SUPPLEMENTARY INFORMATION: GAME-THEORY APPENDIX

Incentivizing Wealthy Nations to Participate in the COVID-19 Vaccine Global Access Facility (COVAX): A Game Theory Perspective

This appendix provides game-theoretic modeling and analysis, as well as additional related discussion, supporting the qualitative findings presented in the main text.

1. The Contribution Game

We begin by analyzing an abstract formulation of the game being played by high- and upper-middle-income countries ("HICs"), modeling HICs as each choosing whether to take an action, called "contribution," to expand and/or accelerate vaccine access for low- and lower-middle-income countries ("LICs"). Contributing adds to the global public good by hastening the end of the pandemic, but HICs may not have sufficient individual incentive to do so. This can lead to a "tragedy of the commons," in which HICs choose not to contribute but all would have been better if they had all contributed. The analysis of this section highlights two ways in which such tragedies can be avoided: first, by enabling groups of players to coordinate their actions; and second, by creating new inducements for HICs to contribute to COVAX.

1.1 Setup

A number N of HICs, indexed as $i = 1,2, \dots, N$, each decide whether to take an action ("contribute") that will expand and/or accelerate vaccine access in LICs. "Contribution" by country i is a zero-one decision that costs $C_i \geq 0$ for country i but benefits all HICs by speeding the end of the pandemic; in particular, each country $j = 1,2, \dots, N$ gets benefit $B_{ij} \geq 0$ if country i contributes.

Contributing makes country i *individually* better off if $B_{ii} > C_i$ and makes all HICs *collectively* better off if $\sum_{j=1,\dots,N} B_{ij} > C_i$. This leaves three basic possibilities.

- *Reliable contributor*: if $B_{ii} > C_i$, then country i has an individual incentive to contribute and hence can be relied upon to do so, regardless of what others do.
- *Unreliable contributor*: if $\sum_{j=1,\dots,N} B_{ij} > C_i > B_{ii}$, then HICs benefit collectively when country i contributes, but country i does not have an individual incentive to do so.
- *Undesirable contributor*: if $C_i > \sum_{j=1,\dots,N} B_{ij}$, then country i's contributions are sufficiently cost-ineffective that HICs are collectively better off when i does not contribute. This does not appear relevant in the case of COVAX contributions.

Some countries are certainly reliable contributors but, to focus on the most challenging case and highlight ideas, we will assume that all HICs fall in the "unreliable contributor" category. In this case, there is potential for a win-win outcome in which all HICs contribute and all benefit, or a lose-lose outcome (specifically, a "tragedy of the commons") in which none contribute and all suffer.

McAdams D, *et al. BMJ Global Health* 2020; 5:e003627. doi: 10.1136/bmjgh-2020-003627

To simplify equations, we also henceforth assume that HICs have symmetric payoffs, i.e., $B_{ij} = B$ for all i, j and $C_i = C$ for all i. (This symmetry assumption is not essential.)

Each HIC seeks to maximize its *net payoff*, denoted Π_i, equal to the sum of all benefits associated with all HIC contributions minus the cost associated with its own contribution (if any). If K HICs contribute, then those that contribute get net payoff $\Pi_\mathrm{i} = KB - C$ while those who do not get $\Pi_\mathrm{i} = KB$. Our assumption that HICs are "unreliable contributors" means that $B < C < NB$.

1.2 The potential for a "tragedy of the commons"

In a Nash equilibrium, each player chooses the strategy that is best for themselves individually, given the strategies of others. If all countries are unreliable contributors, as we have assumed, the unique Nash equilibrium is for all HICs not to contribute—a tragedy of the commons as the global public good associated with bringing the pandemic to a faster end is not realized.

Fortunately, there are many ways to *change the game* that HICs are playing, so as to avoid the tragedy of the commons.[1] Here we focus on the two conceptually simplest approaches, (i) by enabling HICs to act together collectively and/or (ii) by creating new inducements for HICs to contribute.

1.3 Avoiding the tragedy of the commons via collective action

If countries were able to act together as a single player, they would all choose to contribute together and hence each get payoff $\Pi_\mathrm{i} = NB - C > 0$. Achieving such "collective action" can be difficult in practice, and the game theory underlying collective action is too complex to survey here,[2] but the following examples illustrate some of the possibilities.

> *Example: collective action via delegation.* The simplest way to coordinate behavior is to delegate one's decision-making authority to a larger body. That larger body can then make decisions in the *collective* best interest of all members. For instance, 27 EU countries plus Norway and Iceland have joined COVAX together as part of a joint effort termed "Team Europe."
>
> *Example: collective action via community self-enforcement.* The Nobel prize-winning economist Elinor Ostrom documented how real-world players have avoided the tragedy of the commons through community self-enforcement, creating mechanisms to articulate norms of behavior and then incentivize compliance with these norms, such as by shaming those who violate them. COVAX does not have a formal mechanism to sanction HICs that do not contribute but, by publicly naming contributors, COVAX has created an informal mechanism by which contributors (such as China) have gained prestige and non-contributors (such as the United States) have been shamed on the public stage.

[1] "Game-Changer" (McAdams 2014) provides an accessible and exhaustive discussion of all the ways to escape the tragedy of the commons, and the "prisoners' dilemma" more generally.

[2] "Governing the Commons" (Ostrom 1990) highlights how communities have managed to coordinate on mutually-beneficial actions in very different contexts, from fishing to farming. For related game-theory background, see the textbook treatment in "Games of Strategy: Fifth Edition" (Dixit, Skeath, and McAdams 2020).

Example: collective action via conditional commitment. Players can also influence others by conditioning their own behavior on what others do. For instance, consider the game that HICs played early on, when making non-binding expressions of interest to COVAX. Some HICs may have done so *hoping* that their peer nations would also do so, but planning to withdraw their own involvement if their peer nations did not. In this context, HICs have an incentive to express interest—and then followup in stages of increasing mutual commitment—because they know that their peer nations will likely back out unless they also are involved. For instance, suppose that countries i, j are peer nations and that country j understands that country i will only contribute to COVAX if country j also contributes. Country j will then view its contribution as *causing* country i also to contribute. This gives country j an incentive to contribute so long as $2B > C$, rather than just when $B > C$, expanding the range of scenarios in which country j contributes—and likewise for country i, if it believes that country j will back out unless it contributes.

Because collective action can arise in many different ways, and for many sorts of reasons, we do not explicitly model how collective action may arise among different groups of HICs in the COVAX context. Instead, we take as given the *coalitional structure* of HICs, assuming that each HIC either acts on its own as an "individual country" or belongs to a "coalition" with multiple member countries whose contribution decisions are made collectively to maximize the total net payoff of all member countries. For instance, China might act as an individual country, while "Team Europe" countries act as a coalition.

In the unique Nash equilibrium of the resulting game, each individual country contributes if $B > C$ but not if $B < C$, while each coalition with K members contributes if $KB > C$ but not otherwise.

1.4 Avoiding the tragedy of the commons by changing players' payoffs

An even simpler way to avoid the tragedy of the commons is to create new inducements to contribute—effectively lowering the individual cost of contributing. For instance, suppose that each HIC gets a direct benefit of \$100 million when contributing \$300 million to COVAX, due to helping speed the global recovery, i.e., $B = \$100\text{M}$ and $C = \$300\text{M}$. If acting individually, each HIC will choose not to contribute. However, suppose that contributors get access to some extra benefit worth \$250 million. This reduces the *net cost* of contributing from \$300M to \$50M, given which all HICs now have an individual incentive to contribute.

The key question then becomes: what extra inducements can be attached to an HIC's decision to contribute to COVAX? As we note in the main text, there are three natural ways in which contributing to (and collaborating closely with) COVAX can benefit HICs. (Section 2 of this appendix provides a detailed mathematical model that formalizes the following verbal arguments.)

Insurance value: First and most obviously, COVAX provides insurance against the event that all of the vaccine candidates that an HIC has backed turn out to fail. Funding COVAX then allows the HIC to secure at least some doses, so long as there is at least one successful COVAX vaccine.

Quicker repurposing of stranded assets: In the race to develop a Covid vaccine, some investments will be wasted as vaccine candidates are found to be unsafe or ineffective. Collaborating closely with COVAX could allow HICs to reduce such waste. In particular, suppose that an HIC invests to stand up a vaccine-production facility to produce Vaccine A but then Vaccine A is found to be unsafe or ineffective. Because COVAX has close ties with many vaccine manufacturers (including those producing doses on COVAX's behalf), HICs that partner closely with COVAX can move more quickly to re-purpose their stranded investments to produce another similar vaccine that has been shown to be safe and effective.

Wider market for successful domestic vaccines: Vaccine-producing countries can leverage COVAX as a channel to access additional markets for domestic vaccines—both by selling domestically-produced doses through the COVAX Facility and by reaching deals to re-purpose otherwise-stranded COVAX production facilities. For instance, suppose that all of the vaccines in the COVAX portfolio fail but China and the United States each have a successful domestic vaccine. Because China is collaborating with COVAX, China will then be in a stronger position to reach a deal for its vaccine to be produced in stranded production facilities being supported by COVAX.

2. The Capacity-Expansion Game

This section continues our analysis of the Contribution Game, but now in a more specific context. Outside of COVAX, each HIC has two basic options when attempting to secure doses for its own population: lock up existing production capacity ("grab supply"), or invest in additional production capacity ("grow supply").

"Contribution" here corresponds to the decision to grow supply. As in the basic story of the Contribution Game, each HIC naturally benefits when others invest to grow supply, as doing so hastens the end of the pandemic and leads to lower vaccine prices (since there will be more global supply). However, to focus on the most challenging case, suppose that HICs do not care at all about such benefits, i.e., $B_{ij} = 0$ for all i, j. The cost C_i here is country i's *extra* cost associated with growing rather than grabbing supply, with $C_i > 0$ if grabbing supply is cheaper or $C_i < 0$ if growing supply is cheaper.

Without any extra inducements to grow supply, the only countries that do so will be those that, for whatever reason, find it cheaper to grow supply.

Let $0 < \lambda_i < 1$ denote the likelihood that at least one of the "bets" that country i has made will pay off, allowing its population to be vaccinated. Let $0 < \lambda_{COVAX} < 1$ be the likelihood that at least one vaccine candidate in the COVAX portfolio is successful, conditional on country i's bets all failing. With probability $(1 - \lambda_i)\lambda_{COVAX}$, country i's investments will have all been wasted, but COVAX will have a successful vaccine. Country i's options at that point depend on the extent to which it has previously collaborated with COVAX

> *Possibility #1: No collaboration*. An HIC that has not funded or otherwise collaborated with COVAX will be able to vaccinate its population with probability λ_i but otherwise will not be to access a successful COVAX vaccine. Normalize the payoff associated with this case to zero.

Possibility #2: Funding the COVAX Facility. HICs that have funded COVAX get enough doses to vaccinate 10-50% of their populations, enjoying an extra value that we will denote as $V_{10-50\%}$ when their vaccines fail but a COVAX vaccine succeeds.

Possibility #3: Repurposing domestic production. HICs that have chosen to "grow supply" and that have also engaged in close technical collaboration with COVAX may be able to repurpose their own stranded production facilities, if any of those facilities were being prepared to produce a vaccine that is similar to a successful COVAX vaccine. Doing so could allow such a HIC to vaccinate all of its population, creating an extra value that we will denote as $V_{100\%}$ when their vaccines fail but a COVAX vaccine succeeds.

These considerations feed back into the original decision that HICs face whether to grab or grow supply, as follows.

Grabbing supply: A country that chooses to grab supply has two options: (i) disengage completely from COVAX, getting payoff normalized to zero, or (ii) fund COVAX and get expected payoff $(1-\lambda_i)\lambda_{COVAX}V_{10-50\%}$ due to the insurance that COVAX provides to funding countries. Let $F>0$ denote the cost of funding COVAX. A country that grabs supply will therefore choose to fund COVAX if $(1-\lambda_i)\lambda_{COVAX}V_{10-50\%}>F$, but not otherwise. So, a country that grabs supply gets expected payoff $\max\{0,(1-\lambda_i)\lambda_{COVAX}V_{10-50\%}-F\}$.

Growing supply: A country that grows supply faces a similar choice, but now has an additional option to collaborate with COVAX to be able to repurpose the new production capacity under its control. In particular, a country that has grown supply has four options: (i) disengage completely from COVAX, getting payoff $-C_i$, (ii) fund COVAX and get expected payoff $(1-\lambda_i)\lambda_{COVAX}V_{10-50\%}-C_i$, (iii) collaborate technically with COVAX and get expected payoff $(1-\lambda_i)\lambda_{COVAX}V_{100\%}-C_i-T_i$, where T_i is the cost of collaborating technically, or (iv) fund and collaborate technically with COVAX and get expected payoff $(1-\lambda_i)\lambda_{COVAX}V_{100\%}-C_i-F-T_i$.

So long as T_i is small, countries that choose to grow supply have an individual incentive to colloborate technically with COVAX, allowing stranded production facilities to be repurposed more quickly. However, note that countries that collaborate in this way have less need for COVAX doses and hence less incentive to fund the COVAX Facility. One natural way to resolve this tension would be for COVAX to prioritize technical collaboration with those countries that have also provided funding. This would, in turn, create an extra inducement for HICs to fund COVAX.

Overall, we conclude that, by identifying and cultivating ways that HICs can benefit from the COVAX Facility—notably, through supply-chain harmonization and knowledge sharing to make it easier to re-purpose existing production facilities—COVAX can increase HICs' incentives both to fund COVAX and to invest in increasing overall vaccine-production capacity.

McAdams D, *et al. BMJ Global Health* 2020; 5:e003627. doi: 10.1136/bmjgh-2020-003627