# TRUST AND HUMAN CHALLENGE VACCINE TRIALS

# TRUST AND HUMAN CHALLENGE VACCINE TRIALS: EXAMINING THE RELATIONSHIP BETWEEN PUBLIC OPINION AND TRIAL DESIGN SELECTION

By BENJAMIN D. MARSHALL, B.A.

A Thesis Submitted to the School of Graduate Studies in Partial Fulfilment of the Requirements for the Degree Master of Arts

McMaster University MASTER OF ARTS (2022) Hamilton, Ontario (Philosophy)

TITLE: Trust and Human Challenge Vaccine Trials: Examining the Relationship Between Public Opinion and Trial Design Selection

AUTHOR: Benjamin D. Marshall, B.A. (McMaster University)

SUPERVISOR: Professor Ariella Binik

NUMBER OF PAGES: iii, 82

## Lay Abstract

In a challenge trial, "healthy volunteers are intentionally exposed to [diseases] in a controlled environment," to give researchers a better understanding of a disease in order to develop cures or preventative measures for it (WHO 2021, Preface). Many research ethics scholars believe that conducting challenge trials could negatively impact the public's faith in the institution of medical research, but the relationship between public trust and conducting challenge trials is complex and existing literature on the subject does not sufficiently clarify it. This paper begins by exploring whether or not challenge trials can be ethically conducted. Once I show that they can be under particular circumstances, I examine how public trust concerns largely result from the fact that 'public' and 'trust' are not well defined. After defining them, I formulate my own account of how public trust should apply to a risk/benefit analysis for the purpose of trial design selection.

Abstract

In a challenge trial, "healthy volunteers are intentionally exposed to pathogens in a controlled environment, in order to promote understanding of the pathogenesis, transmission, prevention and treatment of infectious diseases in humans." (WHO 2021, Preface). Intentional infection is an uncomfortable concept, and as a result there is a widely held belief amongst research ethics scholars and commentators that a significant ethical concern with challenge trials is their potential to negatively impact the public's trust in the institution of medical research (Eyal 2022, 4). However, the relationship between public trust and the ethics of conducting and assessing challenge trials is complex and existing literature on the subject does not sufficiently clarify it. This paper will begin by examining the ethical permissibility of challenge trials. Once these trials are shown to be ethically permissible under particular circumstances, I will explore how concerns about the way these trials allegedly exacerbate public mistrust largely result from ambiguities in the terms 'public' and 'trust'. After both terms are defined, I will formulate my own account of how public trust should apply to a risk/benefit analysis for the purpose of trial design selection called the community engagement account, which argues that trial design selection policy should focus on demonstrating trustworthiness rather than garnering trust. Because demonstrating trustworthiness requires meeting a set of known expectations, this account identifies local, specific publics as those whose expectations should be of concern when discussing public trust and trial design selection. To examine the expectations of these publics, this account defends community engagement as the measure which should be used to acquire evidence of harmful public mistrust towards the institution of science that could potentially result from conducting a challenge trial.

Table of Contents

## Declaration of Academic Achievement

I, Benjamin David Marshall, wrote this thesis with guidance from my supervisor Dr. Ariella Binik and my friend and mentor Jordan Vaters (M.A.). My novel contributions include identifying the underexplored and underdeveloped nature of public trust as it applies to challenge trial design selection, recapitulating public trust through the lens of trustworthiness, and formulating a novel understanding of public trust as it applies to challenge trial design selection.

## Introduction

There is a widely held belief amongst scholars that challenge trials have the potential to dangerously impact the public's trust in the institution of medical research (Eyal 2022, 4). To conduct a challenge trial, "healthy volunteers are intentionally exposed to pathogens in a controlled environment, in order to promote understanding of the pathogenesis, transmission, prevention and treatment of infectious diseases in humans." (WHO 2021, Preface). Involving healthy volunteers is necessary for accurately testing the safety and efficacy of vaccines, and challenge trials are often regarded as a uniquely beneficial, sometimes even essential, avenue for conducting human vaccine research (Harris 2005, 243). However, there is a consensus that trials involving healthy volunteers must meet a high standard of scientific justification without exceeding an acceptable level of risk to be ethically conducted, and novel experimentation always involves an element of undefined risk (WHO 2021, 6). This not only makes it difficult to know if the risks of a trial are acceptable, but also whether or not the benefits supersede them.

While there is a general consensus that challenge trials can be ethically conducted, it might seem to some as though intentionally exposing healthy individuals to risks in a trial offering them little to no direct benefit violates the most basic responsibilities researchers owe to the wellbeing of those under their care (Ezekiel, Miller, & Grady 2008, 273). If the public holds this view, there are concerns that conducting a challenge trial could irreparably damage the reputation of the institution of medical research. This would impede upon the effectiveness and continuation of future scientific endeavours, a consequence which could prove too harmful to be outweighed by any benefits provided by the trial (WHO 2021, 5) (Hope & Macmillan 2004, 110).

This paper closely examines the relationship between public trust and human challenge vaccine trial design selection. By investigating the existing literature on challenge trials for discussion on the role of public trust in a risk/benefit analysis, I argue that this relationship has not been adequately explored or characterized thus far, and that concerns regarding the effects of backlash stemming from public mistrust are therefore incorrectly and inconsistently applied as a consideration in trial design selection. After examining existing arguments, I provide my own account of public trust which defines the circumstances under which it should factor into a risk/benefit analysis for the purpose of trial design selection, and to what extent. This account argues that public trust should be included in a risk/benefit analysis if and only if evidence gathered by community engagement makes it clear that conducting a challenge trial would hinder the scientific community's ability to continue conducting important medical research.

Chapter one begins with an exploration of challenge vaccine trials and what they entail as they are conducted today, before delving into their history to better understand why their conduct may or may not cause a negative impact on public trust. I then compare the benefits and ethical considerations of challenge vaccine trials and typical field trials, highlighting the consensus view that intentional infection is not *prima facie* unethical. While this makes public backlash against otherwise beneficial challenge trials unjustified, I claim that the consequences of a potentially significant loss in public trust are important enough to consider in certain instances, regardless of their legitimacy. At the end of the chapter, I conclude that public trust as it is currently formulated in the literature is not well-defined enough to be consistently applied as a consideration in trial design selection.

To develop a functional account of public trust, chapter two argues that, because the 'public' refers to multiple groups, we cannot define the relationship between public trust and challenge

trials without understanding which public we are concerned with. I examine different conceptions of public trust as it is discussed in the challenge trial ethics literature, categorizing them by how they characterize public opinion in a risk-benefit analysis. This exercise displays how existing public trust arguments fail to identify a public, and in doing so cannot accurately reflect the impact conducting a challenge trial would have on public trust.

Chapter three argues that community engagement with specific, identifiable publics, rather than general or undefined ones, should be used to determine the role of public trust in trial design selection. Beginning with an examination of trust itself, I argue that there should be a focus on trustworthiness instead of trust because doing so better reveals when public opinion is important to consider, and therefore which publics are of concern when applying public trust in a risk/benefit analysis. Demonstrating trustworthiness demands repeatedly achieving the best results to cultivate a reputation of competence and good intentions, and because an impartial, third party risk/benefit analysis is the most accurate measure available for assessing trial designs, following it in all instances should best contribute to achieving this reputation. However, because trustworthiness demands *repeatedly* demonstrating results, legitimate evidence which indicates that conducting a particular trial would damage public opinion enough to limit future scientific endeavours should nonetheless influence or factor into a risk/benefit analysis. Using community engagement to acquire this evidence not only allows one to accurately determine the impact conducting a trial would have on the trust of an identifiable populace, but it also instructs how to mitigate the effects of any potential loss of trust. Public trust is therefore best served by incorporating public opinion into a risk/benefit analysis only when community engagement indicates that failing to do so would damage the reputation of the institution of medical research enough to eliminate continued opportunities for it to demonstrate trustworthiness.

## Chapter 1: Assessing Challenge Trials and Possible Alternative Trial Designs

### Definition and Introduction

This chapter will examine a number of considerations used in determining the permissibility of conducting challenge trials, particularly focusing on human challenge vaccine trials, to provide context for examining public trust in more depth. I will begin by explaining what vaccine challenge trials entail as they are conducted today. I will then overview their history to display both how they have progressed and the possible causes for public hesitancy towards medical experimentation on humans involving intentional infection. After this, I will discuss the scientific benefits of conducting challenge trials, before examining the ethical considerations they entail to fully explain what a challenge trial involves and why one might choose to conduct it alongside or instead of alternative trial designs. The chapter will conclude with a brief examination of public trust as an especially underexplored yet often utilized ethical consideration and its possible implications on trial design selection.

### Description

The design of a vaccine challenge trial, including the nature of the pathogen being studied, will in large part determine the level of benefit and risk it entails (Rid & Rosenberg 2020, 750). Challenge trials can be used to examine the progression of a disease or the efficacy of a treatment or preventative measure, all of which involve unique design considerations aimed at minimizing risk and maximizing scientific benefit. Regarding human vaccine challenge trials in particular, their ethical conduct necessarily involves obtaining valid informed consent from volunteers, inoculating them with a vaccine candidate, exposing them to a given pathogen, studying the results, and ensuring that volunteers do not experience a level of harm that could be deemed impermissible as a result of infection (WHO 2021, Preface-2). To meet requirements

ensuring the relative safety of volunteers, it can be presumed that a disease studied in a challenge trial of concern is either self-limiting or treatable (Miller & Grady 2001, 1028) (Jamrozik 2021, ix). While it might be possible to conduct a challenge trial for diseases not meeting these criteria, there is not yet a consensus on the permissibility of conducting them or what conditions are necessary to ethically do so (Dawson et. al. 2020, 5), and establishing one extends beyond the scope of the current inquiry.

Challenge trials, for the purposes of this paper, may thereby be classified as trials which intentionally expose consenting volunteers to self-limiting or treatable diseases for the purpose of understanding the pathogenesis, transmission, prevention and treatment of infectious diseases in humans. When using the term 'consenting', I am omitting instances where consensus surrounding what constitutes valid consent is not established. This includes consent provided by minors, parties under potential undue inducement as a result of inequitable financial compensation, etc. (Jamrozik 2021, 64-74). Not only is it unlikely that consent is uniquely implicated in challenge trials, but each of these instances of questionable consent deserve individual consideration beyond the scope of this paper. These stipulations should facilitate a clear analysis of public trust as a characteristic which bears weight in determining the permissibility of specific trial designs without deviating into other important ethical dilemmas.

Jenner Trial

In order to examine challenge trials as they are currently conducted, it is important to understand how and why they have evolved over the years. The first recorded challenge trial was conducted by Edward Jenner in 1796, during the height of the European smallpox epidemic. At the time, smallpox was responsible for the deaths of 400 000 people annually, and for blinding a third of survivors, so finding a cure or preventative measure was a paramount scientific concern

(Riedel 2005, 21). Noticing that the typical facial scarring caused by smallpox was far less prominent amongst milkmaids than amongst the general populace, Jenner hypothesized that being infected with cowpox provided immunization, or at least a significant level of resistance, to smallpox. Cowpox is a disease that, while similar to smallpox, is far less severe (Riedel 2005, 23). As such, if one could achieve immunity or resistance to smallpox by instead suffering a comparably far milder deliberate cowpox infection, and such a practice was adopted by the larger medical community, countless individuals could be spared from grievous injury or death. Seeing an opportunity to rid humanity of a terrible affliction, Jenner designed an experiment to determine the truth of his theory.

Jenner recruited 8-year-old James Phipps to test his hypothesis. He extracted viral material from the lesion of Sarah Nelms, a milkmaid experiencing a cowpox infection, and inoculated Phipps with the infectious material. Over the next two weeks, Phipps would suffer a fever and subsequent chills, ultimately recovering completely when his infection had fully subsided. Once healthy, Jenner exposed Phipps directly to smallpox. When an infection failed to develop, thus demonstrating Phipps's newfound immunity, the experiment was declared a success (Reidel 2005, 24).

This discovery and the resulting breakthroughs in vaccination research not only led to the eventual eradication of smallpox but have immensely contributed to the fight to eliminate numerous other devastating viruses worldwide such as polio, meningitis, and, most recently, COVID-19. However, a closer look at Jenner's experiment exposes multiple ethical issues, many of which were a prominent theme in early medical research involving human subjects. For one, the scientific justification authorizing both Jenner's hypothesis and methods were questionable at best. The anecdotal nature of his evidence supporting the conclusion that cowpox provided

immunity to smallpox would not be sufficient to expose any risks to humans by modern research ethics standards (Bambery et. al. 2016, 93). While his evidence may have constituted adequate justification for further exploring the link between cowpox and immunity to smallpox, there is widespread recognition that more preliminary research would be required in order to ethically conduct a trial of this nature on human subjects.

Additionally, trials must adhere to the obligation held by researchers to balance risks towards the health of present volunteers with the health of a much greater number of potential future individuals benefiting from the results of this research. This obligation, which forms the basis of modern research ethics doctrines, stipulates that volunteers can be exposed to increased risk in studies with a high level of justification and expected benefit, provided that a trial does not entail an unacceptably high risk of experiencing potentially severe effects as a result of such research (Rid & Roestenberg 2020, 750). In other words, there are some risks which may be deemed categorically unjustifiable, and others which are recognized to be both necessary for the proper conduct of research and mild enough to expose volunteers to. This is complicated by the fact that research inherently involves unknown variables, meaning that justification in part relies on the level of certainty we have regarding both the risks and benefits of a particular study. Not only was the risk of death or permanent injury as a result of smallpox well beyond the threshold of what constitutes acceptable risk by any reasonable measure, but the efficacy of variolation was unproven at the time, meaning that the risk of infection was undetermined and therefore unjustifiable to expose a healthy volunteer to. Consequently, the speculatively high risk of both infection and severe or permanent harm would likely render this trial impermissible regardless of any potential future benefit to hypothetical individuals gained as a result. Both acceptable risk

thresholds and uncertainty will be explored in more depth later in this chapter to solidify these claims.

Another characteristic of concern in this trial involves the concept of informed consent. If an individual is participating in research which risks their wellbeing while offering them little to no potential benefit, a high standard of informed consent must be met, and it is unlikely that the Jenner trial would meet this standard. There is an understanding that the level of risk an individual is able to legitimately consent to depends on their ability to accurately comprehend such risk (Jamrozik 2021, 64-5). The harm posed by smallpox infection would likely exceed Phipps' understanding, and the consent obtained would not be truly informed. Not only this, but if an individual passes their infection on to uninformed third parties, this would violate *their* right to informed consent (Lynch 2020), so the possibility of causing widespread harm would at the very least necessitate measures of containing third-party infection that Jenner failed to implement.

Additionally, there are conditions regarding who may provide legitimate informed consent that are generally agreed upon, typically involving requirements surrounding maturity and competency (Jamrozik 2021, 64). In this case, the volunteer, James Phipps, was 8 years old. Conducting a trial that involves obtaining informed consent from minors or proxy consent via their guardians, for the purpose of altruistic medical research participation, requires a high level of justification and a relatively low risk threshold to be considered ethically permissible (Jamrozik 2021, 67-8). In the Jenner trial, this justification could not have possibly been present, as the risks involving intentional infection and the science of variolation were undetermined at the time and therefore could not be accurately predicted. This would mean that enrolling Phipps at this stage of research was likely impermissible. Some might argue that, because of the

endemicity of smallpox at the time and the lack of effective treatments available, acquiring immunity meant that Phipps stood to gain personally by participating in this trial, changing its status as altruistic which could justify reducing our standard of information enough to make his participation permissible (Bambery et. al. 2016, 93). An argument of this nature carries a large burden of proof in light of the severe risks presented by a potential smallpox infection, and as such is beyond the scope of this paper, but in the absence of such an argument it can be presumed that Jenner did not meet the justification necessary to involve a minor in this trial.

Post-Jenner Challenge Trial History

The adoption of Jenner's challenge trial method in later experiments by other researchers only served to entrench a tradition of conducting medical research on a regular basis that, by contemporary standards, are ethically dubious. Such trials include the Walter Reed yellow fever experiments, which would carry a risk threshold considered insufficient by contemporary standards (Clements et. al. 2017), and the Jonas Salk flu vaccine experiments, which involved institutionalized asylum patients, complicating informed consent due to both their presumably diminished mental capacities and their incarceration status (PBS 2011). These abuses stemming from challenge trials culminated in experiments conducted by Nazi Germany, so horrifying in nature that the resulting hesitancy towards biomedical research at large is still felt today (Harris 2005, 242). These experiments prompted the implementation of international guidelines for medical research ethics, such as the Declaration of Helsinki, detailing the parameters within which medical research may be ethically conducted (Hope & Macmillan 2004, 111).

Despite the formation of these guidelines, the following decades continued to see challenge trials which violated them. Experiments such as the Willowbrook hepatitis trial, which involved intentionally infecting children with hepatitis (Goldby et. al. 1971), and the Guatemalan

syphilis and gonorrhea challenge trials, which involved intentionally infecting vulnerable populations with bacteria known to cause the aforementioned diseases (WHO 2021, 2), contributed to the further erosion of public confidence in biomedical research. While contemporary standards of research approval and oversight are even higher than they were in the decades immediately following the Nuremburg trials, the consequences of past experiments on the trust between medical research involving human subjects and the general populace cannot be ignored. This trust is essential to both the continuation of scientific research and the widespread adoption of medical advancements, and as such there must be a commitment to retaining and improving upon this trust whenever it is in jeopardy (WHO 2021, 5).

Additional Concerns

In addition to this history of mistreatment, many believe that there is reason to presume the existence of a general or intuitive inclination[i] amongst the public that the intentional infection present in challenge trials is inherently wrong, which forms the basis of concerns that challenge trials may especially damage public trust (Hope & Macmillan 2004, 110) (Eyal 2022, 4). If the public does hold this perception, this may support a prerogative to conduct alternative study designs whenever possible for the purpose of determining vaccine efficacy and safety, in the name of protecting the reputation of the institution of science (Hope & Macmillan 2004, 110). A belief in the existence of this intuition is made evidently prominent by its role as a foundational principle of international doctrines of medical ethics including the Declaration of Helsinki and the Nuremburg code, which stipulate standards of risk minimization for the protection of the institutions of medical science and research via maintaining public trust (Rid & Roestenberg 2020, 750). The validity of these claims in all instances will be assessed at the conclusion of this chapter. As of now, it remains to be seen whether this is a legitimate ethical

concern, or a mere instinctual reaction based on feelings of discomfort towards the idea of intentional infection.

In light of possible scientific hesitancy resulting from these historical and intuitive concerns, and the possible consequences of this mistrust on the institution of science as a whole, it might seem that conducting challenge trials poses an unnecessary risk towards public trust. The merit of this claim is the motivation for this inquiry. Many would contend that, under certain circumstances, it would be ethically irresponsible *not* to conduct challenge trials (Jamrozik 2021, 33) (Harris 2005, 243). This claim is likely made for two reasons. First, the unique benefits held by challenge trials may provide enough scientific justification to legitimize any discrepancy in risk they may present. Second, the actual, tangible risks of a challenge trial might not in practice violate our threshold of acceptability as one might presume a trial involving intentional infection would. Therefore, if the additional benefits they convey are substantial, it would be imperative to conduct a challenge trial on the basis that it better achieves a balance between those considerations owed to volunteers participating in medical research and those owed to society at large. At the very least, it is generally agreed upon that conducting a challenge trial is not *prima facie* unjustifiable.

Scientific Benefits

In order to arrive at an accurate ethical assessment of challenge trials, the unique scientific benefits and drawbacks afforded by them must first be outlined. The uniqueness of a benefit is inherently relational, and as such, challenge trials are often compared to the most commonly used alternative study designs to determine which benefits they hold over and above them (as well as their potential drawbacks). These alternatives are cell models and animal models in preclinical research, and field trials, also referred to as clinical trials, for research with

human subjects (Kent Scientific 2020) (Miller & Grady 2001, 1029). The first two models, however, are typically used to perform preliminary research on vaccine candidates for the purpose of preparing them for human experimentation (which provides more accurate and generalizable results than non-human testing). Therefore, because the most common human models are either challenge or field trials, field trials are the most commonly cited alternative to challenge trials (Eyal 2022, 1) (Jamrozik 2021, 28). These trials involve inoculating volunteers with a vaccine candidate before releasing them, where they go about their lives as normal (with the addition of attending regular check-ins). Once they contract or interact with the relevant disease in their day-to-day lives, they are brought in for thorough examination to determine the efficacy of the vaccine candidate (Kent Scientific 2020).

1) Number of Volunteers & Use of Resources

Many of the discrepancies between challenge and field trials result from the respective scales of each study design. In order for medical research of any kind to reliably progress, sufficient data must be collected to account for possible outliers and ensure that the results of a given study are as generalizable as possible. Because field trials do not guarantee that a volunteer will achieve infection, and because infection (or evidence of exposure if the vaccine candidate is wholly successful in preventing infection) is necessary to observe vaccine efficacy, these trials require more volunteers to achieve an adequate sample size. To illustrate, while a field trial might require anywhere from 100-10,000 participants in order to achieve a satisfactory sample size, challenge trials may only require the services of 10-40 volunteers to achieve the same results (Bambery et. al. 2016, 93).

By making the acquisition of volunteers more feasible, challenge trials achieve a number of additional scientific benefits. The selection process for volunteers is already quite stringent; in

order to run trials as safely as possible, there are often requirements regarding volunteer medical history, age, and sex, as well as racial or ethnic qualifications where medically relevant (Jamrozik 2021, 62). As such, researchers need to enlist enough volunteers to achieve a sufficient sample size, while at the same time ensuring the volunteers they choose to enlist are not exposed to levels of harm we might consider unacceptable based on pre-existing medical or physical conditions.

On these grounds, conducting a challenge trial instead of a field trial not only alleviates the practical burdens of acquiring enough volunteers to achieve sufficient results, but also allows for researchers to be more selective with their volunteers. This would not only improve volunteer safety, but it could also potentially improve the generalizability of results from a given trial (Jamrozik 2021, 30). By having the ability to focus results on a more specific group of people identified by a set of conditions, we can more accurately extrapolate results within populations encompassed by the selected volunteers. This is especially useful when testing vaccines for diseases which may only affect specific populations.

There are a number of additional practical benefits of requiring fewer volunteers. Firstly, the resources necessary to conduct a trial involving 40 individuals are far more realistic to acquire than a trial with 10,000 volunteers. Producing vaccine candidates is difficult, costly, and does not provide any guarantee that a successful vaccine will be created. As such, limiting the number of resources used in the development and transportation of these candidates is generally preferable, provided other rights are not disrespected. Requiring fewer volunteers to test a given vaccine candidate also provides the advantage of being able to test more individual vaccine candidates in a given trial (Rid & Roestenberg 2020, 750). If, for example, both a field trial and a challenge

trial involve 100 volunteers, a field trial may not even gather sufficient data on a single vaccine candidate, whereas a challenge trial could feasibly test anywhere from 3-10 vaccine candidates.

Because a trial involving a vaccine candidate does not ensure that the candidate will be successful in either efficacy or safety, the efficiency conferred by a challenge trial in testing multiple candidates could potentially allow for a successful candidate to be discovered and produced in exponentially less time than if a field trial was used. For example, if a field trial tests vaccine candidate A, it may take weeks before enough volunteers experience natural infection and researchers are able to determine whether or not A meets ethical standards of effectiveness and/or safety. If A is deemed insufficient at preventing or mitigating infection, then testing candidate B could take another few weeks, and then C, and so on and so forth. Alternatively, if a challenge trial is conducted, infection of volunteers is guaranteed, and the collection of sufficient data on a given candidate is limited only by the acquisition of volunteers and the duration of time researchers deem it necessary to ensure any long-term effects of a given infection or inoculation are sufficiently mitigated, both concerns which are also present in field trials. As such, conducting a challenge trial would likely result in the production of a successful vaccine candidate weeks or months earlier than if a field trial were conducted (Bambery et. al. 2016, 93-4).

Consequently, the ability to test multiple candidates means that challenge trials confer an additional benefit of ensuring that the final vaccine candidate chosen to be developed provides the most benefits to the general population while carrying the least risk. Returning to the previous example, if a field trial is only able to test vaccine candidate A, and A is determined to be sufficiently effective and safe, vaccine A would be produced and distributed. However, by testing more candidates, challenge trials could reveal that candidate B confers the same benefits

as candidate A while also carrying a lower risk of side effects. In this case, vaccine B would be the ethical choice over vaccine A, a choice which would not be known to exist in the absence of research supporting vaccine B. Not only would this support conducting challenge trials over field trials in some instances, but it would also support the use of challenge trials prior to, alongside, or following field trials to ensure the results of a field trial confer the strongest possible benefits and fewest risks to the general populace (Bambery 2016, 93).

The component of efficiency inherent in challenge trials derived primarily via this need for fewer volunteers categorizes one of the leading arguments for their use, which is creating a viable vaccine in a shorter amount of time (Miller & Grady 2001, 1029). In doing so, distribution and the inoculation of the general populace are expedited, and the impact of disease is mitigated with more ease and fewer infections. The actual impact of accelerating vaccine rollout varies based on the prevalence of the disease in question and the efficacy of the vaccine. Regardless of the scale of this impact, however, sooner is always better than later from an ethical perspective, provided everything else remains equal. This is highlighted, albeit in the extreme sense, by the website 1daysooner.org, which was created for the purpose of acquiring volunteers to participate COVID-19 vaccine testing on the basis that accelerating vaccine rollout by even a single day is ethically significant and worth facing the risks inherent in testing novel vaccine solutions (1daysooner.org 2020). While the benefits acquired from sooner producing and distributing a vaccine may not be this extreme in all cases and will need to be weighed against the risks necessary for achieving expedited production, they should nonetheless be considered incredibly significant.

2) Control Afforded by Attenuation

One unique feature of vaccine challenge trials that makes them potentially preferable to field trials from a scientific perspective is the ability they grant to control for numerous additional variables. This control is granted chiefly through the use of attenuated, as opposed to wild type, strains of a given disease. Attenuated disease strains are strains of a given disease which have been altered and/or selected for in order to ensure that the harms suffered by a volunteer via infection are as mild as possible without compromising the accuracy of scientifically relevant data (Jamrozik 2021, 31). Wild-type strains, on the other hand, are those which exist naturally in the wild, and as such, the designation of wild-type covers a multitude of mutations and variants. Being inpatient, challenge trials involve intentional infection in a controlled setting, which affords researchers the option to use attenuated strains to achieve and observe a milder yet still scientifically viable infection. Because field trials require infection to occur naturally, it is not possible for exposure to occur on an inpatient basis, meaning such trials necessarily utilize wild-type disease strains. In this respect, they stand in direct contrast to most modern challenge trials.

The control specifically afforded by attenuated strains over wild-type strains lends well to the repeatability of a given experiment. By utilizing the same strain of a disease on all volunteers being studied, differences in the results of an experiment can be more accurately attributed to deficiencies in the vaccine candidate, or towards certain features which may have been selected for in the testing population (e.g. a negative reaction to infection in, say, men over 40 would be attributed to the characteristics of being male and over 40 with additional certainty, and consequent vaccine candidates may thereby correct for these specific concerns). Utilizing data given by wild-type strains acquired through natural infection, on the other hand, carries the necessary consequence that there will be an element of randomness to the interpretation of results of a given trial. Regarding the previous example, it would no longer be as certain whether

or not a negative reaction in men over 40 has to do with their age and/or sex, or if that particular testing population had collectively encountered a more severe strain or mutation of the disease in the wild compared to other testing populations. In short, the fewer variables left uncontrolled in an experiment, the more a given experiment can be successfully repeated to demonstrate efficacy, which allows researchers to assert their results with greater certainty.

One caveat of utilizing attenuated strains regards the potential generalizability of the results of a given trial. While wild-type strains may imply a reduction in the ability for researchers to assert with certainty which variables are responsible for which results, they are also the strains which researchers attempt to provide immunity against at the conclusion of their research. In other words, the general populace is exposed to wild-type rather than attenuated strains, and as such a vaccine resulting from a given trial must protect against wild-type strains. Though challenge trials do not necessitate the use of attenuated strains but merely afford researchers the possibility of their use, it was previously established that challenge trials may only utilize disease strains which are either self-limiting or treatable, which in turn may require disease strains to be attenuated in such a way that they do not violate this condition (Jamrozik 2021, 31).

If this is the case, then it could be argued that field trials provide unique scientific benefits that a challenge trial cannot provide. Because the infection acquired in a field trial is incidental rather than deliberate, the severity of infection that researchers may ethically observe is far greater, seeing as infection would be deemed more or less inevitable and therefore beyond the scope of normative ethical considerations (though the truth of this claim in all instances will be debated further). As such, one might use this as scientific justification to conduct a field trial rather than a challenge trial. However, this ignores the possibility of utilizing challenge trials prior to or alongside field trials. This would not only confer the benefits of control afforded by

attenuation but would also allow for the acquisition of more generalizable results by letting researchers observe the interaction between a vaccine candidate and a given disease as it exists beyond the laboratory setting. Additionally, utilizing a vaccine candidate in a field trial which has already proven affective against attenuated strains of a disease would allow researchers to either certify the candidate for widespread use much sooner or isolate the causes for its inadequacy against wild-type strains with more accuracy.

3) Control Afforded by Inpatient Observation

In addition to attenuation, challenge trials allow researchers to conduct a study entirely inpatient. The potential inpatient nature of challenge trials confers scientific benefits derived from the ability to observe the course of infection from onset to recovery. Certain challenge trial designs may even focus strictly on observing the effects of a given disease without utilizing a vaccine candidate at all to study the progression of the disease itself (WHO 2021, 1). The conditions for examining the permissibility of such trials vary slightly, yet importantly, from the current inquiry, which is focused specifically on human challenge vaccine trials. While the results of this inquiry *may* apply to the justification of these particular trials, an independent study would be required to determine if this is actually the case and/or in which instances.

By having the ability to consistently monitor the procession of a disease and its interaction with a vaccine candidate from infection to mitigation, researchers are able to determine the efficacy of a given candidate at every stage of infection with a higher degree of certainty. Not only this, but the safety of the vaccine candidate itself can be adequately determined from the moment of inoculation onwards. In a field trial, infected individuals are only able to be examined after infection has already occurred, so observing the procession of a given disease and the consequent efficacy of a vaccine candidate inherently begins at some unspecified time after

infection has been established, largely after symptoms have presented themselves (Jamrozik 2021, 45). As such, both the effectiveness and safety of a vaccine candidate are likely more certifiable when observed through a challenge trial than through a field trial alone.

## Ethical Considerations

Until this point, this paper has focused on defining human vaccine challenge trials, exploring their history, and examining the scientific benefits and potential pitfalls they may entail. While important in and of themselves, the purpose of this inquiry is to investigate how these are implicated in an ethical determination regarding the permissibility of challenge trials. In doing so a conclusion regarding the importance of public trust as a factor in this determination can be arrived at. As such, the current examination will follow the prior analysis of scientific benefits and analyze the potential ethical considerations they entail. This will be done initially through an exploration of the relative risks and benefits entailed by the defining features of challenge trials, with an emphasis on determining what constitutes permissible risk. Following this, there will be a brief discussion surrounding consent, proceeded by an exploration of the ethics of study design, and finally the chapter will conclude with a brief look at the ethics of intentional infection itself.

### 1) Risk

To determine what constitutes permissible risk in vaccination research involving healthy subjects, it is important to first understand what we mean when we say 'risk'. A particular risk can be defined as, "a possible harm in which the probability of harm may be predictable to a certain degree but never controllable with certainty. The chance that pervades the risk to human subjects in biomedical research is also associated with the benefit that may accrue from it" (Van Ness 2001, 4). This understanding of risk thereby defines the concept as a function of probability

and magnitude (Van Ness 2001, 2). The more likely an occurrence, the riskier; the more injurious, the more harmful.

Human vaccination trials of any sort entail inherent harms involved with inoculation and the experimental use of novel vaccines on human subjects, harms which are justified by their necessity and the benefits acquired exclusively from such testing. As a result, there must be certain harms which volunteers may risk experiencing that are collectively accepted as ethically palatable, provided such harms are unavoidable in the pursuit of certain scientific goals (Resnik 2012, 2). What concerns researchers when considering the use of a challenge design as opposed to alternative trial designs is whether a particular challenge trial a) does not entail an excessive risk of potential harms we would consider ethically impermissible and b) does a better job than alternative trial designs to mitigate risks of unacceptable harms or provides unique and substantial benefits which constitute adequate scientific justification for permitting more risk.

If risking certain harms in medical research involving healthy volunteers is permissible, it must be determined which harms are considered unacceptable and at what threshold of risk, and whether or not challenge trials entail such impermissible harms. For example, a high risk of a mild harm occurring might be tolerated (e.g., pain suffered from needle injection) but even a slim chance of harms which could cause permanent injury or death would likely not be. Neither of these determinations, however, allow for the assertion that certain harms at a particular level of risk are unconditionally unacceptable for volunteers to face.

Examining policy does little to clarify what constitutes this impermissible risk or harm. Using American policy as an example, "U.S. law sets no definite limits on the level of risk that healthy volunteers may face in research. Federal research regulations require only that risks be minimized and reasonable in relation to benefits to participants and the expected gain in

knowledge (a social benefit)" (Resnik 2012, 2). What constitutes 'minimized' or 'reasonable' is left ambiguous, but it will be interpreted shortly. Turning to international declarations, the Nuremburg code explicitly states that, "the degree of risk to be taken should never exceed that determined by the humanitarian importance of the problem to be solved by the experiment." (International Military Tribunal 1947, 182). Somewhat conversely, the code also states that, "No experiment should be conducted where there is an a priori reason to believe that death or disabling injury will occur" (International Military Tribunal 1947, 182). The only true points of consensus, referred to as 'the Common Rule', are that additional risks must be justified by additional benefits, and risks should only be undertaken when necessary for genuine scientific discovery (Miller & Joffe 2008, 445).

This discrepancy regarding the validity of imposing a maximum limit on risks towards otherwise healthy volunteers is no clearer when examining the literature on the subject. Some authors, such as Nir Eyal, propose a number of reasons suggesting that such a threshold cannot rightly exist. Chiefly, he suggests that imposing an absolute cap on risk fails to accommodate for situations where infringing on these risks might provide overwhelming benefit such that remaining under this cap would be ethically disingenuous (Eyal 2020, 147-9). This argument is made both from a philosophical perspective as an admonishment of moral absolutism, and a practical perspective which considers the possibility that the weight of potential benefits might justify conducting riskier research under particular circumstances. This rationale is supported by the above policies which solely define acceptable risk as a measure relative to foreseeable benefits without imposing a definite limit and is consequently contradicted by the blanket imposition made by the Nuremburg code not to conduct research which could lead to death or disability. Eyal also discusses an additional concern that limiting research with absolute caps on

risk is excessively paternalistic, provided that volunteers have given free and informed consent (Eyal 2020, 148).

Conversely, there have been a number of researchers who have suggested various upper limits on risk for a variety of reasons. The most apparent of these, referenced in aforementioned U.S. law, is that of minimal risk. While not defined in the strict sense, The Royal College of Physicians provides two interpretations of what minimal risk entails. The first interpretation states that minimal risk may involve a high likelihood of harm occurring, but that the harm itself be mild. The second stipulates permitting no more than a, "very remote possibility of injury or death" (Hope & MacMillan 2004, 111). This interpretation aligns with the intuitive notions of permissible risk described above but does not itself provide a determination of when a harm exceeds 'mild', or what constitutes a 'remote' possibility. Some attempts defining these risks in more detail have stated that minimal harms should not exceed those faced in clinical evaluation, everyday life, or charitable participation (Rid 2014, 74).

Deviating from minimal risk, there are other attempts at establishing a maximum risk threshold along a more definable set of criteria. One such method compares the risks of volunteering in medical research to those faced by a volunteer firefighter, claiming that such an interpretation allows us to expose volunteers to risks beyond those faced in everyday life while still coalescing with demonstrated notions of permissible risks that may be consented to (London 2007). Another interpretation states that comparing volunteering in medical testing with the risks faced in live organ donation not only allows us to exceed minimal risk in a similar manner to the volunteer firefighting example, but that such a comparison also more accurately captures the relationship between patients and medical practitioners (Miller & Joffe 2008). Yet another interpretation disagrees with the notion of comparison to other scenarios whatsoever on the

grounds that social acceptance and morality are distinct concepts, instead stating that a 1% chance of serious injury is a threshold that protects both research participants and the interests of those who stand to benefit from medical experimentation and should therefore constitute an acceptable risk threshold (Resnik 2012).

The debate regarding which of these interpretations, if any, should be relied upon for instituting or rejecting an absolute upper limit on risk faced by volunteers is ongoing. While there are a number of reasons for this, a lack of consensus is partially characterized by concerns regarding public trust. Many ethicists who argue for an absolute limit on risks faced by volunteers do so because they are concerned that the perception of research by the public as excessively risky could interfere with the continuation of research in the future. There is a worry that challenge trials inherently entail such a perception, which in turn invokes assertions that the risks they entail should be limited to placate the public (Bambery 2016, 98). Others argue that, "if HCTs are independently unethical, that alone will be reason enough to oppose them. If that truly is the case, it will not matter whether they exacerbate public mistrust or not, as they should be prohibited due to their independent impermissibility alone." (Eyal 2022, 2). While the discrepancy between perceived and actual risk will be discussed in further chapters, there is widespread recognition that challenge trials can be designed in line with our (admittedly abstract) notions of permissible risk, provided they only expose volunteers to diseases which are self-limiting or treatable.

It was observed in the previous section that challenge trials offer a number of unique scientific benefits which would weigh favourably in a risk/benefit analysis permitting their use. What must be explored, then, are the unique risks entailed by challenge trials which could be mitigated in alternative trial designs, or the ways in which challenge trials uniquely mitigate risks

involved in alternative designs. In other words, setting benefits aside, comparing the risks present in challenge trials to those in field trials will paint a clearer picture of what a risk/benefit analysis involving the two would entail.

As asserted earlier, challenge trials inherently involve utilizing fewer volunteers to complete research. Requiring fewer volunteers to test a given vaccine candidate means that fewer individuals are exposed to the risks of inoculation and possible side effects. As a result, there is an argument to be made that challenge trials might entail less aggregate risk than field trials because the risks of exposing 40 individuals to a self-limiting or treatable disease in a controlled setting are comparable to the risks of inoculating hundreds or thousands of volunteers with a vaccine candidate carrying unknown potential side effects (Bambery 2016, 93-4). An actual determination of whether or not this is the case will be dependent on numerous individual factors within a given trial (e.g., the type of disease being studied, the nature of the vaccine candidate being used, the results of preliminary testing as a measure of foreseeable risk, etc.). At the very least, these considerations suggest that the possible discrepancy in total risk merits careful consideration in a risk/benefit analysis.

While this consequentialist rationale is important to consider, it must be met with considerations for the wellbeing of individual volunteers that ensure they do not suffer impermissible harms in the name of scientific discovery. The nature of intentional infection, which defines challenge trials, also guarantees with complete certainty that volunteers will be exposed to the possibility of experiencing harm from a given disease, meaning that in this regard they always carry more risk to a particular volunteer than a field trial. In a vacuum, if all other considerations between a challenge and field trial are balanced (i.e., they provide similar results along a similar timeframe), the guarantee of otherwise healthy volunteers risking the possibility

of harm stemming from intentional infection is enough to obligate researchers to conduct alternative designs, i.e., field trials. This leads to the assertion that challenge trials should require a stronger scientific justification to proceed than field trials, where the risk of being exposed to the possibility of harm from a disease is determined primarily by the rates of transmission in a given area (Bambery 2016, 99).

However, while challenge trials inherently involve a higher risk of possibly suffering harm resulting from disease exposure, this does not necessarily mean they involve a higher risk of suffering *severe* harms from exposure. Therefore, depending on which conception of risk limitation one adopts (if one is adopted at all), there are a number of reasons why conducting a challenge trial may actually be considered less risky overall. Utilizing minimal risk, for example, a challenge trial acting in accordance with stipulations restricting the use of diseases causing, "irreversible, incurable or possibly fatal infections" (Bambery 2016, 98) could better ensure that harms do not exceed 'mild' due to a number of features inherent in their design. These design features will consequently be explored from an ethical perspective to determine the unique risks and mitigations they provide.

2) Ethics and Study Design

As discussed in the scientific benefits section of this chapter, the two unique design possibilities afforded by challenge trials are the use of attenuated disease strains and the ability to conduct a wholly inpatient study. The goal of this particular inquiry is to determine whether or not either of these aspects of challenge trial design are morally advantageous in addition to being scientifically advantageous.

Starting with attenuated strains, the most apparent ethical benefit they afford over wild-type strains is the relative safety of infection experienced by volunteers exposed to them. Because

sufficient data is necessary for the production of adequate scientific results, a set number of people will necessarily have to face infection, regardless of whether it is incurred naturally or deliberately. As such, there is an argument to be made that intentionally infecting individuals with an attenuated strain of a given disease is ethically preferable to allowing the same number of individuals to face a more severe, yet natural, infection caused by a wild-type strain (Jamrozik 2021, 31). Not only would this be relevant when considering the severity of infection itself, but strains may additionally be attenuated to be receptive to certain treatment options that wild type strains might resist, which would make providing aid to volunteers much more feasible.

Were each trial being compared to the absence of conducting vaccine trials altogether, this might be met with the retort that individuals in a field trial would have faced infection anyways, so there is no ethical dilemma involved in observing their infection, whereas a challenge trial infects individuals who would have been otherwise healthy for the purpose of such observation. However, the clarity of this argument diminishes when one compares conducting one trial directly to conducting the other, as the initial premise (that individuals in a field trial would have faced infection anyways) is no longer necessarily true. Expedited results stemming from a challenge trial could potentially precede the infection that individuals would have acquired in a field trial. In other words, a sufficient sample size achieved through a challenge trial could lead to the production and distribution of an effective vaccine which prevents severe infection before field trials have finished acquiring their results, meaning that those who would have been observed in a field trial are able to prevent their infection before it otherwise would have occurred (Bambery 2016, 93-4). As such, while the objection presented regarding intentional infection is still a cause for consideration, and one that will be explored further, it is one that does not apply categorically or with absolute certainty.

Shifting focus to the inpatient nature of challenge trials, the ability to intervene and mitigate infection the moment symptoms become severe enough to prompt intervention is ethically significant and deserves attention in a risk/benefit analysis. As volunteers are being consistently monitored by researchers, their wellbeing is consistently measurable, and medical attention can be provided the moment concerns for patient wellbeing exceed an acceptable threshold (Jamrozik 2021, xii). Conversely, individuals who acquire infection naturally may not be able to receive necessary medical attention in an appropriately timely manner. In conjunction with the potentially increased level of harm entailed by a wild-type infection, the volunteers infected in a field trial could likely experience a worse state of affairs than they would have if they were involved in a challenge trial. As both trials require sufficient sample sizes of infection in order to proceed, overseeing 40 intentionally infected individuals could be less risky and harmful than allowing 40+ individuals to face unsupervised natural infection.

Unfortunately, inpatient observation is generally considered to be excessively burdensome for volunteers. By requiring volunteers to remain under observation for the duration of infection, researchers are consequently asking volunteers to make a number of sacrifices. Employment, for example, could be jeopardized, and any dependents that volunteers are responsible for would need alternative care (Jamrozik 2021, 53). While the financial burdens incurred in these cases would likely be compensated for, the practical burdens entailed by missing work for an extended period of time, or being unable to care for one's children, are still present, and therefore must be accounted for (Jamrozik 2021, 53-4).

Stemming from the burdens incurred by inpatient trials, there is the additional dilemma of the right to withdraw from a trial which is revoked by requiring volunteers to remain inpatient for the duration of their infection. Intentional infection is largely justified on the basis that a

patient under constant observation would likely receive better care in the event that such an infection becomes unjustifiably harmful. As a result, it is unclear what the ramifications would be if a volunteer recuses themselves from a trial before their infection has concluded, and what the ramifications on culpability would be if they were to consequently suffer impermissible harms from this infection. Additionally, if an individual is infectious at the time of departure, it is possible they would spread their infection to third parties. This third-party risk will be discussed shortly (Jamrozik 2021, 55).

Both of these components of trial design entail possible personal benefits acquired through intentional infection, which itself is ethically significant. If these trials are being conducted in areas where the disease of interest is endemic, it can be argued that the chances of facing natural infection are higher than the odds of not facing natural infection. If this is the case, then the concerns present from infecting an 'otherwise healthy' volunteer do not entirely hold, because it is unlikely that they would have been otherwise healthy (Bambery 2016, 99) In this case, it would likely be preferable to face milder infection from an attenuated strain of the disease and be monitored throughout infection, thereby acquiring a level of immunity either through the vaccine candidate or through the infection itself, than it would be to risk natural infection. As a result, a challenge trial confers personal benefits in ways that a field trial cannot, a fact which is significant in a risk/benefit analysis.

3) Moral Significance of Intentionality

One pervasive question which surrounds intentional infection regards the moral significance of intentionality itself. Comparing the potential state of affairs of a volunteer in a challenge trial and a healthy individual would suggest that being responsible for bringing about the possible harms occurring in the former is intolerable. A volunteer in a field trial merely

consents to the risks of injection and novel vaccination, so they would be considered otherwise healthy, which supports the above assertion that there is a *prima facie* responsibility for researchers to conduct field trials over challenge trials when human testing is necessary. However, the nature of benefits afforded by challenge trials also prompts the comparison of a volunteer in a challenge trial to an individual afflicted by a disease that would have otherwise been prevented, had a vaccine been developed faster as a result of a challenge trial. Because the same number of people have to face infection no matter what trial design is used to achieve an adequate sample size, there is an ethical preference that these infections are attenuated rather than wild-type, and there are also advantages offered by inpatient observation when compared to outpatient infection. All of these advantages devalue culpability as a special moral consideration in a challenge trial. In other words, while responsibility may bear weight in, for instance, assigning blame for negligence, whether or not one is responsible for causing intentional infection would largely be irrelevant as a factor in assessing the strict risks and benefits of conducting a challenge and/or a field trial. Challenge trials would therefore not be considered wrong in and of themselves on the basis of intentional infection (Hope & Macmillan 2003, 115).

Though intentional infection may not be considered inherently wrong, there are two instances where being responsible for causing infection in an otherwise healthy individual could potentially deem a challenge trial unnecessarily risky. The first of these regards the possibility of third-party infection. Third party infection occurs when a volunteer in a challenge trial becomes a transmission vector for a disease, infecting another individual who was not involved in experimentation. This is primarily unethical as a violation of stringent stipulations of informed consent present in any medical experimentation involving humans. While volunteers provide informed consent to face the consequences of infection, third parties do not. Because informed

consent is a necessary condition for conducting medical research, an experiment which causes third party infection is highly unfavourable (Lynch 2020, 924). Additionally, as stated previously, the balance of benefits and harms incurred by both challenge and field trials relies in part on the number of individuals each trial exposes to harm. The more individuals that become infected with a disease as a result of a challenge trial, the less justifiable it becomes to conduct such a trial. However, while this concern is certainly important, an inpatient challenge trial can mitigate any and all risks of third-party infection through proper quarantining procedures, and there is a presumption that the presence of third-party risk would necessitate an inpatient trial, so we can ignore this concern on the assumption that it is entirely preventable.

Public Trust

The second instance where the effects of intentional infection might be inherently wrong is the way it could be perceived by the public, and whether or not the possibility of negative public perception is a relevant risk which could carry potential significant harms towards the institution of medical research. This consideration will be the focus of the paper moving forward. While, as we have seen, challenge trials involve a number of other ethical considerations, for example testing with non-treatable or non-self-limiting diseases and what constitutes valid consent, these considerations have mostly either achieved some level of consensus or have at least been sufficiently explored to a degree that their lack of consensus is established. Public trust, on the other hand, has been asserted as an important concern, but until recently has not explicitly been explored as it relates to challenge trials in particular (Eyal 2022).

As stated at the beginning of this chapter, there are both intuitive and historical reasons for presuming a negative public reaction to harms resulting from medical testing. This presumption is consistent with the assertion made by some that the uncertainty inherent in medical testing

makes it preferable to err on the side of caution when conducting medical research on human volunteers, even in light of potential scientific benefits (Jamrozik 2021, 51). The worry held by researchers is that, should a volunteer experience harm as the result of intentional infection in particular, the historical abuses perpetuated by the medical research community combined with an instinctual revulsion to the concept of intentional infection would lead the public to lose faith in scientific institutions (Bambery 2016, 93). As a result, future medical research may be restricted, and the mistrust fostered as a consequence of harm occurring in these trials could lead to more widespread adoption of anti-science (and particularly anti-vaccine) rhetoric, which itself could lead to future harm caused by a lack of herd immunity or the overwhelming of hospital resources.

These claims appear quite frequently in discussions surrounding the permissibility of challenge trials and will be explored in more depth next chapter. While they seem to be intuitively accurate, the literature on this subject is underdeveloped in spite of the frequency of its invocation. Within this debate, there are two assertions which demand exploration. The first regards the relevance of harms stemming from public mistrust in medical testing and scientific institutions. The lack of clarity surrounding what constitutes the public, trust, and harm towards the institution of science has been both largely speculated on and inconsistently yet regularly utilized. As such, the relevance and application of public trust in a risk/benefit analysis remains unclear; does the risk we refer to in a risk benefit analysis extend to public confidence in science and the possibility of harms stemming from its degradation? It seems as though some ethicists either implicitly or explicitly believe so (Eyal 2022). The second assertion stipulates that the potential harms from public trust are uniquely relevant to challenge trials, over and above other trial designs. If the harms which stem from public mistrust are deemed relevant, they must be

shown to be especially relevant to challenge trials in order to properly qualify as a risk which could independently factor into a risk/benefit consideration between challenge trials and alternative study designs.

This paper has, until this point, explored the larger debate surrounding challenge trials, their relative permissibility, and the instances in which their use would be considered and either endorsed or discouraged. The following chapter will explore these concerns through a careful analysis of the existing and prominent literature on challenge trials. In particular, the following chapter will examine this literature for invocations of public trust and will assess the strength of different justifications used to assert that challenge trials are likely especially damaging to this trust. Ultimately, the goal of this inquiry will be to understand what specifically is meant by the term 'public trust', to what degree it should be considered in trial design selection, and whether or not a definition of the term is consistently adopted by ethicists discussing challenge trials.

---

[i] I think it must be an inclination. Either there is no intuitive inclination in the literature that the idea of intentionally infecting someone is inherently uncomfortable, and therefore the argument that public trust would be especially damaged if something were to go wrong as a result of a challenge trial would not hold or carry nearly as much weight; or, it is an inclination, and we have to address it as such

---

## Chapter Two: Analyzing Uses of Public Trust in the Challenge Trial Literature

At the end of the previous chapter, I examined the idea that public trust in the institution of medical research is more severely damaged by conducting challenge trials than less controversial trials because of a presumed negative public perception of their design. Such a conclusion is one which appears frequently in literature discussing their permissibility (Eyal 2022, 4). When taking this stance, authors often assert the need for challenge trials to adopt special requirements intended to counteract a loss of trust and bolster public confidence in medical research (Hope & MacMillan 2003, 110). Occasionally, it is implied that, wherever possible, alternative trial designs should be utilized to protect public trust (Eyal 2022, 5). Even if this is not stated directly, the assertion that challenge trials pose a risk of harming public trust in ways alternative designs do not makes this implication clear; if true, it would require challenge trials to provide additional scientific justification and/or risk mitigation measures in order to shift a risk/benefit analysis in its favour (Rid & Roestenberg 2020, 762).

Protecting public trust is an important part of conducting medical research. Given the previously explored history of vaccination research, it is apparent that violating ethical criteria (and consequently the trust of the public) in pursuit of scientific gain carries generational consequences (Harris 2005, 242). However, in order to properly protect this trust, it needs to be properly understood. The concern which characterizes the central thesis of this chapter is that public trust, as it exists in contemporary challenge trial literature, lacks an adequate definition that is consistently used in the literature, leading to inconsistent applications of the term. As a result, objections built on public trust concerns as they currently exist are not persuasive.

The chapter begins by briefly analyzing the term 'public trust' as it is used in the broader research context, with a focus on understanding the possible reasons for its inconsistent application in the challenge trial literature. I then construct and examine three possible approaches to evaluating public trust in trial design selection. I call these (1) the intuitive approach, which automatically includes public trust in a risk/benefit analysis under the presumption that challenge trials are more likely than other trial designs to damage public trust, (2) the evidential approach, which aims to implement public trust in a risk/benefit analysis based on evidence revealing the potential backlash conducting a challenge trial would have within particular contexts, and (3) the deontological approach, which proposes that upholding public trust is an intrinsically valuable necessary condition for conducting a trial, over and above a risk/benefit analysis. These approaches will then be contrasted by Nir Eyal's approach, which argues that public trust concerns should be independent from a risk/benefit analysis altogether. The chapter will conclude by arguing that the evidential approach is the most appropriate for developing a functional understanding of public trust as it applies to challenge trial design selection policy, but that existing evidential approaches have thus far been insufficient in this exercise due to the lack of clarity regarding the terms 'public' and 'trust'.

## Scientific Research and the Public Trust

David Resnik's "Scientific Research and the Public Trust" provides a strong foundation from which to examine what exactly is meant when the term 'public trust' is utilized in scientific research. Often, he states, there is a justified assertion that maintaining public confidence in scientific institutions is an important consideration. However, he believes that public trust has not been properly defined in medical research literature. As such, it becomes difficult to consistently adhere to any directives such a concept might influence, especially in the

development of policy surrounding what kinds of scientific research can be ethically conducted in which instances.

One main reason for this lack of consistency, Resnik proposes, are the different interpretations of both 'public' and 'trust' used in the literature. He begins this inquiry by attempting to define trust. While I define trust more thoroughly in chapter 3, a brief overview of Resnik's interpretation provides helpful context for this chapter. Trust, he cites, "is a relationship between or among people" (Resnik 2011, 2). Its purpose is to, "facilitate cooperative social interactions […] which depend on shared expectations of behavior"; it, "involves risk taking", and "can generate ethical and legal duties" (Resnik 2011, 2-3). One chooses to trust another, "because they judge them to be trustworthy" (Resnik 2011, 3), a point which will be particularly salient when discussing the trust placed in the institution of scientific research by the public in the following chapter.

Beyond these general definitions, developing a functional account of trust becomes practically complex and dependent on context; it might be implicit or explicit, it might be concrete or abstract, and it might be between individuals, groups, or individuals and groups. The literature provides examples of more well-defined instances of trust in scientific research, beginning with the trust that exists between researchers and/or research institutions. When collaborating on research, scientists trust that they will receive the credit they are due, and that their research will be treated with fairness. There is also the trust that exists between researchers and volunteers, as volunteers in a trial take a risk by placing themselves in a medically vulnerable position, and they must therefore trust that researchers will act ethically and in their best interest. Additionally, there is the trust that exists between research institutions and institutions which provide funding to them (Resnik 2011, 3-4).

Though important in their own right, Resnik does not believe these trust relationships sufficiently capture the concept of 'public trust', i.e., the trust that exists between the 'public' and the institution of science. He states that authors contributing to the literature on public trust likely equate the 'public' with society as a whole, and that the existence of this trust is, "so obvious that it hardly requires justification" (Resnik 2011, 4). This point is supported by empirical studies referenced by Resnik, which display that the majority of the public expects scientific research to be, "honest, reliable, and objective", i.e., the public trusts science to operate in this manner (Resnik 2011, 6). The concern put forth by scholars is that conducting scientific research which is negatively perceived by the public could lead to the degradation of this existing trust which is so heavily relied upon for the continuation of highly important research.

Resnik points out that the problem with this understanding is that 'society' is made up of a number of different 'publics' which means that deriving ethical or policy implications from public trust in its current formulation entails contradiction (Resnik 2011, 5). He proposes that this is because, "different people in society may have different expectations of science, and therefore place different kinds of trust in science." (Resnik 2011, 5). Some, for example, might trust that science will provide life-saving medication as soon as possible and without unnecessary delay, whereas others might expect researchers to only distribute medication that has been thoroughly tested with the utmost stringency and no sooner. He invokes the point that the 'public' as it has been utilized actually refers to two groups: "the general public" and, "a specific public, such as a particular community or group" (Resnik 2011, 7). Examples of what constitute a particular specific public include 'the immediate community', low middle-income countries, and the scientific public (Jamrozik 2021) (Kerasidou 2018). The conflicting expectations described above would likely exist as a result of deriving public trust from either a

general and a specific public, or two different specific publics, each with contrasting concerns. As such, Resnik suggests that the development of policy which exists on the basis of satisfying the expectations of the public should explicitly reference the public it is concerned with (Resnik 2011, 7-8). As it stands, the continued failure to adopt this measure not only risks developing policy on insufficient grounds, but also risks the further degradation of a shared understanding of what protecting public trust means.

In the process of explaining that public trust has not been adequately defined and thus inconsistently applied, Resnik implies that public trust *should* be considered in the development of policy influencing trial design selection in *some* way. However, he does not explain how such a consideration would function, only that it has not yet functioned adequately. In particular, the implications of his analysis have yet to be considered in the challenge trial literature. Authors discussing challenge trial design ethics have instead relied on a number of different conceptualizations of public trust and how it should function in trial design selection, none of which define a consistent public before establishing their expectations. These accounts range from asserting public trust is inherently at stake when a challenge trial is conducted and thus alternative designs should be used whenever possible, to asserting that public trust should not influence trial design selection whatsoever.

To examine why existing accounts of public trust and trial design selection have thus far been inadequate, I have constructed three ways to categorize public trust as it currently presents itself in the literature. I will begin with the intuitive approach of public trust, which presupposes the legitimacy of the claim that public trust is damaged by the conduct of a challenge trial. I will then move on to the evidential approach, which states that measures should be taken to limit a loss of public trust resulting from conducting a challenge trial only when there is evidence that

such a loss of trust will occur, before examining the deontological approach which designates the maintenance of public trust as a necessary condition for conducting medical research.

The Intuitive Approach

The intuitive approach to public trust and challenge trial research is the most common one adopted in the literature. The approach proposes, "it is widely recognized that CHI studies—although not fundamentally different than other clinical research—require special ethical attention and efforts to maintain public trust" (Rid & Roestenberg 2020, 750). Defenders of this approach make this claim because they believe that human challenge vaccine trials are, "likely to raise more concerns among the public than most other types of medical research" (Hope & MacMillan 2004, 110), and as a result they should be regulated more stringently than other forms of medical research involving healthy human subjects (Hope & MacMillan 2004, 116). Regulations intended to protect public confidence proposed by this approach often include independent expert reviews, publicly available justification for trial design selection, and limitations on challenge trial conduct (Bambery et. al. 2016, 94).

The justification for imposing additional regulations rests on the assertion that "the only sound argument" for limiting healthy, properly consenting, and sufficiently compensated adults from taking on the otherwise acceptable risks of any form of medical research is if such research would have a significant negative impact on public confidence (Hope & MacMillan 2003, 113). To use an extreme example, one author states that, "[i]f significant numbers of people were to die as a result of taking part in medical research, then this would be likely to have the effect of bringing such research into disrepute," which "would […] reduce the amount of research that could take place because of a public reaction against such research" (Hope & MacMillan 2004, 113). Such a violation of public confidence would likely have much larger negative effects

towards the broader medical community, so even if research which weakens public trust could be otherwise ethically conducted (i.e., would not kill significant numbers of people) it should not be or should require stronger scientific justification.

The assertion that challenge trials are more likely to incite public concern than other medical research involving humans is typically predicated on both the contentious history of vaccination research and the intuitive discomfort presumably held by the public towards intentional infection, which will be discussed respectively. Chapter one provided an overview of vaccination research involving human subjects, particularly research involving intentional infection, displaying a long and complicated history of ethical abuses. Many authors argue that this history creates a foundation of public hesitancy towards challenge trial research, and that continuing to conduct this research only serves to entrench the public's perception of the medical community as unconcerned with the safety of volunteers. These authors claim that, due to this history, the public would likely be surprised to know intentional infection for the purpose of medical research still takes place (Bambery 2016, 92).

An extension of this assertion proposes that the public holds a negative view of intentional infection on the basis that the concept itself is 'alien to their expectations', resulting in widespread mistrust towards the institution of science should they be conducted (Hope & MacMillan 2004, 116). Authors proposing this view make the claim that, "because it seems counterintuitive for clinical researchers to deliberately harm participants by causing disease […] CHI studies can face public criticism and potentially cause distrust" (Rid & Roestenberg 2020, 750). This claim, compounded by the previously mentioned history of ethical abuses committed by the medical research community, constitute sufficient reason for authors to assert the necessity of the above proposed regulations to counteract a loss of trust.

To assess the intuitive approach, the truth and impact of this claim must be proven. While the contentious history of vaccine research is established fact, both its impact on current public opinion and the idea that intentional infection is alien to public expectations are debatable. To preface, we are not focused on whether or not intentional infection is unethical (as chapter one established that is not the case) but if a significant portion of the public would wrongly perceive it to be unethical, and how the widespread adoption of this perception would negatively impact the institution of science.

The most evident problem with qualifying the intuitive approach is a lack of evidence supporting the notion that the public is opposed to challenge trial research on the basis of either history or intuitive discomfort. In fact, the very studies which assert the existence of a publicly held aversion to challenge trials also state that, "[facing criticism is] by no means unique to CHI studies, and to date it appears that CHI studies have not exposed participants to excessive risks, led to the unintentional spread of pathogens, or generated excessive public criticism." (Rid & Roestenberg 2020, 750). Studies used in the challenge trial literature actually tend to display public support for the study design, whereas author accusations that challenge trials are alien to public expectations of science are largely unsupported (Eyal 2022, 4). At best, if supporters of the intuitive account reject studies displaying public support for challenge trials for whatever reason, public opinion of their design would be unknown, and we are back to relying on suppositions. At worst, their determinations of public opinion are inaccurate, as would be their conclusions regarding the damage conducting a challenge trial causes to public trust.

The intuitive approach's reliance on supposition becomes evident when the language used by its proponents is examined. For example, the quote used at the beginning of this section stated that it is 'widely recognized' that challenge trials 'can' face criticism. Not only is this

statement relatively meaningless when properly evaluated, as most any trial *can* face criticism, warranted or otherwise, but it serves to imply an increased risk of backlash without actually asserting whether or not this risk of backlash even exists. Justifying these criticisms on the basis that they are 'widely recognized' does not actually assess the legitimacy of public trust criticisms of challenge trials, but merely their popularity, further revealing the suppositional nature of these claims. Additional sources point to challenge trials being a "potential concern" with the "potential to spark backlash", neither of which distinguishes challenge trials from other trial designs in the ways these authors imply (Bambery 2016, 92-4).

The quote above discussing the hypothetical impact of numerous deaths resulting from a challenge trial to justify necessitating additional regulations for them on the basis of preserving public trust makes this supposition all the more evident. It implies causality between actual *harms* stemming from medical research and the deterioration of public trust (i.e., if many people died then the public might mistrust science) rather than claiming that relevant public mistrust results from the *public perception* of otherwise permissible harms. If there were impermissible harms occurring, up to and including multiple deaths, this alone would suffice to condemn nearly any trial (Eyal 2022, 2). Examples which attempt to centralize public trust concerns without invoking harm as the basis for public mistrust, such as the Alder Hey scandal[i], are a) independently unethical for alternative reasons, and b) not challenge trials, or even trials involving healthy human subjects (Hope & MacMillan 2004, 113). As it stands, there seems to be little indication that conducting otherwise safe challenge trials violates the expectations of the public.

One factor complicating this assessment is that it is impossible to assign a set of expectations to an unidentified public. By universally asserting that public trust is more likely

than other trials to exacerbate public mistrust, or by omitting a public of interest in their assertions, authors supporting the intuitive approach imply that intentional infection violates the expectations of medical research held by the general public. To illustrate how this may be problematic, while intentional infection might defy the expectations of certain members of the general public, conducting a field trial that takes longer to acquire results than a challenge trial without providing publicly available justification for this decision (i.e., if a challenge trial is too risky in this instance) could foreseeably violate the trust of a much larger portion of the general public who expects the institution of science to acquire results as soon as ethically possible.

This failure by authors to address which public they are referring to forces them to claim that public trust concerns are applicable only where the possible harms faced by volunteers described above would actually alter public perception and negatively impact the institution of science (Hope & MacMillan 2004, 113), which conversely implies they are discussing a specific public. By stating that this principle only applies when there is an actual risk of public mistrust harming the institution of science, and by consequently claiming that minimal risk can be exceeded with sufficient justification (Hope & MacMillan 2004, 113), these authors do not assert anything beyond the unqualified inclusion of public backlash in a risk/benefit analysis, and only when it is relevant. Because the relevance of such backlash cannot be determined prior to the conduct of a trial without identifying a public of concern[ii], saying that public trust is a concern only when public trust is a concern is far too many words to say nothing at all.

Seeing as we are developing a functional account of public trust, we need to assess not only the legitimacy of using a particular interpretation of public trust as it relates to challenge trials, but why using it could be problematic. Doing so requires understanding the nuances of a risk/benefit analysis. Conducting an analysis for the purpose of selecting a trial design involves

two classes of judgements; relative judgements, in which the benefits of a particular trial are weighted against its risks, and comparative judgements, which measures a trial against alternative designs (Rid & Roestenberg 2020, 762). Regarding comparative judgements, which is the focus of trial design selection policy, "any risk of public controversy" needs to be outweighed by sufficient scientific justification in order to choose a design which entails that risk (Rid & Roestenberg 2020, 762). However, without identifying a public, it becomes impossible to determine what the risk of public controversy is, and therefore the level of scientific justification needed to outweigh it. As a result, it seems likely that policy regulating trial design selection which relies on the undefended supposition that risks of public controversy are both present and sufficiently impactful would unjustifiably favour less controversial trials, leading to their conduct when more beneficial but seemingly riskier trial designs would be more appropriate.

The evidential approach differs from the intuitive approach because it stems from interpretations of public trust which attempt to avoid the suppositional issues of the intuitive approach. Supporters of the evidential approach do not presume the impact of challenge trials on public trust in a risk/benefit analysis, but neither do they propose to ignore it. Instead, they demand evidence on a case-by-case basis for determining 1) if conducting a challenge trial would damage public trust, and 2) how doing so would harm the continuation of research and adoption of public health measures. From this, appropriate policy can be developed which respects volunteer and public interests while maintaining public confidence in the institution of science.

The Evidential Approach

The evidential approach does not imply that challenge trials are more likely than other forms of medical research to disrupt public trust. However, it does rely on the claim that, "[t]he concept

of conducting research in which healthy volunteers are intentionally exposed to pathogens which can cause infection, and in some cases disease, can appear ethically counterintuitive - particularly when natural infections with such pathogens can lead to severe adverse outcomes, including death" (WHO 2021, Preface). This allows authors to claim that intentional infection *might* harbour conditions which could give rise to harmful public mistrust without presupposing this to be the case. As such, research ethicists are able to worry about the possibility of damaging public trust by conducting a challenge trial but only when there is adequate justification for thinking so and when such a loss of trust would result in meaningful backlash, avoiding some difficulties of the previous sections.

One instance where an evidence requirement is apparent is in the, "WHO Guidance on the Ethics of Controlled Human Infection Studies". While deconstructing the ethics of challenge trials, the WHO reports that tens of thousands of volunteers have been enrolled in them over the last few decades, and of these, there have been very few incidents of serious harm or death occurring (WHO 2021, 2). For this and other reasons, they state that there are no "morally compelling reasons" to distinguish modern challenge trials from alternative trial designs, instead stipulating that they only need follow existing ethical protocol for research involving healthy human volunteers (WHO 2021, 4). However, one of these existing protocols regards responsible stewardship of the institution of science, which involves maintaining public trust.

To this point, the guidance references the attention garnered by recent COVID-19 challenge trials as evidence that trust might be in jeopardy, and because of this they are ultimately able to justify imposing additional regulations on them. They cite the global landscape of public trust as 'complex' and use this complexity to assert that challenge trials should 'carefully consider' trust building approaches such as community engagement and the utilization

of independent ethics review boards. The WHO is consequently able to assert that maintaining public trust might very well be relevant in a particular instance of trial design selection, and it is up to researchers to discover if and how much that is the case. This account is therefore valuable in that it suggests a need for more empirical work looking into public perceptions of challenge studies, a need made evident in the previous section. However, a lack of a consistent 'public' utilized in their arguments makes assessing this account difficult. The text references the 'complex global landscape' that influences the acceptability of trials, suggesting a general public, but then brings up 'consultation and engagement activities with the public' (WHO 2021, Preface) as a risk mitigation measure, which implies a more specific, communal public.

Other authors have a similar approach the issue of public trust and trial design selection. They suggest that the implementation of independent ethical review boards and strong community consultation are important measures for preventing any possible deterioration of public trust (Jamrozik 2021, xi). Very rarely, if ever, do they make a determination on the degree that negative public perception acquired through this engagement should affect the implementation of challenge trials; they instead state that, "weighing the potential benefits and burdens associated with HCS requires […] learning from local communities regarding local priorities and the public acceptance of potential research designs" (Jamrozik 2021, 3). While a regular insistence on the use of communal engagement and review boards implies that it is a specific public being referred to throughout the paper, this particular paper also discusses the ability for these trials to undermine trust "domestically and internationally" (Jamrozik 2021, 88), suggesting there may be a discrepancy in publics used.

By not declaring a specific public of concern in trial design selection, the evidential account is able to assert that public trust should potentially, but not automatically, be considered

in a risk/benefit analysis, without actually establishing how this should be done. While authors suggest consulting with the local public to determine what expectations they hold of medical research, they also suggest monitoring online, global attention for evidence that conducting these trials violates public expectations. Additionally, the measures they propose for mitigating a loss of public trust are 1) community engagement and 2) independent ethical review boards. The first is inherently directed towards a local public which forms a community either geographically or situationally. The second is only relevant towards public trust insofar as its goal is to provide an accurate, unbiased account of which trials should be undertaken in which instances, an exercise that is intrinsically valuable and only satisfies public expectations as a result.

The evidence for proving that a loss of public trust is a problem is therefore not guaranteed to be consistent with proposals for mitigating such a loss. For example, engaging with the immediate geographic community does little to influence an online skeptic, so including the opinion of the skeptic as evidence for public mistrust in challenge trials might stand to harm this immediate community who would have otherwise benefitted from acquiring a vaccine sooner. In its current formulation, the evidential approach is simply not useful for developing a functional account of public trust as it relates to challenge trials. This begs the question of whether or not including public trust in a risk/benefit analysis for trial design selection is a useful endeavour at all.

The Deontological Approach

The deontological approach attempts to solve the above dilemma of public trust's role in a risk/benefit analysis by maintaining that, while public trust is an important consideration in trial design selection, it should be excluded from these analyses entirely. Instead, this approach proposes that maintaining public trust is a necessary and independent condition for ethically

conducting medical research. Both the intuitive and the evidential approaches to defining public trust as it relates to challenge trials understand the term in a consequential sense; damaging public trust is bad because it results in poor consequences. The deontological approach, on the other hand, takes the position that maintaining public trust is inherently valuable, and should not be infringed upon or damaged under any circumstances. As such, whether or not a trial damages public trust would entirely supersede a risk/benefit analysis altogether.

In order to make such a claim, this account takes respecting the trust of the public to be the goal of medical research, meaning that attaining scientific benefits and prioritizing the welfare of volunteers is only important insofar as it helps us achieve this goal. Unsurprisingly, this approach is rarely defended in the bioethics literature. If public opinion was infallible, or if a risk/benefit analysis formed by an independent ethics review board was consistently more flawed than public opinion, prioritizing public opinion might make sense, but only because adhering to it would produce the most benefits and involve the least risk, not because it respects the opinion of the public. Without a hypothetical infallibility or superiority over an independent review board, the impracticality of this approach is only made more evident.

Adhering to the deontological approach also entails accepting unfavourable consequences. For one, as observed in the intuitive approach, harm resulting from medical research generally fosters public mistrust in it as an institution. This creates a paradox, where respecting the public's perception and selecting a given trial design by strictly adhering to their opinion rather than consulting an independant risk/benefit analysis will likely result in volunteers facing more risks, causing harm and thus damaging public trust (in instances where public opinion deviates from this analysis). As a result, maintaining public trust in the long term might have to involve damaging it in the short term, a point which will be particularly salient in chapter 3, and which

cannot be reconciled by a strict deontological account. Additionally, this necessary condition of upholding public trust would have to hold for all trial designs, not just challenge trials. As chapter 3 will establish, the public may not always grasp every nuance of medical research, and they may place or withhold their trust in medical research very broadly because of the actions of an individual or local organization (McDonald et. al. 2008, 38). If the public sufficiently and generally distrusts medical research, it would mean that no trial designs could be selected at all, and that the ethical trial design choice would be to simply allow disease to spread unmitigated, which would ironically eliminate any credibility held by the institution of medical research.

Finally, we are still left without a way of determining which public is of relevance, nor a way to determine their opinions. This means the deontological account suffers the problems of the first two accounts, in that the impact challenge trial design selection has on public trust is still suppositional because no public is identified. From a functional standpoint, then, the inclusion of public trust as a concern which influences trial design selection seems too abstract to be considered, regardless of its importance. The account provided by Nir Eyal agrees with this conclusion and defends ignoring public trust in ethical considerations for trial design selection altogether.

Eyal's Response

Nir Eyal's paper, "Research Ethics and Public Trust in Vaccines: The Case of COVID-19 Challenge Trials", takes the position that the need to protect public trust is, "very real, but that anchoring research ethics to this goal is unwise and unsupported" (Eyal 2022, 1). Instead of considering public trust as an independent concern regarding trial design selection, he proposes that a strict risk/benefit analysis should be adhered to, as such a measure best protects the interests of both volunteers and the public. Using COVID-19 vaccine challenge trials as a case

study, he attempts to demonstrate that allowing the protection of public trust to influence a comparative risk/benefit analysis is ethically dubious and functionally unnecessary.

To begin, Eyal uses data from COVID-19 challenge trials in the UK to establish that, from the standpoint of risk faced by volunteers *and* potential societal benefit, they were likely preferable to field trials (Eyal 2022, 1). Even though the risks imposed by these trials were considered safer than those imposed by field trials, and even though conducting a challenge trial in this context provides tremendous scientific benefit, they faced scrutiny from scholars on the basis of protecting public trust. These concerns of 'exacerbating mistrust' as a result of conducting a challenge trial are best represented in what Eyal calls the strong formulation of the mistrust argument, which closely resembles the intuitive approach.

This version of the mistrust argument follows from the weak formulation, which supposes that intentional infection is itself unethical, a conclusion previously shown to be false. If this were the case, Eyal rightly states, "that alone [would] be reason enough to oppose [it]. If that truly [was] the case, it [would] not matter whether [challenge trials] exacerbate public mistrust or not, as they should be prohibited due to their independent impermissibility alone" (Eyal 2022, 2). Like the intuitive account, the strong formulation proposes that relevant stakeholders would wrongly *perceive* intentional infection to be unethical and condemn otherwise ethical challenge trials as a result (Eyal 2022, 2). If worries about their widespread adoption are true, the resulting loss of public trust might negatively impact vaccination rates, volunteer participation in medical research, and the continuation of vaccination research without regulatory roadblocks. The stance taken by defenders of this version of the mistrust argument claim that the actualization of these consequences would create a worse state of affairs than if a

less controversial but riskier and less beneficial trial were undertaken (Hope & MacMillan 2004, 113).

As observed in the exploration of the intuitive account above, it is unclear how the likelihood of the public perceiving challenge trials as unethical could be accurately determined, and how much this perception would lead to the consequences feared by scholars. Expanding on previous accusations regarding the suppositional nature of the intuitive approach, Eyal claims that existing support for assertions surrounding the prominence and impact of public mistrust is largely speculative. He states that the literature on public trust and challenge trials often invokes public trust as a concern, but rarely attempts to qualify it with legitimate evidence. Such attempts are so rare that, "virtually the only evidence cited is in the study by Dawson et al. They mention: (a) a historic article with unclear connection to HCTs and to dissuading recruitment; (b) the case of Jesse Gelsinger, […] and (c) the 2014–16 Ebola outbreak in West Africa, which did not involve any HCTs" (Eyal 2022, 4).

Conversely, Eyal is able to refer to studies proving that 'the public' (referring to a number of different specific publics) actually supports the conduct of challenge trials (Eyal 2022, 4). Though this does not prove the general public supports challenge trials, it does prove that it cannot be supposed that a relevant group of stakeholders are against them. Even if one were to justifiably claim that a relevant subset of stakeholders perceives intentional infection to be unethical, they would still have to determine the existence and magnitude of the impact of their mistrust on the continuation of medical research and public vaccination rates. Public hesitancy towards intentional infection does not necessarily imply that the negative effects of mistrust will be realized, and making such a claim, "exaggerates how much current refusal to get vaccinated is founded on (perceived) trial ethics qualms" (Eyal 2022, 4).

While discerning the impact of merely conducting a challenge trial on public trust is therefore a dubious exercise, there is concern that the occurrence of a severe adverse event resulting from a challenge trial could more tangibly lead to the possible detrimental effects of public mistrust outlined in the previous sections. Often, the case of Jesse Gelsinger and its resulting impact on the field of gene therapy is used as evidence that public mistrust could create unfavourable outcomes and is difficult to rebuild[iii]. However, Eyal believes that even though an adverse event resulting from a challenge trial could pose a risk towards public trust, the same is true of other trial designs. If a particular challenge trial provides a better risk profile than alternative trial designs, the adverse event argument would likely favour conducting challenge trials, returning us to his original assertion that following a strict risk/benefit profile without considering public trust is the most ethically salient option for trial design selection.

Finally, Eyal discusses the implications of making the decision to conduct a given trial on the basis of public opinion rather than empirical evidence. Because we are relying on the strong mistrust argument, "any resulting public mistrust must be, not because the [challenge trial] is unethical or thwarting proof of safety, but because it is wrongly perceived as unethical or as thwarting such proof. This should raise some doubt about letting mistrust concerns dictate our devotions" (Eyal 2022, 5). In short, Eyal does not believe that policy should be developed on the basis of misinformation, which is what public opinion amounts to when it deviates from a presumably accurate risk/benefit analysis. On the one hand, this makes a lot of practical and ethical sense. It situates the well-being of volunteers as the primary factor in deciding which trial to conduct by only exposing them to risks which have been consented to and justified by their necessity in achieving desirable scientific benefits. Because third party reviews presumably operate on more accurate and current available information than the general public utilizes, it

also likely ensures the best results, as risks may be taken that seem controversial but are actually ethically acceptable. As a result, Eyal argues it would be unethical to deviate from a risk/benefit analysis because that would involve either exposing volunteers to unjustified risk or failing to achieve potential scientific benefits out of an unnecessary overabundance of caution.

On the other hand, it is not entirely clear that Eyal's assertions disqualify the potential of usefully including public trust as a consideration in a risk/benefit analysis. While he convincingly argues that creating policy on the basis of *misinformation* is likely immoral, we are concerned with the possibility of creating a functional account of the actual *impact* of public trust which can be used in a risk/benefit analysis that is both practical and ethical. While Eyal displayed that quantifying and utilizing public trust in a risk/benefit analysis is difficult and has not yet been done accurately, it is unclear if he successfully disproves the possibility that conducting a challenge trial could conceivably result in a loss of public trust in science, and that this possible loss of trust could limit future medical research and/or the adoption of vaccines/medications/procedures.

This conclusion is supported towards the end of the paper where Eyal states, "Do we really want to pander to the public when its potential mistrust is based on factual error, or misguided ethics? Perhaps as a compromise we should sometimes do so, when all else fails" (Eyal 2022, 5). Eyal himself is saying that *if* public trust would be damaged by conducting a challenge trial and *if* this damage is enough to constitute all else failing, then maybe a 'compromise' must be reached. This admits that his problem is not with including public trust in a risk/benefit analysis, but with the ways it has thus far been done. An account which identifies a public and provides evidence of a sufficiently negative sentiment towards challenge trials would consequently necessitate the inclusion of public trust in a relative risk/benefit analysis. As such, the next

chapter will defend a version of the evidential approach which accounts for the need to identify a public of interest when implicating public trust as a concern in challenge trial design selection.

Defining Public Trust: Next Steps

Unfortunately, this examination of the evidential approach as it is currently formulated proves that it is not itself sufficient to develop a functional account of public trust. It does, however, provide us with direction. We can conclude that it is possible for challenge trials to negatively affect public trust in particular cases, and that this trust is important for the continuation of vital medical research and the widespread use of vaccines which is necessary for achieving herd immunity. We can also conclude that including public trust as a consideration in a risk/benefit analysis requires an empirical determination of public expectations, which requires identifying a particular public.

In the following chapter I will address each of these points by developing my own account of public trust called the community engagement approach. I believe that it is important to analyze the 'trust' element of public trust to understand what we are actually attempting to protect when we say we are concerned about public trust, and how we can best protect it. Because trust characterizes a relationship between parties, doing so should also reveal who it is we are interested in when we include public trust concerns in a challenge trial design selection risk/benefit analysis. Alongside the guidance provided by the examination of existing public trust accounts in this chapter, having a concrete understanding of the trust that should be used in this context will place us in a position to create a functional understanding of public trust which can be practically applied to trial design selection policy.

---

[i] The Alder Hey Scandal performed research on organs taken from the bodies of individuals who had not indicated consent to donate their organs, or whose families had not provided this consent

[ii] While determining the expectations of an undefined public is not possible unless one guesses them correctly, surveying and assessing the opinions of a particular public might be feasible, an opinion which forms the basis of the evidential argument and the conclusion of chapter 3

[iii] The case of Jesse Gelsinger refers to an experiment conducted in 1999 in which volunteer Jesse Gelsinger died as a result of participation in early gene therapy testing, which resulted in public backlash that set the field of gene therapy research back by a number of years (Eyal 2022, 4)

## Chapter Three: Trustworthiness and the Community Engagement Approach

In categorizing existing literary conceptions of public trust and challenge trials, the previous chapter revealed the confusion surrounding what constitutes the 'public' in public trust. Even the evidential account, formulated in response to the suppositions prevalent in the intuitive account, uses evidence of distrust in one public to justify imposing regulations intended to bolster the trust of an entirely different public. (WHO 2021, Preface). This means there is still no method for accurately assessing the legitimacy and magnitude of public trust as a consideration in trial design selection, nor is there clarification on which publics are of concern in which instances. I believe exploring trust resolves this dilemma. If trust can be defined, it is because a relationship between publics has been properly characterized, meaning these publics are known.

After analyzing trust, this chapter proposes a strong version of the evidential approach, called the community engagement approach, as an improved conception of public trust as it applies to trial design selection. Rather than gathering evidence and implementing measures indiscriminately, this approach requires gathering evidence from specifically identifiable local publics and developing mitigation measures focused on those same publics. I will begin defending this approach by defining trust, highlighting the importance of maintaining it for the continuation of medical research. I will argue that a functional account of public trust asks the institution of science to be concerned with having a reputation of being trustworthy rather than with being trusted. This is because demonstrating trustworthiness is directly within their control, and since doing so garners trust anyways, this approach is both practical and effective, i.e., functional. I will then examine trustworthiness, concluding that it is demonstrated by competence and intention. Consistently achieving the most favourable results displays both of these traits, thereby fostering public trust in the long term, even if it is damaged in the short term.

In many cases, the community engagement approach agrees with Eyal's thesis that following a risk/benefit analysis formed by a third-party ethics review board is the highest available standard of risk assessment and should consequently achieve the best results time after time. However, concluding that public opinion should not factor into ethical considerations does not properly address the concerns of a public trust argument. On the basis of a supposition, it either rejects the possibility that conducting a challenge trial *could* damage public opinion enough to harm the institution of science altogether or denies that anything should be done about it, which is wrong for all the reasons explored in the intuitive account. The community engagement approach instead demands that the contexts in which public opinion is relevant to trial design selection should be defined by gathering *targeted* evidence. Using this approach, I will argue that if and only if this targeted evidence indicates a particular trial design might immediately damage public trust enough to limit the ability of the institution of science to conduct important research in the future, and thereby its ability to continue demonstrating trustworthiness, it should be included in a risk/benefit analysis[iv]. Because this evidence can only be acquired for an identifiable public, I will claim that public trust concerns should be focused exclusively on local publics when discussing the issue in the context of trial design selection.

Trust

Trust is something we all exercise many times on a daily basis. We do it every time we cross a bridge, deposit money to the bank, make a promise, eat at a restaurant, drive a car, or tell a secret; in each of these instances, and countless others, we place our faith and wellbeing in the hands of numerous individuals. Sometimes these individuals are known to us, but oftentimes they are not, and yet we rely on both to varying degrees. Without this trust, society as we know it would be unable to function. The variety of contexts within which we trust, however, makes it

difficult to create a universally applicable formulation of the term, and as such there are various legitimate definitions of trust. It can be understood as cognitive or affective; it can be conceived of as a choice, capacity, or function; it can be a one, two, or three-part relation; it can be a moral, immoral, or amoral consideration, and can be understood as general, specific, thick, thin, rich, deep, or swift (McDonald et. al. 2008, 35) (Hardin 2002).

The blanket context of clinical research does little to clarify how public trust should be characterized as it applies to public trust and challenge trials. For example, while fiduciary trust is often defended as an ideal formulation of trust to apply to conducting clinical research because it appropriately characterizes the asymmetrical relationship between researchers and volunteers, it would likely not be as appropriate for evaluating the relationship between researchers who are collaborating with one another. The relationship between the institution of research and the public, as explored in chapter two, is more complex than either of these because 'the public' in itself does not categorize a particular group. We therefore need to develop an account of trust which not only accounts for this volatility by allowing for context-specific applications of the term but one that also properly characterizes the behaviours of different publics in different contexts.

The account of trust which will therefore be defended in the context of challenge trial design selection is Russel Hardin's trust as encapsulated interest, which will be explained throughout the chapter. This formulation is comprised of establishing trust as a three-part relation and trust as a cognitive, rational faculty. The three-part model of trust expands upon the two-part model, which asserts only that A trusts B, without specifying a domain. Such a model is typically used for explaining the abstract trust an individual has towards others, claiming that if A trusts B, they do so unconditionally, and if A distrusts B this distrust holds in all instances (Hardin 2002,

61). However, it seems unlikely that merely conducting an otherwise permissible challenge trial, even if seemingly contentious, will alone be the difference between the public trusting and mistrusting the institution of science universally. Because the issue at hand demands a focus on the public's trust in science within particular contexts, i.e., conducting a challenge trial, this understanding is insufficient. An even more simplistic formulation of trust is the one-part model, which only asserts that A trusts, but this focuses on a particular individual's capacity to trust generally, and because public trust in specific contexts is the focus of this analysis, this definition does not seem appropriate here; there should not be a focus on building trust generally, but on the effect that conducting a challenge trial has on public trust (Hardin 2002, 61).

In the three-part model of trust used by Hardin, A trusts B within the domain of X (Hardin 2002, 9). By specifying a domain, the three-part model allows for narrowing the scope of trust from universal to a given set of circumstances. One might then be able to say that they trust researchers within the domain of conducting policy-approved research, but not research which extends beyond policy recommendations or is novel enough that policy has not yet been developed. In this case, one not trusting researchers implicitly, but within a specific context, and as such their trust in medical research may be rightly violated (though not necessarily) if novel research is conducted without ethical review regardless of the outcome (Resnik 2011, 6).

Alongside viewing trust as a three-part relation, Hardin interprets trust to be a cognitive function or rational assessment. According to this theory, trust is simply an assessment of trustworthiness, rather than, say, an affective attitude one has towards another. This conception is in part provided in chapter 2 by Resnik, when he states that, "people decide to trust others because they deem them to be trustworthy" (Resnik 2011, 3). This particular conception, however, demands revision. The use of the word 'decide' in this context implies that trusting

another is a choice one makes. While it is true that *exercising* trust is a choice, or *acting* in a trusting manner is a choice, trust as a cognitive faculty presumes a benign assessment (Hardin 2002, 7). One trusts or mistrusts another because they deem the other trustworthy or untrustworthy based on information available to them, and then they choose how to act, but the actual trusting itself is based purely on one's assessment of another's trustworthiness.

To determine whether or not someone is trustworthy, a prospective truster assess two main attributes of the prospective trustee. The first is their motivations, and the second is their competence within domain X to fulfill the truster's expectations (Graham et. al. 2022, 1). The motivational argument begins by claiming that if I believe your interests align with mine, I can trust that you will act in my interests because I trust you will act in yours. Going a step further, there is the encapsulated-interest account of trust proposed by Russell Hardin, in which the trusted party encapsulates the interests of the truster as their own and acts in the interest of the truster in part *because* they are trusted to do so, likely for the purpose of maintaining an ongoing trust relationship (Hardin 2002, 4). In other words, if your wellbeing is contingent upon me doing something, I would take your known interest in it being done as a reason for doing it because I want you to have your interests met[v]. If, in my assessment of your trustworthiness, I determine that you would fulfill my expectations in part because you know I expect you to fulfil them, then I can say that you are trustworthy on the encapsulated-interest account and would therefore be justified in trusting you provided the competence condition is met[vi] (Hardin 2002, 4).

The role of competence in the encapsulated-interest model of trust is more straightforward. Essentially, if I determine that you are competent within a certain domain, I trust that you have the ability to demonstrate your trustworthiness by acting on my trust and I may

therefore justifiably place my trust in you within that domain. Alternatively, if I think you are incompetent within a domain, I will not expect you to act competently within that domain such that whatever I am trusting you to do will be adequately done, and I would therefore not be able to justifiably trust you regardless of how well-intentioned you may be (Hardin 2002, 4). For example, even though I trust my mother implicitly because I know she encapsulates my interests, I would not trust her to perform brain surgery on me because she lacks the requisite skills to do so and would likely fail.

What does all this mean for the implementation of challenge trials? Essentially, it shows that the focus on public trust thus far has been misplaced. According to the encapsulated-interest model, trust is a rational assessment of trustworthiness. As such, trust cannot be built without demonstrating trustworthiness, and if an institution proves itself to be trustworthy, it follows that the public would trust it. The discussion thus far has been focused on whether or not the public trusts the institution of science, when it should instead be focused on whether or not the institution of science is trustworthy enough for people to place their trust in it, and how such trustworthiness can be established within the context of trial design selection.

Trustworthiness

If the goal of defining public trust is to create a functional, useable concept for the purpose of developing policy regarding challenge trial design selection, trustworthiness deserves to be its focus. Trust arguments have thus far concerned themselves with the actions of the public in response to medical research. In addition to wrongly conflating trust as a belief and trust as an action, this places the onus on prompting others to act in a certain way when a feasible account requires that one focuses on their own actions.

The particular motivations of a given public are often diverse and consequently difficult to discern or homogenously influence. As such, it is more practical to successfully prioritize respecting volunteers and conducting meaningful research, thereby demonstrating both motivation and competency over time. Doing so likely resolves public expectations in a much more sustainable fashion, even if not immediately so[vii]. While I will argue that this is often true, I also believe that there are contexts where the potential impact of immediate public distrust, justified or otherwise, warrants consideration in a risk-benefit analysis, and that discovering these contexts requires the use of a functional understanding of trustworthiness.

The previous section defines trustworthiness from the perspective of the truster. In her assessment of trustworthiness, Karen Jones utilizes the three-part model of trust developed by Hardin and recapitulates it from the perspective of the trustworthy. Her reinterpretation of Hardin's model states that, "B is trustworthy in respect to A in domain of interaction D, if and only if she is competent with respect to that domain, and she would take the fact that A is counting on her, were A to do so in this domain, to be a compelling reason for acting as counted on" (Jones 2016, 61). Not only does this directly invert the three-part model of trust, it also explicitly includes competence and encapsulated interest as conditions for trustworthiness. While Hardin includes these as well, he only does so as an implication, wherein A trusting B is dependent on B demonstrating trustworthiness and trustworthiness is consequently defined as demonstrating competent encapsulated interest. Jones also discusses rich trustworthiness, a demonstrably more effective extension of encapsulated interest which involves the trustee accurately signalling instances in which they have encapsulated the interests of the truster because the trustee cares about the wellbeing of the truster as an end in and of itself (Jones 2016, 74).

Applying this model to challenge trials, we would say that the institution of science is trustworthy in respect to the public in the domain of safely and effectively conducting challenge trial research if and only if it is competent with respect to safely and effectively conducting challenge trial research, and it would take the fact that the public is counting on it, were the institution of science to do so in this domain, to be a compelling reason for acting as counted on. This statement is dense but can be clarified by examining its constituent components individually.

First and foremost, competence is demonstrated to the public by achieving results[viii]. By successfully demonstrating the ability to meet the expectations of scientific progress held by the public, the institution of medical research proves their ability to continue to do so in future, or at least provides good reason for the public to believe they are able to. Competence is also concerned with the ability to adequately conduct and mitigate risks without accident. For example, the challenge trial literature often utilizes the case of Jesse Gelsinger[ix] and the resulting consequences on the field of gene therapy as reason to be hesitant to conduct challenge trials, even though this case did not involve a challenge trial (Eyal 2022, 4). The worry held by scholars is that, should an accident occur resulting in severe injury or death, the public might think the institution of medical research incompetent towards the conduct of challenge trials or the assessment of risks in such a trial, both of which would lead the public not to trust them to conduct future challenge trials.

The second component of Jones's applied definition is her extensive account of encapsulated interest, which, "Take[s] the fact that A is counting on her, were A to do so, to be a compelling reason for acting as counted on" (Jones 2016, 74). In this instance, trustworthiness involves being perceived as acting on the interests or expectations of the public *because* the

public is relying on these actions to occur. One way of doing this is, again, to achieve results. Developing vaccines, medicines, and procedures which alleviate and prevent suffering (i.e., successfully conducting medical research) are core expectations the public holds of medical research[x] (Resnik 2011, 6). Providing tangible evidence of having met such expectations is a good way to prove to the public that their interests have been encapsulated. Even negative findings and unsuccessful trials should be understood as achieving results because they are necessary for scientific progress. They show the public that researchers honestly report their failures to ensure the safety of the public, further demonstrating encapsulated interest.

The public also expects that researchers will respect the interests of volunteers, namely by prioritizing their safety (Resnik 2011, 6). An accident resulting in permanent injury or death could be perceived by the public as preventable at best and malicious at worst, thus placing the motivations of medical research in question in the eyes of the public. For example, the historically unethical treatment of medical research volunteers resulting in injury or death for the sake of scientific discovery, as was the case in the Walter Reed yellow fever experiments, are discussed in the challenge trial literature as likely being unethical by contemporary standards (Clements et. al. 2017). It is then asserted that this complicates public perception of the motivations of contemporary challenge trials, which explains existing public hesitancy towards them, at least in part. A reputation of trustworthiness is much easier to lose than it is to cultivate (Hardin 2002, 90), and it is likely that those old wounds still bleed for at least some relevant portion of the public. Given the presumption that public trust must be fragile as a result, any occurrences of impermissible harm are likely to negatively affect the public's perception of the motivations behind medical research much more impactfully (Bambery 2016, 92). Therefore,

limiting risk as much as possible, unless sufficiently justified, is essential for demonstrating encapsulated interest and trustworthiness.

What this establishes is that, both from the perspective of competence and motivation, trustworthiness is best developed by achieving results and limiting risks to volunteers. In other words, following the result of a risk/benefit analysis provided by a third-party review board is likely the most effective path towards cultivating a reputation of being trustworthy at an institutional level, because it is the most likely way to ensure both beneficial results and minimal risk[xi] (Eyal 2022, 5). By this account, even if the optics of the best available trial design in a given context are contentious and cause an initial loss of public trust, conducting this trial should also theoretically provide the best results with the least possible risk. By repeatedly demonstrating successful, responsible, and safe experimentation, the institution of science should prove itself both competent and ethically motivated, strengthening its reputation of trustworthiness long-term.

However, this account entails two important caveats. Firstly, it presumes that a risk/benefit analysis conducted by an independent ethics review board provides the most accurate possible assessment of risk and justification. While this is theoretically true, in practice it may not be, and could result in volunteers facing unjustified risks, which in turn affects the trustworthiness of institutions conducting challenge trial research. This follows the competence argument, in which a preventable accident would give the public reason to believe that the institution of science is incapable of accurately conducting these assessments. Remedying institutional incompetence, however, goes beyond the scope of this paper, and even so it seems unlikely that a third-party ethical review would create a less accurate risk/benefit profile than the public.

The second caveat involves the necessity of *repeatedly* demonstrating competence and encapsulated interest to cultivate long-term trustworthiness. If one damages my trust, I am unlikely to trust them again in future, meaning that I am unlikely to give them further opportunities to demonstrate their trustworthiness (Hardin 2002, 90). What has been argued thus far is that riskier, less beneficial trial designs are most likely to violate the expectations of the public (either by imposing unnecessary risks or unnecessarily limiting the potential for making important discoveries), and the resulting backlash of conducting the 'worse' trial will hinder future medical research and thus future opportunities for demonstrating trustworthiness more than conducting the seemingly contentious but 'better' trial. However, we briefly discussed the possibility that conducting a seemingly contentious trial (i.e., a challenge trial), even if it is truly the best available trial to conduct from an accurate risk/benefit perspective, could initially damage public trust, either in itself or in the event of a serious adverse event. If this initial distrust is severe enough that it alone would limit future research opportunities, and thus the ability to demonstrate trustworthiness, there is an argument for conducting a riskier, less contentious trial design instead. This worry, which will be called the trustworthiness problem, characterizes most strong public trust arguments in the challenge trial literature.

The Community Engagement Approach

The trustworthiness problem identifies *when* conducting a particular trial might sufficiently harm public trust such that it should be included in a risk/benefit analysis but does not identify *which* public we should be concerned about, a task we have already deemed necessary for developing good policy. To properly address this without reverting to the suppositions of the intuitive and evidential approaches, I have developed the community engagement approach. It proposes that determining whether or not conducting a challenge trial will result in the

trustworthiness problem requires relying on evidence acquired by directly engaging with specific, identifiable communities. If doing so proves that conducting a particular challenge trial will result in the trustworthiness problem, an informed decision can then be made on how this distrust can be effectively mitigated, either by continuing a dialogue with this public to address their concerns or, as a last resort, selecting a different trial design[xii]. Consequently, dissent originating from unidentified, different, or more general publics should be ignored because it is unlikely to be helpful in determining if the trustworthiness problem will occur and, as seen in the evidential account, tends to obfuscate the expectations that should be of concern in this study.

I will begin defending the community engagement approach's focus on identifiable publics on the grounds that doing so is practical, and therefore ideal for developing a feasible account of public trust, whereas deriving the expectations of a global or general public is not. I will then re-examine previously explored mitigation measures, displaying how none of them feasibly defend a focus on global public mistrust resulting from the optics of a trial design, and how they instead display an existing, implicit focus on local publics. I will conclude by discussing the advantages of an emotional appeal that a community engagement focus provides.

1) Feasibility

The first reason for focusing strictly on a local public is because doing so allows for the feasible acquisition of useful evidence for identifying when the trustworthiness problem is likely to occur. Essentially, it needs to be clear that people hold a negative view of challenge trials that would lead to relevantly harmful anti-science beliefs should these trials be conducted, rather than anti-science beliefs that include a negative view of challenge trials alongside negative views of all medical research. This is because we are focused on the impact *challenge trial design selection* has on public trust, not the conduct of medical research altogether. Breaking this down

into two components, this inquiry needs to reveal 1) whether the public holds anti-challenge trial beliefs that they hold more strongly than for other forms of medical research, and 2) whether these beliefs are prominent enough among this public to result in the trustworthiness problem.

A strict focus on community engagement allows one to address both of these components. In short, directly engaging with a specific community allows researchers to ask specific questions, receive specific answers, and accurately determine whether or not conducting a challenge trial will result in the trustworthiness problem. While this approach allows one to identify both the direction and magnitude of the expectations of an identifiable public, deriving the expectations of the general or global public only reveals a difference in attitudes amongst diverse local publics. A cursory Google search makes it apparent that finding global support for medical research is just as easy as finding global dissent; on what grounds would more importance be assigned to one than the other? Because the motivations and relevance of general online dissent are inherently vague and contradictory in this way, it becomes impossible to legitimately prioritize focusing on anti-science beliefs, or to separate these more generalized anti-science beliefs from specifically anti-challenge trial sentiments. Without being able to accurately or justifiably discern what the public expects, it cannot be known what kind of impact conducting a challenge trial would have on the institution of science, meaning good policy cannot be developed. Because the focus here is on developing a functional account of public trust, a focus on identifiable local publics is prudent.

2) Mitigation

This feasibility is emphasized when examining the mitigation measures proposed by the intuitive and evidential accounts intended to limit harm to public opinion caused by the conduct of challenge trials. Often, the implication that challenge trials especially harm public trust leads

authors to insist these measures are necessary for conducting them at all (Hope & MacMillan 2004, 113). Notably, this paper has focused on how public trust should factor into a risk/benefit analysis, but mitigation measures tell us how a loss of public trust should be effectively managed, should it occur. While these points are distinct, they are two sides of the same coin, as mitigation is seen as a means for 'cancelling out' any loss in public trust in a risk/benefit analysis (Hope & MacMillan 2004, 113). As a result, insisting that a particular measure should be used alongside a trial reveals not only an attitude that the trial will harm public trust, but *how* conducting it will harm this trust. Examining mitigation measures thereby provides clues on what one is concerned with when discussing the inclusion of public trust in a risk/benefit analysis, so the focus will be on both in this section.

In chapter two, both community engagement and the utilization of independent ethical review boards were proposed as possible measures to mitigate any loss of public trust resulting from challenge trial conduct. Often, this literature also discusses publicizing the rationale behind a trial design selection risk/benefit analysis as a means for educating the public and quelling possible mistrust based on misinformation (Jamrozik 2021, 88-89). It was also shown that the evidential approach's main problem is that the unqualified dissent of one public is used to justify the imposition of measures intended to mitigate the loss of the trust in an entirely different public. A functional account of public trust and challenge trials should address evidence of a relevant loss of trust from a given public with measures intended to mitigate distrust directed towards that same public. It therefore seems logical to examine where proposed mitigation measures are focused to see what they reveal about which publics should be of concern.

Beginning with third-party ethical review, this measure is intended to create an unbiased risk/benefit analysis, and the removal of bias is meant to bolster public trust both locally and

generally. As explored in chapter 2, however, providing an accurate, unbiased account of which trials should be undertaken in which instances is independently valuable, only satisfying public expectations of honesty in science as a result. It is therefore unclear why authors insist that this measure is necessary to bolster trust which may be lost from conducting a challenge trial, rather than insisting on the implementation of this measure in all instances. This confusion is compounded by the fact that an independent ethical review board used for selecting a trial design would have to be instituted before a trial design is selected, meaning it applies in all instances regardless. Suggesting that this measure is necessary to remedy the trust lost by conducting a challenge trial is anachronistic, as it must be conducted before one even knows if a challenge trial will be conducted.

Insisting on the publicization of risk/benefit rationale for selecting a given trial, which also focuses on both local and general publics, incurs a number of its own problems. For one, it is just as likely that public trust is best served when rationale is given for why researchers did NOT use a challenge trial as for why they did. As all medical research either is or is not a challenge trial, publicizing trial design selection rationale should seemingly be standard practice. Additionally, publicizing the rationale of a risk/benefit analysis means having conducted one in the first place. If it has yet to be determined how public trust should factor into a risk/benefit analysis for trial design selection, there is no justification for asserting that conducting a challenge trial requires implementing unique measures (i.e., publicization of rationale) to counteract a purely supposed loss of public trust. If the relevance of public trust in a risk/benefit analysis can be determined, then it is already known how conducting a challenge trial would impact the institution of science, regardless of whether or not rationale is publicized. As such, while insisting on the publicization

risk/benefit rationale might be beneficial for a number of reasons, it does not seem like counteracting the trust lost from conducting a challenge trial in particular is one of them.

There is one exception to this point. If a strict risk/benefit analysis supports conducting a challenge trial but there is evidence that existing negative public sentiment towards challenge trials would result in the trustworthiness problem, and if, in this instance, publicizing risk/benefit rationale would shift public opinion enough to conduct the best trial available, then it may be said that this measure should apply to challenge trials in particular. However, there are multiple reasons why this scenario is not useful for the current inquiry. Firstly, coming to this determination would require isolating the expectations of a public, which presupposes a local public as per the feasibility argument provided earlier. Providing a risk/benefit rationale would therefore be targeted towards this same community. It is unclear whether or not, at this point, the publicization of risk/benefit rationale would be considered its own measure, or it would be part of the informed consent process when undertaking community engagement.

There are also numerous studies which prove that anti-science opinions formed on the basis of misinformation[xiii] are emotionally charged, and that providing legitimate evidence which conflicts with these opinions only serves to make skeptical individuals double down on their false beliefs (Hornsley et. al. 2018). As such, publicizing risk/benefit rationale for trial design selection would likely do little to effectively mitigate a loss of trust in the individuals who should be focused on. Additionally, while the publicization of trial design selection rationale is useful for accountability or respecting the informed consent of the public as they form their opinions on the ethics of challenge trial research, transparency does not build trust so much as it removes the conditions for trusting (Graham et. al. 2022, 1). While it may therefore be valuable to insist on publicizing the results of a trial design selection risk/benefit analysis, it would be for independent

reasons rather than strictly for the mitigation of public trust lost from conducting a challenge trial. Because the goal here is to create a feasible account of public trust, and this measure is not useful as a means for mitigating a loss of public trust, it should not be used to indicate which publics are of concern.

3) Community Engagement and Demonstrating Trustworthiness

The only mitigation measure left to explore is community engagement. As observed earlier, a focus on community engagement implies focusing on a community, which indicates a reliance on local, identifiable publics to determine the relevance of a loss of public trust resulting from conducting a challenge trial. While there has been discussion on the usefulness of community engagement in terms of acquiring an accurate set of expectations from which to develop policy, the following section will explore how community engagement is also a useful tool for building and maintaining a reputation of trustworthiness in the institution of science.

Before proceeding with this analysis, it is important to clarify one major aspect of trust and trustworthiness. As stated earlier, trust is merely a rational assessment of trustworthiness. However, humans are not necessarily rational beings, and can only operate from their own perceptions based on information they have acquired. Because a prospective truster is fallible in this way, building public trust does not necessarily involve *being* more trustworthy, but being *perceived as* more trustworthy, i.e., having a reputation for trustworthiness. While I have argued that repeatedly acting in a trustworthy manner should build this reputation, the trustworthiness problem complicates the ability to repeatedly demonstrate trustworthiness, meaning public perception would remain divorced from rationality if the trustworthiness problem were in effect.

As the previous section indicated, emotionally charged anti-science opinions are unlikely to be remedied by a logical appeal. The inverse of this is that an appeal to emotion can be useful in

positively impacting public trust. This is especially relevant for demonstrating encapsulated interest. Because encapsulated interest, and especially rich trustworthiness, involves acting on another's best interest because *they* are interested in your doing so, rather than because your interests are simply aligned by coincidence, developing relationships with individuals at the communal and individual level is a good way of demonstrating good intentions to the public.

In one instance, the community engagement approach achieves this by making the institution of science feel like a more identifiable public. While there has been a focus thus far on defining the public which trusts the institution of science, I have largely ignored the fact that this institution is itself made up of a number of different publics, which may or may not be perceived in different ways by general and local non-scientific publics. By engaging with a community, the institution of science ceases to be a faceless entity and is instead associated with particular researchers or research organizations who have made themselves known to this community. In terms of encapsulated interest, it is much easier to imagine that a person or group of persons you have personally associated with hold your best interests at heart than a distant and likely profit-motivated organization which will never know you as an individual.

This point is emphasized by a study conducted by Michael McDonald et. al. which surveyed individual's perceptions of medical research. They found that "Participants […] conveyed a general lack of trust in pharmaceutical companies based on their image of such institutions being driven by the pursuit of profit over the pursuit of knowledge and the best interests of subjects" (McDonald et. al. 2008, 38). However, according to one patient surveyed who held this view, her concrete relationships built with individual researchers helped her overcome this distrust (McDonald 2008, 38). While solving the problems of privatized medical research and healthcare

are beyond the scope of this paper, this shows how believing somebody holds your best interests at heart can often occur simply by knowing who this 'somebody' is.

Additionally, trust relationships are not necessarily one-way as they have been characterized up until this point. Reciprocity is an important element of trust; if I am able to prove that I trust you, you are more likely to believe our motivations are aligned, therefore making me seem more trustworthy (Hardin 2002, 17). By engaging with the community and trusting them enough to earnestly consider their opinions, the institution of medical research can display a sense of trust in their community members. Therefore, not only does community engagement indicate which publics we should be concerned about when selecting a trial design, it also functions to improve trust in communities where it has been identified as needing improvement.

Concerns

The trustworthiness problem accurately characterizes the most pressing concern we have when we discuss public trust as it relates to challenge trial design selection. However, this is not the only concern scholars discuss when they talk about public trust. A second prominent concern regularly voiced in the literature is that a loss of public trust will reduce the adoption of vaccines. Because herd immunity is an important part of reducing the impact of a disease, maintaining short-term trust might be more important in a risk benefit analysis than conducting the 'best' trial if this trial would lead to such a reduction.

This concern was not explicitly discussed because, even if it should occur, it does not limit the ability for the institution of science to continue demonstrating trustworthiness. As such, even if the public loses trust in science in the short-term, there is nothing preventing them from eventually regaining this trust should the vaccines resulting from a trial repeatedly be proven safe and effective. Additionally, the link between trial design selection and vaccine skepticism is

weak at best (Eyal 2022, 4). People would likely refuse a vaccine due to a belief that they are ineffective or harmful, rather than because they have reservations about a trial design. Most people will still get it even if they do not trust the institution of science, provided they believe getting a vaccine will be better for them (O'Neill 2002, 11-12). If this existing distrust would prevent them from getting vaccinated it would likely do so regardless of whether or not the vaccine was tested in a challenge trial.

Does this mean we have to suppose global public distrust is not a concern that should be dealt with? I do not believe it does. Because the purpose of this analysis is to reduce or eliminate the role of suppositions in the development of trial design selection policy, it would be imprudent to disqualify this as a concern. While maintaining public trust *is* a concern for a more global public, and while this does apply to a global public opinion on challenge trials in general, it is unclear whether public trust affected by challenge trial design *selection* in a particular instance can truly be considered a concern for a global public, at least to the degree that the trustworthiness problem demands. In any case, there is no way to determine the effect trial design selection has on public trust at a global level. As such, using global public opinion to influence policy focused on determining whether or not to conduct a challenge trial in a particular instance is unwise. However, there are different ways to address concerns of global public mistrust in challenge trials, such as limiting misinformation or eliminating profit-driven medical research, which would likely be more effective and deserve their own considerations.

---

[iv] In the context of trial design selection, public trust as a concern is often framed as part of a utilitarian dilemma. Irrevocably breaking the trust of the public by conducting an otherwise advantageous challenge trial leads to conducting less research overall, leaving cures undiscovered which ends up delaying scientific progress as a whole and costing us exponential future human lives. Supporters of any public trust argument believe that this overshadows the comparatively miniscule impact of delaying a cure or exposing volunteers to higher but less controversial risk by conducting relatively less effective field trials

[v] While this account of trust resembles the fiduciary account, there are important distinctions. The primary difference between these two accounts is that fiduciary trust is built upon the obligation or duty for the trustee to act in the truster's interests, whereas in the encapsulated interest model the trustee encapsulates the interests of the truster as their own and acts to fulfill this trust because they know they are being trusted in this way and/or because they share the interest of the truster fulfilled. In short, a fiduciary obligation does not require me to actually encapsulate another's interests as my own and act on it *because* I want their interest met, but only because I am obligated to meet their interest. Because of the importance of perceived motivations in the realm of developing trust, predicating trust on a legal obligation is not ideal for our purposes; it shows that researchers only protect the interests of the public out of fear of repercussions rather than genuinely caring about the wellbeing of volunteers and the public as their own ends

[vi] It is important not to confuse this understanding of trustworthiness with the rich trustworthiness described by Karen Jones, which involves the trustee wanting to meet the truster's interests because they care about them, and the trustee thus encapsulate the interests of the truster as their own because the truster's wellbeing is an end in and of itself for the trustee.

[vii] The argument here is that seemingly controversial actions which demonstrate competence and encapsulated interest might immediately harm public opinion, but by repeatedly proving oneself to be trustworthy, public opinion will eventually shift and benefit from conducting these actions. In addition to aligning with Resnik's and Hardin's cognitive account of trust, this is consistent with Eyal's belief that actions deserving of public outcry are deserving whether or not public outcry actually occurs, and that undeserved outcry has no justifiable place in assessing the risks and benefits of a given trial design; if a third-party risk-benefit analysis is considered the highest standard of risk assessment, let it remain uninfluenced by unjustified dissent (a claim that will be built upon and contested).

[viii] Though it is difficult to prove what causes the public to trust, we observed earlier that accidents resulting in grievous harm damage trust; regardless of how competent researchers may be in these cases, they will likely be deemed incompetent by a relevant portion of the public

[ix] In 2001, Jesse Gelsinger was a volunteer participating in gene therapy medical testing when they suffered a fatal reaction to a compound used to induce respiratory irritation. The resulting backlash from the immediate community ended up suspending gene therapy research, presumably setting the field back decades, and is often used as an example of the importance of maintaining public trust (Rosenblatt 2020)

[x] While our discussion thus far has largely avoided assigning expectations to the 'public' because we have not identified a public, there are core expectations which we can assign relatively safely. Resnik discusses some of these in his paper, stating that, "Virtually everyone expects science to be honest and trustful, and most people expect that scientists will not intentionally bias their results (Laine et al 2007). Most people also expect that scientists will not undertake experiments on people without their consent, or intentionally harm or exploit human research subjects. A vast majority of the members of society probably also expect that scientists will not abuse animals in research (Shamoo and Resnik 2009)." (Resnik 2011, 6-7). When it comes to challenge trials, it is not that people disagree on these core expectations, but whether or not conducting challenge trials fulfill these expectations or violate them

[xi] This presumes that a third-party ethical review is the highest standard of risk analysis and justification

[xii] Eyal discusses the problems of basing trial design selection on the opinions of the public when he states that, "the practical implication of 'X is perfectly right intrinsically but may disastrously upset the public' is only rarely 'Avoid X'. More often, it is 'Explore whether there might be a particular form of X that avoids upsetting the public so much'. Only once that first attempt fails does it usually become wise to settle for the highly suboptimal 'Avoid X'." (Eyal 2022, 5)

[xiii] As we saw earlier, we are not concerned with public dissent based on informed, legitimate criticisms, but only on misapprehensions of trial design safety by the public which could result in dangerous backlash

## Conclusion

At the beginning of this paper, I discussed the assertion that a number of commentators were concerned about the potential impact of conducting challenge studies on public trust. To examine the legitimacy of this claim, I identified public trust as an underdeveloped and misapplied element of the challenge trial ethics literature. While public trust was regularly mentioned as a reason for limiting the conduct of challenge trials, it was rarely, if ever, defended. By taking Resnik's assertion that assessing public trust demands defining 'public' and 'trust' and applying it to the challenge trial literature, it became clear that the term had been taken for granted up until this point. These beliefs are largely echoed by Nir Eyal, who believes that public trust concerns in the challenge trial literature are largely unsupported and speculative.

After identifying public trust as underexplored, I examined the relationship between the public and the institution of medical research through the lens of trustworthiness, which is the ability to prove one's competence and encapsulated interest. I argued that there should be a focus on demonstrating trustworthiness rather than building trust because doing so is not only feasible but also ethical. Initially, I used Resnik's paper to claim that both 'public' and 'trust' needed to be identified, as they properly categorize this relationship. I then provided an in-depth exploration of trust, showing why a focus on trustworthiness makes more sense for developing a functional account of public trust as it applies to challenge trials.

To conclude, I used the above focus on trustworthiness to develop the trustworthiness problem, which defines a set of circumstances under which conducting the best trial available would cause backlash sufficient to limit future opportunities to demonstrate trustworthiness. I draw on this to develop my own functional account of public trust as I believe it should apply to challenge trial design selection policy known as the community engagement approach.

In summary, conducting a trial supported by a risk/benefit analysis developed by an independent third party is the most likely way to ensure the trial which entails the most scientific benefits for researchers and/or the fewest risks for volunteers is conducted, and in most cases, it should therefore be adhered to. By repeatedly demonstrating trustworthiness and always following an independent risk/benefit analysis, the institution of medical research should theoretically develop a reputation of trustworthiness long-term, even if it harms trustworthiness short-term. The only instance where this would be false is if short-term harm to public trust would be severe enough that the public no longer provides the institution of science with an opportunity to continue conducting trials and achieving beneficial results, thereby limiting the ability to cultivate a reputation for trustworthiness. Relying on community engagement not only allows one to accurately determine whether or not this would be the case, it also mitigates harms caused by possible public mistrust, and could even help develop trust pre-emptively.

This analysis aims to contribute to an improved understanding of the relationship between public trust and the institution of science, namely by identifying the decisions which have a greater impact on public opinion. Perhaps a contentious yet beneficial trial could be conducted which bolsters public trust at the same time by, for instance, shifting away from for-profit medical research, or providing better methods of accountability. These are measures that, while not themselves focused on trial design selection policy, have appeared in this paper as more likely causes for an individual to hold a negative opinion on medical research. Additionally, while global trust is set aside in this paper, it is still important to consider. Hopefully technological advancements one day contribute to accurately assessing the impact trial design selection might have on a broader public. Until then, I believe the community engagement approach provides the best avenue to conducting safe and effective vaccine research.

Bibliography

1daysooner.org. "1Day Sooner." *1Day Sooner* (2020): https://www.1daysooner.org/.

Bambery, Ben, Michael Selgelid, Charles Weijer, Julian Savulescu, and Andrew J. Pollard. "Ethical Criteria for Human Challenge Studies in Infectious Diseases." *Public Health Ethics* 9, no. 1 (2016): 92–103. https://doi.org/10.1093/phe/phv026.

Clements, Alan N., and Ralph E. Harbach. "History of the Discovery of the Mode of Transmission of Yellow Fever Virus." *Journal of Vector Ecology* 42, no. 2 (2017): 208–22. https://doi.org/10.1111/jvec.12261.

Dawson, Liza, Jake Earl, and Jeffrey Livezey. "SARS-CoV-2 Human Challenge Trials: Too Risky, Too Soon." *Infectious Disease Society of America,* Oxford University Press (2020): 1-6

Eyal, Nir, Marc Lipsitch, and Peter G Smith. "Human Challenge Studies to Accelerate Coronavirus Vaccine Licensure." *The Journal of Infectious Diseases* 221, no. 11 (2020): 1752–56. https://doi.org/10.1093/infdis/jiaa152.

Eyal, Nir. "Research Ethics and Public Trust in Vaccines: The Case of Covid-19 Challenge Trials." *Journal of Medical Ethics*, 2022, 1–7. https://doi.org/10.1136/medethics-2021-108086.

Fernandez Lynch, Holly. "Minimal or Reasonable? Considering the Ethical Threshold for Research Risks to Nonconsenting Bystanders and Implications for Nonconsenting Participants." *Bioethics* 34, no. 9 (2020): 923–32. https://doi.org/10.1111/bioe.12725.

Goldby, Stephen, Saul Krugman, M. H. Pappworth, and Geoffrey Edsall. "The Willowbrook Letters, "Criticism and Defense." Reprinted in *Intervention and Reflection: Basic Issues in Medical Ethics*. 5th ed. Ronald Munson ed. (Belmont; Wadsworth 1996). pp 273-381.

Graham, Mackenzie, Richard Milne, Paige Fitzsimmons, and Mark Sheehan. "Trust and the Goldacre Review: why trusted research environments are not about trust." *J Med Ethics,* 2022, 1-4. doi:10.1136/ medethics-2022-108435

Hardin, Russell. *Trust and Trustworthiness.* New York, Russell Sage Foundation, 2002.

Harris, John. "Scientific Research is a Moral Duty." *Journal of Medical Ethics,* 31(4) (2005): 242–248. doi:10.1136/jme.2005.011973

Hornsley, Micheal, et. al. "The Psychological Roots of Anti-Vaccination Attitudes: A 24-Nation Investigation" *Health Psychology,* 37(4), American Psychological Association, 2018, Queensland

Hope, T, and J McMillan. "Challenge Studies of Human Volunteers: Ethical Issues." *Journal of Medical Ethics* 30, no. 1 (2004): 110–16. https://doi.org/10.1136/jme.2003.004440.

International Military Tribunal, "Permissible Medical Experiments." Trials of War Criminals before the Nuremberg Military Tribunals under Control Council Law No. 10. Nuremberg October 1946 - April 1949, Washington. U.S. Government Printing Office (n.d.), vol. 2., pp. 181-182.

Jamrozik, Euzebiusz, and Michael J Selgelid. *Human Challenge Studies in Endemic Settings Ethical and Regulatory Issues*. Cham: Springer International Publishing, 2021: ix-88

Jamrozik, Euzebiusz, and Michael J Selgelid. "Human Infection Challenge Studies in Endemic Settings and/or Low-Income and Middle-Income Countries: Key Points of Ethical Consensus and Controversy." *Journal of Medical Ethics* 46, no. 9 (2020)

Jamrozik, Euzebiusz, and Michael J. Selgelid. "Ethical Issues Surrounding Controlled Human Infection Challenge Studies in Endemic Low-and Middle-Income Countries." *Bioethics* 34, no. 8 (2020)

Jones, Karen. "Trustworthiness." *Ethics,* 123(1) (2012): 61–85. doi:10.1086/667838

Scientific, Kent. "Kent Scientific Corporation." *Kent Scientific*, Kent Scientific, 23 Oct. 2020, https://www.kentscientific.com/blog/how-is-a-vaccine-created-animal-research-plays-a-key-role/.

Kerasidou, Angeliki. "The Role of Trust in Global Health Research Collaborations" *Bioethics* (3) (2019): 495-501 10.1111/bioe.12536

Kraft, Stephanie A., Devan M. Duenas, James G. Kublin, Kelly J. Shipman, Sean C. Murphy, and Seema K. Shah. "Exploring Ethical Concerns about Human Challenge Studies: A Qualitative Study of Controlled Human Malaria Infection Study Participants' Motivations and Attitudes." *Journal of Empirical Research on Human Research Ethics* 14, no. 1 (2018): 49–60. https://doi.org/10.1177/1556264618820219.

London, Alex. "Reasonable Risk in Clinical Research: A Critique and a Proposal for the Integrative Approach." *Statistics in Medicine* 25(17) (2006):2869-85

McDonald, Michael, Anne Townsend, Susan M. Cox, Natasha Damiano Paterson, Darquise Lafrenière (2008). "Trust in Health Research Relationships: Accounts of Human Subjects." *Journal of Empirical Research on Human Research Ethics: An International Journal*, 3(4) (2008), 35–47. doi:10.1525/jer.2008.3.4.35

Miller, Franklin G., and Christine Grady. "The Ethical Challenge of Infection-Inducing Challenge Experiments." *Clinical Infectious Diseases* 33, no. 7 (2001): 1028–33. https://doi.org/10.1086/322664.

Miller, Franklin G; Steven Joffe. "Limits to research risks." *Journal of Medical Ethics,* 35(7) (2009), 445–449. doi:10.1136/jme.2008.026062

O'Neill, Onora. *A Question of Trust.* Cambridge, Cambridge University Press, 2002; vii-100

O'Neill, Onora. "Linking Trust to Trustworthiness." *International Journal of Philosophical Studies* 26(2) (2018): 293-300. https://doi.org/10.1080/09672559.2018.1454637

NewsHour. "Report: Medical Experiments Conducted on U.S. Prisoners, Patients." *PBS*, Public Broadcasting Service (2011): https://www.pbs.org/newshour/health/medical-slideshow-code.

Resnik, David B. "Scientific Research and the Public Trust." *Sci Eng Ethics*, 17(3) (2011) September; 17(3): 399–409. doi:10.1007/s11948-010-9210-x.

Resnik, David B. "Limits on Risk for Healthy Volunteers in Biomedical Research" *Theor Med Bioeth*. 33(2) (2012): 137–149. doi:10.1007/s11017-011-9201-1

Rid, Annette. "Setting risk thresholds in biomedical research: lessons from the debate about minimal risk." *Monash Bioethics Review,* 32(1-2) (2014): 63–85. doi:10.1007/s40592-014-0007-6

Rid, Annette, and Meta Roestenberg. "Judging the Social Value of Controlled Human Infection Studies." *Bioethics* 34, no. 8 (2020): 749–63. https://doi.org/10.1111/bioe.12794.

Riedel, Stefan. "Edward Jenner and the History of Smallpox and Vaccination ." *Baylor University Medical Center Proceedings* 18(1) (2005): 21-25.

Roestenberg, Meta, Marie-Astrid Hoogerwerf, Daniela M Ferreira, Benjamin Mordmüller, and Maria Yazdanbakhsh. "Experimental Infection of Human Volunteers." *The Lancet Infectious Diseases* 18, no. 10 (2018): 1–24. https://doi.org/10.1016/s1473-3099(18)30177-4.

Rosenblatt M. "Human challenge trials with live coronavirus aren't the answer to a Covid-19 vaccine." *STAT News*, 2020.

Shah, Seema K., Franklin G. Miller, Thomas C. Darton, Devan Duenas, Claudia Emerson, Holly Fernandez Lynch, Euzebiusz Jamrozik, et al. "Ethics of Controlled Human Infection to Address COVID-19 ." *Science* 368, no. 6493 (2020): 832–34. https://doi.org/10.1126/science.abc1076.

Shah, Seema K., Franklin Miller, and Holly Fernandez Lynch. "The Role of Community Engagement in Addressing Bystander Risks in Research: The Case of a Zika Virus Controlled Human Infection Study." *Bioethics* 34, no. 9 (2020): 883–92. https://doi.org/10.1111/bioe.12806.

Van Ness, Peter. "The Concept of Risk in Biomedical Research Involving Human Subjects." *Bioethics*. 2001 Aug; 15(4): 364–370. doi: 10.1111/1467-8519.00244

WHO guidance on the ethical conduct of controlled human infection studies. Geneva: World Health Organization; 2021. Licence: CC BY-NC-SA 3.0 IGO.