lambda \cdot r \mid j, r \leq x] + (\Lambda_j(x) - \alpha)x \frac{1 - G_j(x)}{G_j(x)} \\ &+ \mathbb{E}[T_{\text{ex}}^j(r) \mid r \leq x], \end{aligned} \tag{5}$$

and

$$\mathbb{E}[V_i(r) \mid r \geq x] = \mathbb{E}[\lambda \cdot r \mid i, r \geq x] + (\alpha - \Lambda_i(x))x + \mathbb{E}[T_{\text{ex}}^i(r) \mid r \geq x]. \tag{6}$$

By [Result 3](#), for groups in J to receive absolute priority allocation, it must be that the value of [expression \(5\)](#) is uniformly higher (over $j \in J$ and x) than the value of [expression \(6\)](#) (over $i \notin J$ and x). [Expressions \(5\)](#) and [\(6\)](#) 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 [expression \(5\)](#) is the increase in revenue gathered from types above x when the allocation probability of types below x decreases; x in [expression \(6\)](#) 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 this discussion and [Result 3](#), providing absolute priority allocation to groups in J is more likely to be optimal when (i) these groups are associated with high welfare weights λ , (ii) the designer is not too concerned about revenue (α is relatively low), and (iii) groups in J have high externalities. This has a few implications. First, although a high welfare weight λ raises the value of [expression \(5\)](#), it is never a sufficient force on its own. This is because $\mathbb{E}[\lambda \cdot r \mid j, r \leq x] = 0$ 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 nonprioritized groups. Indeed, [expression \(6\)](#) is lower bounded by $\alpha(\max_{i \notin J} \{\bar{r}_i\})$, where \bar{r}_i could be on 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 [expression \(5\)](#) being high uniformly over x . It suffices that the label j is highly predictive of $a = \text{Safe}$ and v_{ex}^j is high, or that the label j is highly predictive of $a = \text{Risky}$ and h_{ex}^j is high.

For instance, consider j to be the group of front-line health workers. As we already argued in [Section IV](#), this group is likely to be associated with high welfare weights λ , and a high health externality h_{ex}^j . Because this label reveals that $a = \text{Risky}$ with high probability (by definition, these agents work directly with COVID-19 patients), we can think of $\mathbb{P}(a = \text{Risky} \mid j, r)$ as being approximately 1 (in particular, almost constant in r). If, moreover, the designer does not place a very high weight α on revenue, then the assumptions of [Result 3](#) are likely to hold—indicating that this group j should be prioritized.

If the designer does place a high weight α on revenue, the conclusion must be modified. Indeed, as we formalize in the next result, when α is high, the designer could benefit from selling early access to vaccines to wealthy people with high WTP.

RESULT 4. Suppose that it is optimal to use a market allocation within group i and a free allocation within group j (with Result 2 providing the supporting assumptions on the primitives). If $V_i(\bar{r}_i) > \bar{V}_j > V_i(0)$, then it is optimal to start vaccinating agents in group i first, then later vaccinate all agents in group j , and then vaccinate the remaining agents in group i after that.

The intuition for Result 4 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 in 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 vaccines.

When the weight on revenue α is high, Result 4 could apply. The designer first offers vaccines at high prices to the general population. Then, high-externality groups (e.g., health workers) are vaccinated free of charge. Finally, the vaccines are again allocated using a market mechanism, with prices gradually decreasing over time. We can further calculate the threshold price at which the first stage should stop: if J denotes the high-externality groups, then the threshold WTP p^* at which the allocation process switches to the second phase is determined by

$$(7) \quad V_{I \setminus J}(p^*) = \mathbb{E}[V_J(r)],$$

assuming that $V_{I \setminus J}(r) = \sum_{i \notin J} V_i(r)$ is nondecreasing. The left-hand side of equation (7) can be approximated by

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

if p^* 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^* can be charged approximately its WTP if it is close to the maximal WTP). The right-hand side of equation (7) is equal to

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

which is a special case of [expression \(5\)](#) 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

$$p^* \approx \sum_{j \in J} (\mathbb{E}[\lambda r \mid j] + \mathbb{E}[T_{\text{ex}}^j(r)]).$$

Since both the welfare weight λ and the expected externality in groups including health workers are high, p^* should be much higher than the average WTP of health workers, and higher than the social value of their expected externality.³³

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, as the derivation is unchanged even if

$$\mathbb{E}[\lambda \mid r \geq p^*] = 0.$$

Another explanation is that the weight on revenue α is low, at least in developed economies. However, unless the weight is zero (which seems unlikely), in the case of the extremely wealthy, $\alpha \bar{r}_i$ could still be large. The most likely explanation, in our view, is that a policy of selling vaccine access to the wealthy would have some degree of “repugnance”; we return to this point in [Section VII](#).

In our discussion, we described the last stage as using a market mechanism to allocate remaining vaccines to the “rest” of the population. However, even when using prices, the designer could still rely on labels to guide the allocation. The next result casts light on optimal allocation across groups when a market allocation is used.

RESULT 5. Suppose that it is optimal to use a market allocation within groups i and j . Then, it is optimal to vaccinate group i entirely before group j if and only if

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

33. In practice, because relatively few people would be able to pay such a high price, the first and second phases could happen simultaneously without affecting the welfare properties of the allocation in a substantial way.

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 vaccinate agents in both groups for some time, and then vaccinate the remaining agents in group j .

Result 5 stands in sharp contrast to **Result 1**. Under market allocation, there is in general significant overlap in the vaccination periods 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 a higher WTP on average than agents in group j , then simultaneous vaccination can only be achieved if agents in group i face higher prices for the vaccines.

Summarizing our findings, we emphasize again that the optimal allocation within each group influences the optimal allocation across the groups. If $V_i(r)$ is nondecreasing 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 in each group by clearing the group-specific market (independently of all the other groups). As a result of group-specific market clearing, prices will depend on labels; in particular, market-clearing prices may be zero for groups receiving absolute priority allocation (**Result 3**) while being high (at least initially) for the general population (**Result 4**). In contrast, if $V_i(r)$ is nonincreasing 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 IV**.

Finally, we illustrate a “mixed” case with an example whose additional purpose is to further clarify the role of prices.

1. *Illustrative Example.* Suppose there are three labels— $I = \{1, 2, 3\}$ —and the distribution of marginal values of vaccination in groups is such that $\bar{V}_1 > \bar{V}_2 > \bar{V}_3$. Suppose for now that prices cannot be used. Then, in line with **Result 1**, the optimal mechanism first randomly allocates vaccines to all members of group 1, then to all members of group 2, and then to all members of group 3. The marginal (flow) value for society of this allocation is depicted in the left panel of **Figure I**. Note that these marginal values are constant in each group.




FIGURE I

An Example of the Optimal Mechanism with Three Groups When Prices Can (Right) or Cannot (Left) Be Used

Now suppose prices can be used. Let us assume that $V_1(r)$ and $V_2(r)$ are nonincreasing and $V_3(r)$ is nondecreasing, and that $\bar{V}_1 > V_3(\bar{r}) > \bar{V}_2$, meaning that the marginal value of vaccinating a member of group 3 with the highest WTP is more than the average marginal value of vaccinating a member of group 2, but less than the average marginal value for group 1. In this case, the optimal vaccination schedule proceeds as follows (see the right panel of Figure I): We first vaccinate all members of group 1 for free in random order. Then we use a price schedule for group 3 with prices decreasing over time at a rate that ensures assortative matching between the highest-WTP agents in group 3 and vaccines that become available between t_1 and t_2 . When the price is such that the marginal social value of an agent in group 3 that would purchase at that price is equal to \bar{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 principal benefit of the price mechanism in this example is that it allows us to vaccinate the high-marginal-value people in group 3 earlier than the low-marginal-value people in group 3, resulting in a modified priority across groups under which some agents in group 3 are vaccinated before group 2. By contrast, when free allocation is used within group 3, all agents in group 3 are vaccinated after group 2 because $\bar{V}_2 > \bar{V}_3$ and there is no

way to identify the high-marginal-value people in group 3. This analysis also illustrates how the use of prices can lead to strictly higher welfare overall, since more social value is unlocked earlier in the distribution process.

VI. OPTIMAL ALLOCATION UNDER A “PURE-HEALTH” OBJECTIVE FUNCTION: A PARADOX

In earlier sections, we focused on a standard (at least in economics) welfare function that aggregates all agents’ utilities (directly and indirectly through externalities and the value of revenue). However, in popular discourse and the medical ethics literature, it is common to consider objective functions that focus solely on the health aspects of the problem. In this section, we apply our methods to a “pure-health” objective function and derive a paradoxical conclusion.

To analyze how our findings about the optimal mechanism change under this alternate objective, we begin by zeroing out all terms not related to health outcomes from the social value (2); that is, we eliminate v , v_{ex} , and $(\alpha - \lambda)p$. This results in the “pure-health” value function

$$(8) \quad V(v, v_{\text{ex}}, h, h_{\text{ex}}, \lambda, t, p) = \delta(t)\mathbf{1}_{\text{Risky}}(\lambda h + h_{\text{ex}}).$$

All of our formal results continue to hold under the alternative objective (8), with $V_i(r)$ redefined as

$$(9) \quad V_i(r) = r \cdot \mathbb{E}[\lambda | i, r, \text{Risky}] \cdot \mathbb{P}(\text{Risky} | i, r) + h_{\text{ex}}^i \cdot \mathbb{P}(\text{Risky} | i, r).$$

The first component of the health objective (9) corresponds to the private health benefit and is a product of three terms. The last term $\mathbb{P}(\text{Risky} | i, r)$ is the best estimate of the probability of an agent choosing $a = \text{Risky}$ conditional on information available to the designer. The second term $\mathbb{E}[\lambda | i, r, \text{Risky}]$ is the best estimate of the social welfare weight, which now conditions on the fact that the agent chose $a = \text{Risky}$. Finally, the first term reflects the agent’s WTP, which—conditional on choosing $a = \text{Risky}$ —corresponds to her health benefit. Thus, looking at the first component alone, WTP is actually closely aligned with the pure-health objective. In particular, if the designer is not too concerned about inequality in group i (i.e., if λ does

not vary with r conditional on i) and the assessed probability $\mathbb{P}(\text{Risky} | i, r)$ does not depend on r (e.g., because i already reveals that $\mathbb{P}(\text{Risky} | i, r) = 1$), then a market allocation is optimal.

However, a free allocation may be preferred under the health objective if the second component of the [objective \(9\)](#)—the health externality multiplied by the probability of choosing $a = \text{Risky}$ —is large and strongly decreasing in the WTP r . The previously considered group of ride-share drivers may serve as an illustration. Because the poorest drivers 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. Meanwhile, the health externality is large because ride-share drivers come into close contact with many people. Hence, it may be optimal to use a free allocation for ride-share drivers under the health objective.

To determine the optimal priority of groups, we compute the analogs of [expressions \(5\)](#) and [\(6\)](#):

(10)

$$\mathbb{E}[V_j(r) | r \leq x] = \mathbb{E}[\lambda r \mathbf{1}_{\text{Risky}} | j, r \leq x] + h_{\text{ex}}^j \cdot \mathbb{P}(\text{Risky} | j, r \leq x),$$

(11)

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

A key determinant of prioritized groups under the health objective is the label-revealed probability of choosing $a = \text{Risky}$. Thus, the health objective supports even more strongly the idea that health workers should receive absolute priority allocation. These groups have a high probability of choosing $a = \text{Risky}$, and a high health externality. More generally, [Result 3](#) 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.

In contrast to the utilitarian objective, the health objective is less likely to support absolute priority allocation to groups such as teachers—especially if they can teach remotely ($a = \text{Safe}$ with high probability). This is because the health objective attaches no value to socioeconomic externalities such as the benefits of in-person instruction for children and their parents.

More generally, the pure-health objective gives lowest priority to groups who are most likely to choose $a = \text{Safe}$.

The last observation implies that the pure health objective yields a somewhat paradoxical insight about vaccinating vulnerable populations (e.g., the elderly) versus people in high-transmission settings (e.g., college students).³⁴ Because people who are at high risk when infected are much more likely to choose $a = \text{Safe}$, their contribution to the health objective function is low. In contrast, healthy young people are likely to choose $a = \text{Risky}$, which means that their contribution to the health objective is large. Thus, under the pure-health objective (so long as the vaccines do indeed block transmission of the virus—see [Bubar et al. 2021](#)), potential spreaders of the virus should be vaccinated before people who are at greatest risk.³⁵ This is despite the fact that we reached the opposite conclusion under the utilitarian objective. The paradox is that *if the goal is just to maximize overall population health, then individuals whose observables indicate high private health risk (and hence a high likelihood of choosing the safe action) get lower priority than agents whose observables indicate low private health risk (and hence a high likelihood of taking the risky action)*.

This finding highlights a very strong ethical assumption—made implicitly by the pure-health objective—that no weight should be placed on the (potentially substantial) cost borne by those who choose the maximal level of precautions. To see that sharply, imagine a person who will die for sure if they choose $a = \text{Risky}$ (h is extremely high) but who suffers from being forced to stay at home (v is high). Under the pure-health objective, the value of vaccinating such a person is zero because that individual receives the health benefit h even when they are not vaccinated.³⁶ In contrast, vaccinating such a person would be highly desirable under the utilitarian objective.

34. See [Gans \(2022\)](#) for a related point.

35. In practice—unlike in our abstract model—vaccinating people who choose to self-isolate is valuable because exposure to the virus cannot be reduced to zero. Indeed, [Bubar et al. \(2021\)](#) found that total virus incidence is minimized by vaccinating the most likely spreaders but that vaccinating the elderly was still important for minimizing mortality.

36. Again (see note 35), in practice, the value is not exactly zero because no level of precautions can completely eliminate the risk of infection; nevertheless, the value is small relative to h and v .

VII. CONCLUDING REMARKS

Our baseline framework focuses on the main trade-offs associated with the choice of vaccine prioritization scheme. 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. In this section, we discuss additional points and extensions.

VII.A. *Vaccine Hesitancy*

In the baseline model, WTP was assumed to be nonnegative. But in practice, certain agents may have a negative WTP (see, for example, [Kutasi et al. 2021](#); [Gans 2023](#)). To capture vaccine hesitancy in our framework, we would simply allow for negative health benefits h , reflecting a belief by an agent that a vaccine is harmful to her health. Because of externalities, it may sometimes be optimal to provide monetary incentives for agents with negative WTP in exchange for them agreeing to getting vaccinated.

Specifically, fixing group i , suppose that the objective function $V_i(r)$ is nondecreasing. Then, the optimal allocation is to (eventually) vaccinate all agents with $r \geq r_i^*$, where $r_i^* < 0$ is the threshold WTP at which $V_i(r_i^*) = 0$, that is, at which the costs (private disutility from vaccination and the revenue loss) are equal to the benefits (positive externality from vaccination). Prices start out positive and gradually turn negative (the decline must be slow enough that agents with high WTP prefer to get the vaccine early on, rather than wait and collect a monetary payment). On the other hand, if $V_i(r)$ is decreasing (and nonnegative in expectation), then all agents in group i should be vaccinated in random order. In this case, all agents collect a lump-sum payment for getting vaccinated (even the ones whose WTP is positive), with the payment chosen to convince the most negative-WTP agent to participate. A decreasing $V_i(r)$ may indeed arise when low (negative) r is related to skepticism about the pandemic that results in disobeying 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.

Paying people to get vaccinated may raise ethical concerns. For example, [Satz \(2010\)](#) and [Sandel \(2012\)](#) challenged the classical economic principle that a voluntary transaction can never

leave the involved parties worse off. They point out that the monetary payment could introduce an element of coercion if agents lack a sufficient understanding of the medical consequences (e.g., potential negative side effects of the vaccine) or are “forced” to accept due to their financial situation. It is also known that negative prices can alter agents’ perception of the value of the transaction (Roth 2007). Our framework does not incorporate these additional considerations. That being said, if vaccine hesitancy results from low WTP combined with additional costs of getting vaccinated (e.g., the costs associated with finding a vaccine provider, or of missing days at work), our framework predicts that lowering these costs would be socially beneficial—and this conclusion does not appear to raise any ethical concerns.³⁷

Offering monetary compensation to overcome vaccine hesitancy has been used in practice.³⁸ More commonly, however, governments instead imposed restrictions on those not vaccinated (such as limiting access to public spaces and buildings). Although we do not study interactions with other types of policy interventions, our framework may nevertheless be helpful in thinking through the effects of other policies on the optimal allocation of vaccines. Forbidding certain activities prior to getting vaccinated can be modeled as reducing the socioeconomic benefit v . If an agent’s v is reduced by $d < v$, and that agent had a negative h , then their WTP becomes $h + v - \max\{h, v - d\} = h + d$. At the same time, the WTP for those choosing the safe action remains unchanged. If the restrictions are severe enough (d is large), then WTP may become positive, and our baseline analysis applies. If $V_i(r)$ were decreasing to begin with, such a shift in WTP would typically maintain its monotonicity, and the optimal allocation would be to offer vaccines at a zero price. Thus, policies restricting socioeconomic activities prior to getting vaccinated may serve as a substitute for offering monetary compensation.

Finally, we note that a well-documented phenomenon producing behavior somewhat observationally similar to vaccine

37. For empirical evidence that small monetary compensation can lead to increased vaccination rates, see Campos-Mercade et al. (2021). The medical literature is split on the question of whether paying people to be vaccinated is morally acceptable—see Largent and Miller (2021) for arguments against negative prices, and Persad and Emanuel (2021) for the counterarguments. However, there seems to be some degree of consensus that compensating people for vaccine-related expenses is morally permissible.

38. See, for example, Oza (2021) and Terrell (2021).

hesitancy is that demand for health services often drops discontinuously when prices become strictly positive, especially among poorer populations (see [Newhouse 1993](#)). Our model can capture this phenomenon with an atom at $r = 0$ in the distribution of WTP. Exactly as we might expect, when the social value $V_i(0)$ of vaccinating agents with $r = 0$ is particularly high, $V_i(r)$ will be decreasing in the neighborhood of $r = 0$, and hence the optimal vaccination schedule will have a stage with free allocation to agents in group i .

VII.B. *Elastic Supply of Vaccines*

Our model takes the supply of vaccines as given, abstracting away from the potential relationship between the price of vaccines and total vaccine supply. While this assumption is stylized, it allows us to show that prices can be useful even if they do not increase supply. Assuming fixed supply maps directly to the vaccine allocation problem in at least two ways. First, from the perspective of most countries, the supply of vaccines during the COVID-19 pandemic was essentially fixed in the short run—and certainly did not depend on the within-country allocation mechanism. Second, in general, countries constrained prices to be low—often zero—and instead subsidized supply, effectively choosing the supply directly ([Castillo et al. 2021](#); [Kominers and Tabarrok 2022](#)).

Having said that, our framework can indirectly model the supply effects via the weight on revenue α . Especially for developing economies, monetary costs may be a 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—and this favors a market allocation. The most likely outcome is the coexistence of a public market where vaccines are allocated at low or zero prices to groups with the highest externalities, and a private market with relatively high market-clearing prices that generate substantial revenue.³⁹

VII.C. *The Choice of Labels*

For simplification, we treated the set of labels as given. From a purely theoretical perspective, the more information the designer can access—even if the information is not directly related

39. Such a solution was advocated, for example, by [Rajagopalan \(2020\)](#).

to the problem at hand—the better the performance of the optimal mechanism (versions of this argument can be recognized in the [Akerlof 1978](#) idea of tagging for tax purposes, and in the [Holmström 1979](#) informativeness principle for moral hazard problems). In practice, however, using a larger set of labels leads to trade-offs. First, verifying eligibility based on detailed characteristics can be administratively costly, especially when intermediaries (such as pharmacies) lack incentives to exert effort or if it is easy for agents to “forge” labels. Second, some characteristics could be directly relevant and easy to certify but such that agents would be hesitant to reveal them because of privacy concerns (e.g., their detailed health status). Third, using certain labels (such as gender or race) may raise ethical or fairness concerns. Finally, a system based on many labels would necessarily be complex.

We have argued throughout that the additional role of prices in screening for welfare weights and externalities relies on the existence of residual correlations (conditional on labels) between WTP and the variables of interest. Consequently, as labels get more informative, the role of prices in the optimal mechanism diminishes. Yet the discussion suggests that there are practical limits on how informative the labels can be, implying that prices can be useful as a screening device even if the underlying information is, in principle, accessible to the designer. While formalizing the trade-off is beyond the scope of this article, practical experience suggests that labels should be relatively coarse, easy to verify, and difficult to manipulate; frameworks such as the NASEM guidelines ([National Academies of Sciences, Engineering, and Medicine 2020](#)) advocate for using only a few broad categories based on age, occupation, and documented health conditions,⁴⁰ so it is possible that the marginal screening improvement achieved by employing prices could be substantial.

VII.D. Across-Country Allocation

While we focused on within-country allocation when interpreting our results, our framework can also inform the optimal allocation of vaccines across countries. It is interesting to point out that while a vast majority of countries decided not to use prices for the internal allocation of COVID-19 vaccines, the across-country

40. In addition, [Wrigley-Field et al. \(2021\)](#) argue in favor of using geographic location as a label, mainly as a proxy to obtain a more equitable vaccine allocation along racial and ethnic characteristics.

allocation was almost entirely market-based, with high-income countries receiving a disproportionately large supply of vaccines early on; attempts to ensure a more equitable distribution of vaccines across countries, such as the one that COVAX tried to implement (see [Budish et al. 2022](#)), were largely unsuccessful in mitigating global vaccine inequality. As our framework makes clear, neither of the extremes—pure priority allocation or pure market allocation—is likely to be optimal.

Setting aside political considerations, it is instructive to use our framework to consider the problem of optimal allocation of vaccines at an international level. In this case, an agent's country is simply part of their label. If agents from developing economies receive higher welfare weights as a result of their lower expected wealth (or higher health need in the absence of a well-developed health care system), then the optimal allocation of early vaccines would be less skewed toward rich countries than the allocation that arose in practice in the COVID-19 pandemic. Even with equal welfare weights, there may be reasons to favor equitable vaccine allocation across countries; for example, there is an enormous positive health externality associated with widespread global vaccination to limit future (potentially dangerous) mutations ([Gilbert and Hatchett 2021](#)). Finally, our model highlights the social value of vaccinating specific target populations such as the especially vulnerable—potentially motivating approaches like the “fair priority” model of [Emanuel et al. \(2020a\)](#), which aims for vaccines to be allocated across countries to prevent greater and more urgent harms on the margin.

VII.E. *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 $\alpha = 0$ because the designer does not benefit from this (inherently wasteful) activity. All of our mathematical results continue to hold in such a model; however, they must be reinterpreted accordingly. For example, a “market allocation” now means that people who spend the most time in the queue get the vaccine first. Meanwhile, a “free” allocation means that there is no queue and a lottery is used to determine priority. In such a context, some of our assumptions may naturally be reversed. For example, we argued that poorer agents may have a lower WTP r , all else being

equal. When values are measured in terms of disutility from waiting in a line, it may be the case that poorer agents are associated with higher r , which now becomes “willingness to queue.”⁴¹

Analogously to how prices bias access in favor of the rich, queueing biases access in favor of those most “able to queue.” For example, younger and healthier people may find it relatively less costly to stand in line, and these characteristics may correlate negatively with the social value of vaccination. Moreover, the exact context and nature of the “costly activity” matters: standing in a physical line will lead to different sorting than waiting on hold to make an appointment over the phone. Although our framework captures these effects by modeling the joint distribution of the effective private cost (captured by the parameter r) and social values of vaccination, the detailed economic analysis of optimal queueing is beyond the scope of our article. For related discussions of using “ordeals” in determining access to health care (including moral ramifications), see [Eyal, Romain, and Robertson \(2018\)](#), [Rose \(2021\)](#), or [Zeckhauser \(2021\)](#).

It is natural to ask whether the biases introduced by either pricing or queueing could be avoided if agents were instead paying with “vaccine tokens” issued and distributed by the government. A time-varying number of tokens required to get the vaccine would serve as the “price.” Dependence of the allocation on labels could be achieved by issuing different numbers of tokens to different groups. The variable r —the WTP for vaccines in tokens—would be determined by their opportunity cost, which would depend on the allowed uses of tokens. At one extreme, if tokens were non-tradable and only exchangeable for a vaccine, their opportunity cost would be zero, and hence they would not serve any screening purpose. At the other extreme, if tokens were exchangeable for money or fully tradable (like in the proposal put forward by [Kang and Zheng 2023](#)), then they would implement a market allocation of vaccines augmented with label-contingent lump-sum payments, which could help relax the dependence of WTP r on ability to pay. By varying the degree of tradability of the tokens (e.g., by implementing a “tax” for exchanging them for money, or making tokens exchangeable only for some other forms of publicly

41. Poorer agents may have higher willingness to queue because of lower shadow cost of time relative to money—although note that this is not always the case, for example, because poorer agents may have less ability to take time off from work to get vaccinated.

provided health care), the designer could to some extent engineer the distribution of r . That being said, vaccine tokens would have problems of their own. Primarily, they could reduce vaccination rates in some critical groups: among high-priority groups due to the high opportunity cost of tokens (high resale price) in early stages of the pandemic, and among vaccine skeptics due to an increase in their effective value of not getting vaccinated. In addition, tradable tokens would grant access to early vaccination for very rich people willing to buy out enough tokens in the secondary market, without the benefit of raising revenue for the designer.

VII.F. Decentralized Implementation

In our approach so far, we have focused on the optimal allocation (and potential payments); we have not discussed how they can be implemented in practice. A pure priority system, like the one described in [Section IV](#), 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 require pharmacies (and other entities administering vaccines) to verify personal information in order to determine the price paid by the patient, in a way similar to how identity of the patient is used to determine copayment under a health insurance scheme.⁴²

VII.G. Prices and Ethics

[Roth \(2007\)](#) has discussed the idea that repugnance acts as a constraint on market design. In our context, the use of monetary payments may indeed raise various moral concerns. Selling vaccines at high prices necessarily creates correlation between wealth and health. [Walzer \(1983\)](#) argued that access to care in

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

the health sphere should be governed by the principle of need, and thus should not depend on the ability to pay. Relatedly, [Satz \(2010\)](#) pointed out that the role of markets should be limited if their operation could lead to deepening inequality along dimensions such as health outcomes. A further objection has been raised by [Sandel \(2012\)](#), who observed that the nature of the good being allocated may be altered by the fact that a monetary payment is involved—in this context, the social value of the scientific discovery (an effective vaccine) may be diminished by its commodification. Finally, [Kass \(1997\)](#) argued that the instinctive reaction of disgust to the idea of “selling vaccines to millionaires” should be taken seriously, as it is likely an indication of a deeper moral intuition.

A stripped-down version of our model may shed light on some of these concerns. Suppose that a designer faces a population with an equal fraction of “vulnerable” and “nonvulnerable” agents. The designer controls treatment that is effective when applied to vulnerable agents but has small social value when applied to nonvulnerable agents. Suppose further that half of the agents are “poor” and half are “rich,” and that wealth is unrelated to vulnerability. To simplify the argument, suppose that the designer has the capacity to treat 75% of the population but 25% will have to be denied treatment.

If the designer directly observes who the vulnerable agents are, then all previously discussed ethical and economic principles will be in agreement that we should treat all vulnerable agents (the only disagreements could arise over which nonvulnerable agents to treat). However, this outcome is no longer feasible when vulnerability is each agent’s private information. Without access to vulnerability data, the designer can at most hope to reach 75% of the vulnerable population by assigning treatment randomly.⁴³ Prices may help restore the desirable outcome in this case. Suppose that the rich will always pay the price (within some reasonable range); the poor nonvulnerable will never pay it; and the poor vulnerable will be willing to pay a small price. Then, charging a small price ensures that 100% of vulnerable agents are treated. Therefore, in this example, any social theory that

43. Of course, in practice, the designer will have some information about vulnerability (captured by labels in our formal model). Our argument here applies to the residual information asymmetry between the planner and agents that is likely to persist even if the designer has access to some data. The argument has more force when that residual uncertainty of the designer is substantial.

evaluates outcomes based solely on the vulnerability criterion must prefer the price allocation. Any social theory that evaluates outcomes based on the final allocation of treatment and puts positive weight on treating the vulnerable agents can reject the price allocation only if it puts negative value on the fact that disproportionately many rich agents are treated.

Although it is extremely stylized (in a way that overstates the effectiveness of prices in more realistic cases), the preceding example is representative of the general trade-offs associated with using prices. Relative to [Walzer \(1983\)](#), our framework illustrates that adding payments to the system may help identify the agents most in need of the health treatment; however, due to the imperfect nature of WTP as a screening device, pricing will induce some degree of undesirable correlation between health outcomes and ability to pay. Our analysis thus shows that there is an inherent trade-off in Walzer's principle, the resolution of which may depend on the extent to which WTP correlates with need rather than wealth. One corollary is that the principle may be consistent with moderate prices but would probably rule out very high prices. Addressing Satz's argument, we note that while selling vaccines to the rich has the immediate consequence of raising inequality in access to health care, the resulting monetary revenue could be subsequently used (e.g., by increasing the supply of vaccines) to provide free access to the poor, thereby actually achieving more health equity in the end. Meanwhile, economic arguments are ineffective against direct moral criticisms of prices such as the commodification argument of Sandel and the "wisdom of repugnance" idea of Kass. It is an open question how policy makers should trade off these moral sentiments—assuming they are indeed shared by a significant fraction of the population—against the efficiency gains of using prices.

APPENDIX

Our model can be solved using techniques developed by ADK.⁴⁴ Even though ADK used a different objective function,

44. These techniques can 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 \(2012\)](#) applied it to multiple goods with heterogeneous quality. [Loertscher and Muir \(2022\)](#) relied on similar techniques to solve a problem of a revenue-maximizing seller in the presence of

their method applies for any objective function that is linear in the quality of the good allocated to every agent. As explained in [Section III](#), 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. Proofs of [Lemmata 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 $q(\min\{v, h\}) - p + \text{const}$ (see [expression \(1\)](#)). Thus, by defining $r = \min\{v, h\}$, we obtain that r is the WTP 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 by [Jehiel and Moldovanu 2001](#); [Che, Dessein, and Kartik 2013](#); and [Dworczak](#) [Kominers](#) [Akbarpour 2021](#)). 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.

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 [expression \(2\)](#))

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

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 $Q_i(r)$ is the expected quality and $t_i(r)$ is the payment charged to an agent with label i and

resale; [Ashlagi, Monachou, and Nikzad \(2020\)](#) showed that these methods can be also used in designing the optimal dynamic allocation in a multigood environment by optimizing over how much information is disclosed about different types of objects; finally, [Kleiner, Moldovanu, and Strack \(2021\)](#) demonstrated 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.

WTP r .⁴⁵ The expected payoff for the designer from using such a mechanism is

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

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

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

ADK showed that in an incentive-compatible mechanism with nonnegative transfers (as is assumed here),

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

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

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

45. 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.

where $V_i(r)$ is as defined in [expression \(3\)](#). Moreover, the relevant results of ADK apply to our setting by substituting $V_i(r)$ in for $\tilde{V}_i(r)$.

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 $\mathbb{E}[V_i(r)]$, which is equal to

$$(12) \quad \int_0^{\bar{r}_i} (\Lambda_i(r)\gamma_i(r) + \alpha(r - \gamma_i(r)) + v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} \mid i, r) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} \mid i, r)) dG_i(r).$$

A simple calculation (using integration by parts) shows that under random allocation, the first term in the integrand of [expression \(12\)](#) integrates out to $\mathbb{E}[\lambda \cdot r \mid i]$; the second term disappears (revenue in a mechanism without prices is zero); and the last two terms integrate to $v_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Safe} \mid i) + h_{\text{ex}}^i \cdot \mathbb{P}(a = \text{Risky} \mid i)$.

B. Proofs of [Results 1–5](#)

To derive [Results 1–2](#), we restate the results of ADK in our context. To identify an optimal mechanism with prices, we proceed in two steps:

- (i) First, vaccines are allocated optimally “across” groups: F is split into $|I|$ CDFs F_i .
- (ii) 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 refer to these two steps as the “within problem” and the “across problem,” respectively (see ADK for formal definitions of these optimization problems).

For the setting without prices ([Section IV](#)), 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 in each group (the designer cannot use prices), the value from allocating a unit of quality to group i is simply \bar{V}_i , as defined in [Lemma 3](#).

Therefore, the across problem can be formally written as

$$(13) \quad \max_{(F_i)_{i \in I}} \left\{ \sum_{i \in I} \left(\mu_i \bar{V}_i \int_0^1 q dF_i(q) \right) \right\},$$

$$(14) \quad \text{s.t. } \sum_{i \in I} \mu_i F_i(q) = F(q), \forall q \in Q.$$

It follows immediately that in any optimal solution, $\max(\text{supp}(F_i)) \leq \min(\text{supp}(F_j))$ whenever $\bar{V}_i < \bar{V}_j$, which corresponds to the statement that all agents in group j receive a weakly higher-quality vaccine (i.e., all agents in group j are vaccinated earlier) than any agent in group i . (If $\bar{V}_i = \bar{V}_j$, 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](#).

To prove the remaining results, we restate two theorems from ADK. The first one describes the solution to the within problem (with prices), and the second one describes the solution to the across problem. The statements differ slightly from ADK because of 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 nonnegative lower bound \underline{r}_i).

THEOREM 1 (ADK). Define

$$\Psi_i(x) := \int_x^1 V_i(G_i^{-1}(y)) dy.$$

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

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

where $\text{co}(\Psi_i)$ denotes the concave closure of Ψ_i . An optimal solution is given by an expected-quality schedule

$Q_i^*(r) = \Phi_i^*(G_i(r))$, where Φ_i^* is nondecreasing and satisfies

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

for almost all x .

THEOREM 2 (ADK). Let $\widehat{V}_i(x) \equiv -\text{co}(\Psi_i)'(x)$. There exists a nondecreasing function $V^{\min}(q)$ such that for all i and q , the optimal solution $(F_i^*)_{i \in I}$ to the across problem satisfies

$$\begin{cases} F_i^*(q) = 0 & \text{if } \widehat{V}_i(0) > V^{\min}(q), \\ F_i^*(q) = 1 & \text{if } \widehat{V}_i(1) < V^{\min}(q), \\ F_i^*(q) \text{ solves } \widehat{V}_i(F_i^*(q)) = V^{\min}(q) & \text{otherwise.} \end{cases}$$

Moreover, $V^{\min}(q) = \min_{i: F_i^*(q) < 1} \{\widehat{V}_i(F_i^*(q))\}$.

PROOF OF RESULT 2. The proof follows directly from [Theorem 1](#). When $V_i(r)$ is nondecreasing, Ψ_i is concave; hence, $\text{co}(\Psi_i) = \Psi_i$ everywhere. Thus, $Q_i^*(r) = F_i^{-1}(G_i(r))$, corresponding to a market allocation. When $V_i(r)$ is nonincreasing, Ψ_i is convex; hence, $\text{co}(\Psi_i) > \Psi_i$ on the interior of the domain.⁴⁶ Thus, $Q_i^*(r) = \int_0^1 qdF_i(q)$, corresponding to a free allocation.

PROOF OF RESULT 3. The first condition for absolute priority allocation to groups J , $A(0) \geq \sum_{j \in J} \mu_j$ is clearly necessary, because if it does not hold, then it is not feasible to vaccinate all agents in groups in J immediately (at time 0). If that condition holds, then a sufficient condition for absolute priority allocation to groups in J can be deduced directly from [Theorem 2](#). Indeed, all agents in groups in J receive the vaccines before any other group if, for all $j \in J$ and $i \notin J$, the slope of $\text{co}(\Psi_j)$ at 1 is lower than the slope of $\text{co}(\Psi_i)$ at 0; moreover, this condition is necessary when $A(0) = \sum_{j \in J} \mu_j$. The slope of $\text{co}(\Psi_j)$ at 1 is equal to $\min_x \{\mathbb{E}[V_j(r) | r \leq x]\}$, and the slope of $\text{co}(\Psi_i)$ at 0 is equal to $\max_x \{\mathbb{E}[V_i(r) | r \geq x]\}$, by direct calculation, proving the result.

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

PROOF OF RESULT 4. This result follows directly from [Theorems 1](#) and [2](#). For a group j with free allocation, $\text{co}(\Psi_j)$ is linear, so that the slope $\widehat{V}_j(x)$ is constant in x , equal to \overline{V}_j . For a group i with market allocation, we have $\text{co}(\Psi_i) = \Psi_i$, and so $\widehat{V}_i(0) = V_i(0)$, while $\widehat{V}_i(1) = V_i(\bar{r}_i)$.

PROOF OF RESULT 5. This result follows directly from [Theorems 1](#) and [2](#), given that for a group i with market allocation, we have $\widehat{V}_i(0) = V_i(0)$ and $\widehat{V}_i(1) = V_i(\bar{r}_i)$.

C. An Algorithm to Determine the Optimal Vaccine Allocation Scheme

We explain how the optimal vaccine allocation mechanism can be computed, in several steps:

- (i) **Computation of social values from the primitives of the model:** Compute the function $V_i(r)$, given by [expression \(3\)](#), for each group i . In practice, this step requires estimating the distribution of agents' characteristics entering the social objective function conditional on label i and WTP r .
- (ii) **Determination of the mode of allocation (prices versus randomization) within each group, via ironing:** For each group i , compute the function $\Psi_i(x)$, as defined in [Theorem 1](#); concavify that function and differentiate it to obtain $\widehat{V}_i(x) \equiv -\text{co}(\Psi_i)'(x)$, which can be thought of as the "ironed" social value function. Intuitively, random allocation of quality will be used on intervals of WTP (in the quantile space) on which $\widehat{V}_i(x)$ is constant (due to ironing); assortative allocation of quality (via prices) will be used on intervals of WTP on which $\widehat{V}_i(x)$ is strictly increasing.
- (iii) **Determination of the optimal across-group allocation:** Starting from the distribution F of vaccine quality (obtained from the dynamic supply of vaccines, as described in [Section III](#)), apply the greedy procedure described in [Theorem 2](#) to obtain the distribution of vaccine quality F_i^* that is optimally made available to each group i . Intuitively, the greedy procedure of [Theorem 2](#) allocates vaccine qualities "sequentially" (from the lowest quality to the highest) to the group i that has the lowest marginal social value $\widehat{V}_i(x)$ (computed at the optimal within-group

- allocation) when a fraction x of agents in group i were already matched to a vaccine quality.
- (iv) **Within-group matching:** For each group i , having obtained the distribution of vaccine quality $F_i = F_i^*$ in the previous step, apply [Theorem 1](#) to characterize the optimal expected quality $Q_i^*(r)$ allocated to an agent with WTP r in group i . Intuitively, this step combines steps (ii) and (iii) by matching agents in group i to available vaccines, where vaccines are ordered from lowest quality to the highest, and agents are ordered according to their ironed social values (that is, the ordering is random among all agents with the same ironed social value).
- (v) **Pricing:** Finally, using the envelope formula, compute the supporting transfers:

$$t_i^*(r) := Q_i^*(r)r - \int_0^r Q_i^*(\tau)d\tau.$$

The collection $(Q_i^*(r), t_i^*(r))_{i,r}$ describes the expected vaccination time of an agent with label i and WTP r , along with the associated price. Alternatively, the allocation pins down the price of each quality level q of the vaccine offered to agents with label i . To obtain the price schedule as a function of time, we can then map qualities of vaccines q back to the timing of the vaccine t through the mapping $q = \delta(t)$.

STANFORD UNIVERSITY, UNITED STATES

UNIVERSITY OF CHICAGO, UNITED STATES

NORTHWESTERN UNIVERSITY, UNITED STATES, AND GROUP FOR RESEARCH IN APPLIED ECONOMICS, POLAND

HARVARD BUSINESS SCHOOL, HARVARD UNIVERSITY, AND A16Z CRYPTO, UNITED STATES

REFERENCES

- Akbarpour, Mohammad [®] Piotr Dworzak [®] Scott Duke Kominers, “Redistributive Allocation Mechanisms,” *Journal of Political Economy* (forthcoming).
- Akerlof, George A., “The Economics of ‘Tagging’ as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning,” *American Economic Review*, 68 (1978), 8–19. <https://doi.org/10.1093/oso/9780199253906.003.0006>.
- Ashlagi, Itai, Faidra Monachou, and Afshin Nikzad, “Optimal Dynamic Allocation: Simplicity through Information Design,” Working Paper, Stanford University, 2020. <http://dx.doi.org/10.2139/ssrn.3610386>.
- Athey, Susan, Juan Camilo Castillo, Esha Chaudhuri, Michael Kremer, Alexandre Simoes Gomes, and Christopher M. Snyder, “Expanding Capacity for Vaccines against Covid-19 and Future Pandemics: A Review of Economic

- Issues,” *Oxford Review of Economic Policy*, 38 (2022), 742–770. <https://doi.org/10.1093/oxrep/grac037>.
- Bibbins-Domingo, Kirsten, Maya Petersen, and Diane Havlir, “Taking Vaccine to Where the Virus Is—Equity and Effectiveness in Coronavirus Vaccinations,” *JAMA Health Forum*, 2 (2021), e210213. <https://doi.org/10.1001/jamahealthforum.2021.0213>.
- Boettke, Peter, and Benjamin Powell, “The Political Economy of the COVID-19 Pandemic,” *Southern Economic Journal*, 87 (2021), 1090–1106. <https://doi.org/10.1002/soej.12488>.
- Bown, Chad P., “Covid-19 Vaccine Supply Chains and the Defense Production Act,” *Oxford Review of Economic Policy*, 38 (2022), 771–796. <https://doi.org/10.1093/oxrep/grac026>.
- Brito, Dagobert L., Eytan Sheshinski, and Michael D. Intriligator, “Externalities and Compulsory Vaccinations,” *Journal of Public Economics*, 45 (1991), 69–90. [https://doi.org/10.1016/0047-2727\(91\)90048-7](https://doi.org/10.1016/0047-2727(91)90048-7).
- Bubar, Kate M., Kyle Reinholt, Stephen M. Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H. Grad, and Daniel B. Larremore, “Model-Informed COVID-19 Vaccine Prioritization Strategies by Age and Serostatus,” *Science*, 371 (2021), 916–921. <https://doi.org/10.1126/science.abe6959>.
- Budish, Eric, Hannah Kettler, Scott Duke Kominers, Erik Osland, Canice Prendergast, and Andrew A. Torkelson, “Distributing a Billion Vaccines: COVAX Successes, Challenges, and Opportunities,” *Oxford Review of Economic Policy*, 38 (2022), 941–974. <https://doi.org/10.1093/oxrep/grac038>.
- Bulow, Jeremy, and John Roberts, “The Simple Economics of Optimal Auctions,” *Journal of Political Economy*, 97 (1989), 1060–1090. <https://doi.org/10.1086/261643>.
- Campos-Mercade, Pol, Armando N. Meier, Florian H. Schneider, Stephan Meier, Devin Pope, and Erik Wengström, “Monetary Incentives Increase COVID-19 Vaccinations,” *Science*, 374 (2021), 879–882. <https://doi.org/10.1126/science.abm0475>.
- Castillo, Juan Camilo, Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Tasneem Chifty, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Greg Larson, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Wiecek, “Market Design to Accelerate COVID-19 Vaccine Supply,” *Science*, 371 (2021), 1107–1109. <https://doi.org/10.1126/science.abg0889>.
- Chakravarty, Surajeet, and Todd R. Kaplan, “Optimal Allocation without Transfer Payments,” *Games and Economic Behavior*, 77 (2013), 1–20. <https://doi.org/10.1016/j.geb.2012.08.006>.
- Che, Yeon-Koo, Wouter Dessein, and Navin Kartik, “Pandering to Persuade,” *American Economic Review*, 103 (2013), 47–79. <https://doi.org/10.1257/aer.103.1.47>.
- Condorelli, Daniele, “What Money Can’t Buy: Efficient Mechanism Design with Costly Signals,” *Games and Economic Behavior*, 75 (2012), 613–624. <https://doi.org/10.1016/j.geb.2012.02.018>.
- , “Market and Non-Market Mechanisms for the Optimal Allocation of Scarce Resources,” *Games and Economic Behavior*, 82 (2013), 582–591. <https://doi.org/10.1016/j.geb.2013.08.008>.
- Delacrétaz, David, “Processing Reserves Simultaneously,” Working Paper, University of Manchester, 2020.
- Doğan, Battal, and Madhav Raghavan, “Equitable Allocation of Vaccines in a Supply Network,” Working Paper, University of Bristol, 2022.
- Dworczak, Piotr [®] Scott Duke Kominers [®] Mohammad Akbarpour, “Redistribution through Markets,” *Econometrica*, 89 (2021), 1665–1698. <https://doi.org/10.3982/ECTA16671>.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 97 (2007), 242–259. <https://doi.org/10.1257/aer.97.1.242>.

- Emanuel, Ezekiel J., Govind Persad, Adam Kern, Allen Buchanan, Cécile Fabre, Daniel Halliday, Joseph Heath, Lisa Herzog, R. J. Leland, and Ephrem T. Lemango et al., "An Ethical Framework for Global Vaccine Allocation," *Science*, 369 (2020a), 1309–1312. <https://doi.org/10.1126/science.abe2803>.
- Emanuel, Ezekiel J., Govind Persad, Ross Upshur, Beatriz Thome, Michael Parker, Aaron Glickman, Cathy Zhang, Connor Boyle, Maxwell Smith, and James P. Phillips, "Fair Allocation of Scarce Medical Resources in the Time of Covid-19," *New England Journal of Medicine*, 382 (2020b), 2049–2055. <https://doi.org/10.1056/NEJMs2005114>.
- Eyal, Nir, Paul L. Romain, and Christopher Robertson, "Can Rationing through Inconvenience Be Ethical?" *Hastings Center Report*, 48 (2018), 10–22. <https://doi.org/10.1002/hast.806>.
- Gans, Joshua S., "Optimal Allocation of Vaccines in a Pandemic," *Oxford Review of Economic Policy*, 38 (2022), 912–923. <https://doi.org/10.1093/oxrep/grac027>.
- , "Vaccine Hesitancy, Passports and the Demand for Vaccination," *International Economic Review*, 64 (2023), 641–652. <https://doi.org/10.1111/iere.12609>.
- Gilbert, Sarah, and Richard Hatchett, "No One Is Safe Until We Are All Safe," *Science Translational Medicine*, 13 (2021), eabl9900. <https://doi.org/10.1126/scitranslmed.abl9900>.
- Goldstein, Joshua R., Thomas Cassidy, and Kenneth W. Wachter, "Vaccinating the Oldest Against COVID-19 Saves both the Most Lives and Most Years of Life," *Proceedings of the National Academy of Sciences*, 118 (2021), e2026322118. <https://doi.org/10.1073/pnas.2026322118>.
- Grigoryan, Aram, "Effective, Fair and Equitable Pandemic Rationing," Working Paper, University of California, San Diego, 2021.
- Hartline, Jason D., and Tim Roughgarden, "Optimal Mechanism Design and Money Burning," in *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, (New York, NY: Association for Computing Machinery, 2008), 75–84. <https://doi.org/10.1145/1374376.1374390>.
- Holmström, Bengt, "Moral Hazard and Observability," *Bell Journal of Economics*, 10 (1979), 74–91. <https://doi.org/10.2307/3003320>.
- Jehiel, Philippe, and Benny Moldovanu, "Efficient Design with Interdependent Valuations," *Econometrica*, 69 (2001), 1237–1259. <https://doi.org/10.1111/1468-0262.00240>.
- Jehiel, Philippe, Benny Moldovanu, and Ennio Stacchetti, "How (Not) to Sell Nuclear Weapons," *American Economic Review*, 86 (1996), 814–829.
- Kang, Mingshi, and Charles Z. Zheng, "Optimal Design for Redistributions among Endogenous Buyers and Sellers," *Economic Theory*, 75 (2023), 1141–1180. <https://doi.org/10.1007/s00199-022-01442-4>.
- Kang, Zi Yang, "Markets for Goods with Externalities," Working Paper, Stanford University, 2020.
- Kass, R. Leon, "The Wisdom of Repugnance: Why We Should Ban the Cloning of Humans," *New Republic*, 216 (1997), 17–26.
- Kleiner, Andreas, Benny Moldovanu, and Philipp Strack, "Extreme Points and Majorization: Economic Applications," *Econometrica*, 89 (2021), 1557–1593. <https://doi.org/10.3982/ECTA18312>.
- Kominers, Scott Duke, and Alex Tabarrok, "Vaccines and the Covid-19 Pandemic: Lessons from Failure and Success," *Oxford Review of Economic Policy*, 38 (2022), 719–741. <https://doi.org/10.1093/oxrep/grac036>.
- Kutasi, Kristóf, Júlia Koltai, Ágnes Szabó-Morvai, Gergely Röst, Márton Karsai, Péter Bíró, and Balázs Lengyel, "Can Hesitancy be Mitigated by Free Choice across COVID-19 Vaccine Types?," Working Paper, Rice University, 2021.
- Largent, E. A., and F. G. Miller, "Problems with Paying People to Be Vaccinated against COVID-19," *Journal of the American Medical Association*, 325 (2021), 534–535. <https://doi.org/10.1001/jama.2020.27121>.
- Loertscher, Simon, and Ellen V. Muir, "Monopoly Pricing, Optimal Randomization, and Resale," *Journal of Political Economy*, 130 (2022), 566–635. <https://doi.org/10.1086/717894>.

- Makhoul, Alan T., and Brian C. Drolet, "A Reserve System for the Equitable Allocation of a Severe Acute Respiratory Syndrome Coronavirus 2 Vaccine," *Chest*, 159 (2021), 1292–1293. <https://doi.org/10.1016/j.chest.2020.08.2042>.
- Milgrom, Paul, *Putting Auction Theory to Work* (Cambridge: Cambridge University Press, 2004).
- Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, 6 (1981), 58–73.
- National Academies of Sciences, Engineering, and Medicine, *Framework for Equitable Allocation of COVID-19 Vaccine* (Washington, DC: National Academies Press, 2020).
- Newhouse, Joseph P., *Free for All?: Lessons from the RAND Health Insurance Experiment* (Cambridge, MA: Harvard University Press, 1993).
- Ostrizek, Franz, and Elia Sartori, "Screening while Controlling an Externality," *Games and Economic Behavior*, 139 (2023), 26–55. <https://doi.org/10.1016/j.geb.2023.01.011>.
- Oza, Anil, "Cash for Shots? Studies Suggest Payouts Improve Vaccination Rates," *ScienceInsider*, August 21, 2021. <https://www.science.org/content/article/cash-shots-studies-suggest-payouts-improve-vaccination-rates>.
- Pai, Mallesh, and Philipp Strack, "Taxing Externalities without Hurting the Poor," Working Paper, Rice University, 2022.
- Pancs, Romans, "A Vaccine Auction," Working Paper, ITAM, 2020.
- Pathak, Parag A., Govind Persad, Tayfun Sönmez, and M. Utku Ünver, "Reserve System Design for Allocation of Scarce Medical Resources in a Pandemic: Some Perspectives from the Field," *Oxford Review of Economic Policy*, 38 (2022a), 924–940. <https://doi.org/10.1093/oxrep/grac034>.
- Pathak, Parag A., Tayfun Sönmez, and M. Utku Ünver, "Improving Ventilator Rationing through Collaboration with Experts on Resource Allocation," *Journal of the American Medical Association Open*, 3 (2020), e2012838. <https://doi.org/10.1001/jamanetworkopen.2020.12838>.
- Pathak, Parag A., Tayfun Sönmez, M. Utku Ünver, and M. Bumin Yenmez, "Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value behind in Health Care Rationing," Working Paper, Massachusetts Institute of Technology, 2022b.
- Persad, Govind, and Ezekiel J. Emanuel, "Ethical Considerations of Offering Benefits to COVID-19 Vaccine Recipients," *Journal of the American Medical Association*, 326 (2021), 221–222. <https://doi.org/10.1001/jama.2021.11045>.
- Persad, Govind, Monica E. Peek, and Ezekiel J. Emanuel, "Fairly Prioritizing Groups for Access to COVID-19 Vaccines," *Journal of the American Medical Association*, 324 (2020), 1601–1602. <https://doi.org/10.1001/jama.2020.18513>.
- Rahmandad, Hazhir, "Behavioral Responses to Risk Promote Vaccinating High-Contact Individuals First," *System Dynamics Review*, 38 (2022), 246–263. <https://doi.org/10.1002/sdr.1714>
- Rajagopalan, Shruti, "The Best Way to Vaccinate Most Indians in the Least Time," *Mint* (September 14, 2020). <https://www.livemint.com/opinion/columns/the-best-way-to-vaccinate-most-indians-in-the-least-time-11600100665869.html>.
- Ray, Debraj [®] Arthur Robson, "Certified Random: A New Order for Coauthorship," *American Economic Review*, 108 (2018), 489–520. <https://doi.org/10.1257/aer.20161492>.
- Rose, Julie L., "Rationing with Time: Time-Cost Ordeals' Burdens and Distributive Effects," *Economics & Philosophy*, 37 (2021), 50–63. <https://doi.org/10.1017/S0266267120000103>.
- Roth, Alvin E., "Repugnance as a Constraint on Markets," *Journal of Economic Perspectives*, 21 (2007), 37–58. <https://doi.org/10.1257/jep.21.3.37>.
- Saez, Emmanuel, and Stefanie Stantcheva, "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106 (2016), 24–45. <https://doi.org/10.1257/aer.20141362>.
- Sandel, Michael J., *What Money Can't Buy: The Moral Limits of Markets* (New York: Macmillan, 2012).

- Satz, Debra, *Why Some Things Should Not Be for Sale: The Moral Limits of Markets* (Oxford: Oxford University Press, 2010).
- Schmidt, Harald, "Vaccine Rationing and the Urgency of Social Justice in the Covid-19 Response," *Hastings Center Report* 50 (2020), 46–49. <https://doi.org/10.1002/hast.1113>.
- Schmidt, Harald, Parag Pathak, Tayfun Sönmez, and M. Utku Ünver, "Covid-19: How to Prioritize Worse-Off Populations in Allocating Safe and Effective Vaccines," *British Medical Journal* 371 (2020a), m3795. <https://doi.org/10.1136/bmj.m3795>.
- Schmidt, Harald, Parag A. Pathak, Michelle A. Williams, Tayfun Sonmez, M. Utku Ünver, and Lawrence O. Gostin, "Rationing Safe and Effective COVID-19 Vaccines: Allocating to States Proportionate to Population May Undermine Commitments to Mitigating Health Disparities," Working Paper, Massachusetts Institute of Technology, 2021.
- Schmidt, Harald, Utku Ünver, Michelle Williams, Parag Pathak, Tayfun Sönmez, and Lawrence Gostin, "What Prioritizing Worse-Off Minority Groups for COVID-19 Vaccines Means Quantitatively: Practical, Legal and Ethical Implications," Working Paper, 2020b.
- Sönmez, Tayfun, Parag A. Pathak, M. Utku Ünver, Govind Persad, Robert D. Truog, and Douglas B. White, "Categorized Priority Systems: A New Tool for Fairly Allocating Scarce Medical Resources in the Face of Profound Social Inequities," *Chest*, 159 (2021), 1294–1299. <https://doi.org/10.1016/j.chest.2020.12.019>.
- Terrell, Kenneth, "These Companies Are Paying Employees to Get Vaccinated," *AARP.org*, September 10, 2021. <https://www.aarp.org/work/companies-paying-employees-covid-vaccine/>.
- Varian, Hal R., "Position Auctions," *International Journal of Industrial Organization*, 25 (2007), 1163–1178. <https://doi.org/10.1016/j.ijindorg.2006.10.002>.
- Vellodi, Nikhil, and Joshua Weiss, "Optimal Vaccine Policies: Spillovers and Incentives," FRB Richmond Working Paper no. 21-06, 2021a.
- , "Targeting Interacting Agents," Working Paper, Paris School of Economics, 2021b.
- Walzer, Michael, *Spheres of Justice: A Defense of Pluralism and Equality* (New York: Basic Books, 1983).
- Weitzman, Martin L., "Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need it Most?," *Bell Journal of Economics*, 8 (1977), 517–524. <https://doi.org/10.2307/3003300>.
- Wrigley-Field, Elizabeth, Mathew V. Kiang, Alicia R. Riley, Magali Barbieri, Yea-Hung Chen, Kate A. Duchowny, Ellicott C. Matthey, David Van Riper, Kirrthana Jegathesan, Kirsten Bibbins-Domingo, and Jonathon P. Leider, "Geographically Targeted COVID-19 Vaccination is More Equitable and Averts More Deaths than Age-Based Thresholds Alone," *Science Advances*, 7 (2021), eabj2099. <https://doi.org/10.1126/sciadv.abj2099>.
- Zeckhauser, Richard, "Strategic Sorting: The Role of Ordeals in Health Care," *Economics & Philosophy*, 37 (2021), 64–81. <https://doi.org/10.1017/S0266267120000139>.