Modeling Global Policy for Managing Polioviruses: An Analytical Journey

Kimberly M. Thompson, Radboud J. Duintjer Tebbens
Harvard School of Public Health
677 Huntington Ave., 3rd Floor
Boston, MA 02115
(617)432-4285, fax: (617)432-3699
kimt@hsph.harvard.edu, rduintje@hsph.harvard.edu

The success of the Polio Eradication Initiative promises to bring the world the benefits of sustained improvements in quality of life (i.e., cases of paralysis and deaths avoided) and saved costs from cessation of vaccination. Obtaining these benefits requires that policy makers manage both the transition from the current massive use of oral polio vaccine (OPV) to a world without OPV and the risks of potential future reintroductions of polioviruses. In 2001, we began a case study on retrospective polio risk management to demonstrate the importance of using a dynamic disease model to correctly estimate the cost-effectiveness of vaccines. Discussions with the CDC about the case study led to an opportunity for us to develop a large model to support the prospective decision making process. This paper tells the story of our journey, emphasizing insights about the requirements for analysts to create tools that really help high-level decision makers.

Key words: polio eradication, decision analysis, risk analysis, dynamic disease model, outbreak, process, variability, uncertainty

Introduction

"All decisions are based on models… and all models are wrong" *John Sterman* (Sterman, 2002)
"All models are wrong but some are useful" *George Box* (Box, 1979)

In theory, modeling decisions provides the opportunity to analyze options systematically, which allows explicit consideration of the possible alternatives given the decision makers' preferences and constraints. For complex problems, models can help provide a shared vision of the system, show the different components and how they interact, and synthesize the existing information. Models can also offer insights about important sources of variability (i.e., real differences between individuals that matter in the context of the decision) and uncertainty (i.e., imperfect information) and their implications. Thus, models can theoretically help decision makers choose more wisely, but how much do models really help in practice?

Behind the story of any model lies a story about the analysts and process that created it, yet these stories rarely find their way into the literature. Unfortunately, this means that valuable lessons learned do not help other analysts who would benefit from strategies that might help reduce the amount of "muddling through." This paper introduces the complexity of the decisions brought by eradication of polio and tells the story of the first four years of our experience developing a decision analytic modeling tool to support global policies for managing the risks of polio after eradication. By writing such a paper while the modeling process continues to evolve and expand, we recognize that we run the risk of providing a perspective prior to knowing the final

outcome. We believe that the benefits at this point outweigh the risks, and we hope the reader will agree. We organize the remainder of this paper by providing the context about polio risk management a then offer the story of the project structured according to the 5 requirements that we believe drove its early success.

Context

On April 12, 2005, we celebrated the 50-year anniversary of the publication of the largest and first clinical trial for a vaccine. In 1955, researchers demonstrated the effectiveness of the Salk polio vaccine (Francis et al., 1955) and the news media that day exclaimed exciting themes: "The vaccine works." "It is safe, effective and potent." "Polio is conquered." The promise of a vaccine that would end the terror caused by polio brought hope and led to long lines of people wishing to get the vaccine.

Thompson and Duintjer Tebbens (2005) provide a retrospective analysis of the polio vaccination history for the U.S. and document the incredible story of success, starting from a peak of over 21,000 cases of paralytic polio in 1952 to 0 cases in the U.S. since 1999. During that time, the U.S. began vaccination in 1955 with the Salk Inactivated Polio Vaccine (IPV), switched in the early 1960s to the Sabin Oral Polio Vaccine (OPV) due to its relatively lower cost and easier administration, and finally returned to an enhanced IPV (eIPV) in the late 1990s at a relatively very high cost to avoid the burden of Vaccine-Associated Paralytic Polio (VAPP) (Miller et al., 1996). By the late 1990s, use of OPV led to the only cases of paralytic polio (less than 10 cases per year with an annual birth cohort of approximately 4 million children). The existence of the eIPV option combined with concern about the perception of risks of vaccination exceeding the benefits weighed in favor of the shift to eIPV for routine vaccination, and other developed countries similarly moved toward IPV use. However, the Polio Eradication Initiative (PEI) and most of the countries in the world continue to rely on OPV, because it provides population or herd immunity and it remains much cheaper and easier to administer.

In 1988, with an estimated 350,000 global cases of paralytic polio annually, the World Health Assembly committed to global eradication of polio. In 2003, wild polio viruses only continued to circulate in 6 countries (Nigeria, India, Pakistan, Niger, Afghanistan, and Egypt), but recent challenges (political and logistical) led to outbreaks, exportations, and reintroduction of polio in a number of previously eradicated African countries with outbreaks in 2004-2005. These recent reintroductions and outbreaks demonstrated the rapid ability of polioviruses to spread in susceptible populations, and provide important insights about the need to maintain vigilance in vaccination at least until the successful global eradication of wild polio viruses.

With the success of eradication approaching, many questions emerge and world leaders face a large range of complicated decisions. For example, while the World Health Assembly leaders who committed to eradication in 1988 expected that this would mean complete cessation of polio vaccination (as occurred following the eradication of smallpox), current perceptions by some countries about the risks of potential bioterrorism and the current use of IPV for routine vaccination lead to the logical question of whether developed countries will stop using IPV. For countries now using OPV, the complicated set of options includes switching to IPV (at high cost and presumably at the expense of investing in other public health measures), continuing to use

OPV, or stopping vaccination. Continuing to use OPV emerges as problematic for 2 reasons: (1) OPV use leads to a small, but finite and reasonably predictable number of cases of VAPP, and (2) recent circulating Vaccine-Derived Polio Viruses (cVDPV) events occurred in countries that reduced their coverage with OPV resulting in outbreaks (caused by the OPV viruses reverting to a neurovirulent form and circulating among susceptible individuals). The possibility of cVDPVs means that low levels of OPV coverage present a significant risk, and that continued use of OPV should involve efforts to maintain very high coverage, most likely in the form of the supplemental immunization activities now used in the polio eradication initiative (i.e., National Immunization Days). The risks for outbreaks after eradication range from the potential for cVDPVs, potential breaches in containment, intentional reintroduction, and the little known possibility that one of the very small number of immunocompromised individuals who can excrete polioviruses that they never cleared after they received OPV (immunocompromised VDPVs or iVDPVs) might also at some point represent a potential source. The small risks represent important concerns when we consider the impacts of reintroduction of polioviruses as the population susceptibility increases.

As national and world health leaders consider the options, they recognize that the opportunity to stop OPV vaccination necessitates development of a strategy (i.e., response plan) and the tools (i.e., vaccine stockpile and/or on-going production) for responding to a future outbreak. They also remain interested in understanding the risks, costs, and benefits of their various options. With the success of national eradication of polio and the consequent reduction in the burden of disease, decision makers tend to want to reprioritize resources away from polio risk management and toward other issues. However, they must recognize that this might not represent the optimal strategy in the longer term.

Modeling can provide an opportunity to explore the trade-offs, and we currently provide analytical support to help answer key questions working collaboratively with a large team at the CDC and WHO. Over the course of this work, we identified several requirements that allowed us to play a role in the development of national and international policy, and that we believe transcend the application to polio.

Requirement 1: Vision, hard work, and timing

Our work on this project reflects some good luck, but more importantly it reflects vision and preparedness. In 1998, the lead author began working with a doctoral student to review the published, peer-reviewed pediatric cost-effectiveness and cost-benefit analyses in an effort to assess the state of the literature. That review yielded several key insights, including the preponderance of analyses for vaccine interventions, little recognition of variability and uncertainty, little consideration of changes over time, and a lack of dynamic disease modeling in the context of assessing the benefits of vaccines. The lead author also appreciated the opportunity that existed to bring together multiple analytical tools into integrated policy analyses. For example, in the case of cost-effectiveness modeling for vaccine interventions, the use of system dynamics tools for modeling outbreaks and the importance of factoring in the benefits of herd or population immunity seemed like an obvious and important combination to demonstrate quantitatively. In 2000, the second author contacted the lead author looking for a master's thesis project and expressed interest in working on this topic. Performing a

retrospective cost-effectiveness analysis emerged as the best strategy for demonstrating the importance of considering time and using a dynamic disease model, and we selected polio as a case study because it offered a long history (but not too long) and we did not find any evidence of other analysts already working on it. We also recognized that the approaching eradication of polio meant that this would be a good time to reflect on the aggregate benefits of historical interventions, and while other analysts apparently viewed polio as no longer interesting, we suspected that it would continue to be of interest for the foreseeable future.

Throughout most of 2001, we reviewed the literature on polio and built a model that allowed us to perform a first-cut retrospective cost-effectiveness analysis for polio interventions, which we presented at the December Society for Risk Analysis (SRA) meeting. Figure 1 shows our first draft model results demonstrating the importance of using a dynamic model to capture the population immunity benefits of OPV. The lines in the figure demonstrate the expected number of annual paralytic cases over time in the absence of vaccination, with vaccination using a static model (one in which only the recipient derives benefits from the vaccine), and with vaccination using a dynamic model. In the process of reviewing the literature and building our model, we contacted a number of polio researchers, including leaders at the U.S. Centers for Disease Control and Prevention (CDC). In addition, the lead author collaborated with another vaccine economics researcher at Harvard to develop the Harvard-CDC Joint Initiative on Vaccine Economics (JIVE) project, which created a mechanism for us to develop research collaborations with the CDC. We traveled from the SRA meeting to the CDC and presented our insights to leaders of the CDC's Global Immunization Division and Polio Eradication Branch. We hoped that our presentation would yield better data and information about some of the big uncertainties that we brought to the discussion, and we found at the end that the people we met with immediately appreciated the opportunity that our model offered with respect to modeling decisions prospectively. The vision that integrated policy modeling is essential to provide the best possible information about the risks, costs, and benefits of decisions and our efforts to both dig in to the problem and learn the history made it possible for us to be in the right place at the right time. Our collaboration with CDC's polio team really began at that meeting.

Requirement 2: Everyone on the team must recognize that the process matters and manage expectations

As we started to discuss how we might support the CDC's efforts, we quickly learned about the size of the community involved in polio eradication and the complexity of the decisions. We also recognized that the CDC group included a wide range of scientists, both laboratory experts and epidemiologists, but that it did not include many economists, modelers, or decision analysts. At the time, the only economist in the group included a newly hired person working on a paper that focused on estimating the costs of the five options that the CDC and WHO identified as the major decisions. We started by developing a proposal for our effort that focused on modeling these five options and using our dynamic model to estimate the impacts of future outbreaks that might occur in the context of this set of policies. This generated a lot of questions from everyone involved. As we learned more, we began to suggest that the model would need to deal with many complexities not by ignoring them but by characterizing them. We expanded the model to explicitly consider the risks that could lead to outbreaks and we started to talk about how

different countries might bring very different perspectives to the discussion. We also began to delve into the uncertainties and to sift through a huge amount of information.

In October 2002, with the project in its infancy, we traveled to WHO to participate in an informal consultation on economics research on post-certification immunization polices. Figure 2 shows a picture of our model that we presented at that meeting. We felt pleased that we could fit a simple schematic of the model on a single page, but this proved too complex for some of the people attending the meeting. That meeting led to three important insights: (1) WHO took the lead on policy with input from the CDC and other leaders, so influencing the actual policy meant we ultimately needed to work with and help WHO, (2) the WHO policy makers appeared very pessimistic about the actual utility of economic models in the context of influencing policy, and (3) the economists attending the meeting (all with experience working on prior economic analyses with WHO) expressed frustration about the policy makers asking the wrong questions or changing the questions at the end of their analyses. We expected to attend that meeting to give the group an introduction to our efforts. Instead, we realized as we listened to the discussion that we needed to make the case for why the type of analysis we were doing would be helpful. One of the most important last-minute additions to our presentation came from extracting the following quote made by one of the key WHO decision makers who noted that: "Though the [World Health Assembly] was the appropriate forum for discussing and debating the merits of an eradication initiative against polio, it has been argued that the delegates may not have had sufficient information to make a truly informed decision." The fact that World Health Assembly committed to the eradication of polio without data on the estimated human or financial resources and with no clear statement on the strategies that would need to be pursued or timeframe surprised us, and this helped us make the case that it was not too late to give decision makers good information.

Listening to the other analysts there, we also recognized that we would need to deal with the complexities, even if the decision makers did not care about them at the time and would willingly scope them out. More than one of the analysts talked about being told not to consider an issue for an analysis and then later finding the results rejected because the analysis did not consider that issue. We listened carefully to what prior analysts said about their experiences and we observed the interactions between our CDC collaborators and the WHO. Based on this, we decided that we needed to make the model as simple as we could and expect to model in the context of discussions in which some of the participants might prefer to play down the complexities. However, we also still needed to model the complexities rigorously and expect that the information might change. This made it essential for us to be flexible and clear about what others could expect from us and what we needed from them. At this point, we focused on process and emphasized in our discussions and writing that decision makers that rely on analytical tools consistently identify the process of developing the model as valuable, particularly in clarifying the key issues and sorting through the options with the best available information. We worked to make it clear that instead of focusing on providing a model with "the" answer, we offered a rigorous process and tools for the entire team that would provide defensible and helpful insights to help them evaluate the options and make more informed choices.

Requirement 3: Effective communication (essential!) and beer (helpful)

The enthusiasm about the model within the group of collaborators at CDC helped tremendously in our efforts to collect and synthesize information. We found one of the most unique aspects of the CDC polio team came from the impressive collaboration between the lab and field scientists. Even though the lab lies within one division of the CDC and the field team in another, nearly 10 years ago the leaders of these two groups started a tradition of everyone meeting at Moe's (a local dive) on Friday after work to talk over beer and encourage collaborative solutions. We truly became part of the team when one of the meetings we attended in Atlanta ended with a trip to Moe's. From our perspective, the receptivity of our work and recognition of its potential impact by members of the CDC polio team stemmed initially from a commitment by a high-level decision maker at the CDC to give us a chance to work on this project following our first meeting there. However, we believe that the relationships that we developed with our collaborators as partners in the modeling process became the most important driver for the expansion and continuation of the project. While analysts cannot always choose to enter a project with a highly functional team, we appreciated that our efforts on this project did not include fixing a dysfunctional team as a first step, and that this represented an important advantage. We also took the approach of working with everyone that we could at all levels, from the folks at the top to the new people, and we viewed all interactions as opportunities to learn, as well as to answer any questions that others might raise about modeling.

With respect to working with WHO, we encountered a few bumps in the communication process. After our first presentation and initial enthusiasm expressed by WHO at the informal consultation, some of our CDC collaborators attended a follow up meeting where they presented the model. After that meeting, they returned not with information about some of the key questions we raised, but with a message of skepticism on the part of the WHO that investment in our complex model and the process would yield helpful information. Fortunately, our CDC collaborators remained firm in their commitment to the process and we focused on building the model working with them. One of our first tasks centered on enumerating the complete list of the actual decisions and their interactions, which we achieved using the decision tree shown in Figure 3, and doing so in a way that recognized the different starting points for individual nations (Sangrujee et al., 2003). We decided that we would need to stratify the world to capture the highest-level of variability in the decision makers and risks, because we clearly saw the differences in starting points and that countries might rationally prefer different strategies depending on their risks and available resources. In particular, we knew that highly developed countries like the U.S. would probably continue routine inactivated polio vaccine (IPV) use into the foreseeable future, even with the global eradication of wild polioviruses. At the same time, competing demands for resources would probably lead developing countries looking at the opportunity to stop oral polio vaccine (OPV) to prefer to stop polio vaccinations completely as soon as possible after global eradication. We presented our framework of the decision options when we gave an update of our model at another consultation meeting at the WHO in September 2003. The presentation proved very useful to the WHO given discussions at the time about the post-eradication risks and its qualitative framework for characterizing them, and at that point the WHO took a strong interest in our efforts to characterize the risks. We began to see evidence of our work in some of the discussions and in the development of the PEI strategic plan.

Requirement 4: Recognition of the responsibility of analysts to organize, demand rigor, and sustain confidence in the vision

Finally, with the decision options identified, we turned our attention to refining our dynamic disease model and the other components of the overall model (Duintjer Tebbens et al., 2005a,b,c) with the simplified diagram of the components shown in Figure 4. The dynamic disease model represents one of the most significant aspects of our efforts. Remarkably, no similar model existed, even with all of the historical outbreaks that occurred with polio. Although we started with a relatively simple SIR (Susceptible-Infected-Removed) model, the model quickly grew in complexity to account for important aspects of the model. For example, we recognized that the underlying population immunity structure reflected a mixture of individuals fully susceptible to disease, and those partially infectible, with important differences between those who had recent infection with the wild virus or vaccine, those with historic live virus infection, and those protected by IPV vaccination. Figure 5 shows a simplified schematic for the first age group (with newborns entering the model, but not showing the similar schematic for the second age group or the arrows connecting them). We ultimately decided to model 25 separate age groups with different initial immunity profiles to capture the differences in historical vaccination and to maintain or ability to include different target population groups in the context of modeling a wide range of potential outbreak response strategies. We went through several rounds of iteration on the model as we talked with the experts at the CDC and WHO about the model inputs. Fortunately, our colleagues also offered to work with us to apply the model to real outbreaks so that we could learn about how the model predicted some retrospective cases. In the context of informing future policies, we face the challenge of trying to characterize the impacts of potential outbreaks prospectively, and given uncertainty about the actual conditions we must make a significant number of assumptions and model the uncertainty about the many possible futures. Figure 6 provides an example of a potential outbreak and demonstrates the expected impacts of different strategies for responding to that outbreak. We emphasize that once specific information about any real outbreak becomes available the model can and should use that information to provide a more-informed estimate (e.g., the model uses as a starting point characteristics of the different types of polioviruses "averaged" for types 1, 2, and 3, but clearly once we know the type of virus responsible for any given outbreak the model should use inputs specific for that type).

Early on in the process, we decided that we would develop the various components of the model and write papers to submit for peer review and publication. We recognized that our model would need to withstand critical reviews by technical experts in addition to answering the high-level questions posed by policy makers. This expectation on our part raised the bar with respect to the rigor and increased the amount of iteration and discussion, but it also led to significant delays associated with obtaining required clearance on manuscripts with collaborating authors from CDC and WHO. Fortunately we started the project well in front of the discussions about post-eradication policy and unfortunately the Polio Eradication Initiative faced significant challenges in interrupting transmission of wild polioviruses, which also delayed the timing of the global discussions and negotiation. Our commitment to build the complex model also meant that we took a longer path than what we expected at the beginning. However, we now find ourselves with a useful tool at the point in time when it can make a difference.

In September 2004, we presented some of the first key insights from our model at WHO and found the decision makers at multiple levels very engaged. Our demonstration of the potential outbreaks and outcomes from one of the key risks in the first few years (i.e., the risks of circulating vaccine-derived polio viruses) provided strength to the case that eradication of polioviruses would require eradication of OPV as soon as possible after the confirmed eradication of wild polioviruses. We also emphasized that policy makers should prepare the world for the relatively high probability of at least one outbreak occurring after eradication, and for the concept that we must prepare to respond to such an outbreak and not consider such an event to signal the failure of global eradication. Our model included key placeholders for a stockpile and response strategy, and we continue to emphasize the importance of planning and preparation. After the September 2004 meeting, our CDC and WHO collaborators began to work with us very intensively to review the model components, investing a significant amount of time in the final iterations of the model. In the end, our persistence (made possible by strong collaboration and support from the CDC) prevailed. During 2004, we also began to work with the vaccine policy leaders of the Pan American Health Organization (PAHO) on a regional model, which we presented at a Technical Advisory Group meeting in November 2004. At this point, our efforts continue to expand and utility of the process and the model continue to grow. However, we continue to also learn more about the multiple feedback loops and stakeholders with interests in the process.

Requirement 5: Curiosity and good questions

As we look to the future, we see several tests of the model and process that will ultimately determine its fate. We hope that as long as we continue asking good questions we will provide valuable contributions to the team. By asking the questions that need asking, as opposed to the ones with easy answers, our role as analysts has thus far included challenging others on the team to think about the system more broadly, and to more fully consider the interactions of the various components and changes over time. Some of the biggest challenges that we anticipate will come in providing answers to key questions in ways that meet the needs of different audiences and reflect the best available scientific evidence. While we talk about variability, national leaders talk about the options as they see them. Our discussions about the future and highly uncertain risks (e.g., bioterrorism) reveal very different perceptions of the risks and preferences for managing uncertainties. Successful elimination of paralytic polio will require the elimination of OPV, yet coordinating the process of cessation will mean obtaining global agreement with the policy and compliance with its implementation.

Modeling the risks suggests a relatively high probability of at least one case of paralytic polio occurring after OPV cessation and emphasizes the need for preparations to prevent re-establishment of circulation, and we emphasize that analysts must consider the dynamic nature of the choices and use appropriate models. This comes of no surprise to analysts who concentrate their efforts on system dynamics, but we note that dynamic models remain underutilized in many policy modeling contexts.

Coming full circle to the retrospective analysis, we make a strong analytical case for economic analysts to use dynamic models in cost-effectiveness and we highlight the changes that occur in the cost-effectiveness of vaccines over the course of their lifecycles. We believe that the polio

experience provides a wealth of lessons, most notably the need to appreciate the major changes that occur over time.

Conclusions

Analysts can contribute significantly to real policy decisions, but only if they commit to a process for doing so and recognize the dynamic nature of both policy and science. We hope that providing a review of the requirements for our success (at least success to this point) yield some useful insights for other modelers and similar processes. We recognize that a model only helps when people use it, and in this regard we aspire to create and maintain a living model that will support the decisions of real high-level policy makers as they continue to face the challenges of managing the global risks of polio over time.

Acknowledgments

The authors used resources from unrestricted grants to the Harvard Kids Risk Project to support their efforts on this paper. Much of the work described in this paper derives from the authors' collaboration supported by CDC grant number U50/CCU300860, TS-0675. The authors gratefully thank their manly collaborators at CDC, WHO, and around the world who continue to contribute to the team effort required to build a successful model to support real policy.

Figure 1: First results from our retrospective model, presented at the CDC in December 2001, showing paralytic polio incidence without vaccine (top-blue), with vaccine static (middle-green), and with vaccine dynamic (bottom-red) transmission model

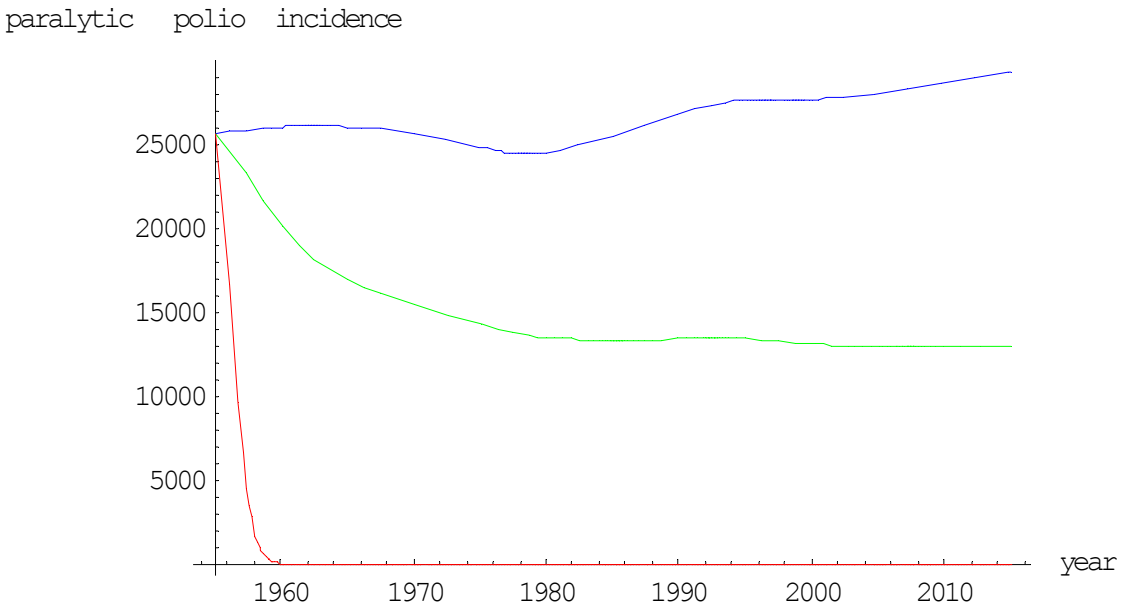

Figure 2: Initial influence diagram of the polio decision model

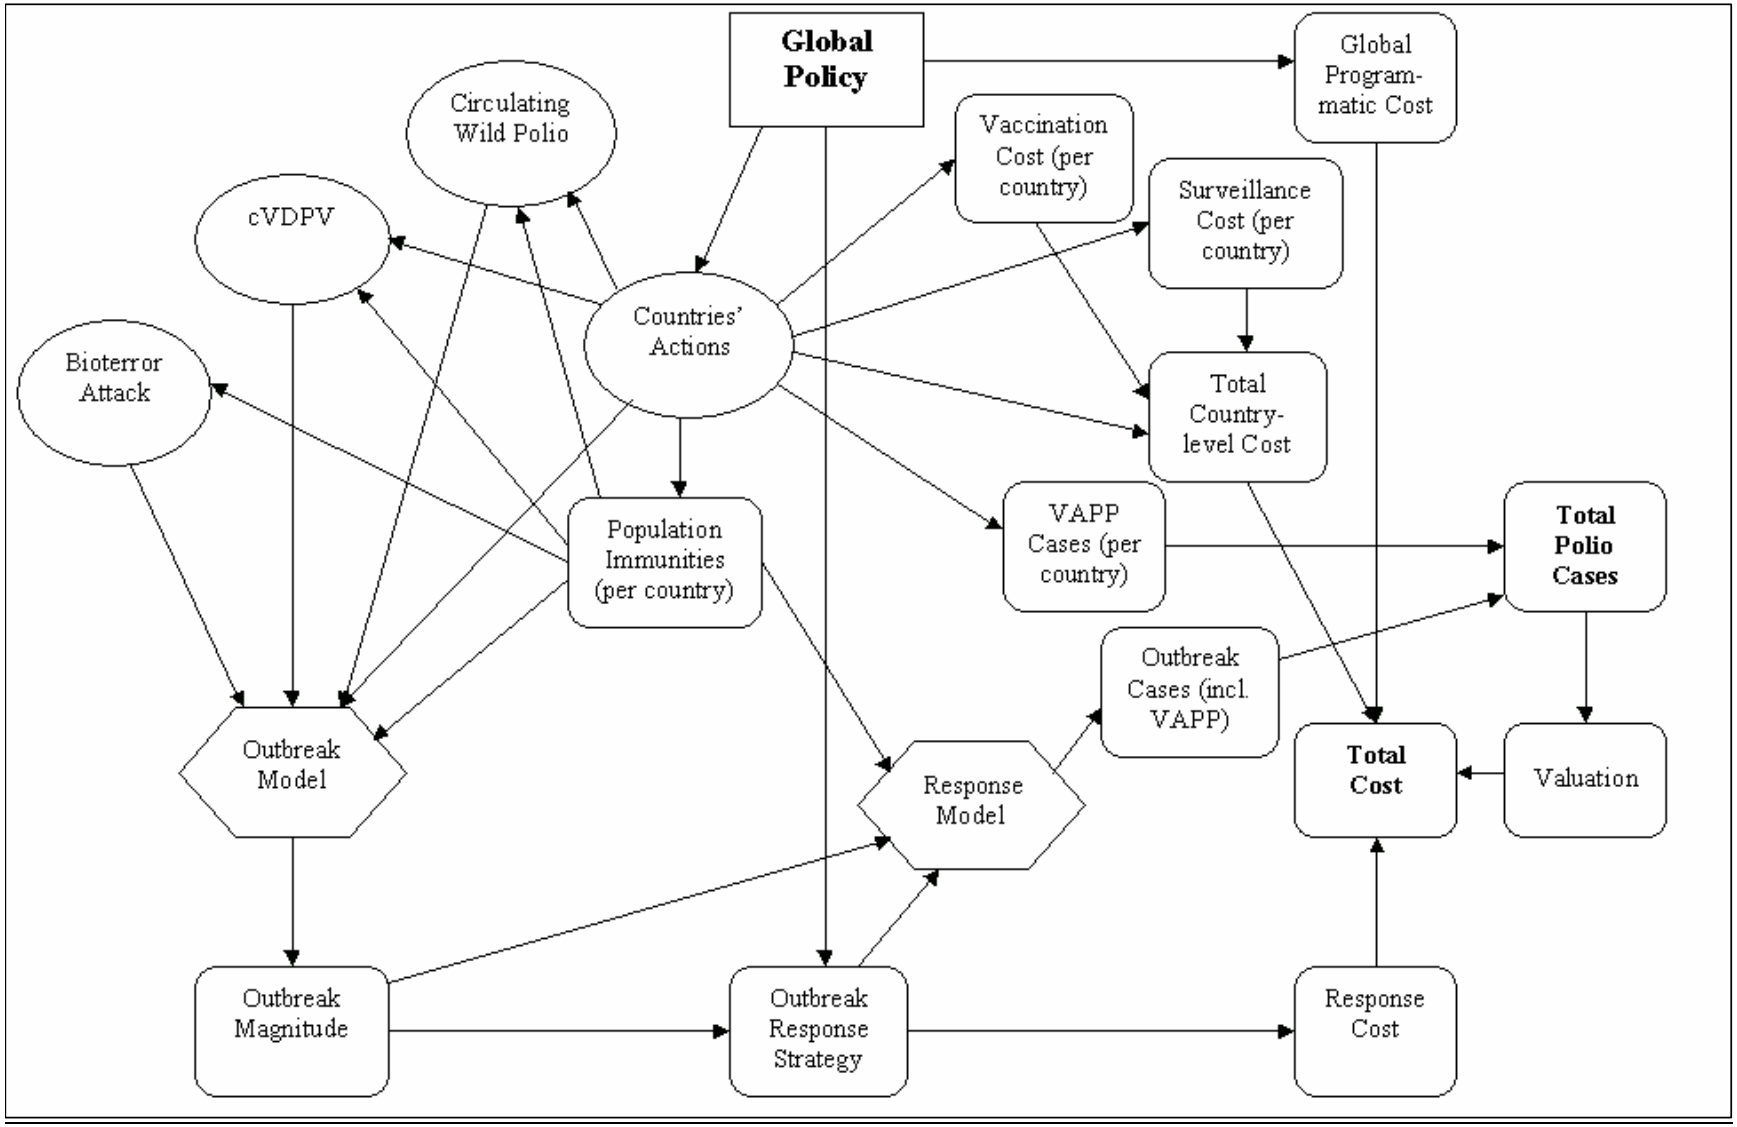

Figure 3: Major decision options for all countries – first five years after certification of the world as free of wild polioviruses

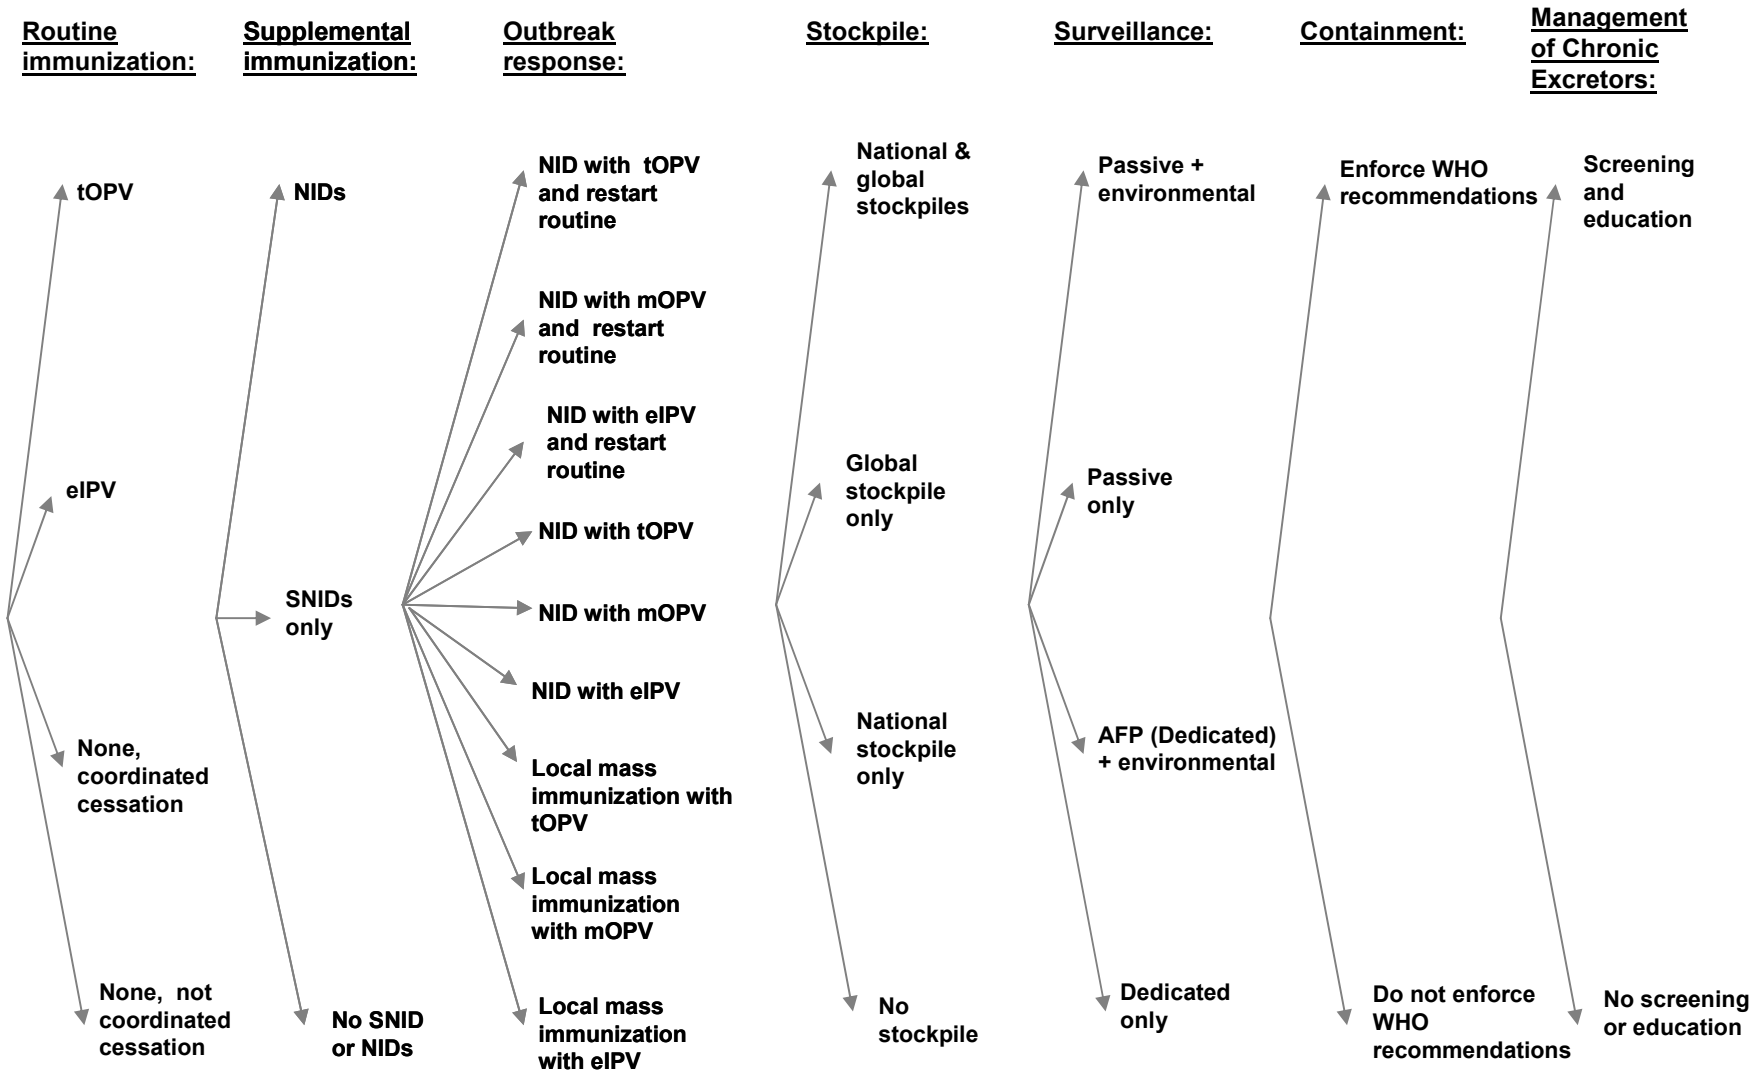

Figure 4: Simplified diagram of model components

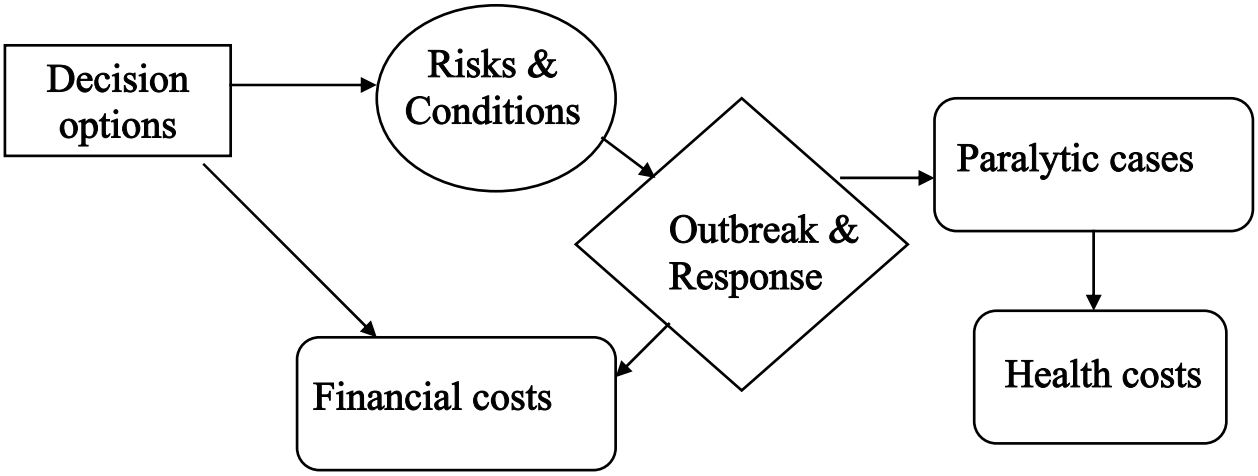

Figure 5: Schematic of dynamic disease model components for the first of 25 age groups extending on basic SIR though inclusion of 3 different groups with partial immunity (due to prior recent or historical vaccination or exposure to wild poliovirus)

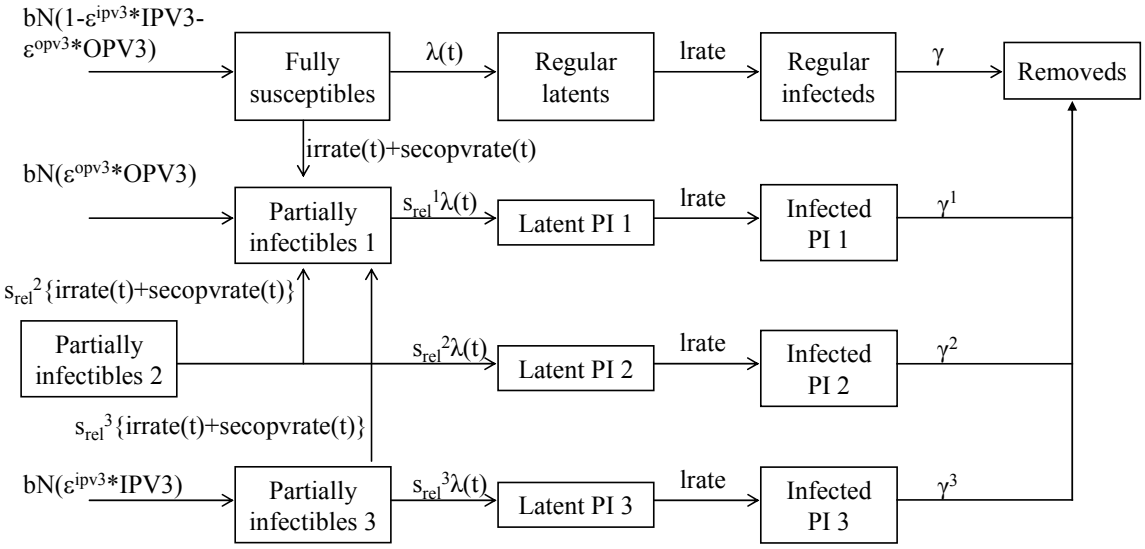

$$\lambda(t) = (\gamma R_0/N)*\Sigma^{age}[RI(t)+ \Sigma^{\text{partially infecteds}}\{i_{rel}*IPI(t)\}]$$

Figure 6: Example of a prospectively modeled outbreak after cessation in a hypothetical country*

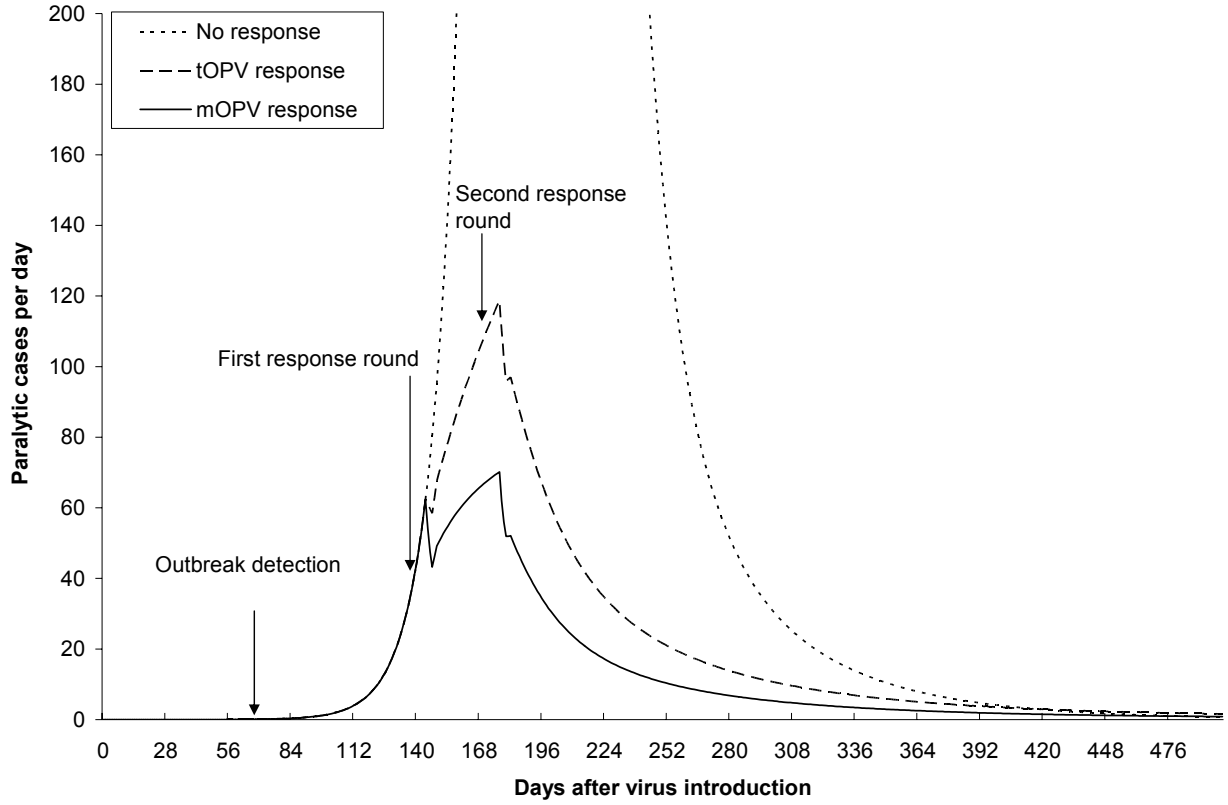

*Source: Duintjer Tebbens et al. (2005a). We assume the outbreak occurs in a low-income country with R_0=13 and a population of 100 million people 5 years after cessation of all polio immunizations and 10 years after stopping supplemental immunization activities. Detection occurs as soon as the cumulative incidence reaches 1 paralytic case and we assume a delay from detection to response of 70 days. The response scenarios assume two immunization rounds at a 30-day interval covering 90% of all children younger than 5 years of age in 3 days. The "no response" curve reaches a peak of over 1,700 cases on day 197, but ultimately the failure to respond would lead to exportations of cases into other parts of the world and to global recirculation of polioviruses signaling the failure of eradication.

References

Box, GEP. "Robustness in the strategy of scientific model building." In Robustness in Statistics, Launer RL and Wilkinson GN (Eds.), New York, NY: Academic Press, Inc., 1979, page 202.

Duintjer Tebbens RJ, Pallansch MA, Kew OM, Cáceres VM, Sutter RW, Thompson KM. "A dynamic model of poliomyelitis outbreaks: Learning from the past to help inform the future." American Journal of Epidemiology, 2005a; In press.

Duintjer Tebbens RJ, Pallansch MA, Kew OM, Cáceres VM, Jafari, H, Cochi, SL, Sutter, RW, Aylward B, Thompson KM. "Risks of paralytic disease due to wild or vaccine-derived poliovirus after eradication." Submitted to Risk Analysis, 2005b.

Duintjer Tebbens RJ, Sangrujee N, Thompson KM. "The costs of future polio risk management policies." Submitted to Risk Analysis, 2005c.

Francis T Jr, Korns RF, Voight RB, Boisen M, Hemphill FM, Napier JA, Tolchinsky E. "An evaluation of the 1954 poliomyelitis vaccine trials." Am J Public Health. 1955;45(5, Part 2):1-63.

Miller MA, Sutter RW, Strebel PM, Hadler SC. "Cost-effectiveness of incorporating inactivated poliovirus vaccine into the routine childhood immunization schedule. Journal of the American Medical Association." 1996;276(12):967-971.

Sangrujee NK, Duintjer Tebbens RJ, Cáceres VM, Thompson KM. "Policy decision options during the first five years following certification of polio eradication." Medscape General Medicine 2003(December 19);5(4). (Available at: http://www.medscape.com/viewarticle/464841).

Sterman JD. "All models are wrong: Reflections on becoming a systems scientist." System Dynamics Review 2002;18(4):525.

Thompson KM, Duintjer Tebbens RJ. "Retrospective cost-effectiveness analyses for polio vaccination in the United States." Submitted to Journal of the American Medical Association.