More from Less: Optimising Vaccines in a Constrained World

WITOLD WIĘCEK

Abstract

This paper argues that more focus on how vaccines are dosed can deliver large benefits. The current practice of vaccine development and licensure delivers vaccines that are safe and effective, but can lead to doses that are higher than is socially optimal. After a vaccine has been approved, there are often few commercial incentives to make adjustments. As a result, while there are many successful examples of adjusting dosing, this process can take decades.

This is especially important in pandemics and under supply and fiscal constraints, such as those currently facing Gavi. A simple analysis suggests that there are both large health and fiscal benefits that public health decision makers and vaccine buyers can get from optimisation. This is shown in three illustrative but concrete cases: accelerating the switch to single-dose HPV vaccination, adjusting PCV dosing, and dose-sparing for a COVID-19 vaccine.

To get to optimal vaccines faster, all stakeholders in vaccine development need to contribute. Researchers and pharmaceutical companies should improve the science of vaccine dosing, especially through greater focus on model-informed development. Regulators should push for optimisation data as part of approvals and consider clinical trial designs that enable optimisation. Global health funders and vaccine buyers should identify vaccines where more evidence is needed and be proactive about generating it, possibly by creating new funding models and pull incentives for optimisation.

POLICY PAPER 367 • OCTOBER 2025

More from Less: Optimising Vaccines in a Constrained World

Witold Więcek
Center for Global Development (non-resident fellow)

The Center for Global Development is grateful to the Gates Foundation for contributions in support of this work.

I am indebted to nine experts who agreed to be anonymously interviewed on vaccine optimisation.

Thanks to Javier Guzman, who proposed working on this topic at CGD, and Janeen Madan Keller, who proposed and supported writing this policy paper. Also thanks to Tom Drake, Saloni Dattani, Enlli Lewis, and Jacob Trefethen for comments and conversations about the paper. Sarrin Chethik, Siddhartha Haria, and Sebastian Quaade from the Market Shaping Accelerator helped me with conclusions on market shaping.

Witold Więcek. 2025. "More from Less: Optimising Vaccines in a Constrained World." CGD Policy Paper 367. Washington, DC: Center for Global Development. https://www.cgdev.org/publication/more-less-optimising-vaccines-constrained-world

CENTER FOR GLOBAL DEVELOPMENT
2055 L Street, NW Fifth Floor
Washington, DC 20036

1 Abbey Gardens
Great College Street
London
SW1P 3SE

www.cgdev.org

Contents

Figures

Tables

Boxes

Introduction

This policy paper makes the case for shifting global vaccine policy to optimise the doses and schedules of vaccines for maximum public health impact. While much attention has been given to incentives needed to develop vaccines to meet the world's needs, far less effort has gone into improving the vaccines we already use. Real-world experience and counterfactual analyses both suggest that the public can reap large health and economic rewards from optimising vaccines, but this is not done consistently and tends to happen slowly. The most compelling recent success is the human papillomavirus (HPV) vaccine, where switching from a three-dose to a single-dose regimen is expanding coverage and drastically reducing costs, especially in low- and middle-income countries (LMICs).

There are simple reasons why the existing system tends to deliver vaccines that are far from maximally beneficial. Optimising vaccine regimens is hard to do during clinical development, and producers often reap rewards from being first to market. Regulators judge efficacy and safety, not broader public health benefits. Once approved, it is costly to change the label—and if the change involves reducing the number of doses, it runs counter to the interests of the drug makers.

If changes do eventually happen, they can be slow and costly—for the HPV vaccine, it took 16 years and several clinical trials. I argue that changing the guidelines for HPV and pneumococcal conjugate vaccine (PCV) five years earlier could have saved hundreds of thousands of lives and hundreds of millions of dollars. And not every opportunity is realised. For example, adopting a dose-sparing approach to COVID-19 vaccines could have saved hundreds of thousands of lives.

This issue matters uniquely right now. Consider three major recent headlines in global vaccination: (1) cuts to Gavi's funding, (2) reductions in public funding for vaccine research in the US, and (3) the announcement that the Gates Foundation will sunset in 2045. The case for Gavi to optimise the ways in which vaccines are used is stronger than ever, but optimisation itself relies on research funding that is becoming increasingly uncertain. Within the current context, there is an even greater premium on figuring out solutions to this optimisation issue.

It doesn't have to be like this. There is a positive, proactive research agenda on optimisation we could be pursuing. By improving the foundational science of dosing, we could optimise much earlier in the research and development (R&D) process. But to make progress on this issue, we have to understand and align the incentives of different stakeholders. The main contribution of this short paper is to characterise the problem—which I think is understudied—and consider the perspectives of different entities on which optimisation depends, such as producers, regulators, purchasers, R&D funders, and others.

The vaccine that makes it to market is not optimal

To set the stage, let's consider the chain of decisions that determines choice of vaccine regimens—number and timing of doses and amount of vaccine in each—during vaccine development. While this will overgeneralise across dozens of different vaccines, there are certain common themes that will matter here.

BOX 1. Different meanings of optimisation

This paper purposefully avoids casting a wider net, as I believe the specific issue of improving vaccine dosing is particularly neglected. However, there are many other meanings of vaccine optimisation, each of which is an important problem in its own right. They include:

- Choice of vaccination "portfolio," i.e., which diseases to immunise against with limited country budgets
- Choice of which vaccine to use against each given disease to optimise for cost and effectiveness
- Optimising combined vaccine schedules (since vaccines are often delivered together during visits)
- Improving timing of vaccination (e.g., aligning malaria vaccinations with seasonal patterns)
- Improving delivery to reduce cost and wastage (e.g., choice of vial size)

The first step is choosing the initial regimen. This is often done based on precedent or some simple heuristics specific to each vaccine platform. For example, HPV vaccine developers chose a three-dose schedule because the other virus-like particle protein vaccine for hepatitis B administered doses at 0, 1, and 6 months. It was widely assumed that three doses would be needed.[1] For COVID mRNA vaccines, developers simply tested the ranges that were used for candidate mRNA vaccines against other pathogens (e.g., Zika, cytomegalovirus) and opted for a gap of 21–28 days, a "canonical" gap between the first two doses, which also allowed for quick evaluation.[2]

Once a candidate regimen is proposed, how do developers decide what progresses into Phase 2 and Phase 3 trials? Unlike many drugs, a vaccine developer usually doesn't know which markers predict protection and has to rely on "probable surrogates." In simple terms, they can measure that the vaccine is doing something good, but may not yet know how much of that "good thing" is enough to

1 The HPV vaccine will be a recurring example, so it is worth saying more. Why not simply test one or two doses in the initial Phase 1–2 trials? HPV vaccine clinical trials, similar to many others, looked at antibody titres. This was expected to correlate with protection from cancer, but at the time, the researchers didn't know how durable these antibodies would be, and the consensus was that three doses would be needed. With the benefit of hindsight, it would still have been valuable to also include a single dose in the Phase 1 and 2 trials. Although the data would not have been conclusive at the time, it could have shortened the evidence generation process in the 2010s.

2 Why is a 3–4 week gap between doses so common in vaccines? B-cell formation process takes about one month, so dosing again too soon can lead to "competition" between doses rather than optimal stimulation of the immune response; also see Table 2.4 in Chapter 2 of Plotkin's Vaccines.

prevent disease. Are ten times more neutralising antibodies better? A hundred? And how long will they last?[3]

In the end, the high cost and uncertainty will often lead developers to run efficacy trials with the maximum tolerated dose and a "canonical" schedule.[4] Since Phase 2–3 trials account for most vaccine development costs, this risk aversion is understandable; investigating alternative dosing regimens has a guaranteed large downside and highly uncertain upside. It is likely that developers prefer to risk a vaccine dose that is too high (as long as they are convinced that the side effects will not be bad) than lose even a fraction of the effect; testing multiple regimens in large-scale trials, meanwhile, is prohibitively expensive.

At the development stage, producers are not only focused on efficacy and safety but are also considering commercialisation of the product. For example, the vaccine maker may want to slot the new vaccine into existing schedules for other vaccines, even if the response is suboptimal. Some of the decisions at this stage are exacerbated by regulatory inflexibility: what you test in the large-scale trial is what is filed for licensure. Regulators and vaccination advisory bodies would refuse to approve doses and schedules that haven't been directly tested in large efficacy trials.

Nowhere in this process are broader social benefits explicitly considered. The focus for both the regulator and the developer is to maximise individual efficacy. Even if data are indicative of a lower dose or fewer doses providing comparable protection, the brief of the regulator does not extend to public health benefits.

3 What are these probable surrogates? They are read-outs that are biologically plausible but not yet proven, such as parasite-inhibiting antibodies, polyfunctional $CD4^{+}$ T-cell frequencies, or growth-inhibition assays. In practice, developers pre-specify a short list of these surrogate markers and carry forward the lowest dose whose confidence interval overlaps with the top-dose group on each marker, provided safety is acceptable. This statistical "non-inferiority" approach, combined with human challenge studies for some vaccines, guides which regimens progress to larger efficacy trials.

4 Maximum tolerated dose is the highest dose that was *tested*. Typically (but not always), we expect the immune response to follow an S-shaped dose-response curve: past a certain point, the benefits plateau and there is no point in further increasing the dose. But this shape is hard to establish; you'd need at least five different doses, with at least two in the lower part of the S and two in the upper part, which is more than what developers would typically do.

BOX 2. Optimising vaccines versus optimising drugs

The vaccine development story is often in stark contrast with drug development, where (1) developers often have a much better idea of the target for the drug (e.g., if designing a new medication for cardiovascular disease, we can easily measure cholesterol), (2) they can measure the concentration of the drug in the body, and (3) decisions can be guided by a suite of pharmacological tools, which allow a lot of design to happen *in silico*.

These tools are familiar to the FDA and other regulatory agencies, which play the role of a partner in the choice of regimen—important, given that the FDA is increasingly concerned with dose justification for some drugs (e.g., a high-profile FDA initiative for dose optimisation in oncology). More broadly, we can also see a push for wider adoption of modelling tools among other inventions—such as organ-on-chip technology, the move away from animal testing, and wider adoption of AI. Support for model-based drug development has been part of FDA's mandate, covered under the most recent Prescription Drug User Fee Act.

That is not to say that developing effective drugs is easier than developing vaccines, but rather, drug developers today have many more tools to predict dose-response relationships. Comparable tools are not available in vaccine development. However, the idea of developing better models has gained more traction from vaccine developers in recent years (e.g., see this article by Desikan et al. for a summary).

Most vaccines are optimised—eventually

When vaccines first enter the market, their regimens often fall short of maximising either individual or social benefits. But the good news is that nearly all vaccines eventually go through some form of optimisation. Just as there are some disincentives during a vaccine's initial clinical development, there are also many reasons to optimise over the longer term.[5]

The three common triggers for optimisation are: (1) acute shortage, (2) cost and access considerations, and (3) the search for higher efficacy and durability. We will illustrate this with some concrete examples, but first, let us discuss how optimisation needs are addressed more generally.

5 For the remainder of this paper, I focus on cases where doses can be reduced, as this is most relevant in the case of constrained budgets and/or supply. However, sometimes an increase in doses is needed. Two notable examples of such optimisation occurred with live-attenuated vaccines, where breakthrough cases of measles (1980s) and varicella (2000s) led to guideline changes from using a single dose to two doses. These second doses are aimed at maintaining herd immunity and target older children (4–6 years-old) to help with the waning of antibodies.

BOX 3. Recent examples of dose and schedule optimisation

More details for all of these are provided in Appendix A: Table A1.

- While the HPV vaccine was initially approved (2006) in three doses, eventually WHO SAGE recommended a single-dose (2022), which has phenomenal efficacy; coverage has increased massively following the switch.
- During the mpox outbreaks in 2022–2023, regulatory bodies approved reducing mpox vaccine doses by 80 percent and delivering them intradermally, based on preexisting Phase 2 data.
- In March 2025, WHO SAGE reviewed evidence on dosing of the pneumococcal conjugate vaccine. They concluded it may be possible to drop one dose in settings where coverage is consistently high, but this should only be done in conjunction with surveillance (see Appendix A for details).
- During shortages of polio vaccine in the last decade, fractionation was recommended and used in many countries.
- During the 2016–2017 yellow fever outbreaks in the Democratic Republic of the Congo (DRC), Angola, and Brazil, doses of the yellow fever vaccine were decreased by 80 percent based on a WHO recommendation. This allowed for emergency vaccination of 7 million people in Kinshasa in August 2016 and was shown to be effective.
- For COVID-19 vaccines, some national immunisation advisory bodies recommended that the gap between doses was adjusted from 3–4 weeks (tested in randomized controlled trials) to 12–16 weeks. Despite some promising evidence, there was no willingness to decrease the amount of antigen in each dose, but in 2021 Moderna dropped the mRNA dose from 100 mcg to 50 mcg for boosters.

All of these initiatives have been successful in the sense of delivering effective vaccines quickly or expanding coverage.

How does optimisation happen? Let's return to the story from the previous section. At the end of successful vaccine development, a regulator grants a pharmaceutical company a license to market their product, with a label specifying dosage, number of doses, and other details. If public health necessitates a change, what are the options for optimisation of already-approved vaccines? There are two possible routes: (1) label changes approved by regulators, or (2) off-label programmatic recommendations by national immunisation advisory bodies (NITAGs) or the World Health Organization (WHO).[6]

When efficacy or durability must be demonstrated, sponsors or public funders may run large trials. In some specific cases, there is an incentive for the vaccine maker to sponsor such trials, which may

6 There are also emergency use authorisations (EUAs), which do not themselves revise the product label, but are issued by the regulators.

lead to label changes—option (1) mentioned above. The most common case is expanding the vaccine to new age groups, where there is an obvious profit motive for the vaccine maker. However, it can also include optimisation: for example, Merck changed the HPV vaccine label to a two-dose regimen in the mid-2010s, which made sense in the context of a competitive market and low coverage.

However, not all changes need large RCTs. Sometimes a combination of immunogenicity data and immunobridging or modeling[7] can be used, for example, when updating seasonal vaccinations. During the 2022 mpox response, for instance, fractional-dose intradermal use through regulator EUA drew on immunobridging evidence generated by both public-sector studies and sponsor data. Whatever the evidence base, label changes typically will require initiative from the vaccine's owner.

However, the sponsor for optimisation research is more often non-commercial, with major funders being governments or the Gates Foundation. The switch to a single-dose HPV vaccine regimen is a good example here. This will typically lead to route (2) mentioned above, an off-label recommendation by national immunisation advisory bodies or the WHO.[8]

Once an off-label recommendation is given, any decision by national-level decision makers to follow it is not a straightforward question of acting on evidence, especially in LMICs. Gavi and purchasers (individual countries, UNICEF Supply Division, the Pan-American Health Organization) can also have indirect influence on these decisions through their role in shaping vaccine markets and uptake. I will say more on the perspectives of buyers and relevant decision makers, but first, let's complete the picture by discussing the magnitude of potential benefits.

Benefits of vaccine optimisation can be huge

Optimisation can save lives, reduce side effects, improve coverage, ease supply constraints, and provide cost savings. But what is the scale of these potential benefits? Predictably, this question has to be assessed on a case-by-case basis, and calculating counterfactual benefits is difficult and often model-dependent.

Various model-based estimates are available for all of the examples cited so far, but they can often overstate benefits due to certain simplifying assumptions which are not tractable. Therefore, I believe that simple illustrative examples are more convincing in this case.

7 Immunobridging may simply mean looking at the amount of antibodies generated by optimised doses and comparing it to longer-term follow-up data to make predictions about durability of protection, without waiting years for follow-up. It can also mean more complex extrapolation, such as across age groups. In our HPV example, an RCT was conducted on 15–20 year-old girls to show efficacy, and immunogenicity was also measured on 9–14 year olds. These two results combined suggested that 9–14 year-olds would be sufficiently protected by a single dose.

8 Although not my focus in this paper, it's also worth pointing out that, alongside the evidence on safety and efficacy, epidemiological and economic modelling can (and should) also aid these decisions. Unlike route (1), off-label recommendations consider public health impacts.

I consider three counterfactual scenarios with details, including underlying references, of each in Appendix A:

- **HPV:** It took 16 years between licensure of the first HPV vaccine and the WHO SAGE recommendation to use a single dose. If the switch to a single-dose happened five years earlier, 150,000 deaths could have been averted per Gavi's model of vaccination benefits. This likely overstates benefits if later catch-up campaigns were possible, but even then, the benefits would be very large.
- **Pneumococcal disease:** PCV is effective, but it is a major cost driver for both Gavi and countries transitioning out of support; in 2026–2030, it is projected to account for 8.5 percent of Gavi's total budget and cost over $1 billion. Emerging evidence (discussed by WHO SAGE in March 2025) suggests that switching to a prime-boost regimen in settings with good coverage could be beneficial. If such a policy was implemented by Gavi countries with coverage above 80–90 percent in the previous five years (2020–2024), it could have already generated savings of $120–250 million.
- **COVID:** In 2021, many experts agreed that the 100 mcg dosing of Moderna's COVID-19 vaccine was set too high. Moderna subsequently reduced the dose to 50 mcg for boosters, but only beginning in autumn 2021. It was programmatically feasible for countries to deliver 50 mcg doses, which would unlock supply for up to 400 million more people to be vaccinated in 2021—but to be conservative, I divide that by half. Scaling up available estimates of vaccination benefits in 2021, between 250,000–500,000 lives could have been saved that year.

To be clear, my main consideration here is to show how large the potential benefits from further optimisation can be. These examples are simplistic on purpose, but I think they represent a best-guess on the magnitude of potential benefits, or at least they don't extend beyond what seems plausible in typical modelling of the benefits of vaccines.

Additionally, I do not mean to detract from monumental efforts made by non-commercial sponsors, clinical researchers, and vaccine makers in finding more optimal doses. There are many success stories.

Lastly—and most importantly—none of the decisions I tentatively propose here would be straightforward, nor should they be. All of them would require faster generation of additional evidence, closer coordination among the main stakeholders, and careful consideration of trade-offs, as they should—the essence of public health is to make decisions that balance uncertain benefits with uncertain risks. However, I believe all these decisions could have feasibly been arrived at in the timeframes I suggest, without a crystal ball.

Understanding incentives and challenges for different stakeholders

Let us now think more granularly, breaking down the problem into perspectives of different stakeholders. The discussion so far has provided some examples of incentives that may pull them away from optimisation. For example, we saw that the early clinical research is slanted heavily towards finding the maximally effective vaccine, as opposed to understanding the dose-response profile. Meanwhile, the objective of the pharmaceutical companies is to sell as many doses as possible, while keeping R&D costs low. However, even if a vaccine R&D is done by a nonprofit, ultimately the overarching ambition is to get to a vaccine that works—one that meets a target product profile set by regulators or the WHO—which is also the main perspective for these organisations. Moreover, the WHO and global partners like Gavi want to ensure equitable access to safe and efficacious vaccines.

To get a more complete view, let us start by mapping out the range of entities involved in optimisation (Figure 1). Roles, challenges, and successes of different entities involved in vaccine optimisation are then summarised in Table 1. I will highlight issues that each entity is facing, but also note where and how they contribute to optimisation.

FIGURE 1. Entities involved in vaccine optimisation

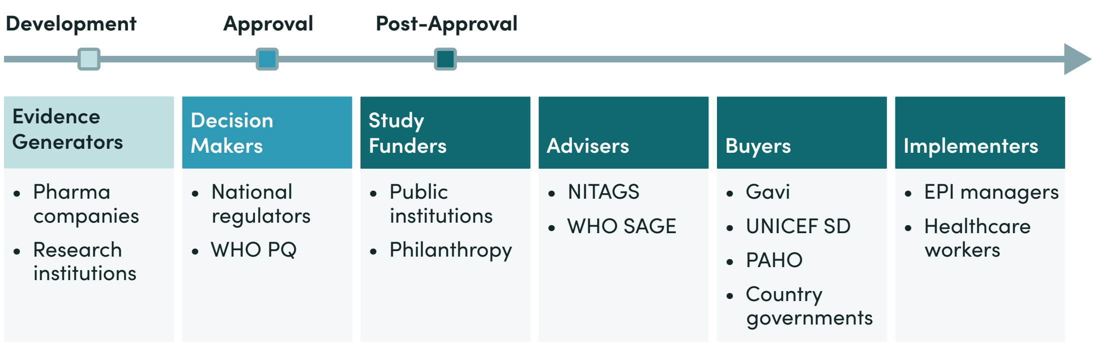

TABLE 1. Entities involved in vaccine optimisation with details on their functions, challenges, and impacts

Stage	Actor	Role	Challenge	Example
Evidence	Early-Stage Research Institutions	Develop vaccine candidates	Often insufficient resources to comprehensively test doses	Development of low-antigen R21 malaria vaccine by Jenner Institute
	Pharmaceutical Companies	Originate some candidates, sponsor clinical trials, manufacture, and market vaccines	High R&D cost for large trials (necessary for approval and label change); pressure to reach market fast	Switch from 100 mcg to 50 mcg booster by Moderna
	Philanthropic and Public Funders	Conduct post-approval research on vaccine optimisation	Limited funding; fragmented system of entities that may not coordinate their research	Gates Foundation creating agenda for PCV 1p+1/fractionation and single dose of HPV vaccine
Decision	Regulators	Give approvals for different dosing, schedules, or mode of delivery	Usually slow and typically requires initiative from manufacturer (request for label change), unless EUA; aims to maximise individual efficacy and safety, not population benefits	FDA/EMA/MHRA approval for 1/5 dose of mpox vaccine in 2022
	WHO/WHO SAGE	Prequalification and recommendations on use of vaccines, including off-label use	Risk-averse; equity considerations may prevail over cost-benefit analysis	1/5 dose of yellow fever vaccine; fewer doses of HPV and pneumococcal vaccines
	National Immunization Technical Advisory Groups (NITAGs)	Recommend off-label use at national level	Area of competency and willingness to make recommendations varies country-to-country; rare in LMICs	Increasing gaps between two doses of COVID-19 vaccine
Rollout	Buyers (Gavi/UNICEF SD, PAHO, individual countries)	Can shape optimisation decisions through procurement and rollout campaigns, especially advance market commitments	No direct role in evidence generation	Gavi market shaping role for PCV
	Implementers (immunisation programme leads, HCWs)	Determine programmatic feasibility of modifying regimens (e.g., intradermal delivery or splitting vials)	Strong status quo bias, especially where there is low capacity	Successful intradermal delivery of fractional-dose Inactivated Polio Vaccine (fIPV) in India

Building on the figure above, let's contrast some situations in which incentives are and aren't sufficient by highlighting them on the entity-by-entity basis, expanding on the examples provided. Let us also provide a more narrative summary that builds on Table 1.

For **small-scale originators of new vaccines**, such as university spin-off labs or small biotech firms, it may not be feasible to have complicated factorial designs at first.[9] For big pharmaceutical companies, which often take over development and run efficacy trials, optimisation would be beneficial, but they are typically racing against the clock and trying to limit expenses. While optimisation may drive down the cost of goods, the main consideration for a pharmaceutical company may be getting to the market first to guarantee volume of sales, rather than maximising profit per unit. As mentioned, there is an uncertain upside if the lower dose works, but an even larger downside if it does not, and the overarching goal for most developers (big and small) may be to develop a vaccine that works, rather than optimising. In some cases of research conducted by public bodies, optimisation can take more of a center stage, as was the case with the R21 malaria vaccine with Matrix-M adjuvant.[10]

On **the producer** side, there can be strong market incentives to optimise. If supply is constrained, producers may choose to reduce the number of doses or amount of antigen in order to reach more people. Even if supply is not constrained, the pressure can also be brought on by competition. For example, by reducing the HPV vaccine to two-doses and then a single dose, Merck was able to reach more girls and fend off competition from GSK's Cervarix vaccine (and eventually others). Moreover, supply of Merck's vaccine was constrained for several years, so fewer doses allowed them to reach more girls.

However, it is important to remember that most of the evidence relevant to change of regimen in this case was generated by **non-commercial sponsors**, which suggests that market incentives were insufficient to justify R&D investment in optimisation. Non-commercial entities can bridge the gap between private and social benefits, but there is a risk that this process will be inefficient if coordination is lacking. It also requires a long-term commitment and competes for funding with other common goods in global health. The decision to move to a single dose HPV vaccine relied on at least four separate RCTs, mainly sponsored and coordinated by the Gates Foundation.

Regulators also play a crucial role in situations where market incentives are insufficient for producers to optimise. Their basic mandate is limited to assessing efficacy and safety of the vaccine regimen that producers seek to license. While the FDA's approval of the intradermal 1/5th dose of Jynneos proves that regulators can play an active role during emergencies, they also indirectly create

9 It's difficult to find primary literature quantifying this issue, but Rhodes et al. offer a valuable perspective in their 2020 paper about dose finding, which highlights some shortcomings of current approaches and calls for development of better modelling techniques.

10 The R21 malaria vaccine was initially developed at the Jenner Institute with funding from European and Developing Countries Clinical Trials Partnership. From the beginning, it has been tested with a broad range of doses and researchers found that increasing the amount of antigen did not improve response. Consequently, the vaccine that was approved has five times less antigen than the RTS,S malaria vaccine and was initially a few times less expensive.

incentives for optimisation in routine vaccine development. First, by collaborating with producers, regulators can bring focus to dosing decisions. Second, regulators can request post-licensure studies that seek to optimise doses.[11]

Unlike regulators, who focus on establishing individual benefits and safety, **national immunisation technical advisory boards (NITAGs)** can make recommendations for off-label use that take public health benefits into account—and they can make them under considerable uncertainty. This is crucial in public health emergencies, and especially during pandemics, when the social benefits of vaccination (and consequently, from optimisation or any other means of accelerating vaccine availability) can be massive, as we attempted to model with colleagues in the case of COVID-19 vaccines.[12]

However, the example of COVID also shows that countries that could have benefitted most from optimisation did not adopt it. While the gap between two doses was increased in the UK, Canada, India, Thailand, and some EU countries, this was not the case in LMICs where supply was much scarcer.

Where NITAGs do not have capacity to make off-label recommendations, they will defer to **WHO Strategic Advisory Group of Experts on Immunization (SAGE)** position statements. As seen in the examples in Box 3, changes to vaccination practices were usually precipitated by WHO SAGE recommendations. However, the bar for a SAGE position statement is set very high, and WHO has to weigh a complex set of priorities, such as. vaccine equity, stakeholder acceptability, and unassailable scientific standards for such recommendations. In my opinion, this may preclude some viable optimisation recommendations from materialising.[13]

As mentioned earlier, we should also distinguish between recommendations and **public health decision problems at national level**. To name a few issues facing decision makers:

- Accountability, authority, and striving for certainty: Ministers or public health bodies may not want to assume sole responsibility, especially if decision has to be made under uncertainty
- Unclear priorities: For example, how to weigh cost-effectiveness against equity considerations?

11 In the future, regulators could also require pharmacological modeling and dose justification before approval—more on this discussed in the recommendation section. Of course, any requirement to justify dose or conduct more dosing studies post-approval has to be balanced by the realistic impact on R&D budgets and timelines. If the regulators' requirements are too stringent, they risk having a net adverse effect.

12 For example, in Castillo et al. (2021), we estimated the social benefits of COVID-19 vaccines in early 2021 to be $5,800 per course of vaccination, and about $500–$1,000 for marginal courses, based on GDP benefits alone. In other words, any intervention which could expand vaccine supply would yield a ten- to hundred-fold return on investment. In Więcek et al. (2022), we extended this argument directly to fractional dosing of COVID-19 vaccines.

13 To be more exact, in making its recommendations, SAGE assesses quality of evidence, scale of potential benefits and cost-effectiveness, acceptability to stakeholders, implementers' attitudes, equity, and more. A short example of the SAGE framework applied to HPV vaccine is available online and worth reviewing.

- Contractual issues: Pre-existing agreements may specify particular dosage
- Indemnity: If doses change, who assumes responsibility for adverse events?
- Equity and messaging: Even if lower doses have public benefits, using them in LMICs can create a perceived double standard (i.e., low doses for poorer countries)

Note that when it comes to national-level decision making, three of the successes in Box 3—COVID-19 dosing gaps, PCV 1p+1, lower mpox dose—are attributable to decisions made in well-resourced settings: the first two in the UK, and the US, UK, and EU in the case of mpox. This is unsurprising, but it's worth keeping in mind that changes to vaccination practices are usually not initiated in the settings that would benefit the most.

Also in this category, I reserved a separate box for **implementors: Expanded Programme on Immunisation (EPI) managers and healthcare workers**. The most salient example in their case is when intradermal delivery is required, which in turn may require additional training, procurement of different syringes, and use of different vial sizes. While it's difficult to assess individual claims about how challenging such programs are to implement, it is probably fair to assume that implementers have a bias towards maintaining the existing practice.[14]

In addition to regulators and public health bodies, **vaccine buyers** can more directly shape the incentives of producers. Gavi's 6.0 strategy for the 2026–2030 period promises to "optimise" and "prioritise" vaccines in face of new financial constraints . This suggests that optimisation as a policy lever is well suited to the current environment. I will attempt to outline the ways in which buyers could create better incentives for optimisation in an upcoming paper.

What **other factors** may impact optimisation decisions? In various interviews, experts also raised the following concerns: unclear legal liability in the case of off-label recommendations, lack of cooperation from producers (e.g., providing instructions/leaflets for alternative delivery methods), and unclear areas of competence in lower income settings. These constraints tend to be more binding during public health emergencies.

The future of vaccine optimisation

The discussion of incentives and disincentives in the previous section provides some hints about where potential solutions to optimisation issues may lie. Before moving to conclusions, it is useful to consider which particular vaccines could potentially benefit from these solutions (see Box 4).

14 For example, when mpox intradermal delivery was rolled out in the US, some predicted that the method would be unfamiliar to healthcare workers. However, these problems have not materialised, as far as I am aware. Similarly, the rollout of fractional IPV in India shows that it can be feasible in LMICs.

For an illustrative consideration of programmatic factors in the case of mpox, see questions 10–12 in WHO's recent mpox fractional dosing FAQ document.

BOX 4. Current and future candidates for optimisation

- Tuberculosis*: The M72 candidate is currently in a Phase 3 trial with 20,000 participants, using two doses given a month apart. Both the number of doses and the gap between them should be investigated further.[15]
- RTS,S and R21 malaria vaccines: Supply is limited and the rollout has been slow. Open questions include the number of doses, fractionation of the last dose, gaps between doses, and aligning the vaccination with seasonal risk.
- Dengue vaccine*: Currently given as two doses, but a single dose may offer sufficient protection.
- Group B Streptococcus vaccine*: Shows promising response at lower doses.
- Mpox* and cholera vaccines: Face acute and persistent shortages. Data already supports reduced dosing in emergencies, but there are questions about implementation.
- PCV: Recent trials show the viability of 1p+1 and fractionation, but implementation questions in LMICs still need to be resolved

**Listed in Gavi's most recent Vaccine Investment Strategy longlist of eight vaccines*

A conventional conclusion to a policy paper is to offer concrete policy recommendations. At the end of this section, I will propose three ideas for advancing the vaccine optimisation agenda, but first a dose of realism is warranted. I believe the problem described in this paper is somewhat understudied and not easily tractable. To some extent, the central problem I describe here is part of the broader issue in vaccine research and development: high costs of testing, low rates of success, and sometimes weak market incentives. Thus, some of the best solutions to the problem of optimisation are the ones that target the overarching problem of vaccine R&D more broadly.

In particular, any solution that can bring down the cost of evidence generation would be a boon for the vaccine optimisation agenda. As we've seen, even when there is no incentive for vaccine makers to generate evidence, non-commercial sponsors are willing to step in. The main barrier to more relevant research on optimisation is how onerous—both costly and slow—evidence generation can be. Broadly, solutions can take the form of innovations to reduce costs of large clinical trials, but they can also be about improving clinical surrogates and model-based techniques.

Even more broadly, any innovation that drives down the expected cost of vaccine R&D—such as new tools that improve vaccine candidate design and selection, yielding better vaccines and higher trial success rates—would also indirectly benefit the optimisation agenda.

Having made this general point, let's now ask how improvements can be made at the margin under the status quo. In short, these include investing more in dosing trials earlier on in development,

15 This may be especially valuable considering that, for earlier TB vaccine candidates, higher doses could lead to worse responses.

developing better modelling tools, encouraging regulators to push for stronger dose justification, and market shaping.

Recommendation 1. Improve the science of vaccine dosing

Developing better in silico modelling techniques to guide dosing decisions in early clinical development would benefit all future vaccine development. Currently, these models are understudied and lag far behind what is done in drug development (see Box 2). Addressing this challenge requires joint efforts by academic researchers, pharmaceutical companies, and regulators. Developing better models—a major, long-term research effort—is only worth it if pharmaceutical companies use them and regulators accept them.[16] Nevertheless, there are promising parallels with improved dosing decisions for other therapeutics.

Recommendation 2. Regulators should push for optimisation data as part of approvals

Regulators and public health bodies, such as the US Food and Drug Administration (FDA), European Medicines Agency (EMA), and WHO prequalification program, could encourage greater focus on optimisation in early clinical research by requiring dose and schedule justification as part of licensure. There is also an additional opportunity through regional regulatory capabilities, for example, the African Medicines Agency could take a leading role in optimisation.

These requirements do not have to be onerous and could take the form of dose de-escalation or additional dose ranging in early clinical trials. As demonstrated by the EUA approving a 1/5th dose of the mpox vaccine in 2022, pre-existing immunogenicity data can make all the difference in emergencies.

Recommendation 3. Identify vaccines where more evidence is needed and invest in generating it now, not later

Building on the previous recommendation, pre-existing data is crucial because it allows for timely decision making. As discussed, the current model often delays optimisation research until licensure is complete. Instead, we should systematically review the cost-benefit of generating optimisation evidence for all vaccines currently in the pipeline and sponsor the most promising ones.

For example, given the huge potential benefits of developing a TB vaccine, investing in dosing trials for M72 (see Box 4) now may turn out to be highly cost-effective. Putting optimisation on the agenda today could unlock significant benefits earlier.

16 Pharmaceutical companies can also contribute by sharing immunological data that are currently siloed. However, most of the relevant data for immunomodeling still has to be generated in future research projects. On a related note, there have been positive signs of adopting modeling techniques by some vaccine developers, like GSK and Moderna.

Recommendation 4. Understand incentives and develop funding models

Optimisation will happen faster if we create stronger incentives to do it. Currently, it relies on non-commercial entities, especially the Gates Foundation, to sponsor research into alternative regimens. But what if alternative mechanisms could be used to speed this process up? Perhaps producers could be incentivised directly through some form of pull or push funding? Could buyers signal preferred vaccine characteristics early in development? These are open questions that will need to be answered in the future. Developing a "market shaping toolset" for optimisation of future vaccines would be especially valuable in future pandemics.

We should think concretely about PCV. As per WHO SAGE recommendations, the rollout of an alternative two-dose (1p+1) regimen depends on surveillance capabilities of countries, rendering it impossible for less well-resourced countries. Given potential cost savings from reducing the number of doses, it may be rational to explore funding new surveillance networks.

Additionally, in taking an incentives approach to optimisation, it is especially relevant to consider the viewpoint of Gavi and vaccine buyers since they have a direct ability to shape the market, such as through advance commitments. I will explore this topic in a follow-up note.

Appendix A. Three case studies of potential benefits

HPV: Switching to single-dose vaccination earlier averts 150,000 deaths

HPV (especially types 16 and 18) is a major cause of cervical cancer, which was estimated to have killed 350,000 women in 2022. Fortunately, HPV vaccines are incredibly effective at reducing cancer risk.

The first HPV vaccine approval was given in 2006, with WHO prequalification following in 2009 for Gardasil (Merck) and Cervarix (GSK). Gavi began providing routine support in 2013, and by the end of 2022, 16 million girls were fully immunised. Initially, three doses were given, but beginning in 2014 (following a WHO recommendation), mainly two doses were used.

Supply constraints, cost, and the need for multiple doses continued to pose challenges, resulting in a slow rollout of HPV vaccines in Gavi countries during the 2013–2022 period. However, after WHO SAGE issued a single-dose recommendation in 2022, Gavi immunised an additional 14 million girls in 2023 alone, almost doubling vaccinations in a single year. Gavi's own estimates of future needs and benefits project approximately one death averted per about 60 girls vaccinated in the target population of 86 million by 2025 (these numbers are available in Gavi's 2023 progress report and Vaccine Investment Strategy).

Let's assume the recommendation to switch to a single dose had been made earlier than 2022. If 3–4 year efficacy trials started shortly after Phase 3 trial results were completed, they would likely have reported findings before 2015. Let's be conservative and consider a hypothetical recommendation to switch to a single dose was made in 2017, ten years after the completion of Gardasil's Phase 3 trial and FDA approval, more than twice the time it takes to run a large trial. This would shift the rollout of single doses halfway through the 2013–2022 period mentioned earlier, allowing for at least 8 million additional doses.[17]

What are the health benefits of expanding coverage? According to Gavi, the HPV vaccination program up to the end of 2023 averted an additional 600,000 future cervical cancer deaths from about 30 million vaccinations. This estimate is relative to the counterfactual of no vaccination, so following Gavi's model, we would claim 150,000 deaths averted. This may sound high, since many cohorts could still have been reached later, but consider two facts. First, as mentioned before, there was a substantial unmet need, so it's likely that extra supply would have gone to cohorts which would

17 This estimate is, again conservative, as HPV vaccine rollout likely accelerated over the years. It's also worth noting that until about 2014, three doses were still being used in some settings. On the other hand, I assume here that increased supply would translate directly to increased coverage. This is a convenient assumption, since there has been a lot of opposition to the HPV vaccine. But the counterfactual we are discussing here still allows for roughly five years of rollout in Gavi countries before the switch to a single dose, presumably enough time to make large-scale implementation feasible at least in some countries.

have otherwise remained unvaccinated for a long time. Second, a 2014 study showed that once women became sexually active, the risk of cervical HPV infection was about 25 percent per year, even in countries with considerably lower risk. Therefore, it's not unrealistic to assume something close to the counterfactual of no vaccination.

PCV: Optimising doses saves Gavi $120–250 million over five years (2020–2024)

Current PCV dosing practice varies:

- **3p+1:** Three priming doses in infancy and a booster at 9–15 months, common in the US
- **2p+1:** Two priming doses in infancy plus a booster at 9–15 months, used in UMICs and HICs; the switch from 3p+1 to 2p+1 was precipitated by decisions in Canada, the UK, and later EMA between 2004–2008
- **3p+0:** Three priming doses in infancy without a booster, typically in LICs (synced with DTP vaccination at 6, 10, 14 weeks). This approach is most common in Sub-Saharan Africa, where it's hard to get toddlers back in for a booster shot, but also because risk is high in the first year of life. However, a booster dose is seen as essential for achieving herd immunity in the long term.
- **1p+1:** One priming dose in infancy plus a booster at 9–15 months, used in the UK since 2020. This approach is feasible once higher levels of protection are established. Immunogenicity studies and modelling indicate no tangible increase in disease burden, a 33 percent cost saving, and potentially stronger immune responses than 3p+0.

Coverage has been increasing quite rapidly, reaching up to 65 percent worldwide according to the latest WHO statistics, largely thanks to Gavi's support for routine immunisation campaigns. PCV has been a major source of spending for Gavi (see Figure 15 here). For example, between 2009–2020, Gavi contributed $3.3 billion to PCV costs, with more recent years seeing an annual volume of 150–200 million doses.[18]

At a recent WHO SAGE meeting in March 2025, a proposal was made to switch from 3-dose to 2-dose regimens for infants, or use 40 percent dose for PCV13 only. The current WHO proposal is to implement this only (1) in high coverage settings and (2) alongside specific pneumococcal surveillance that can flag any problems.

By 2020, many Gavi countries were already approaching good coverage. As many as 18 Gavi countries—including Bangladesh, Kenya, and Sudan—had over 90 percent coverage, with 11.5 million

18 In 2019 alone, procurement reached 161 million doses within Gavi's Advance Market Commitment (AMC) (similar volumes were seen between 2015–2021), with current projections rising to over 200 million doses. The cost per dose has steadily declined: $3.50 initially, $3.30 in 2017, and $2.90 in 2019, reflecting Gavi's success in lowering prices. Costs fell further to about $2.00 per dose following SII's PNEUMOSIL joining Gavi's AMC in 2020. Countries transitioning out of Gavi support have also been able to access PCV at AMC prices.

children vaccinated (WHO data). Another 13 Gavi countries—including Pakistan and Tanzania—had 80–90 percent coverage, with 13 million children vaccinated.

Let's consider **a hypothetical recommendation to switch to 1p+1 over the 2020–2024 period in countries with high coverage.** This could translate to both large cost savings and higher efficiency through dropping one dose. If countries with >80 percent coverage switched to 1p+1 starting in 2020, this would mean savings of about $120–250 million over the last five years ($2 per dose x 5 years x approximately 12–25 million doses per year) depending on whether coverage of 80–90 percent is needed to switch to 1p+1. These funds could then be redirected to another effective vaccination program, such as malaria.

I am not suggesting that this switch would have been wise given the state of evidence in 2020, nor that it would have been programmatically feasible in these particular countries at that particular time. I also only look at the last available year of data and do not account for COVID-19 impacts and surveillance considerations. The purpose of this calculation is to point out that there are large benefits, *conditional on* sufficient evidence and guidelines being available to these countries.

COVID-19: Halving dosage of the Moderna vaccine in the first year of vaccinations could have saved hundreds of thousands of lives

In a paper written during the COVID pandemic, my colleagues and I argued that the evidence available at the time for three of the first highly effective COVID vaccines—developed by AstraZeneca, Moderna, and Pfizer—was sufficient to justify lower doses. Here, I focus on the case of Moderna's mRNA-1273. In 2020, Moderna tested and obtained emergency approval for a 100 mcg formulation of their mRNA vaccine. While Pfizer and Astrazeneca vaccines also had high efficacy, and it's likely that lower doses of these vaccines may have worked well, Moderna dosed their mRNA vaccine four times higher than Pfizer's. There was agreement among many scientists that 50 mcg could have achieved similar efficacy. Moderna eventually dropped the dose for boosters to 50 mcg, starting in autumn 2021.

In Więcek et al. (2022), we suggested that fractionation in 2021 could have expanded supply by 500–1,500 million doses, with large potential reductions in infections and mortality. But such estimates are highly model- and assumption-dependent (e.g., we used sterilising immunity, which proved incorrect). Here I opt for a much simpler argument, focusing on Moderna only.

Moderna delivered about 800 million doses in 2021 before updating their formulation, 25 percent of which went to LMICs. By that point, as many as 5 billion people had been vaccinated, so Moderna's share was about 7–8 percent of the market. If doses were halved, we could hypothetically vaccinate an additional 400 million people in 2021, but to be conservative, let's assume only 200 million extra vaccinations.

Watson et al. estimate 14–20 million deaths averted by vaccination in 2021. Scaling this estimate proportionally to the number of doses, we could be looking at an additional 500,000–1,000,000 lives saved. Some subsequent publications have called that into question, suggesting that the impact of vaccination may be half of these initial estimates. Conservatively dividing by two again, **250,000–500,000 additional lives** could have been saved. This estimate is arguably "assumption-heavy," but still quite conservative; most crucially, it applies to the Moderna vaccine only, which was a small portion of all effective vaccines administered throughout 2021.

Accelerating widespread vaccination also means shorter lockdowns. In this regard, the value of extra pandemic vaccine supply is marginal. In a 2021 Science paper, my colleagues and I estimated the social benefits of additional COVID doses. Even under the most pessimistic assumption (expanding capacity from 4 billion to 4.2 billion courses), the economic benefits were estimated at **$180 billion**. This sounds shockingly large, but consider that estimates of excess mortality suggest 10–15 million deaths in 2021 alone. On the economic side, the IMF projected pandemic-related GDP losses of $500 billion per month in 2021, not accounting for other harms.

TABLE A1. Recent examples of vaccine optimisation

Vaccine	Proposed Change	Rationale	Key Evidence that Triggered the Change	Who Ran the Pivotal Studies	Formal Recommendation (Body, Year)
HPV	Cut primary series from 3 doses to **1 dose** for 9–14 year-old girls	Fewer injections lowers cost and improves coverage	Non-inferiority immunogenicity studies and several efficacy RCTs conducted between 2007 and 2022	Sponsored by Gates Foundation (mainly) and US public funding	WHO SAGE recommended single dose in 2022 (following recommendation to use two-dose schedules and license variation circa 2015)
mpox (JYNNEOS)	Give **0.1 mL intradermal (ID)** instead of 0.5 mL sub-cutaneous (SC) (1/5 dose)	Stretch finite stock during 2022–23 health emergencies	Phase 2 RCT (2015) comparing ID vs SC; post-2022 data showed comparable neutralising titres	Bavarian Nordic with US NIAID; follow-on NIH studies	US FDA emergency-use authorisation & CDC interim guidance, August 2022 (a few months after outbreaks); similar recommendations in the EU and the UK
COVID-19	Reduce **booster** from 100 µg to **50 µg** (half dose) for mRNA-1273 (Moderna)	Maintain immunogenicity, reduce reactogenicity	Moderna Phase 2 and 3 immunogenicity data	Moderna and NIH	US FDA and EUA approved 50 µg booster, October 2021 No recommendation or approval for lower doses of primary series.
PCV (Pneumococcal conjugate)	Move from 3p+1 (infant series + booster) to **fewer doses (1p+1) and fractional dosing**	Maintain herd immunity with fewer doses; programme cost & logistics, especially for LMICs	Cluster RCT in Gambia; trials in Fiji; post-introduction surveillance in South Africa showing carriage & Invasive Pneumococcal Disease (IPD) control with reduced schedules	MRC Unit Gambia at LSHTM; Fiji MOH & Murdoch Children's; NICD South Africa (mainly Gates Foundation funding)	WHO SAGE recommendations for 2p+1 or 3p+0 issued in 2009; new recommendations for 1p+1 and fractionation discussed in March 2025 (25 years after initial licensure)
Yellow fever	Use **0.1 mL (1/5 dose)** during outbreaks	Acute 2016 shortage in Angola and DRC; outbreaks in Brazil; earlier trials showed long-lasting seroconversion with 1/5 dose	Randomised immunogenicity trials in Brazil between 2009–2013; 2016 campaign monitoring data from Angola	Public health bodies with WHO support	WHO emergency recommendation and WHO SAGE endorsement in 2016
Polio (fIPV)	Give **two 0.1 mL ID fractional IPV doses** instead of the standard 0.5 mL IM dose	Shortages during switch from oral polio vaccine	RCTs in India between 2012–2014; Cuban infant studies; programmatic rollout data	ICMR and WHO-supported investigators; Cuban MoH; CDC	WHO SAGE recommendation and GPEI guidance in 2016–2017