Learning from the pioneers:

Lessons about data platforms
drawn from the WWARN experience

Elizabeth Pisani and Stella Botchway
September 2016

Executive Summary

Researchers are increasingly being encouraged to share individual patient data from clinical trials, but there's remarkably little experience about how that can be done equitably, ethically and efficiently, especially for diseases where most research takes place in low and middle income countries. One pioneer in this area is the WorldWide Antimalarial Resistance Network (WWARN), conceived almost a decade ago by malaria researchers from across the globe to bring together clinical, in vitro, molecular, pharmacological and (later) medicine quality data. It was hoped these data, analysed together, would allow for the efficient tracking of drug-resistant malaria.

WWARN facts and figures (at October 2016)

Individual patients included in database:	135,000
Clinical trials included in database:	186
% of all published trials of artemisinin combination therapies:	80%
Molecular studies included in database	103
Number of data contributors and other collaborators	268
Published individual patient meta-analyses using WWARN-standardised data	17
Principal funder:	Bill and Melinda Gates Foundation
Additional support:	ExxonMobil Foundation, Ministry of Foreign Affairs, France
	European & Developing Countries Clinical Trials Partnership
	UK Medical Research Council, UK Department for International Development

Researchers and those that fund them, together with the journals that publish research results, are all moving towards a more "open" model of science, of which data sharing is an important part. To contribute to discussions about how best to share data, the Wellcome Trust -- on behalf of the Public Health Research Data Forum -- commissioned a study to capture the learning provided by the experience of WWARN, which is both rich and still rare. WWARN founders, staff and collaborators were generous in collaborating with the study, which was carried out by independent researchers. The study, which was designed as a learning exercise rather than a formal evaluation, was based on a comprehensive records review, in-depth interviews with 47 people involved with WWARN (including some who have chosen not to contribute data to the platform), and a witness seminar which yielded sometimes divergent views about the genesis and development of the network.

The aims of WWARN shifted over time but there's no doubt that the collaboration has contributed to a better understanding of malaria treatment efficacy. Pooled analyses based on data compiled and standardised by WWARN have informed changes to international guidelines on antimalarial treatment and dosage (see box, below). The collaboration, housed at Oxford University since 2009, is now under the umbrella of the Infectious Disease Data Observatory (IDDO). Supported by the Wellcome Trust, Médecins Sans Frontières, the WHO-affiliated Special Programme for Research and Training in Tropical Diseases (TDR), Drugs for Neglected Diseases Initiative and Foundation for Innovative New Diagnostics, IDDO is adapting the platform for Ebola, visceral leishmaniasis and other diseases.

Seven key lessons were identified around the factors that make sharing individual patient and pathogen data feasible and useful – as summarised below. We hope that WWARN's

experience will help inform the efforts of other researchers and, particularly, those of funders, policy-makers, companies and other communities that support and benefit from scientific enquiry, as they work to share knowledge to improve health.

Lesson 1: Data sharing platforms work well when the motivations of platform funders, developers, contributors and users are considered at the outset, and the incentives driving each of them are aligned.

WWARN initially hoped to provide national policy-makers with real-time information about the geographic spread of drug-resistant malaria. This goal was not realised because incentives were misaligned. There was little demand for global analysis from national policy makers. More importantly, clinical trialists in endemic countries (who still advance through publication in peer reviewed journals) did not want to share data with other scientists before they had published their own results.

The platform took off after WWARN switched its focus to pooled analyses. WWARN now requests data from principal investigators conducting efficacy trials of antimalarial drugs to answer specific research questions. Data contributors are invited to participate in analysis and paper writing, and are credited appropriately on resulting publications. This motivated researchers to contribute their data to the curated resource, and has produced world-class science: new methods have been developed and pooled analyses have led to changes in international malaria treatment guidelines. Finally, potential users, including drug developers and the World Health Organisation, are approaching the platform with specific questions.

Lesson 2: Both disease experts and data scientists are core to the design of a successful data sharing platform with public health aims

The malaria specialists who conceived of WWARN wished to minimise barriers to sharing and maximise flexibility of potential outputs. While they put in countless unpaid hours debating which information was most useful and developing standards, they resisted pre-defining the end uses of the database. Information scientists, who know how to develop shared resources to achieve network effects, pressed for greater clarity of purpose at the outset. In retrospect, defining the core purpose of the database more clearly at the start -- a task for disease experts -- would have resulted in a more efficient development process.

Lesson 3: Data curation is expensive, but essential if shared data are to be useful

WWARN invested a great deal of time and effort developing the standards and the tools to allow datasets to be standardised and combined across time and location. It is these investments that make the shared data useful. The value of a well-curated data set grows over time but the costs are front-loaded; investors should not expect a quick return. However the cost of curation tools developed for one disease platform may be apportioned across future platforms, because instead of developing new systems from scratch, many existing structures can be adapted.

Lesson 4: Data sharing platforms should be underpinned by clear, equitable governance structures that can evolve in line with changing community norms.

Less visible than the curation tools, but just as important, are the procedures WWARN has developed to govern contribution and use of data shared through the platform. From the start, data use agreements were clearly worded and not overly legalistic. Developed in uncharted territory in the face of widespread scepticism, the early terms of submission required permission from investigators for every use of shared data -- a huge administrative burden for the platform. Community confidence and norms supporting data sharing have since grown, and with them the possibility for more open models of sharing. Terms which default to greater sharing while allowing contributors to exercise more control if desired may help balance trust-building with the flexibility to evolve along with the data sharing zeitgeist.

New partnerships lead to rapid, policy-relevant analyses

More than a decade ago, researchers from the Liverpool School of Tropical Medicine collaborated with the WHO-housed Special Programme for Research and Training in Tropical Diseases (TDR) to show that weight-for-age translations were inaccurate in some parts of the world, meaning that children may be getting inappropriate doses of medication. Their re-calculated rates were used by pharmaceutical firm Sanofi in formulating fixed dose combination of the antimalarial artesunate - amodiaquine.

In late 2013, WWARN published an analysis based on data from 7,072 patients, pooled across 26 studies. They showed that young children taking dihydroartemisinin-piperaquine seemed to be getting a lower intake dose than adults, and were at higher risk for recrudescence. This suggested that children may need higher doses, but it wasn't known if higher doses would be safe. This evidence was independently reviewed by WHO, which changed its dosing recommendations for children taking the drug. The Liverpool group quickly switched the focus of a planned dosing study, looking instead at the efficacy and risk of cardiotoxicity of a higher dose of DHA-piperaquine in children. In short, a hypothesis derived from pooled analysis of well-curated, shared data led rapidly to a targeted clinical trial. An urgent question was answered, and policy quickly changed. The result should be fewer treatment failures and more healthy children.

Lesson 5: Institutional arrangements have important implications for data sharing; these should be considered with care at the outset of a data sharing venture.

Initiated by a small group of relatively well-resourced malaria specialists in consultation with senior researchers from many malaria-endemic countries, WWARN was conceived as a loose network of professional researchers. For practical reasons, driven in part by the needs of funders, it was then embedded within an academic department at Oxford University. This arrangement has had important consequences. Staff may have career expectations which are not advanced by their functional roles, particularly if they are subject to traditional publication-based measures of academic productivity. In addition, in a culture where "ownership" of data is still contested, siting a data sharing platform in a northern academic institution may create the perception of a "data grab" and limit the willingness of key partners, including endemic country researchers, to participate fully.

Lesson 6: More interaction with national authorities and participation of endemic country researchers in setting research questions may contribute to more locally relevant and actionable research results.

WWARN's current strength is now in producing high quality peer reviewed papers that make the best use of a large, multi-country database to yield learning of global significance. The platform has not yet found a ready audience among national policy makers, however.

A granular understanding of local data will become increasingly important as disease prevalence falls. Involvement of endemic country researchers in analysis and interpretation should thus also grow, but right now, endemic country researchers rarely use shared data resources. This is in part because those who have the skills to perform complex analyses are too senior to spend time on such tasks. Scientific collaborations, including data sharing networks, must be funded to work with partners to reinforce skills at more junior levels. Support should include financial and professional incentives that encourage endemic country researchers to conduct secondary analysis that answers questions raised by local policy makers.

Lesson 7: As data sharing platforms mature, institutional arrangements may shift

As lesson 2 suggests, scientists who specialise in disease areas are the critical drivers of platform development, and specialists in informatics design are essential to its conception. Disease specialists are also best placed to define early research questions, and will thus be key in the "proof-of-concept" phase, when the utility of the platform is demonstrated. Universities or other research bodies are thus ideal incubators for new disease platforms, but because of incentive structures and perceptions of conflict of interest, they may not be the best home for data platforms in the long term. Since there are so few models to draw on, it's impossible to recommend alternative models with confidence. However as data platforms mature, it is possible that day-to-day data management may be more cost-effectively handled by a neutral public health entity that employs a few specialist advisors and enforces standards and transparent governance structures developed and agreed by the broader scientific community. University-based disease specialists are likely to remain the most important users of a platform. They must continue to be involved in shaping its direction, without necessarily controlling it.

1 Background

In recent years, many funders of biomedical research have begun to encourage -- and in some cases oblige -- grant-holders to "share" their data with other researchers. The belief is that data shared will become data reused; the combination of data generated by different studies will allow for the reproduction of initial analysis and for more powerful meta-analysis of complex questions; these analyses will in turn translate into better policies and practice and thus to better health, at limited additional cost. Squeezing more knowledge out of existing data is considered especially important for the often neglected diseases that affect people in poorer countries, where most research is funded by public agencies or charities.

One organisation that has brought scientists, their data and their brains together successfully is the WorldWide Antimalarial Resistance Network (WWARN), first mooted in a poster session at the Molecular Approaches to Malaria conference in Melbourne in 2004, actively planned since 2007 and formally established in 2009. The WWARN database has since its inception included data covering four areas of importance in tracking drug-resistant malaria: clinical, in vitro, molecular and pharmacological. Data on the quality of antimalarial medicine was added in 2010. Data are extremely granular, relating to individual patients or parasites. The collaboration has so far assembled information from 186 clinical trials including 135,000 patients. This includes a remarkable 80% of all published clinical trials reporting on the efficacy of artemisinin combination therapies. A further 103 molecular studies are included in the database. These data have been used to produce a number of pooled analyses, published in the scientific literature, which have informed changes to international guidelines on antimalarial use and dosage.

WWARN's achievement is particularly notable in that the data sharing initiative was conceived and instituted before funders, academic institutions or scientific journals had begun actively to encourage or require such initiatives, when most researchers were generally hostile to the very concept of allowing other analysts access to the information they had collected. As one interviewee for this study put it: "Seven years ago, data sharing was a swear word." In addition, there were few models to learn from, so WWARN proceeded in part by trying things out, expanding on things that worked and adapting things that didn't.

In the intervening years the tide has begun to turn in favour of sharing data that might be combined with other information to lead to more rapid gains in public health. The failure to share research and surveillance data quickly during the West African Ebola outbreak in 2012, which may have contributed to much avoidable suffering, put wind in the sails of the open data movement. However the many institutions now actively promoting data sharing give very little guidance on how data should be shared in order to maximise health gains, in part because they have little systematic information on what works, what doesn't, and why.

1.1 Purpose and methods of this study

Aiming to begin to fill that information gap, the current study analyses the extraordinarily rich body of experience built up over the past decade by WWARN, which is unquestionably a pioneer among those sharing individual patient level data from clinical trials conducted in low and middle income settings. The goal is to provide a solidly evidence-based analysis of the building blocks of a successful data sharing platform. Our aim is to identify which elements of WWARN's experience are fundamental to meaningful data sharing and which may have been artefacts of timing, particular political or funding constraints, disease-specific characteristics or other factors perhaps less pertinent for future platform development. We hope to draw lessons that will help inform the members of the Public Health Research Data Forum as they continue to develop their policies and support for data sharing.

This study, funded by the Wellcome Trust, is not intended to provide an evaluation of WWARN. The staff and collaborators of WWARN have been extraordinarily generous with their time and resources, providing full access to project records as well as many hours of interview time. While they provided comments on a first draft of this paper, WWARN staff were not involved in the analysis of the data, nor in the formulation of lessons or conclusions.

Methods

This study is based on three major sources of information: a records review, a series of in-depth interviews and a health histories seminar. We also conducted an analysis of publications by WWARN and its collaborators, the results of which will be reported elsewhere.

Records review

We reviewed records dating back to the earliest presentations of the concept that became WWARN. These records included: conference and meeting presentations; papers published in academic journals; grant proposals and grant-related reporting forms; progress reports by scientific module heads, internal strategic plans and business plans; minutes of all board and Scientific Advisory Committee (SAC) meetings; iterations of terms of submission, memoranda of understanding and contracts with data providers; correspondence with the World Health Organisation (WHO) and other key partners; analytics covering use of the WWARN website, downloads of tools and social media reach; surveys of user and stakeholder attitudes; the current contents of the WWARN website.

The minutes of all board and SAC meetings, the progress reports from scientific module heads and key academic papers proposing or describing WWARN were entered into NVivo software and coded thematically from a codebook based initially on the original study protocol, and developed iteratively. For other documents, notes were taken and those were in turn coded as above.

In-depth interviews

The principal investigator conducted interviews, generally of 60-90 minutes in length, with 47 individuals, purposively selected to give a wide variety of perspectives on WWARN's evolution. Table 1 shows how the interviewees relate to WWARN.

Table 1: Characteristics of people interviewed for this study

Relationship to WWARN	Number of interviewees
Current WWARN staff/consultant	5
Former WWARN staff/consultant	7
Current or former scientific leaders or group head*	5
WWARN board members	3
WWARN scientific advisory committee members	2
Data contributors from industry	2
Other data contributors	3
Malaria researchers who do not contribute data	3
Secondary user or analyst	4
Global health organisation	11
Policy maker	2
TOTAL	**47**

*These are all also WWARN founders, and frequently referred to as such in the text

Of the 47 people interviewed, 21 are women and 11 are from low or middle income countries. Some 46% are/were based in Europe at the time of their involvement with WWARN, with the remainder evenly split between Asia/Australasia, Africa and North America. There were no interviews with people based in Latin America, a region that has not engaged extensively with WWARN to date.

Most of the science group heads, board members and scientific advisory committee members have also contributed data to the resource. In the text of the report, interviewees who fulfil more than one role are indentified by the role most relevant to the context of the quote.

Interviewees signed consent forms for the interviews. Where consent was given for recording, interviews were audio-recorded, and notes were simultaneously taken. One interviewee refused consent for recording; recording was not logistically possible in three cases, and equipment failed for two interviews. Transcripts and/or interview notes were entered into NVivo software and coded thematically as above.

Health histories seminar

On June 22, 2016, we brought together a number of individuals involved with the inception and evolution of WWARN in a "witness seminar" format. The seminar was chaired by Rob Terry, of the WHO-affiliated Special Programme for Research and Training in Tropical Diseases (TDR), and João Nunes, a lecturer in international relations from York University;

they were provided with a detailed issues brief by the study PI. The four-hour seminar allowed for spirited discussions of the genesis and development of WWARN. Alternative views were aired and in some cases reconciled, and additional records were unearthed for the records review.

The seminar was transcribed by a third party; the transcript was entered into NVivo software and coded thematically as above.

2 The WWARN story: a classic pivot

Like many pioneering organisations, WWARN has evolved significantly since its inception; in tech start-up language, it has performed a "pivot".

The current narrative is that WWARN was from the start intended as a collaboration between scientists in search of new knowledge. [1] However in early conference talks, published papers and grant proposals, WWARN was envisaged as a tool for surveillance. In 2006, WWARN co-founder Carol Sibley and WHO drug resistance surveillance expert Pascal Ringwald proposed: "An open, public database [which] would record the data in real time so that the data could be accessed, analysed and productively used by all interested parties". (C. H. Sibley and Ringwald 2006) Those "interested parties" were depicted predominantly as national policy makers and programme managers.[2]

It was only later that WWARN evolved into a successful scientific collaboration producing sophisticated, high quality pooled analyses of drug efficacy outcomes in individual patients. These analyses have generated important, testable hypotheses, advanced methodology, and informed malaria dosing and treatment decisions at the global level.

The course taken by WWARN is a common one among tech start-ups, in particular data-driven platforms. These ventures typically go through four major phases: conception, construction, proof of concept, and finally expansion and widespread use. Some 75% of all tech start-ups fail, typically because the idea is ahead of its time -- consumer demand or supporting technology don't develop as rapidly as the founders hoped -- or, conversely, because the idea is overtaken by unanticipated changes in technology or consumer preferences. (Gage 2012) Some start-ups, however, survive and thrive even when technology or the anticipated consumer base don't behave as expected. They do this by adapting their offering so that it is better suited to the actual circumstances. Legendary pivots include the micro-blogging site Twitter, and operating system Android. Twitter started off as a staff communications channel for podcast directory platform Odeo. Odeo died because consumers preferred to manage their podcasts using RSS feeds, but the Twitter by-product caught on; it now has 320 million active users worldwide. Android began life as an operating system for the camera market, which shrank as people used their phones to take pictures. Android switched its focus to the growing mobile phone market, and now ships with 80% of all new phones worldwide.(O'Hear and Lomas 2014)

[1] In its earliest incarnation, the proposed network was known as WARN: the World Antimalarial Resistance Network. The second W, which makes the organisation more easily identifiable in electronic searches while distinguishing it from the West African Research Network, was added by 2010.

[2] A review of PowerPoint presentations, minutes of meetings, published papers and grant proposals up to the main WWARN grant proposal of 2009 found 436 references to a database, 281 references to data collation or collection, and 187 mentions of surveillance or monitoring. There were 104 mentions of research/researchers/scientists and just 33 references to collaboration/collaborators.

WWARN's evolution fits this model. By examining the mismatch between expectations and actual events that arose in the conception and construction phases, and the later pivot to a model better suited to the evolving realities, we hope to shed light on some of the underlying determinants of successful data sharing. This section outlines the broad issues relating to motivations, incentives and relationships which influenced the path taken by WWARN. Lessons around more specific issues such as informatics, tools development, human resources and governance structures will be discussed in section 3.

2.1 Conception phase: a "no brainer" endeavour

WWARN was conceived to fill an important gap. Towards the end of the 20th century, great strides were made in the global fight against malaria, largely because of the introduction of a new class of antimalarials based on artemisinin. This responded to the failure of the common drugs chloroquine and sulfadoxine-pyrimethamine due to widespread resistance. By 2005, 11 million courses of artemisinin based drugs were being procured annually.(World Health Organisaton 2016a) In northwestern Cambodia -- the cradle of resistance to all other major classes of antimalarials -- patients treated with artemisinin were, however, beginning to take longer than expected to clear parasites from their blood.

This set alarm bells ringing among researchers. But it also sparked frustration. A great deal of research on the efficacy of various artemisinin-based combination therapies and other antimalarials was taking place, and routine surveillance data were also being collected by many countries using WHO protocols. WHO, which as an institution supports the concept of data sharing, had created a publicly-accessible database providing aggregate data from therapeutic efficacy studies. However clinical signs of treatment failure are the final signal of potential resistance. At least in theory, identifying and tracking genetic markers of resistance and resistance in vitro would flag the possibility of resistance earlier on, while pharmacological studies could help distinguish treatment failure due to resistance from that due to under-dosing, poor absorption or other mechanisms preventing an adequate blood level of the active drugs.

Losing artemisinin to resistance would represent a massive defeat in the war against malaria; to prevent that, it would be necessary to adapt drug formulation and treatment choices as soon as there was reliable evidence that resistance was emerging in a given geographical setting. In other words, funders and other policy-makers working with malaria globally needed to be able to see the "big picture" of emerging resistance. And national policy makers needed to contextualise their local information and act quickly to adapt treatment guidelines and procurement practices as necessary.

A straightforward way to meet these needs was to bring all of the existing clinical research and surveillance data together with genetic and in vitro data that provided earlier signals of resistance. Pharmacological data which checked for under-dosing as an alternative source of treatment failure would also improve understanding. As several of the people involved with the conception and early support of WWARN said in interviews, an integrated, openly accessible database aimed at informing policy makers in their choices about antimalarial use was a "no brainer". In 2007, the body that became WWARN received a grant of close to US$100,000 from the Bill and Melinda Gates Foundation (BMGF) for an initial exploratory conference, held in Hinxton, UK. Scientists from both northern and malaria endemic countries discussed the potential platform and a "coalition of the willing" was formed. The following year, the founders were granted a further US$ 990,000 to cover the year-long

planning of a major grant application. As a result, two grants, of US$20.6 million to Oxford University and US$7.5 million to WHO's Global Malaria Programme, were approved in 2009. But the real-time surveillance database they were originally intended to support did not materialise in the form that founders first anticipated. Why not?

Incentives not aligned

The first reason that data weren't shared as easily as anticipated is that there wasn't enough overlap in the interests of the two groups that WWARN identified in its 2009 grant application as its "stakeholders": "*those who gather primary data for local use and policy makers and funders who need timely information on antimalarial drug efficacy on a far wider scale.*"

The platform users defined by WWARN split largely along the lines of supply and demand. Researchers collecting primary data were to be suppliers. While programme managers might supply surveillance data, the role of policy makers and funders was largely seen as being on the demand side. However it was never clear what expressed need the platform would fulfil for either side.

The supply side

Table 2 summarises what interviewees said about the incentives and disincentives to contribute individual patient data from clinical trials and other therapeutic efficacy studies to the open access, real-time database initially envisaged by WWARN.

Table 2: Incentives and barriers to contributing data to the early WWARN model

Data generator	Incentive for data contribution	Barrier to data contribution
National programmes	Weak: interests are national, not global	Moderate: data quality may be questioned
Academic researchers	Moderate: possibly contribute to better health	Strong: may undermine possibility for publication
NGOs	Strong: public health benefit highly valued	Weak: minor ethical concerns over privacy
Industry	Moderate, in companies which value transparency	Strong: reanalysis may undermine regulatory filings

Academic researchers in southern institutions consider themselves as having most to lose. By depositing data into a database, and especially one housed in a well-known Western university, they allow for the possibility that analysts who may have better skills, better ideas, better software and more time away from the rigours of the clinic or the lab will use data collected in difficult conditions in endemic countries. Until the "proof of concept" papers that came much later, interviewees also worried that weaknesses or errors in the data would be made public or that data would be misused in ways that discredit their community, research institution or nation. Though WWARN provided data standardisation services, these were not seen as a sufficient benefit to outweigh the disincentives to share.

Pharmaceutical companies saw some value in contributing data if it then allowed them access to individual patient data from studies they did not sponsor. It was also good PR for companies that wished to present themselves as transparent. On the downside, however, drug manufacturers with recently registered drugs are concerned about reanalysis of data; they fear different methods may yield different results, which in turn cause potential headaches with regulators.

The demand side

As subsequent events have shown, the threat of artemisinin resistance arising and spreading was very real and the initiators of WWARN were quite right in identifying the potential utility of a comprehensive database which could be used for global surveillance. But potential utility is not the same as demand. In this case, the target users of the database -- programme managers and other policy makers -- were not actually demanding it, in part because they tend to be more interested in local specificities than in global overviews.

Multilateral programme managers and other policymakers at the global level were certainly demanding the "big picture" on resistance that the platform would have provided. But they rarely underwrote that demand with the funding for data collection that might have encouraged researchers to contribute data to a platform that met the funders' broader information needs.

The one actor who most easily straddled the supply and the demand side of the equation was the WHO. To the extent allowed by its always stretched resources, the global body supports the clinical surveillance of antimalarial resistance in fulfilment of its mandate from member states. Because it provides guidance to countries on malaria treatment, it is also a potentially important user of "big picture" data that include markers of resistance preceding therapeutic failure. Recognising the obvious potential synergy between WHO's role as supplier and user of data, BMGF supported its antimalarial resistance monitoring efforts to the tune of US$7.5 million. The grant to WHO, planned in conjunction with the WWARN grant, was intended to support the "supply side" through therapeutic efficacy studies in a number of countries. The funder expected data from those studies to be contributed to the WWARN database to allow for analysis that would be useful to WHO and its member states on the demand side. Institutionally, WHO favours data sharing. However it has always been sensitive to the member states' concerns about ownership of data. Although WHO is generally considered to be influential at the country level, it says it was unable to persuade malaria control programmes to contribute data from WHO-supported surveillance efforts to the WWARN platform.

The pros and cons of opportunism

The need for resistance data to be collated and consistently interpreted was identified by a group of malaria researchers who regularly came together for scientific conferences. University of Washington-based Carol Sibley had a particular interest in genomics, Ric Price of Oxford University and Menzies School of Health Research in Darwin and UCSF's Grant Dorsey were focused on clinical data, Jacques Le Bras at the Paris Descartes University specialised in in vitro analysis, Karen Barnes of the University of Cape Town kept a lens on pharmacology, while Chris Plowe, based at the University of Maryland, specialised in molecular markers of drug resistance. Pascal Ringwald and Peter Olumese from WHO's Global Malaria Programme brought in the perspective of a major global health

organisation, while Piero Olliaro and Olumide Ogundahunsi, from the Special Programme for Research and Training in Tropical Diseases (co-sponsored by several UN family agencies and hosted by WHO) provided views from an organisation that focuses on building capacity in endemic countries. First over coffee and beers, and then slightly more formally at side meetings at conference venues that involved "anyone who might want to come along", this loose group began to define what a data platform might look like, and how it might be developed. The drive was personal, organic and opportunistic. Scientists took it on because they believed it was worth doing, not because a funder had issued a request for proposals. This led to an extraordinary level of commitment which was vital to pushing the project off the coffee table, on to the drawing board and eventually into reality. But it did mean that the concept was largely shaped in the formative stages by like-minded academic researchers and malaria specialists at the more senior levels, with limited input from policy makers, programme managers or others who might be active users of shared data or analysis based on those data. The ad-hoc nature of the conception phase led to missed opportunities for buy-in from potentially crucial groups, especially on the policy side.

2.2 Construction phase: a victim of culture clashes

The small group that dedicated most time and energy to making WWARN a reality were tenaciously dedicated to using science to improve health. To them, the platform was a public good. For junior researchers and those in endemic countries, the idea of giving other researchers access to unpublished data "for the public good" was anathema. "The problem with the public good is that it doesn't feed my family," commented one data contributor in an interview. This was especially the case when it became clear that WWARN would not be able to offer any financial support for new data collection. Even comfortably tenured senior scientists from wealthy nations did not contribute data to the public good before they had published it, a process which can take a year or more.

Meanwhile, tensions arose between WWARN and WHO's Global Malaria Programme. Though key WHO staff were involved in the conception and planning of WWARN, they nonetheless felt its surveillance goal duplicated the global health body's mandate. Their reticence in backing the initiative fully rubbed off on their partners in national malaria programmes. Presented by WHO with exceedingly legalistic terms (see page 23), governments refused to contribute surveillance data to the platform. Together, these constraints made it singularly unlikely that a core goal of the platform -- to provide resistance information in real time to anyone who could use it -- could be achieved.

Founders and funders

WWARN's founders [3] recognised that the public good alone would not overcome the disincentives to data sharing. As one put it: "My own principle is whenever someone gives you something you have to give back, and more… The psychology is important too; you're not contacting people because you demand their data or require their data, it's a network with give and take and there are things that you have to give, too."

[3] We use "WWARN founder" to refer to the original science group heads as well as its first directors.

In the conception phase, three major approaches to "giving back" were planned:

- cleaning and curating contributed data and generating automatic reports which could form the basis of a publication
- providing researchers with protocols and other tools to help them improve their work in the field or lab
- helping endemic country scientists upgrade their own skills or that of their junior colleagues through training.

All of these were to be supported by WWARN staff based in endemic countries, working with local "centres of excellence" to provide training opportunities for researchers in the region.

Worthwhile in their own right, these approaches were also designed to increase the quality of data in the repository. The BMGF was not, however, persuaded that skills development was essential to useful data collation. As one interviewee from a global health organisation put it: "One thing you learn right away is that the Gates Foundation doesn't do capacity building. It may be what people need, they may crave it, it may also be essential to get what you need in terms of results, but they won't support it unless you can dress it as something different." When grant managers asked WWARN to scale back its support for capacity building in malaria endemic regions, many potential collaborators were further disappointed. "We were not exactly flooded with data," commented a science group head.

Institutional incentives

WWARN was conceived as a database designed and managed by scientists in various institutions around the world for the public good. However as they approached the more practical construction phase, WWARN's founders were obliged to make decisions about where the database would sit, and who would manage it. Three options were considered -- it could be set up within WHO, as a stand-alone non-profit organisation or embedded within an academic institution. The first option, which would have tied the organisation to WHO's recruitment practices and salary scales, would have been unacceptable to the funder, according to BMGF staff. The second would have required significant investment in building unfamiliar management and governance systems which would have been time-consuming and which were of little interest to the drivers of the project. That left academia: among academic institutions, Oxford was chosen because it was accessible geographically and had strong links to malaria research, especially in southeast Asia where resistance to artemisinin first emerged. In addition, Oxford was the home base of MalariaGen, a consortium of malaria researchers in endemic countries. MalariaGen maintained a genomics database which, it was hoped, would provide an informatics skeleton on which WWARN could be built.

The choice of Oxford had many advantages, but a "start-up" mentality was not one of them. The cultures of flexibility, rapid response and information sharing are not well established: "The walls are high in Oxford, and they extend outwards," observed one interviewee from a global health organisation. "They put up lots of barriers and things people would have to agree to if they wanted to share data, and it completely wrecked the spirit of data sharing."

In the early part of the construction phase, the news that the physical database and WWARN management would be housed at Oxford very quickly became conflated with

ownership and use of the data. "We certainly thought of it as just one big Oxford data-grab," said a malaria researcher working in an endemic country.

Major progress on the basics

Though neither researchers nor the WHO's surveillance partners contributed data that would have allowed for real-time tracking of resistance as originally planned, WWARN did have a significant body of post-publication data contributed by network founders and their associates, the "early adopters" in the data sharing process. These included some 25,000 individual patient records from 25 countries -- certainly enough with which to build the systems and processes needed for a workable database. There was some disagreement between the informatics team and the founding researchers about the shape of the database. But extraordinary levels of time and energy invested by the scientific group heads and their technical advisory groups had resolved important questions about what should be included in the database; they also extended standards developed by WHO for therapeutic efficacy studies to ensure compatibility while allowing for the inclusion of different data types. With these standards as a foundation, the WWARN team built a functional, semi-automated database curation tool which successfully took in heterogeneous individual patient files, standardising the format so that data could be pooled over time and location. Visualisation functions were included, so that users could see simple summary data displayed on colour-coded maps.

In parallel with the IT development, WWARN had worked hard to develop data submission, access and other governance structures which were workable in the prevailing climate of hostility to data sharing. It also invested significantly in developing newsletters, a website and other channels which would allow the group to communicate progress to potential users.

With this phase of construction complete, a few things had become clear:

- Curating heterogeneous data across four major scientific areas was both feasible and potentially useful
- Given the reality of incentive structures and funder preferences, the database was unlikely ever to attract pre-publication data that could be used for real-time monitoring of antimalarial resistance
- The inability to establish a good working relationship with key individuals within WHO had created a barrier to direct interaction with national policy-makers that was unlikely to be overcome.

In summary, by 2010 it was apparent that WWARN was unlikely to fulfil the role originally conceived for it by its founders. However the venture had a core of dedicated collaborators, well-functioning informatics and communications infrastructures, and three years of funding in hand. Very sensibly, it reinvented itself to build on its strengths.

2.3 The Proof of Concept phase: Pivot to success

Once it was clear that WWARN would not develop as originally planned, its founders, management, board and scientific advisory committee began to re-evaluate. Their principal challenge was to attract more data to the platform by providing incentives that better met the

needs of potential contributors. With real-time data for the surveillance of antimalarial resistance no longer an option, the product, too, would have to be redefined.

WWARN was conceived, driven and managed by academics, and academics from Northern and endemic countries were also the principal source of data for the platform. In retrospect, it thus seems evident that it would function best not as a public health surveillance tool, but as an academic collaboration focused on methodological development, hypothesis generation and the mining of pooled data for new insights about therapeutic efficacy and emerging resistance. Recasting WWARN as an academic collaboration made it possible for contributing researchers to add their names to publications, still the measure of academic productivity. It allowed for the full benefits of a well-curated, standardised database of information collected from tens of thousands of individual patients and parasites to be realised. Far more than just monitoring and mapping markers of resistance, the WWARN database allowed for complex analysis of drug dosing and other issues of importance in malaria treatment in small but important sub-groups such as pregnant women, the malnourished and age groups under-represented in clinical trials.

The structure that drove the pivot was the Study Group, an idea that arose from a WWARN meeting held in Senegal in June 2010, after the WHO representative suggested that WWARN should focus on research rather than surveillance. As the potential for pooled analyses became evident, scientific group heads and their colleagues began to identify new scientific questions, and proactively to seek specific data that would help answer them. The first was pharmacology group leader, the University of Cape Town's Karen Barnes, who suggested a pooled analysis of lumefantrine pharmacokinetics in patients with malaria. It was clear the data in the existing WWARN database weren't adequate to answer the question, so she wrote individually to the senior authors of 31 publications on the subject, asking them to contribute specific variables to the WWARN database to answer a specific question, and assuring contributors of authorship on any resulting publication. "Suddenly, eighty percent of people wanted to be part of something they couldn't do on their own, because they saw the real value there." Eventually, 26 research groups shared their data with WWARN for pooled analysis. As Barnes reported to the WWARN SAC in 2011: "We experienced how much easier it is to get contributions when we present clearly what we will do with data."

The principal rewards for contributing to a study group are of the sort that appeal primarily to academics globally -- publications and the generation of new scientific knowledge. But the model was also instrumental in encouraging contributions of data from industry. In interviews, industry representatives said they saw the benefit of collaboration. "We knew that other teams, MSF and academic groups, were doing studies using our drugs, because we supplied the drugs. We didn't have access to those data." However, they initially worried that depositing patient data in a database would lead either to unhelpful re-analysis or to pooled analysis of data that are not comparable. These fears were put to rest by the transparency of the study group process and the ability to participate in discussions about methods and data inclusion. "Finally, we thought it was in the best interests of science to bring those all together, to pool the data and have people use it who have a good understanding of the strengths and weaknesses of the data, people who are respectful of ownership and who take a scientific approach to formulating questions."

The study groups have created a virtuous circle: once researchers have contributed data for a stated purpose, a one-time investment is made in curating and standardising the data, which are added to the platform. With permission from contributors, these data can now easily be reused for other study group analyses. The approach has increased trust. "Now there's this

proven model, we've published several [study group papers] and … people are getting real added value out of their original clinical trial… So now when you approach people they know about the model and the concept, and that has helped tremendously in getting people to share," said a WWARN staff member. As more global-level programme managers have become aware of the potential inherent in pooled analyses, demand for specific, policy-relevant analysis has also increased.

The more academic focus of the new model also reduces friction for WWARN staff. Instead of aiming to contribute to a global public good through surveillance, they now dedicate their energy to conducting research that is reported in publications in high impact journals. This aligns their output with the interests of the university that officially employs them.

2.4 Maturity and widespread use: attracting new partners, facilitating access

The study group model has led to the expansion of the database; increased trust among both contributors and potential users; increased the visibility of WWARN; and demonstrated the range of methodological and analytic issues that might be addressed using the database. It has also raised the bar on what we should expect from meta-analyses as we move into an era of Big Data. "It's a no-brainer that individual patient data meta-analyses should be the gold standard methodology," commented a WWARN founder. "As soon as you've done it, you realise you shouldn't do anything else."

As a result of these successes, and supported by broader shifts in the data sharing landscape, WWARN is now being approached by potential collaborators, including academic researchers from outside the WWARN "family". Mathematical modellers and others who have not contributed data to the platform but who are interested in using it to answer specific questions as well as to inform new, policy-relevant study design are approaching WWARN and enquiring about access to the data. WHO asked WWARN to participate in an expert review group, work which led to the initiation of a study group on cardiotoxicity of antimalarials. A number of interviewees involved with drug development said they would like WWARN to expand its database and analyses to allow for routine pooled analyses of drug safety.

Summary lesson from WWARN's evolution and pivot

For a data-based platform to succeed, the motivations of platform funders, developers, contributors and users must be considered, and the incentives driving each of them aligned.

Box 1: New partnerships lead to rapid, policy-relevant analyses

More than a decade ago, researchers from the Liverpool School of Tropical Medicine (LSTM) and the WHO-housed Special Programme for Research and Training in Tropical Diseases (TDR) grew concerned that inaccurate weight-for-age translations meant that children in some parts of the world were being given inappropriate doses of medicine. They used data from publicly-available datasets as well as data contributed by research groups working on nutrition to calculate weight for age in different populations of children.(Taylor et al. 2006) Their findings were used by Sanofi in the formulation for the fixed dose combination of artesunate - amodiaquine. (Sanofi-Aventis 2006) The same dosages have been used by all producers of WHO prequalified formulations of ASAQ since then.(World Health Organisaton 2016b)

In late 2013, WWARN published a pooled analysis showing that young children taking dihydroartemisinin-piperaquine seemed to be getting a lower intake dose than adults, and were at higher risk for recrudescence.(WWARN DP Study Group 2013) The study recommended considering higher doses for children; the likely need for this was underscored by the findings of a WWARN-led pharmacokinetic study.(Tarning et al. 2012)

The studies had been unable to investigate whether a higher dose would be safe and well tolerated. With the support of the European & Developing Countries Clinical Trials Partnership, the Liverpool group quickly switched the focus of a planned dosing study, looking instead at the efficacy and risk of cardiotoxicity of a higher dose of DHA-piperaquine in children. Though their detailed results are not yet published, they were able to make them available for review by WHO expert panels. The sum of all of this research led the WHO to change its dosing recommendations for children taking the drug. (World Health Organisaton 2015, 39) The result should be fewer treatment failures and more healthy children.

"That is a clear example where WWARN guided with their pooled analysis what the next steps should be for relevant research," said Malawi-based LSTM researcher Anja Terlouw, who led the cardiotoxicity research. "That helped a team like mine to do a very targeted study that could be used to give exactly the information needed to feed into recommendations. It worked for us, it worked for WWARN, and it worked for the policy-makers who needed the information. It was a great win-win situation."

3 The benefit of hindsight -- lessons in specific areas

Detailed discussions, contestations, experiments and decisions in a number of key areas shaped the broad outlines of WWARN's evolution into a successful and useful platform. This section looks in greater detail at a few key areas: informatics, tools development, governance, and policy relevance. WWARN's experience in these areas may help inform future choices around data sharing initiatives.

3.1 Informatics

Oxford was chosen as a site for WWARN in part because it was the home of MalariaGen, a database that was the product of a collaboration between scientists working on the genetics and genomics of malaria. It was hoped that MalariaGen could provide, if not a blueprint, at least a master plan for the development of the WWARN repository. The genomics project's founder, Dominic Kwiatkowski, became the scientific head of the informatics module of WWARN.[4]

What should go in to the database?

As Kwiatkowski said during a WWARN meeting in 2010, "curation is the process of turning a data repository into information." Databases become useable when their contents are clearly described and easily accessed. They become more useful when the data they hold can be easily analysed, for example because variables have been standardised so that they are comparable between different studies and across time and place.

Purpose

For this curation to take place, important decisions have to be made about which variables will be included, and how they will be treated. Those decisions, in turn, can most easily be made if there is a clear vision of who will use the data, and for what purpose. An overly prescriptive initial plan may have a downside, however, in that it forecloses potential uses that are not foreseen at the outset. "There was a constant struggle between: let's put it all in there and then figure out later what we can do with it, against starting with questions and output in mind," recalled one interviewee. "The idea that you build a database to achieve certain goals was something that some of us didn't really get at the start, because we'd never done it before."

While flexibility is necessary to meet emerging user needs, a failure to define end users or the desired output of a shared data resource clearly from the start leads to a less streamlined construction process. "WWARN built a tool, and then we built another tool, and then we changed it, and then we modified it, and we went back and forth many times. If you did it again for another disease, you'd want to more clearly define your end product so you could start working towards that at the very beginning," said one former WWARN employee involved with data management. Defining outputs with great clarity was especially difficult for WWARN, because there were so few existing platforms to learn from and the universe of possibilities was thus poorly defined. While the informatics team said they recognised the need for flexibility, their frustration with what they considered an amorphous brief from the malaria scientists was palpable.

[4] That WWARN was structured with an informatics module in parallel to the other scientific modules (clinical, in vitro, molecular, pharmacological and, later, medicine quality) is telling: it suggests that the information sciences were seen as an activity in their own right, rather than a cross-cutting foundation upon which the utility of the venture as a whole rests. Communications, in contrast, was always conceived as a service that would underpin the goals of the enterprise as a whole.

Consultations at the design stage

A well-designed platform brings together two or more groups of users with different but complementary objectives, and facilitates interaction between them by making life easier for both sides. (Eastwood 2016; Evans and Schmalensee 2016) The greatest challenge at the design stage is to identify the needs and interests of all the user groups. At a 2010 WWARN meeting, participants worried that consultations with potential users had not been widespread enough. Ric Price, group head for the clinical module, said: "What we decide is important may be completely irrelevant to our target stakeholders. We do need to determine stakeholder needs, to ensure that we don't just build a resource for ourselves." The same meeting noted that WWARN's approach had so far been one-sided, focusing only on the contribution side. Policy users -- the other major group of stakeholders in the original WWARN design, were expected to be satisfied with maps showing the results of the studies included in the database.

Pooled analyses were indeed envisaged at the design stage, and a great deal of time and effort was put in to defining core variables that would allow for these analyses. However the original terms of reference implied that analysis would be conducted by WWARN staff, and little thought was given to the needs of other users. Studies shared with WWARN were visible through an interactive mapping tool but looking for metadata was cumbersome.[5] A freely-available metadata dictionary which will allow other potential users to find out what variables are in the WWARN dataset is only now being prepared; the process would have been far less resource-intensive if the needs of a wider group of potential platform users had been taken into consideration at the start.

Content and format

Clearly, the more heterogeneous data are in either content or format, the more difficult they are to curate. Coping with heterogeneity was one of the biggest challenges facing WWARN, and doing it successfully has been one of the group's greatest achievements. But it was the cause of considerable friction in the construction phase.

The informatics team pressed the scientists to define a small set of core variables that could form the basis of the dataset and, further, to put the onus on contributors to submit them in a standard format. The WWARN scientists strongly resisted this pressure, because they thought it would deter people from sharing their data. "It was a decision taken for pragmatic reasons to reduce the barrier to sharing to an absolute minimum," said a WWARN employee. "Dominic [Kwiatkowski -- head of the informatics module] didn't think that was a sensible way to go because you were setting yourself up for absolute mayhem trying to deal with it all, and I understand his point of view."

The friction reveals two important differences between the MalariaGen model and WWARN. Firstly, genomics data are inherently much more homogenous than clinical data. One former WWARN employee recalled: "I went to the informatics group, thinking that we should automate the data intake. Then I gave them all those clinical trials things, those definitions of intention to treat, per protocol, survival analysis, Kaplan Myers: the poor guys

[5] It is possible to compile a list of included studies from data on the WWARN Explorer web page, either by hovering over 289 (?) individual pins on a map and noting the data that appears in a pop-up box for each, or by copying data that appears in 58 separate drop-down menus. The variables thus discoverable for clinical studies are: study location, date, number of patients, drug, lowest efficacy rate on day 28 and 95% confidence interval.

were lost, and you can't blame them." Secondly, MalariaGen was offering free genetic sequencing to any data contributor -- a very large carrot for which researchers are prepared to invest a bit of time and effort in data management. WWARN had no such carrot to offer.

Using the WHO therapeutic efficacy study variables as a starting point, WWARN went on to define core variables for all its data areas. Some interviewees expressed concern that an overly standardised approach reduces the richness and potential utility of the data. However this concern must be offset against the extraordinary time and effort that goes in to curating each new variable. Interviewees said that to make it worth standardising data in advance, they had to be able to foresee that it would be used in several different analyses. Since adopting the study group model, this tension has resolved itself somewhat, as one WWARN employee explained. Whole data sets are contributed and retained, but variables now only get thoroughly cleaned, standardised and taken in to the database when they are needed for a specific analysis. "That way, you don't spend time putting the data into the database unless you are actually going to use it for something."

Some contributors provide WWARN with datasets that include many variables that are not immediately curated and included in the repository. There is currently no system for logging these unused variables so that they might easily be discovered and identified for curation should the interest arise. As the database expands its user base, this sort of logging may well be an investment worth making.

Prospective versus retrospective data

Many interviewees talked dreamily of reducing the curation burden by collecting data using common protocol templates. With prospective data collection, this is becoming more likely. With retrospective data, it's not even an option. And yet given the prevailing publish-or-perish culture, most data that are available, at least at the construction phase of a new data platform, are likely to have been collected before the widespread use of any shared tools or common standards.

Data managers must decide how much time and effort to invest in trying to curate old datasets that are poorly documented or of uncertain quality. For some variables, such as toxicity data, subjectivity or lack of standards may render retrospective standardisation virtually impossible.

Data quality

Databases such as that created by WWARN aim to compile data that are of high quality. Since data managers are unable to audit data collection, they are obliged to accept most contributions at face value. "Mostly, any data we get, you assume it is right, they've recorded what was observed," said one WWARN staff member. Data managers do, however, perform rigorous checks for missing data, out of range values such as pregnant men, and other anomalies. They check obviously problematic data, consult with contributors and correct where necessary. This curation process itself contributed to increasing the quality of prospectively-collected data. Communication with WWARN staff about anomalous data increased researchers' awareness of potential pitfalls affecting data quality at the data collection and entry stages. The interaction also allowed WWARN to promote tools to support more standardised and more reliable data collection.

Intensive data checking is, however, very time-consuming. As data in a database reaches a certain critical mass and norms and distributions are established for different variables and populations, it becomes possible to use informatics tools to automate many quality checks.

Cost and value

Data curation is an expensive business, especially for heterogeneous data. However the costs are front-loaded, because the grunt work of database design, standards setting and platform construction is done early on. Tools and systems must be revised or developed over time to meet emerging needs, so curation costs never vanish entirely, but they diminish considerably over time as data management becomes increasingly automated. Contributions of data collected prospectively by well-trained researchers using tools that match the standards in the database reduce curation costs even further over time.

The benefits, on the other hand, increase over time. Once a dataset is cleaned, well curated and added to a database, it can be used repeatedly by an infinite number of analysts. Pooled analyses based on standardised individual patient data can be re-run in a matter of minutes as new data are added to the platform. A classic Cochrane-style meta-analysis of aggregate data, in contrast, requires far more work to update. "When someone looks at the cost they say oh, this must be expensive! But if you look at all of the things that it's saving at a very practical level around data management ... the economies of scale of having data curation, data management, long-term secure data storage done in an ethical way that would normally be carried by each of our individual research programmes; then [the worry about expense], well, it disappears," said one data contributor, a member of a large research group. Individual research groups benefit from a well-curated, independently managed database not just because it relieves them of the costs of curation and storage but because they will not have to manage repeated access requests.

Funders who approach data platforms as a long-term investment should see the return on that investment increase over time. Investors who expect a quick return, on the other hand, will likely be disappointed.

Summary lessons: Informatics

A well-defined purpose is the bedrock of a successful database. Too much rigidity in the design may, however, undermine the ability to adapt to emerging opportunities.

Widespread consultation with potential contributors and users at the conception stage will greatly increase the efficiency of database design and construction

The burden of data curation and the potential utility of the platform are both greatest when:

- data are heterogeneous rather than homogenous.
- barriers to data contribution are lowered by accepting a wide variety of formats.
- retrospective data are included
- variables are curated on submission (general) rather than on user demand (specific).

Data platform designers need to consider the trade-offs involved, and match their decisions to their resources.

The value of well-curated data sets grows over time. Investors can not expect a quick return on investment.

3.2 Tools

WWARN's experience with tools calls to mind a truism of the open source software movement: "Every good work of software starts by scratching a developer's personal itch".(Raymond 1999) When WWARN collaborators made tools that met some need they themselves encountered, those tools were very often adopted by others in the malaria community, going on to become industry standards. According to a WWARN founder: "It's difficult to do survival analysis. And so the idea is to empower people to do it in a standardised way. In a way, we were trying to sort our own problem out. But in doing so, that resource becomes useful to other people." WWARN's most popular tool, the Parasite Clearance Estimator, provides a good example. Designed by statistician Kasia Stepniewska to estimate the rate at which parasites are cleared from the blood following treatment, it is now the "go-to" tool for the job among malaria researchers.[6]

When WWARN collaborators tried to second guess what users wanted, however, as was the case with tools for data visualisation and automatic report generation, these tools on offer were less likely to be used. By the time researchers submitted their data to the WWARN platform, they had already published their headline results. "The report didn't help them, it just said the data were slightly different from what they had published. So that didn't go down well," recalled a WWARN science group head.

Tools for data collection and processing

Many WWARN tools were offered through the website to anyone wishing to use them, regardless of whether they contributed their data to the platform or not. These included standard protocols for clinical and in vitro studies which, if used, have an added benefit: more standardised data collection methods increase the comparability and perhaps the accuracy of the data submitted to any platform, thus reducing the burden of curation. Uptake has been mixed: standard protocols have been used most enthusiastically in newly-emerging areas of enquiry such as the use of the ring-stage survival assay.(Witkowski et al. 2013) Where practices are well entrenched, such as with clinical trials, established research groups have tended to stick with their own methods.

More niche tools supporting quality assurance among laboratories measuring in vitro susceptibility and pharmacokinetics have also proven popular. Some 60 labs in 30 countries now receive free antimalarial reference materials from WWARN or use the group's proficiency testing and other external quality assurance services. WWARN collaborators say that this has led to a measurable improvement in data quality.

Endemic country researchers interviewed were unanimous in saying that they would find the tools more useful if they came with more support. "I have one PhD student who got help using the [parasite clearance estimator] tool from WWARN. They gave the output but he could not understand what the output meant. Now he is going back to the old way, because

[6] Of 145 references including "parasite clearance estimator" listed by Google Scholar, 104 refer to the WWARN tool. Since the publication of the WWARN parasite clearance estimator in late 2011, the respective figures are 91/134. The majority of the 91 are studies which refer to the use of the WWARN tool in the methods section. Since web metrics were first collected in May 2013, the parasite clearance estimator has had 1,265 unique views, half of them in Europe. Data on country of access is available for the top 10 user countries, which together account for over 80% of use. Among that sub-section, over 70% of users are in northern countries or Australia, and a further 23% are in Thailand. The remaining two countries in the top 10 -- Kenya and Mali -- account for 5% of use between them. Data at 16 June 2016.

he at least he knows what it means." WWARN did provide more training and site-specific mentoring in its earlier years, but that tailed off around 2012, after funding for regional offices dried up. The user support circulation lists, chat fora and mentoring sessions run by some of the scientific coordinators also dried up around this time.

WWARN staff continue to welcome queries from users. "The tool is online, then we have an email address, they can contact us with their questions," said one. Those that do have been very impressed by the level of support they received from WWARN, saying that it has improved the quality of their science. "We emailed them a list of the studies we wanted to use that matched our criteria, and they told us that we'd missed a couple," said an external data user in a northern institution. "But they also gave lots of useful comments on our analysis plan… That kind of scientific input from WWARN was really good; I wasn't expecting so much of their time." However as a WWARN staff member reflected, field researchers in endemic countries may be intimidated by the relatively anonymous, web-based approach which became the default in the absence of a regional presence. "We are thinking that we can just put tools on a website and expect people to use them. We're communicating in a very Western way, not a Vietnamese or Ghanaian way."

Standards wars

Any organisation that publishes tools runs the risk of being accused of playing God. "Who's the arbiter of the correct method? Who the hell are WWARN to say this is the best? That's a question we faced over and over," said one WWARN founder.

In WWARN's case, the offer of tools made freely available to any researcher wishing to use them heightened existing tensions between the group and key individuals at WHO, who felt the organisation's normative function was being threatened. The overwhelming majority of interviewees, including those working with malaria control programmes, disagreed with the WHO's reading of the situation. National malaria control programmes continue to use the WHO's standardised tools for therapeutic efficacy studies for their clinical surveillance of antimalarial resistance; the tools are simple, functional and well-liked. However they don't meet the needs of clinical researchers who need to collect a wider range of data. WWARN filled the gap by making template protocols for more complex research trials available to this second group.

The development of tools does not inevitably lead to tensions -- they were avoided, for example, in the case of the 2015 microscopy guidelines published under the auspices of TDR and WHO but developed in collaboration with WWARN and other groups. But implementers of initiatives such as WWARN should remain sensitive to the fact that an offer of tools to the community may be perceived by key individuals as a territorial encroachment. Meanwhile funders would do well to accept that "impact" can sometimes best be achieved by contributing to collaborative ventures which will not carry their brand.

In constructing the repository, WWARN built up considerable expertise in determining which aspects of clinical trials needed to be recorded in standardised ways. This expertise has been recognised by the Clinical Data Interchange Standards Consortium (CDISC), an open standards organization which functions as a sort of Wiki for metadata and is the FDA-mandated data standard for all regulatory submissions as of 2017. In 2015, CDISC asked for WWARN's help in developing standards for clinical trials in malaria.

Summary lessons: Tools

When tools fulfil a specific, expressed need, they get used.

Tools are more likely to be used if they come with ongoing support that users feel comfortable accessing. Concerns raised during the support process should be fed back into improvement in tools.

Tools that provide methods for new research areas are more likely to be used than those that cover areas where practices are already entrenched.

New tools and standards can put noses out of joint. Before developing them, consider whether the uptake of the tools is likely to justify possible disruption of relationships.

3.3 Governance

In the context of a data sharing platform which aims to support better health in low income settings, governance can be divided broadly into two major areas. The first is the governance of the data platform itself, including terms of access and use. The second broad area covers the institutional structures through which decisions around data-platform governance are made.

WWARN has seen both evolve to an extent, providing lessons in both areas.

Terms of submissions, terms of use

One of the major challenges of developing information-based platforms (or indeed any platform) is to write the rules of the game in a way that provides all players with the clarity they need, without tying them up in so much red tape that they decide not to play, or so that it becomes difficult to adapt to changing circumstances. WWARN rose to the challenge impressively, although it did take quite a long time to reach agreement on the terms under which contributors would allow their data to be stored in the WWARN database. This was in significant part because of the legal implications of storing individual patient data. The legal department of Oxford University, concerned with protecting the institution from liability, produced an impenetrable seven-page document that complied with UK regulations on data protection but that would have been daunting to even the most enthusiastic prospective contributor of data. WWARN management fought hard, and ultimately successfully, to simplify the agreement and make it more user-friendly. Feedback on an early draft was sought from 25 potential contributors before a final three-page document was published in March 2011, laying out in plain English which data may be submitted, how it may be used, and how it will be acknowledged in any publication.

Fenced, gated or open?

The terms of submission were finalised only after the "pivot" to the study group model, and were thus aimed principally at academic researchers. With open access to the primary data off the table, founders discussed two other possibilities: a "gated" model, in which the database could be accessed by people outside the collaboration through a formal access

process but without reference to the original contributors, and a "fenced" model in which data could only be used with the express permission of the contributors, which must be accorded for each new analysis.

Given the data sharing climate of the time, the misgivings introduced by the early rhetoric about open access, and the fact that WWARN had so little to offer contributors in the early days in exchange for their data, the network opted for the more restrictive model. The decision created an immense administrative burden as WWARN staff were obliged to track down investigators for each new analysis of data that were sometimes over a decade old. Some scientists changed institutions, retired or died after first submitting data to WWARN, complicating the task. "I never had a refusal but on a study with 100 authors it could take a long time, sometimes two or three months," recalled one former WWARN employee. This process obviously greatly slows the pace at which new knowledge can be generated, and the speed at which a data platform can demonstrate its value and attract a critical mass of users.

On the upside, the process of recontacting all contributors is one way of ensuring that they have a chance to contribute their expertise, including their intimate knowledge of the specificities of the data they collected, to ongoing analyses. Interviewees from private companies and global health organisations focused on drug discovery also favoured the approach. "The WWARN model is quite sophisticated, it is very respectful of ownership, so the scientists who collected the data always have a voice in how it is used. It may seem a bit burdensome, but it works," noted one.

As the database expands and existing contributors see the benefits of participating in study groups, the approval process is speeding up. One external user of WWARN data said that the process was actually time-saving, because it meant that her university could negotiate its own cumbersome data transfer agreements with a single organisation, WWARN, knowing that due diligence would have been done with multiple contributors.

Because WWARN was a leader in collating patient-level data from low-income settings, the clumsy "fenced" model was necessary in order to build contributor trust. However WWARN's board has now approved a shift to a "gated" model, using an independent data access committee managed by TDR, which has been working actively towards increasing the sharing of data related to tropical diseases. Those who have already contributed data can decide to delegate decisions on data access to the data access committee and be notified of data requests with an option to opt out within a set time period, or choose instead the original agreement, that they be contacted to give individual approval for each new request for use. Similar choices are being written into terms of submission for new contributors.

Ideally, WWARN's proof of concept in malaria would act as a model for other disease communities, allowing them to skip the resource-intensive "fenced" stage and opt directly for a model in which data can be more easily reused. However as an interviewee working at a global health organisation to develop other disease platforms noted: "Humans are not very good at learning from other people's experience".

Ownership, authorship, and fair reward

The Terms of Submission require contributors to confirm they have the authority to submit the data. However in endemic countries, researchers often work in close collaboration with national malaria control programmes. "For the national programmes, the question of data ownership is very real. More than once I heard people say: why are we giving African data to England?" said one former WWARN employee. This has not proven a significant barrier

in countries where relationships between national programmes and research institutes are strong, but it does underline the need to take researchers' wider partnerships into consideration when drawing up terms of submission and communicating them to interested groups.

Pharmaceutical firms removed all possibility of misunderstandings by demanding legally binding material transfer agreements from the start. "The paradox was that ultimately it was much easier working with Pharma than with [academic] researchers, because with Pharma everything was agreed beforehand, and we just stuck to the agreement; it was so simple," said a former WWARN employee.

In contrast, though the language in WWARN's terms of submission for academic contributors was admirably clear, the commitment it described was sometimes felt to be fuzzy. Contributors were promised that "your contribution to any WWARN publication will be acknowledged in accordance with the guidelines of the International Committee of Medical Journal Editors." In early pooled analyses, all data contributors were listed as authors. Later, in closer accordance with the ICMJE guidelines, both data contributors and WWARN staff were accorded acknowledgement but not authorship, a cause of some dissatisfaction on behalf of data contributors and of some WWARN staff. The six most recent pooled analyses have been published in the name of a WWARN Study Group, with all individuals involved acknowledged in an A-Z listing at the end of the paper. This gives equal weight to the contributions of those whose idea underlies the study, those who performed the analysis and wrote the paper, and those who contributed data on which the analysis was based.

This system was adopted in part to counter the unhelpful impression, reinforced by some of the earlier group analyses, that WWARN unfairly favoured Oxford-based scientists. However it has created an inequity of its own. It is true that, before the A-Z system was introduced, eight out of 11 papers based on data in the WWARN repository listed Oxford-affiliated researchers as both first and last authors and a further two had a first author from Oxford. However it is also true that those researchers, many of whom were WWARN staff, spent many months developing ideas, compiling data, contributing to standardisation, performing analysis and engaging with large numbers of data contributors. In the A-Z system, their contribution is accorded the same weight as that of a researcher who contributed six data points. "A lot of [WWARN] staff missed out because they did all the work and then everyone was an author," noted a data contributor from a major research group.

The importance of language

WWARN worked hard to develop language and terms of submission that are straight-forward and not intimidating. They stand in stark contrast to the terms proposed by WHO to national governments. When BMGF gave WHO US$ 7.5 million to support strengthened therapeutic efficacy studies, WHO agreed to encourage countries to contribute the resulting data to the WWARN database. The template letter sent out by WHO contained language not unlike that originally proposed by Oxford's legal department.[7] Faced with such forbidding

[7] A sample sentence: "It is understood that WHO will accept no liability or responsibility for any loss or damage arising out of any failure on the part of Oxford to comply with the restrictions on access to data set forth in Annex 2 hereto in respect of data (including data generated from blood samples) and/or the provisions contained in Annex 2 relating to the use of blood samples, which the Government has, pursuant

language, most health ministries refused to share their data. The organisation has been more successful in promoting data sharing in other disease areas such as HIV, however, and has recently called for greater sharing of data in public health emergencies.(World Health Organization 2015) Perhaps, then, other disease platforms will find it easier to include national governments.

Communications

Very early on WWARN hired communications specialists who worked to support the platform as it reached out to potential users and other supporters. Some of the scientists interviewed said they had not understood the need for this at the outset. "But it turned out to be quite forward thinking," said one. Because the platform was driven at first by scientists who took a "supply side" approach to developing what seemed to them an obvious public good, it was important to work to generate demand. "You can have the best Explorer on the planet, but if no-one knows about it, it's useless," explained one WWARN staff member, referring to the data visualisation tool that was intended to draw in policy users.

The communications team was central to important steps -- such as the simplification of the terms of submission -- that made data sharing more feasible for potential users. A well planned, solid approach to communications may at first seem less pivotal to the development of a data sharing platform than informatics or the grittier aspects of data governance. However the WWARN experience suggests it is critical in facilitating the discussions with potential data contributors and users that can help to align incentives between different platform users. "The IT experts want to build a great data platform, but they are less concerned with what happens next," said a WWARN employee. "The scientists want the data, the policy-makers need the research evidence, but neither spend as much time worrying about how they will receive it. Communications is the glue that binds the whole enterprise together."

Summary lessons in database governance and communications

Clear terms of engagement agreed up front increase efficiency and reduce mistrust

Overly legalistic or disproportionately cautious language discourages participation

Once set, rules of engagement can be hard to change. Terms which default to more sharing while allowing contributors to exercise more control if desired may help balance trust with the flexibility to evolve with the data sharing zeitgeist.

It's important to invest in systems that support consultation and allow platform developers to communicate effectively with potential user groups

More active discussions within the scientific community about the benefits of sharing data through well-governed repositories might reduce fears still extant among researchers.

---

to the terms of this Agreement, agreed to share with Oxford for use within the WWARN Project, including for the purpose of inclusion in the WWARN Project database."

Institutional arrangements

A platform's terms of submission and use are materially influenced by the platform's institutional governance, both legally and because of user perceptions and preferences -- as seen in the drop in use of US based data services after Edward Snowden revealed the extent of the US government's electronic "snooping".

WWARN's founders were dispersed among many institutions, including universities in Seattle, Baltimore, Darwin, Paris and Cape Town. In its early years it also supported regional centres in Kenya, Senegal, Thailand and (briefly) São Paulo. Despite the geographical and institutional breadth, however, embedding WWARN at Oxford University certainly created perceptions that a small group of well-known and well-resourced northern academics were trying to access data from around the world. This has discouraged some researchers from contributing data to the platform.

The decision to embed WWARN in an academic institution also removed the possibility for truly independent governance structures. WWARN's Board did not hold the executive accountable for key management functions such as fund-raising, in part because ultimate responsibility for hiring, firing and other key management decisions lay with the university. Many interviewees, especially those in other global health organisations, said the structure has resulted in a lack of transparency and accountability. While virtually no-one ascribed this to any ill intention on the behalf of WWARN or its board members, they did say that the muddled governance undermined the perceived independence of the platform, reducing its acceptability to the malaria community beyond academia.

The Oxford base has attracted talented staff who are interested in an academic career, sometimes appointing them to posts including data management in which it is hard to achieve the publications that are the measure of productivity in that environment. Meanwhile, staff outside the academic stream say that university management is ill-equipped to value their work correctly. Those who do work at producing publications are, under the new A-Z study group model, also not rewarded for their work in the classic way. And because WWARN is a stand-alone project within an academic department, there is no job security for staff. As one remarked: "If you're working on six-month [contract] extensions, you can't get a mortgage, your wife can't join you… "

Partnerships

Institutional arrangements can foster or foreclose particular types of partnerships, a point underlined by an African policy maker who asked: "What is the mechanism for WWARN, which is seated in a private university in the UK, to engage with the government?" National governments find it easier to interact with international organisations such as WHO or bilateral development organisations than they do with non-government groups, private companies or academia. In the field of global health, the WHO remains the single most important partner for those wishing to influence policy and practice in low-income countries. It can function as a bridge between national governments and other groups of actors; equally, it can act as an effective drawbridge, closing off the route to effective communication and collaboration.

Some of WWARN's founders underestimated the importance of WHO as a source of technical expertise in the eyes of national malaria control programmes. Others were aware of WHO's importance politically but were unconvinced that the global body could deliver high levels of technical support with the minimal resources at its disposal. These factors

combined with poor personal relations to pull up the drawbridge. But the core problem of a mismatch between the WHO's broad mandate and its slim resources has produced information gaps in any number of health areas. Other data sharing platforms may also be perceived as an existential threat by WHO staff mandated to produce disease information, but not properly resourced to do it well.

Successful collaboration with "make or break" partners such as WHO greatly extends the potential utility of a disease data platform by facilitating the participation of national programmes as both contributors and users. By thinking of the organisation as a key user of a platform at the conception stage, and by consulting widely at country office and regional levels as well as with Geneva, designers can aim to build in outputs or facilitate interactions that meet the WHO's needs.

Summary lessons: institutional governance

Perceptions of the interests of the institution hosting a platform are as important as reality in shaping user preferences

Aligning staff expectations with the functional requirements of a job and ensuring that performance can be rewarded adequately may increase staff satisfaction, reduce staff turnover and improve platform functioning.

If a particular institutional set-up is likely to exclude "make or break" partners, choose a different institutional set-up.

3.4 Policy relevance

There are two major drivers behind the data sharing movement in the health sciences. The first, particularly relevant to industry-funded clinical trials, is transparency. The second, of greater importance to infectious diseases of poverty for which much research is funded by taxpayers or foundations, is the promise of better health. The process of translation of shared data into longer, healthier lives depends to a large extent on the questions that are asked of the data. The more actionable they are, the more likely it is that, once answered, they will be acted upon. But policy action is about more than just technical feasibility. It's most easily achieved when research meets specific, articulated needs. These can't always be anticipated at the outset, but they can be built in to the mechanisms which promote the use of the data platform.

Demand-driven research

Some of the WWARN study groups have had immediate implications for policy, leading, for example, to changes in WHO treatment guidelines. But to date, all but two of the study groups have been led by WWARN staff or collaborators answering questions of their own devising. As the platform expands and attracts new users, other groups, including WHO, are beginning to suggest research questions. However, other than researchers from the University of Cape Town (where the WWARN pharmacology group is based), there have not yet been any proposals from collaborators from endemic countries who are perhaps best

placed to meet the needs of policy makers in those countries. According to one African SAC member: "Countries simply don't have the capacity, the time or the money to trawl through, get data off the net and do analysis, even on their own regional data."

He appreciated WWARN's efforts to involve collaborators in study groups, but found them tokenistic. "In the end it's a researcher at Oxford who does the analysis and asks us to participate, including in writing the paper. But truthfully that's really hard to do if you don't really understand how they did the analysis." He suggested a more genuine collaboration. "We need to involve the contributors much more in the analysis plan, to brief them properly on the methods used, on the potential interpretations. We don't need a formal training, we're researchers, so with a good briefing we understand. But it's by no means evident that we have even that kind of entente."

Others think that's unfair. Those leading and coordinating study groups acknowledge that they generate the research idea, but they always circulate a detailed analysis plan when requesting permission to use data. They have often tried hard to involve collaborators both at that stage and throughout the analysis and interpretation phases, but report being regularly frustrated by the limited feedback they receive. "Some of the guys we've worked with, the minute they go home with a PhD they get made head of a research institution and that's effectively the end of their research career, they don't have time for anything any more," said a senior analyst at WWARN.

Researchers in well-established regional hubs in countries such as Kenya and Senegal have indicated that they would like to take the lead on using the WWARN repository more. But in many other endemic countries, researchers will need more support if they are to invest significant time in leading study groups.

At least one regional network, the Asia Pacific Malaria Elimination Network (APMEN), has successfully promoted policy-relevant research by convening meetings of researchers, programme managers and advisors from international organisations and global malaria programmes. Small grants are provided to help researchers answer the policy makers' questions. "I got 50K [Australian dollars] to do a study that came up in discussions with policy makers," said a researcher in an endemic country in Southeast Asia, who has not contributed data to the WWARN platform. "Even though it is not big grant, it is quite targeted, helping us to kick-start new areas of research."

While WWARN has never had funds to support externally-generated study groups financially, the network welcomes research ideas from any source. Many researchers in endemic countries, including WWARN collaborators, remain unaware of this, and the website provides no clear information about the process for proposing or setting up a study group. "My own feeling is that the communication has not been clear, how to be part of a study group other than just contributing data," commented one endemic country researcher. "What is missing is information, and when information is missing… people become suspicious…. Then because of a simple lack of information, a good thing gets killed."

Data contributors as analysts

"Big data" analyses have an undeniable power to identify norms, patterns and trends. But as malaria control efforts gain ground and the emphasis shifts from control to elimination, local variations and specificities which are typically dismissed as "noise" in large pooled analyses become more important: in other words the noise may become the signal of interest. In these circumstances, the local knowledge which resides with the people who

collected the data becomes even more valuable. "We're talking about people who are on the front line, who are seeing patients on a daily basis, if there's anything that is emerging, these are the first people who will see it," said one endemic country researcher, who has chosen not to contribute data to WWARN because he believes the network undervalues local knowledge.

He and others believe that front-line investigators are best placed to pose research questions about important emerging issues, and that they are also key to the analysis process. "Analysts in the north don't understand what is happening in the villages; all the biases are not even imagined," said a northern academic working in an endemic country. "So data are being crunched a bit blindly, giving results that aren't really reliable in the end. We need more of a link between those two worlds."

In-country researchers are also the most effective communicators of important findings to the national programmes who must ultimately use them. National programmes rarely change their practices based solely on papers in international journals. "Apart from anything else, that's all in English," said an Asian researcher, another non-contributor who works closely with local health staff. "In this region, NMCP staff can't read it." Local studies usually take precedence, and local experts are almost always consulted. Failure to involve them in analyses leads to lost opportunities. "A policy-maker finds something on the net; the conclusions are [perhaps] based on good analysis," said an African SAC member. " but no-one from the region can explain them, because no-one participated in the analysis. That's a problem for us."

Summary lessons: Policy relevance

More interaction with national authorities and participation of endemic country researchers in setting research questions may contribute to more locally relevant and actionable research results. Institutional arrangements should be structured to support this interaction, including through strong local links.

A granular understanding of local data will become increasingly important as disease prevalence falls. Involvement of endemic country researchers in analysis and interpretation should thus also grow.

Endemic country researchers currently rarely use shared data resources. Platform supporters should consider incentivising secondary analysis by endemic country researchers, perhaps through active mentorship schemes coupled with funding for pooled analyses by endemic country institutions that answers questions raised by regional or local malaria programme managers.

4 Data sharing models: WWARN and beyond

Besides having brought together data from an astounding array of individual patients, parasites, labs and clinics around the world, curated them in ways that make them not just useable but useful, and produced world-class new methods, tools and pooled analyses, WWARN has built up a great wealth of experience in the science of data sharing itself. This includes a range of informatics tools and data standards, as well as terms of submission and

other legal and governance standards which continue to evolve to meet changing needs. Most interviewees pointed to different ways in which WWARN's tools, experience and communications approaches could be applied to other diseases. "Just the legal framework alone is an essential piece that we can't underestimate," said one interviewee in a global health organisation now working on a data platform. "The right to use the data, how the data are curated, how they can be published, all these aspects are very sensitive questions, and I trust that WWARN has already gone through a lot of problem solving. We can also benefit from that; that's not trivial it is huge." Groups working on several neglected tropical diseases have recognized this potential, and have asked WWARN to help adapt these basic approaches to data gathered by their collaborators. This work, now in its early stages, is the foundation for the development of the Infectious Diseases Data Observatory (IDDO), an umbrella organisation of data platforms including WWARN. The group has now begun the task of replicating WWARN's success and learning from its mistakes in facilitating the curation and analysis of individual patient data for different diseases, including Ebola and visceral leishmaniasis.

4.1 When will the WWARN model work best?

Of course in the complex world of global health, and especially in the shifting sands of data sharing, exact replication of the WWARN model won't be possible. What follows is a suggestion of the factors that should be considered when assessing the appropriateness of the model, and the adaptations that may be needed.

Similar purpose?

WWARN was set up to create an open access database that would allow all-comers to track markers of antimalarial resistance in real time. Neither the data sharing landscape nor the institutional set-up favoured that outcome. They did, however, favour the collation of individual patient data in a format that allowed for the production of high-quality academic papers based on pooled analyses.

WWARN altered its purpose to suit its structures. If a data-sharing venture hopes to fulfil a different purpose -- if it intends to provide data for real-time surveillance, hopes to facilitate drug discovery, or seeks to promote transparency in science as its principal purpose, for example, then it will probably have to alter its structures: the WWARN model as it now stands may not work. It may well be successfully adapted, but significant up-front work will be needed to map and plan necessary modifications in the governance structures that shape incentives and make the new goal achievable, especially in the current data-sharing climate.

It's also important that people implementing a data sharing effort and those supporting it financially should agree about its core purpose. If investors think they are buying a curation service while implementers aim to establish a scientific collaboration, at least one side is likely to be severely disappointed by the outcome.

Similar incentives?

Incentive structures are fundamental in determining what sort of model is most appropriate in achieving a given goal. The MalariaGen model did not translate neatly to WWARN because data collectors had a strong incentive to contribute in the former case, but not, initially, in the latter. A scan for "deal-breaker" differences in incentives early in the planning phase of a data resource may point to the need for significant modifications in terms of data submission, use, or other governance and institutional arrangements.

Incentives are importantly determined by institutional choices, but also by the current data-sharing zeitgeist. If funders and journals require data to be deposited in highly curated, disease-specific repositories, platforms such as WWARN will acquire an additional value-added in the eyes of potential contributors.

Similar diseases?

WWARN's active participation in the development of CDISC standards has given staff a great head-start in defining meta-data and other standards that are common across the clinical trials process, and across some diseases. However, no two diseases are the same.

Some differences will simply require the development of new, disease-specific standards. "There's a lot of potential for the IT platform to be shared across diseases, but you're always going to need a group of people who have the content knowledge, experience, and background across different aspects [of a particular disease], because these things are so complex," said one science group head. Many people with different skills have contributed to the success of WWARN, but it has perhaps been most heavily dependent on the personal commitment of extremely highly qualified malaria researchers. If similar individuals can be incentivised to contribute their expertise in other diseases to build a WWARN-type platform, and if key WHO staff and other members of the disease community participate in framing the chosen standards, this process should be relatively straightforward. Those are, however, two big "ifs", worth assessing before embarking on extending a WWARN-type platform to each new disease.

Some diseases will certainly require the inclusion of quite different types of information, sometimes with significant ethical or logistical implications. For example to be really useful, databases for Ebola and Chagas disease must include registries, and thus identifiers which allow data collected from specific individuals to be linked over time. This will require modification of both informatics and governance procedures. Again, such modifications may well be possible -- indeed IDDO has, in its Ebola platform, already addressed successfully the security issues required when patient identifiers must be retained along with the patient’s clinical record. However such issues require careful consideration before commitments to a particular model are made.

Figure 1: Adapting the WWARN model for other disease: decisions and action

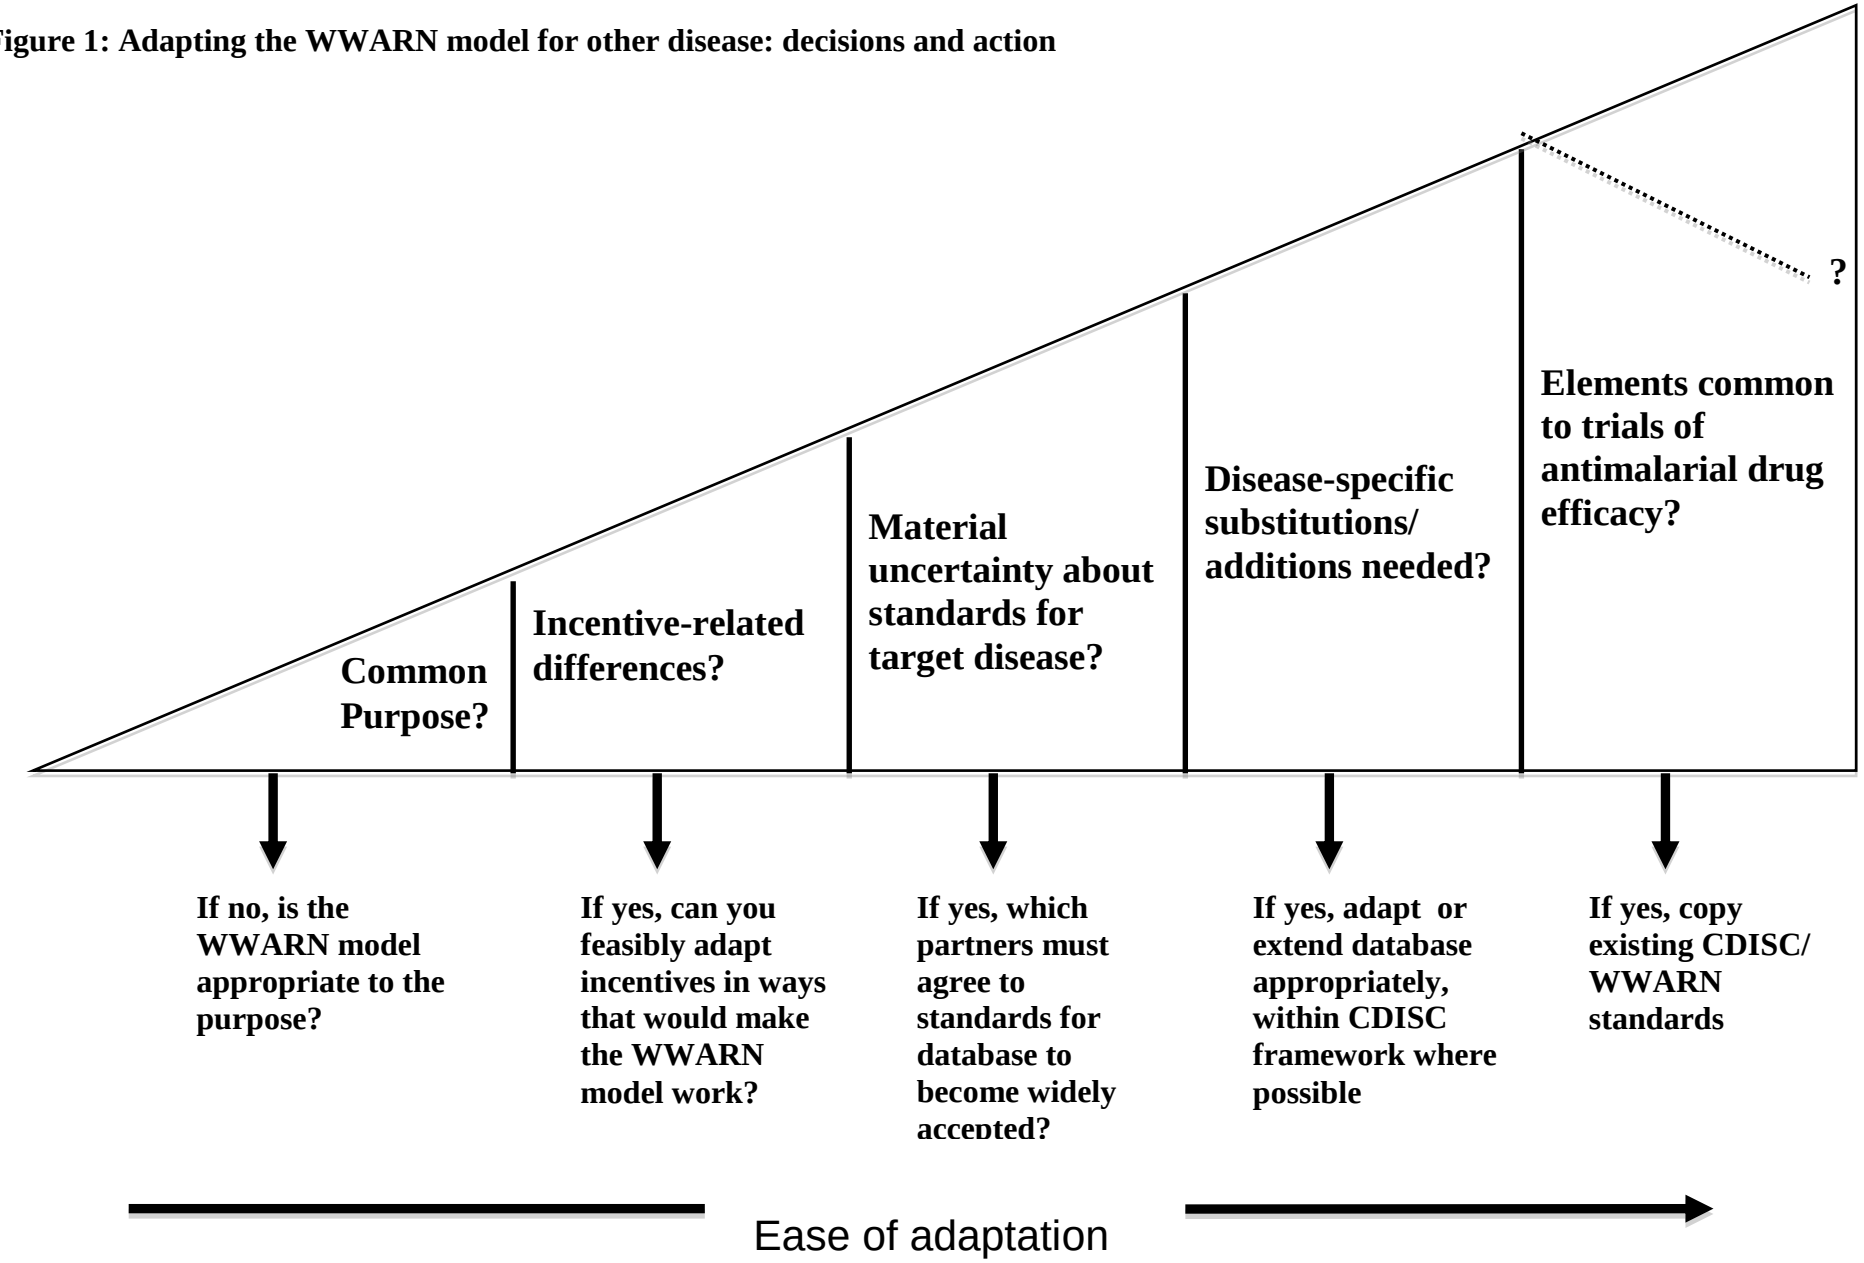

Another disease-specific difference relates to potential costs and benefits. The scientific benefit of pooled analysis is likely to be greatest where data are sparse and individual studies are small. A disease may fit this pattern because it is neglected, or because it is rare, or both. In either case, the cost of data curation is likely to fall heavily on global health funders. The economies of scale achieved by sharing informatics, governance and communications tools across diseases should bring costs down over time. But where data are very heterogeneous, few prospective (and thus more standardisable) data are collected, and prevalence is low, funders will have to consider the opportunity cost of the likely gains in actionable knowledge that pooled analyses would provide.

4.2 Adaptive models

The WWARN experience reminds us that the engagement as well as the management needs can be quite different at different stages of platform development. They are also altered by advances in technology and shifts in the consumer and user culture.

Different needs for different phases of platform development

Table 3 describes the roles of various groups at different stages of development. The groups are not mutually exclusive, of course. People who are collecting data in the field and contributing to the construction phases may also be community leaders involved in standard setting. Senior disease specialists are often also involved in policy-making, sitting on national or international advisory boards.

The WWARN case underlines the extent to which roles and community norms in data sharing have shifted through phases of the initiatives development, as well as over time. In the same way, management and institutional structures are also likely to shift, with the importance of different actors waxing a waning according to the maturity of a platform. Multi-tasking individuals are likely to see the demands on their own expertise shift, too, as described below.

Table 3: Roles of different groups as platform develops

	Data contributors	Disease experts	Data users	Informatics
Conception	Consulted about incentives, terms of submission and use	Actively conceptualising purpose, utility, user constituencies	Consulted about needs, desires, incentives for use	Data scientists focus on overlap in user & contributor interests, guide design
Construction	Early adopters provide data	Nail down contents and standards	Limited consultation	IT engineers implement platform design
Proof of concept	Provide data, analysis ideas	Conduct early analyses, meeting user demand	Wider consultation to define demand	Data managers keep platform on track
Expansion and use	Provide data, contribute analysis ideas, interpretation	Support analysis	Active participation in question-setting and interpretation	Data managers keep platform on track

The central role of scientific input

Disease experts clearly have a central role to play in the conception phase, when they are defining the purpose and utility of a database or platform. They are aided in this by data scientists, whose expertise lies in "matchmaking" the needs, desires and incentives of contributors and users, and in identifying the "sweet spot" of overlap in interests that leads to network effects and thus a successful platform.

Where the goal of a platform is to improve public health outcomes, input from the widest possible community of potential contributors to and users of a platform is critical at the outset. If the needs and desires of either group are not met, the platform cannot succeed. Scientific specialists in both the disease area and informatics are driving the process at this stage, so they must also feel incentivised to invest in the process. Hosting the venture in an academic institution at the early stage may facilitate this by allowing them to contribute as part of their "day job".

Disease experts continue to be very heavily involved in the construction phase, because they must define which data will be included, and how they will be standardised. On the informatics side, engineers take over much of the work as it becomes less conceptual and more practical. "Early adopters", often data contributors who are also part of the disease expert group, contribute data. Engagement with the wider research community continues, though perhaps less intensively. As construction continues and there are standardised data to work with, the core group of experts, still in the driving seat, consult more widely to determine which questions might provide the best "proof of concept" for the platform. They are also likely to drive the early use of the data.

Platform as a service

A well-conceived platform that has taken user needs and desires into account from the start can expect to expand rapidly after the proof-of-concept phase, attracting contributors and users of data and analysis from a pool far beyond the platform's founding "inner circle". It is at this stage that it really transforms itself from a database to a platform -- an entity that exists to facilitate interactions between different user groups by making life easier for all of them. The disease experts who are so central to the development of a successful platform become a user group like any other. While they continue to contribute to discussions about standards as the disease and community needs evolve, their role in day-to-day decision making around the database itself becomes limited.

This is a critical inflection point for a data-driven platform. The informatics work is by now relatively routine -- largely a matter for professional data managers. Important decisions around governance and communications have been taken and procedures established. Platform maintenance becomes largely a service operation, although because science advances and needs mutate, a limited number of dedicated disease specialists will always be needed to guide development and interface effectively with the wider user community. "Early on there are lots of decisions that need to be made by scientists, but once you've programmed that into a robust tool, there's no reason [the database management] shouldn't be outsourced," said one academic-track employee of WWARN. "I think [WWARN] would be better as a separate non-profit in a science park somewhere, it would change the perception that it's a data grab and you'd get a different type of employee, better suited to the work."

WWARN's founders foresaw this possibility. The 2009 grant proposal said: "This arrangement between Oxford University and WARN is understood by both parties as optimal for the initial phase of WARN. Our long term plans include the real possibility that we will evolve into an independent NPO [non-profit organisation] … likely by the end of the second year." A later suggestion involved shifting the platform to an endemic country, most likely South Africa.

Neither transition occurred, the latter because of concerns over the informatics infrastructure. Many interviewees saw an academic base as an appropriate "incubator" for WWARN and similar disease-specific platforms in the conception, construction and proof-of-concept phases. WWARN has not yet reached full maturity -- demands for the addition of drug safety data and significant changes in the terms of use and metadata presentation needed to support a "gated" access model still require very significant input from the malaria community. However several interviewees, including WWARN employees, suggested that once a data platform does reach maturity, it would be better able to serve contributors and users if it was hosted outside a university. "It comes down to having perceived neutrality, credibility and lack of any agenda other than being a public good," said a WWARN founder.[8]

If the day-to-day maintenance of a mature data platform were shifted to an institution in which professional advancement did not depend on publishing papers,

[8] Others, also including WWARN staff, believe that in biomedical research settings "maturity" is an unattainable goal. In their view continual evolution will demand continual incubation at a level best provided in an academic setting.

it would not in any way reduce the importance of active engagement with disease specialists. "You still have to find a way of networking a critical mass of people with content knowledge to be able to use a shared platform effectively," said another WWARN founder. "They know what's worth knowing". Nor would it diminish the need to support researchers in endemic countries in defining and answering questions that are of importance to local communities. It would simply allow for better use of the comparative advantages of different groups. Instead of funding a repository and all the things that make it worthwhile in a single packet in perpetuity, it may be possible to establish coalitions on both the funding and the implementation side. Working towards a single agreed purpose -- the use of existing data to generate more improvements in health through effective curation and relevant analysis of shared data -- different actors could take different parts. Funders that aim to develop skills among scientists in endemic countries could fund capacity building efforts while academics who have strong relationships with southern institutions could help deliver those efforts. Funders who see a shared database as an important resource for their own decision-making could fund its ongoing maintenance as well as supporting disease specialists to contribute to further development; these services could be delivered by a non-profit organisation dedicated to maintaining disease platforms of importance to health in low income settings.

From impact factor to public health impact

It will always be necessary to engage the groups listed in Table 3 in the development of a successful platform. But the amount of time and energy put into engaging and securing the support of different groups depends very much on the broader transformation of the data sharing landscape, including the ways in which academic and other researchers are incentivised in their careers.

For over two centuries, the work of academic scientists in all countries has been recognised principally through the publication of work in peer-reviewed journals. The process is slow, inefficient, expensive and largely opaque, and it encourages data hoarding. Papers that present only outputs from analysed (and therefore aggregated) data eliminate the possibility for exactly the type of pooled analyses that WWARN has demonstrated are so useful. The obsession with publication is especially pronounced in southern institutions. "For us, its not even about authorship in general," commented a WWARN board member who is a professor at an African research institution. "It's about being first author. When I sit on promotions boards, we don't even consider papers if the applicant isn't in the first three authors." This focus puts academic researchers at odds with the policy and programme implementation communities who don't give a fig for citations but do care about having access to actionable information as soon as possible.

Over the years since WWARN first took the bold step of beginning to collate and standardise data from individual patients and parasites, funders have begun to encourage data sharing, though not yet to incentivise it coherently. In 2010, several funders of public health research signalled that they expected grantees to make data available in the scientific commons. In the six years since then, none of them have, to our knowledge, begun to consider scientists' data sharing records when assessing grant applications. Nor have universities, when hiring health researchers. But they

have at least begun to lay the foundations on the road to data sharing, and medical journals are on the construction team.

From March 2014, PLoS journals began to require authors to make the data underlying their publications openly available.(Silva 2014) Earlier this year, the International Committee of Medical Journal Editors published the draft of a similar policy.(Taichman DB et al. 2016) None of these policies requires the careful curation and standardisation that makes the WWARN database really useful -- a tragic lost opportunity if ICMJE policy is implemented as drafted -- but they have contributed to a gradual but increasing acceptance on the part of academic researchers that they will have to make their data more openly available to other scientists. And they have kick-started important work on the development of metadata standards that make the data understandable and useable, tools that make data more discoverable, and citation standards that will allow for easier professional recognition of data sharing efforts.

The road towards more open science is a one-way track. "Standardised data so that everyone can learn and analyse with the same tools: that is going to be the future, however much people are conservative and reject it and are afraid of it," said an interviewee in a global health organisation. "This will happen, so let's make its useful -- at least for research."

Others, including from the global health funding agencies which to a considerable extent influence the terms of trade in public health research in low income settings, are blunter still. "We have to change the conversation to: data sharing is a given. If we [switched to an opt-out system], people would quickly find data sharing is not as bad as they thought and we'd make progress much faster," said one.

In a world in which "data sharing is a given," the business of building a platform is altered dramatically, as are the possibilities it offers. If the contribution of high quality data to a "high impact" repository were considered on a par with published analysis when rewarding researchers with grants or promotions, it would no longer be necessary for platform designers to go through such clumsy contortions to "bribe" researchers into contributing data. The months and years of trying to achieve buy-in from potential contributors would be short-circuited. And the hand-wringing that goes into discussions around authorship would disappear. The sort of database initially envisaged by WWARN's founders would become a possibility and the design and governance, currently skewed by the necessity to incentivising data contribution, could become more balanced. Researchers, no longer haunted by being scooped, would get high quality support for data cleaning, management and storage. "That would be so great," said one researcher who has contributed data to WWARN. "When I think back to what I used to do 15 years ago, with 17 versions of the same data set on my hard drive, all called something like "malaria_final" -- it's really embarrassing!"

That sort of supply side incentivisation of sharing would be a very good start; it could be supported by impact metrics similar to those used for journals. The more data from a given repository gets used (a logical proxy for its utility, useability and quality) the higher the repository's impact factor and the greater the incentive to contribute data to it. But generating data alone, however high quality and however widely shared, is of little use. Ultimately, the value of a data sharing platform in public health is reflected in the extent to which it is used to generate knowledge which leads to changes in policy and practice. Data fingerprinting tools under

development allow the use of data from a well-curated dataset to be tracked. That could lead to the equivalent of a "citation index" which would provide researchers with professional credit for each policy-relevant re-use of the data they collected and shared.

As APMEN has shown, policy-led use can also be directly incentivised through research calls. It should be. Funding for secondary analysis that answers policy-relevant questions is a potentially important way of ensuring that the expertise of those who collected the data is honoured and added to. Data contributors are by definition most familiar with the specificities of a local situation. While many interviewees mentioned that the gulf between researchers on the one hand and programme managers and policy-makers on the other remains unacceptably wide, endemic country researchers certainly have better access to local policy makers and a better understanding of their context and concerns than collaborators in northern organisations.

Analyses that are more relevant to specific local or regional concerns would supplement, rather than replace, the more global methodological work and pooled analyses for which WWARN has become well known. Peer reviewed journals and the paper-as-metric are likely to survive for some time in parallel with structures that favour more open science, so platform users who produce analyses of global importance will continue to be rewarded by papers in high impact journals.

The business model

The costs of developing a platform are front-loaded; they generally diminish rapidly after the construction phase, though they creep up again as a platform increases in size, complexity and functionality. Community engagement and disease-specific standards development must be repeated for each new disease platform, but some economies of scale will be achieved through replication of informatics, governance and communications structures. Though their potential value grows over time, data repositories are by nature long-term ventures, and their costs need to be met over the long term.

Despite all the talk about business models in the early years of WWARN, and the appointment of staff to develop resources, attempts to develop an alternative to grant funding by philanthropic or public bodies were not successful. As data sharing efforts gather steam and the volume of disease areas covered increases, this may become difficult to sustain. "There is less and less appetite for committing [to funding] for long periods of time," said one interviewee from a global health funding organisation. "No funder wants to be blackmailed for life into having to carry on funding the platform... You need to come up with alternatives, to involve the thinking that venture capitalists have."

For neglected diseases of poverty, where the demand from the pharmaceutical industry is limited, it is likely that at least some of the core funding for "public good" data platforms will always be met from public or philanthropic sources. This study did not investigate alternative models for funding data platforms in any detail; however some ideas were raised in interviews and discussions.

One former WWARN employee suggested building a "per-patient" data curation cost into all clinical trials budgets. "Pharma puts at least 200 or 300 pounds [per

patient] into data management to ensure that the data are useable for ever.[9] But in many countries they don't even put a penny on that. If we were to put 10 or 20 dollars per study participant into making sure the data are useable for ever, that could probably work prospectively… but data management needs to be there from the very beginning." The curation fee would then get passed on to a designated database that met certain standards for curation and data standardisation. The solution could not, of course, be applied to retrospective data.

A variation on this idea is to charge a fee similar to the "open access" fee charged by journals. Funders (including pharmaceutical firms) who require data to be made available for use by other scientists would pay for curation of data from the studies they fund. Even if research funders and academic institutions fail to reward data sharers with grants or better jobs, many researchers will be incentivised to deposit their data in approved repositories as more journals follow PLoS's lead in requiring such a move. Such requirements raise the possibility of publisher subscriptions. Could a coalition of companies such Springer and Elsevier (which in 2014 reported a profit margin of 37% on revenues of £2.1 billion) contribute to maintaining community-supported platforms in which authors of papers they publish can usefully deposit underlying data?

Data platforms can increase their value to prospective users by proactively ensuring, as WWARN is doing, that they meet the needs of the journals or funders that require data to be shared. In January 2016, WWARN registered with the Registry of Research Data Repositories; it now includes an "open access" badge on its website.(Registry of Research Data Repositories 2016) In July 2016, the editors of PLoS Medicine agreed that researchers investigating malaria in pregnancy could fulfil their data sharing obligations by depositing the data underlying the study in the WWARN database. Complying with the terms of this agreement will require WWARN to accelerate existing plans to modify its terms of submission and set up an independent data access committee, but the acceptance of WWARN as data repository by a journal dedicated to open science is a big step forward. The next big step facing WWARN and other data repositories that invest significantly in curation is to get users to pay to have their data actively and ethically curated, safeguarded and shared.

The discussion of both impact factors and business models spotlights a black hole in the public health community's knowledge around data sharing platforms. We simply have no adequate standards by which to measure the value of these investments, nor even any idea about the time-frame over which any return can be expected to accrue. Tangible health benefits such as reduced mortality are notoriously difficult to link to research outputs; apportioning the contribution of data curation services to such outcomes would be harder still. And what of less tangible benefits? One malaria researcher in an endemic country said that collaborations based on shared data, including through the WWARN platform, had led to new partnerships, more professional grant applications, and thus more funding for his research institution. Those benefits can't readily be calculated on a profit and loss statement, but that does not make them less real.

[9] Data underlying studies used for product registration purposes must be kept available for 15 years.

5 Summary

WWARN has been a pioneer in bringing together individual patient and parasite data collected in low income settings. The group has developed standards that are applicable beyond malaria, making important contributions to global efforts such as CDISC. It has demonstrated both the feasibility and the value of rigorous data curation, including of heterogeneous data. And it has produced world-class science, developing important new methods and tools, together with pooled analyses which have led to changes in international guidelines, and which should eventually reduce illness and deaths caused by inadequately treated malaria.

The road to this success has not been entirely straightforward. Above all, the WWARN experience demonstrates the importance of early and wide-ranging consultation with potential users of a data platform. To design a useful platform efficiently, you need to start with a clear idea of who will want to use the platform, and what for. That means understanding the often quite different needs of each potential user group, and meeting them, in a way that, incidentally, also meets any needs you may have.

A data platform that aims to contribute to maximising health in low income settings must also maximise the likelihood that analysis of data in the platform is both relevant and adequately communicated to decision-makers in those specific settings. This reduces the relative value of papers in academic journals, increases the relative value of local knowledge, and puts researchers in countries most affected by the disease in question in pole position. To succeed substantially in this goal, platforms must take advantage of the knowledge and expertise of those researchers, match-making them with other researchers who have skills and capacities that may be in short supply in their own setting. Both sides must be adequately rewarded for the work they put in to identifying relevant questions, answering them and communicating those answers. That's not going to happen as long as peer reviewed publications continue to overwhelm all other possibilities as measures of achievement in health research.

Scientists who specialise in disease areas are the critical drivers of platform development, and specialists in informatics design are essential to platform conception. Disease specialists must remain engaged as users -- without them, the diligently curated data will languish, shared but unused. A small number of disease specialists may also have the interest and talent to focus on database management. However, as the platform matures, many university-based academics with an interest in life-saving research may wish to refocus on using, rather than managing, bio-informatic resources. A neutral public health entity that employs a few specialist advisors and enforces standards and transparent governance structures developed and agreed by the broader scientific community may prove a more efficient home for such a resource, as well as being broadly acceptable to a wider range of users.

WWARN continues to demonstrate its value with new uses of its growing database. It has inspired many others to set off down similar paths. We hope that this review of its experiences encourages them on that path, while perhaps charting a way around some of the obstacles they are likely to encounter.

Acknowledgements

We'd like to thank everyone who has participated in this study for their time and the insights they have provided. WWARN has been admirably open with its records, and dozens of WWARN staff and associates have been unstintingly generous with their time. Very special thanks are due to Carol Sibley, Philippe Guerin, Laura Merson, Andrea Stewart and Sam Smith, who went above and beyond in facilitating this work. They certainly do not share all the views expressed in this report, nor any of the blame for errors or omissions.

We thank the Wellcome Trust, and in particular Dave Carr and Katherine Littler from the policy team, for financial support for this work, and for their ongoing interest in finding equitable ways to get more out of investments in health research.

References

Eastwood, Brian. 2016. 'The Return of Platforms (and How to Not Fail at Building One)'. *MIT Sloan School of Management*. May 3. http://mitsloan.mit.edu/newsroom/articles/the-return-of-platforms-and-how-to-not-fail-at-building-one/.

Evans, David S., and Richard Schmalensee. 2016. *Matchmakers: The New Economics of Multisided Platforms*. Boston, Massachusetts: Harvard Business Review Press.

Gage, Deborah. 2012. 'The Venture Capital Secret: 3 Out of 4 Start-Ups Fail'. *Wall Street Journal*, September 20. http://www.wsj.com/articles/SB10000872396390443720204578004980476429190.

O'Hear, Steve, and Natasha Lomas. 2014. 'Five Super Successful Tech Pivots'. *TechCrunch.* May 28. http://social.techcrunch.com/gallery/five-super-successful-tech-pivots/.

Sanofi-Aventis. 2006. 'Application for Inclusion of Artesunate/Amodiaquine Fixed Dose Comobnation for Paediatric Tables in the WHO Essential Drugs List'. Sanofi-Aventis. http://archives.who.int/eml/expcom/children/Applications/OMSCoarsucam_malaria.pdf.

Sibley, Carol H., and Pascal Ringwald. 2006. 'A Database of Antimalarial Drug Resistance'. *Malaria Journal* 5 (1): 48.

Silva, Liz. 2014. 'PLOS' New Data Policy: Public Access to Data'. *EveryONE*. February 24. http://blogs.plos.org/everyone/2014/02/24/plos-new-data-policy-public-access-data-2/.

Taichman DB, Backus J, Baethge C, and et al. 2016. 'Sharing Clinical Trial Data: A Proposal from the International Committee of Medical Journal Editors'. *JAMA* 315 (5): 467–68. doi:10.1001/jama.2015.18164.

Tarning, J., I. Zongo, F. A. Some, N. Rouamba, S. Parikh, P. J. Rosenthal, W. Hanpithakpong, et al. 2012. 'Population Pharmacokinetics and Pharmacodynamics of Piperaquine in Children with Uncomplicated Falciparum Malaria'. *Clinical Pharmacology and Therapeutics* 91 (3): 497.

Taylor, Walter RJ, Dianne J. Terlouw, Piero L. Olliaro, Nicholas J. White, Philippe Brasseur, and Feiko O. ter Kuile. 2006. 'Use of Weight-for-Age-Data to Optimize Tablet Strength and Dosing Regimens for a New Fixed-Dose Artesunate-Amodiaquine Combination for Treating Falciparum Malaria'. *Bulletin of the World Health Organization* 84 (12): 956–64.

Witkowski, Benoit, Nimol Khim, Pheaktra Chim, Saorin Kim, Sopheakvatey Ke, Nimol Kloeung, Sophy Chy, et al. 2013. 'Reduced Artemisinin Susceptibility of Plasmodium Falciparum Ring Stages in Western Cambodia'. *Antimicrobial Agents and Chemotherapy* 57 (2): 914–23. doi:10.1128/AAC.01868-12.

World Health Organisaton. 2015. 'Guidelines for the Treatment of Malaria. Third Edition'. Geneva: WHO. http://www.who.int/malaria/publications/atoz/9789241549127/en/.

———. 2016a. ‘Q&A on Artemisinin Resistance’. *WHO: Malaria*. July. http://www.who.int/malaria/media/artemisinin_resistance_qa/en/.
———. 2016b. ‘WHO List of Prequalified Medicinal Products’. *World Health Organization*. August. http://apps.who.int/prequal/query/productregistry.aspx.
World Health Organization. 2015. ‘Developing Global Norms for Sharing Data and Results during Public Health Emergencies’. WHO. http://www.who.int/medicines/ebola-treatment/blueprint_phe_data-share-results/en/.
WWARN DP Study Group. 2013. ‘The Effect of Dosing Regimens on the Antimalarial Efficacy of Dihydroartemisinin-Piperaquine: A Pooled Analysis of Individual Patient Data’. *PLoS Med* 10 (12): e1001564. doi:10.1371/journal.pmed.1001564.