FINAL VERSION TO BE PUBLISED IN *SCIENCE & ENGINEERING ETHICS*
doi:10.1007/s11948-017-9917-z

The Ethics of Anti-Aging Clinical Trials

1. Introduction

Recently, biomedical research into the effectiveness of anti-aging interventions (AAIs) has reached the clinical trial phase. This area of research is promising for several reasons, not the least of which is that anti-aging interventions can significantly improve the well-being of millions of people, even the lives of those who don't suffer from a fatal disease. If effective, they could cause people to live longer, healthier, and happier lives, and there is great value in this effect. For this reason alone the development of this line of research is worth pursuing. However, it is not the mere extension of life that is the outcome variable of interest in clinical trial of AAIs. In order for the extension of life to be valuable, the AAIs need to extend valuable life. Interventions that merely extend the number of years during which humans suffer through diseased lives contributes no value to those lives, and perhaps have significant disutility for society.

Currently there are at least two clinical trials of interventions on the aging process. One trial is being conducted in the U.S., and is a phase IV clinical trial involving the off-label use of metformin, which is a drug that for many years has been commonly used to treat type 2 diabetes (Metformin in Longevity Study) There is some evidence (Anisimov et al. 2008; Martin-Montalvo et al. 2013) that metformin extends the quality life span of mice, and so researchers believe that it may also do so for humans. The most recent clinical trial is a phase I trial planned to occur in Japan (UMIN-CTR Clinical Trial). The trial is evaluating the pharmacokinetics of nicotinamide mononucleotide (NMN), which has been shown to slow or reverse the aging

doi:10.1007/s11948-017-9917-z

process (Yoshino et al. 2016; Mills et al. 2016). Other trials, including further trials on the use of metformin, are in the planning stages.

Clinical trials of AAIs, if they are true AAIs, have as their dependent variable some quantification of the extension of quality life span compared to some other standard, such as the mean or median age of death for the population from which the subject is drawn or the subject's expected life span given what else is known about the subject's health and behavior. Researchers may conduct clinical trials of other interventions that can indirectly affect the extension of quality life span. A clinical trial evaluating the effectiveness of a new diabetes drug may indirectly indicate its effectiveness in preventing diabetes deaths. But unless the dependent variable in the trial is measuring the drug's ability to extend a person's life, it isn't evaluating the effect it has on extending quality life span.

One of the biggest challenges of conducting a clinical trial of an AAI is that this variable, the extension of quality life span, is dependent on a large number of other variables. Such is not the case with most other clinical trials. For example, a clinical trial investigating the effectiveness of a new drug to treat cardiovascular disease would have as its dependent variable the blockage of blood vessels. The blockage of blood vessels, as compared to extension of quality life span, is dependent on a much narrower range of variables. A person's disposition to avoid risk is not closely associated with cardiovascular disease, but a person who is indisposed to avoid risk is more likely to live a shorter life. But all of the variables upon which cardiovascular disease depends are also variables upon which extension of quality life span depends.

When the value of a dependent variable is determined by such a wide range of independent variables, it is more difficult to eliminate confounders. In the case of clinical trials

doi:10.1007/s11948-017-9917-z

of AAIs, the elimination of confounders implies ethical difficulties. Specifically, clinical trials of AAIs cannot produce reliable data unless they resolve a dilemma related to the elimination of confounding variables. But the resolution of this dilemma requires ethically problematic practices, as I argue below.

The argument for this conclusion starts with the premise that in order for AAIs to be made available to the public, researchers need to demonstrate (a) that the intervention is effective in extending quality life span; (b) that the intervention is acceptably safe; and (c) the extent to which the intervention is effective. The focus of the present paper is on clinical trials that could potentially demonstrate (c), the extent to which the intervention is effective. That is, in order for an AAI to safely be made available to the public, researchers need to know the full effect of the intervention.

Acquiring data to support inferences regarding how (not merely whether) the AAI affects the extension of quality life span introduces the dilemma. The horns of the dilemma are: either researchers use subjects in clinical trials who have already experienced significant aging, or they use subjects who have not aged significantly. If researchers use subjects in clinical trials who have already aged significantly, then the subjects' pre-intervention health and behavior will confound inferences regarding the full effect of the AAI on the extension of quality life span. But if researchers use subjects in clinical trials who have not aged significantly, then the subjects' post-intervention health and behavior will confound inferences regarding the full effect of the AAI on the extension of quality life span. Thus, either pre-intervention health and behavior will confound the inferences, or post-intervention health and behavior will confound the inferences. In what follows, I argue that eliminating one of these confounders requires ethically problematic

doi:10.1007/s11948-017-9917-z

practices, such as violations of norms governing informed consent or foregoing significant improvement in the well being of society.[1]

To my knowledge, there is no ongoing clinical research on AAIs that is expected to significantly increase the quality life span of subjects. So what follows is not directed at current research, or any other research that doesn't aim to significantly extend the quality life span of humans. But it is likely that this line of research will continue to develop and that at some point there will be clinical trials of AAIs aimed at significantly increasing quality life span. As this research develops, it is critical to consider the ethics particular to research on AAIs.

2. Inferring Effectiveness

If society is going to allow the scientific community to pursue research on AAIs, the intention of the research must be to contribute to an AAI that will be made available to the public, even to those without disease. If the research aims to merely develop an AAI that will be available only to those suffering from a chronic disease such as cardiovascular disease or diabetes, then the research may be incentivizing the acquisition of these diseases.[2] It could be that the extension of quality life span from the developed AAI when administered to those suffering from chronic diseases still does not meet the quality or length of the expected life span of someone not suffering from chronic disease, in which case there would be limited utility of the AAI. But if the AAI extends the life span for those suffering from chronic diseases beyond

[1] The problems with anti-aging clinical trials are not problems with current anti-aging clinical trials. The clinical research is not yet developed enough to be measuring the full effect interventions have on the extension of quality life span.

[2] Unless aging itself is considered a disease, in which case everyone would suffer from it. Indeed, for AAIs to get through regulatory approval, it may be necessary to classify aging as a disease.

that of those who don't so suffer, then those who don't so suffer have incentive to acquire chronic diseases. Incentivizing the acquisition of chronic diseases is unethical for a variety of obvious reasons. Thus, if the scientific community is going to develop AAIs, it must be so that they extend the life span of anyone, not simply those who suffer from chronic diseases.

Prior to an AAI being made available to the public, researchers must have reliable information regarding the full effect of the intervention, including how long it increases quality life span and what those years are likely to be like. Researchers need this information for all populations of people to whom the AAI will be available. Thus, if the AAI will be made available to any healthy adult, researchers need to know what the effect will be on healthy adults.

Failing to understand the full effects of an AAI on healthy adults poses a serious risk to the public, especially if the AAI will be made available to them. One way such a failure can cause harm is that making it available to the public without this knowledge increases the risk of the public suffering serious adverse events. Another way is that even if there are no adverse events associated with the intervention, if it is not as effective as advertised it can cause users of it to believe that they will live longer than they actually will, which can cause significant harm not only in how this belief influence their behaviors but also when they come to find out that they will not live as long as expected.[3]

3. Confounding Variable: Pre-intervention Behavior

[3] Some researchers may think that the research into AAIs need continue only up to the point where it can be demonstrated that the AAI is better than alternative treatments. Notwithstanding the controversy over equipoise, it's not clear that an AAI made available to healthy subjects will ever be demonstrably "better" than not taking an AAI, assuming it's a treatment of aging. It could be that for many individuals there are good reasons to not want to live an extended life. Also, it's not clear that it is, in fact, better to live longer. Certainly the potential for increases in well being is great, but it's difficult to know that, all things considered, it is better.

doi:10.1007/s11948-017-9917-z

The development of AAIs must aim at interventions available to the public, which requires that researchers know the full effect of the intervention. It is insufficient to know whether the intervention extends quality life span; researchers must know to what extent it increases quality life span. In order to know the full effect of an AAI, researchers must be able to infer conclusions about how it is likely to affect the public from clinical trials on a sample population of the public. It is this generalization that creates tension between, on the one hand, the need to know the full effect of the intervention and, on the other hand, the ethical conduct of clinical trials. In short, acquiring the data sufficient to make such a generalization requires introducing ethically dubious practices into the clinical trial.

Consider first the current clinical trial evaluating the effect of metformin. This study is using subjects who are over sixty years old and pre-diabetic, but have not yet developed the disease. This study may yield reliable information about how metformin prevents the development of diabetes, which is one way in which a person's life could be extended. But the sample population will not provide reliable information on what the full effect of the drug is on the extension of quality life span. One reason is that the population is already pre-diabetic. Thus, it is impossible to infer what the effect will be on those who are not pre-diabetic. Second, the age of the subjects is sixty or older. This means that all of the subjects will have already aged significantly. Whatever the outcome is, the value of that outcome could be partly attributed to the health and behavior of the subjects prior to participating in the study.

This second reason generalizes to any clinical trial of an AAI. If the subjects have aged significantly prior to participating in the trial, then the value of the outcome (extension of quality life span) could be in part caused by the subjects' pre-intervention health and behavior. In other

doi:10.1007/s11948-017-9917-z

words, once the results are in, there are multiple potential explanations for the values observed. One explanation of the values is that they are what they are due to a causal relation between the AAI and the values. Another explanation is that they are what they are due to (a) a causal relation between the AAI and the values and (b) the pre-intervention health and behavior of the subjects. The problem is that there is no way to eliminate (b), so it's impossible to rule out the conjunction of (a) and (b).

There are techniques that can be used to minimize the effect a potential confounder such as (b) can have on the results, such as the randomization of the subjects. But there is no way of knowing what pre-intervention health or behavior may be influencing the value of the outcome variable, so these techniques will not help in eliminating the confounder. There are myriad behaviors or health conditions that can influence the quality life span of a person, and many of these occur much earlier in life, far removed from the outcome of interest. Many of these may be behaviors or conditions of which the subjects are completely unaware. For instance, suppose researchers randomize the assignment of significantly aged subjects into two groups, an experimental and a control. No matter the randomization, for each individual value—the measurement of the extension of a subject's quality life span—it will be impossible to say whether the value is what it is because of the intervention, or partly the intervention and partly the subject's pre-intervention health, unless the variables that contribute to the length of quality life span are controlled across all of the subjects. But this is impossible, given the many variables that influence the length of quality life span. There is thus no way to control for these conditions or behaviors across the study population, even with randomization, which means it will be

doi:10.1007/s11948-017-9917-z

impossible for a clinical trial to generate results that can support merely (a) rather than the conjunction of (a) and (b).

This problem is particular to clinical trials testing AAIs, because the outcome of interest shares a causal relation with so many independent variables, not all of which can be controlled across a subject population consisting of people who have lived adult lives in diverse environments and behaving in different ways. This is not the case for outcomes related to chronic diseases such as diabetes, stroke, or heart disease, because the range of conditions and behaviors to which they bear a causal relation is comparatively much smaller.

Moreover, when the development or severity of these chronic conditions is the outcome of interest, it may be possible to avoid clinical trials altogether, and instead use observational research to evaluate different exposures and how they influence the outcomes. But observational research will never tell us how effective an AAI is, even if that intervention is already routinely used for the treatment of another condition, because the treatment, in the case of an observational study, wouldn't be given to subjects who lack the condition for which the intervention is intended (healthy subject wouldn't be exposed). This means that the effectiveness of the treatment cannot be determined, as all of the observations would be of people who entered the study with a chronic condition.

If the subjects in the study tend to be older or have aged significantly prior to participation, it may be possible to determine whether an AAI is effective, but not to what extent it is effective, as the subjects' pre-intervention behavior and health confound such an inference. Eliminating these confounders requires eliminating the influence they could have on the outcome, which requires that the subjects' pre-intervention behaviors and conditions are neither

doi:10.1007/s11948-017-9917-z

diverse nor likely to have significant impact on the age of death. This requires that the subjects not have aged significantly.

4. Confounding Variable: Post-intervention Behavior

If the subjects in a clinical trial have aged significantly, then there is no way to rule out the influence of pre-intervention health and behavior on the observed values of the outcome, namely the extension of quality life span. And if this is impossible, then it is also impossible to make inferences regarding the extent of an AAI's influence on the dependent variable.[4] Given that the confounding variable is inherent to the subject population, the only way to eliminate it is to use a different subject population, one which hasn't aged significantly at the time of the intervention, which generally requires that it be composed of young subjects.

However, if a clinical trial uses younger subjects, subjects who have not aged significantly, then a different confounding variable is present, the subjects' *post*-intervention health and behavior. If the intervention begins at a young age, say, thirty, but the outcome of interest doesn't occur until decades later, then there are many years for subjects to behave in ways that could confound the results.

Specifically, the confounding variable is that subjects will behave differently if they believe they may live longer than they otherwise would have. The belief that one may live longer than he or she otherwise would is likely to have significant influence on how one lives his or her life. As things are now, most people in developed countries likely believe that they will live into

[4] A trial using subjects who have aged significantly could still indicate whether an AAI is effective, just not how effective it is. And it is the extent of its effectiveness that is at issue here.

their sixties, seventies, or eighties, barring an unlikely illness or accident earlier in life.[5] This belief must have an influence on how a person lives, though it may be nearly impossible to measure this. People who believe they will live well into their hundreds instead of their seventies or eighties may behave in ways that are more risk averse and overall healthier, or they may exhibit riskier and unhealthier behavior, thinking that his or her life expectancy is determined by the intervention.[6]

Beliefs of this sort are commonly controlled by a placebo, so that no subjects know whether they received the intervention or the placebo. But in the study population of interest, even a blinded study may fail to prevent the confounder. The reasons for this are twofold. First, the intervention occurs over a long period of time (because the subjects are young and the outcome decades after the intervention's introduction) and, second, the intervention is central to how one lives one's life. At a minimum, subjects in the study will withhold belief about how long they will live, which itself can cause changes in a person's behavior.

To be clear, I am suggesting that a young subject's belief that he or she is enrolled in a clinical trial testing an AAI will influence his or her behavior such that he or she wouldn't have behaved that way had he or she not been enrolled in the trial, regardless of whether they believe they received the intervention or a placebo. [7] It is admittedly difficult to evaluate the truth of this

[5] I use this life expectancy merely as an example because it represents the life expectancy of men and women in most democratic, developed nations. There are of course communities, large and small, who have different life expectancies.
[6] 'Life expectancy,' as I use the phrase here, refers to a person's mental state, namely an expectation of how long one will live. This is distinct from 'quality life span,' in that a quality life span is independent of a person's mental state.
[7] The mechanism of confounding is slightly different between the two variables. Pre-intervention health and behavior is an effect moderator, as it plausibly moderates the effect of the dependent

doi:10.1007/s11948-017-9917-z

counterfactual, but there are some reasons to think that beliefs about one's life expectancy influence one's behavior.

Suppose for example that subjects at the age of thirty are told that they are enrolled in a study on an AAI. They are also told that they are getting either the intervention or a placebo and that the odds of getting one rather than the other are 1:1. They are told that if they get the intervention, their life expectancy goes from eighty years old to one hundred and ten.[8] They are told that if they get the placebo, then their life expectancy will not change. A subject in this clinical trial now believes that he or she has approximately a fifty-fifty shot at living to one hundred and ten, and approximately a fifty-fifty shot at living to eighty, for instance.[9] It's hard to see how this belief *wouldn't* influence the subject's behavior. The subject may live the next few decades eating a healthy diet and getting regular exercise while abstaining from drugs and alcohol and avoiding risky behaviors like driving over the speed limit or riding a bike without a helmet. Or, thinking that her life expectancy is determined, she may instead engage in unhealthy behaviors and welcome risk.

---

variable, but the moderating variable is not brought on by the independent variable. Post-intervention behavior, however, is a mediator, as it mediates the values of the dependent variable, and is brought on by the independent variable. Sometimes people use 'confounder' to refer to any variable that causes one to make a type I error. And in the case of clinical trials of AAIs, that's the type of error one would likely be making. I adopt this use here.

[8] I discuss this below, but researchers could choose to withhold the life expectancy. Doing so may be unethical. Alternatively, researchers may instead inform subjects that their life expectancy will be significantly extended. Presumably, being told one will live many years or decades longer and being told one will live significantly longer would have the same effect on the subjects' psychology, as the subjects would think that the extension of live by many years or decades is a significant extension.

[9] A reasonable person will not assign a subjective probability of .5 to either possibility, as one would recognize that there is a small chance that one dies much earlier.

doi:10.1007/s11948-017-9917-z

Further, it is widely accepted that a person's beliefs about his or her health influences his or her behavior (Carpenter 2010; Rosenstock, Strecher, and Becker 1988; Janz and Becker 1984). A person's beliefs about the likelihood of an outcome, its severity, its benefits, the time between belief and behavior, and susceptibility to the outcome can all influence health-related behaviors. Thus, it is reasonable to think that a person's belief about how long he or she is likely to live will influence health-related behaviors, especially since age of death is one of the more salient health outcomes for a person.

When the trial is over, whatever the results are, one potential explanation will be that the observed effect bears a causal relation with the AAI. Another potential explanation is that the observed effect is due to (a) a causal relation between the observed effect and the AAI *and* (b) the post-intervention behavior of the subjects that ensues from their belief that they may live longer. There would be no way to rule out (b).[10] And if it is impossible to rule out (b), then it will confound the results. That is, the observed effect could be due in part to the AAI and in part to the subjects' beliefs that they may live longer—or their withheld belief about life expectancy—which means that it is not possible to infer conclusions regarding the full effect of the AAI.

Manipulating subjects' mental states by providing vague or incomplete information to the them will not help to eliminate the confounding post-intervention behavior. If instead of providing subjects complete information on the possible extension of life expectancy, researchers merely informed subjects that they may be getting a drug that may increase life expectancy, then

[10] It would be possible to rule out (b) if researchers had complete information regarding all of the subjects' behavior over the course of the trial, because then the results could be adjusted according to these observations. But it is not logistically possible to collect this information in a way in which it would be trustworthy.

the post-intervention behavior may confound the results to an even greater degree than it would if complete information was provided. Doing so would fail to prevent the subjects' beliefs and expectations of life expectancy from influence post-intervention behavior. But by providing vague or incomplete information to the subjects, researchers also allow for greater variability in subjects' interpretation of the information. Since the subjects' psychology is likely to influence their behavior and hence the outcome, leaving the interpretation of the information provided to them uncontrolled introduces an additional uncontrolled variable into the trial, which increases the potential for the subjects' post-intervention behavior to confound the results.

Furthermore, if researchers have more detailed information that may help subjects decide whether they wish to participate in the research but fail to provide it, then providing merely vague or incomplete information would fail to contribute to acquiring informed consent. Below, I argue that one way to avoid the dilemma is to not inform subjects that they are participating in the research, but in that case the potential for confounding is avoided. In the case of researchers providing merely vague or incomplete information, researchers may both fail to acquire informed consent as well as fail to avoid the confounding by post-intervention behavior.

It may be that researchers will be genuinely ignorant of the approximate effect an intervention may have on a subject's life expectancy, and so are in no position to provide anything but vague or incomplete information. Providing vague or incomplete information may have a more pernicious, but equally confounding, effect. Providing vague or incomplete information to subjects about how long they can expect to live may induce uncertainty about one's life expectancy, which may cause psychological harm to the subjects. Of course, no one is certain about how long one will live, but most of us are confident that, barring accident or

doi:10.1007/s11948-017-9917-z

unlikely illness, we will live to our life expectancy, whatever it is, and not much longer. We can be confident in the upper limit of our life span. Subjects in a blinded clinical trial won't be so certain, because they would believe that they will likely live to their unmanipulated life span *or* much longer. They would lack confidence in the upper limit of their life span. This lack of confidence in the end point of a person's life can cause psychological harm, harm that has no benefit to the person or society and which subjects wouldn't experience if they weren't enrolled in the research.

Inducing subjects' uncertainty about their life expectancy may also confound the results. Suppose at the age of thirty a subject is told he or she will be taking a drug that may increase life span, or it may not. This information defeats any justification the subject has to believe that he or she is likely to live to their unmanipulated life expectancy, whatever that happens to be. The appropriate epistemic response to this information is to withhold any belief about life expectancy. What is epistemically inappropriate is the confidence that most of have that we will live close to our life expectancy, barring accident or serious illness. The subject now has to live the remainder of his or her life agnostic about how long he or she is likely to live, including whether he or she will die before or after partners or children or friends. Living a life of such uncertainty seems likely to be a life full of psychological stress. And stress significantly increases the chance of developing chronic, life-threatening diseases, especially when experienced consistently over many years.

For the same reason, it may be unethical to introduce a placebo to blind subjects to which groups they are in. Doing so induces uncertainty about life expectancy, which may cause harm

doi:10.1007/s11948-017-9917-z

and stress in the same way that the uncertainty that results from vague or incomplete information causes harm and stress.

As is the case with a clinical trial using significantly aged subjects, randomization into experimental groups and control groups will also not eliminate the confounding variable. The source of the confounding variable is the subjects' belief that they are participating in a clinical trial and its likelihood to increase their quality life span. This is a belief that all subjects will have, whether or not their assignment to a control or experimental group is randomized.

Eliminating the confounder of pre-intervention behavior requires introducing the intervention earlier in the subjects' lives. But doing so introduces another confounder, the subjects' post-intervention behavior, which is mediated by their knowledge that they are participating in a clinical trial and their belief that they may live longer. The use of a placebo, which is the typical method of eliminating this type of confounder, is not likely to work. Even if it were ethical to blind subjects in an clinical trial evaluating the effectiveness of an AAI, subjects' health-related behaviors would still be influenced by their beliefs induced by their participation in the research, which would confound the results.

Thus, a clinical trial using such subjects will not be able to eliminate the effect believing one may live longer has on the outcome of interest, since this belief is likely to influence one's behavior over the course of the trial. The only way to eliminate the confounder of post-intervention behavior is to ensure that the subjects lack the belief that they may live longer than they would if they weren't participating in the research.

The first horn of the dilemma for conducting clinical trials of AAIs is the confounding variable of pre-intervention health and behavior. Resolving it impales one on the second horn,

the confounding variable of the post-intervention health and behavior. Below, I suggest a way between them, but the path is ethically problematic.

5. Requirements for a Clinical Trial

The only way to avoid both horns of the dilemma is to make sure that the young, healthy subjects lack the belief that participation in the clinical trial may cause them to live longer than they likely otherwise would.[11] There are several ways researchers could control the subjects' belief that they may live longer than they otherwise would.

5.1. Option 1: No Informed Consent

One of these ways of controlling the belief is to simply fail to inform subjects that they are participating in a clinical trial of an AAI. Informing human subjects of their participation in biomedical or behavioral research is a cornerstone of biomedical ethics, and codified in statutes (45 CFR 46), declarations (WMA Declaration of Helsinki), and reports (The Belmont Report) that govern biomedical research throughout the world. Under almost any interpretation of these, it would be unethical to not inform subjects that they are participating in a clinical trial evaluating the effect of an AAI.

However, there are some interpretations that may permit an exception. I believe a case could be made that the benefits of the research to the subject and to society are so great that an exception would be warranted, especially since the disclosure of participation would make it impossible to achieve those benefits. Whether this is a strong case would depend on many other

[11] One might think that it is possible to simply constrain the behavior of the subjects. In such a case the subjects would have the belief that they may live longer than they otherwise would, but the range of behaviors available to them is restricted. For several reasons this won't solve the problem. One is that it would require researchers to significantly limit subjects' liberty over the course of their lives. Another is that having such liberty restrictions could plausible affect the outcome of interest.

factors, such as the number of subjects required, available alternative research strategies, and the risks to the subjects. But it's possible that such a case could be made, if the benefits are great enough. It's just unlikely, and difficult to imagine an IRB ever approving this.

5.2. Option 2: Non-confounding Belief

Assuming that the subjects in a clinical trial of an AAI will be informed of their participation in the research along with the details of that participation, for their beliefs about how long they will live to not confound the results, it must be that they would have had those beliefs even if they weren't participating in the research. If it were the case that the subjects came into the research with certain beliefs about life expectancy, and the research didn't influence these beliefs in any way, then the subjects wouldn't behave any differently post-intervention than they otherwise would have.

In order for informed subjects' beliefs to not confound the conclusions, it must be the case that they are informed that the intervention will not significantly alter their life expectancy. Assuming that this is not a lie, this means that the subjects' beliefs about life expectancy are similar to the expected effect of the AAI. For example, if the subjects generally believe that they will live into their seventies or eighties, then they ought to be informed that the expected effect will be to prolong life, but not considerably so, perhaps by only a few years. If they are told that it would prolong life more than that, then researchers run the risk of this belief influencing the subjects' behavior.

If researchers pursue this option of making sure that the subjects' beliefs about life expectancy are congruent with the expected effect of the AAI, to make significant gains in the effectiveness of an AAI, subjects' beliefs about life expectancy will have to change. Since people

doi:10.1007/s11948-017-9917-z

get their beliefs about how long they are likely to live from a variety of sources, including common conceptions of life expectancy, the life spans of family and friends and other people in their communities, in addition to published data on life expectancy, for people to believe that their life expectancy with the intervention may only be slightly longer than it otherwise would be, the belief needs to trickle down from these sources, however that happens.

Let us suppose that it does, and we have a community of people who think they will live a long time, and they believe this based on experience of family members living that long and common conceptions of life expectancy. There is a further problem with this population. Either the proposed intervention is similar to the their life expectancy or it isn't. If it is, then the resulting data will not show how effective the intervention is, because the subjects would have lived approximately that long without it.[12]

But if the intervention purports to significantly extend the subjects' life expectancy, then researchers run into the same problem as above--either the results will be confounded by post-intervention behavior or the researchers will have to refrain from informing subjects that they are participating in a clinical trial of an AAI. So, the only way to ensure that the subjects' beliefs about their own life expectancy will not confound the results, while informing them that they are participating in a clinical trial, is to ensure that the difference between how long they would expect to live if they weren't participating and how long they expect to live given their participation is small.

[12] Moreover, since the effect size of the AAI will be small, many more people will need to be enrolled for the research to be sufficiently powered.

The result of this is obvious: it will take generations to produce enough data to show that a given intervention will significantly extend quality life span over current expectations of quality life span. The alternative is to do it in one generation, but not tell any of the subjects that they are enrolled in a clinical trial testing an intervention designed to significantly increase quality life span.

So, which is ethically optimal? The research can be conducted quickly and efficiently and produce results that may demonstrate the full effectiveness of an AAI. The utility of this is that, compared to the alternative, the well-being of many more people may be improved, given that the intervention may be available to the public sooner. The cost is that for the subjects, they will live much of their lives deceived as to how long they are likely to live (though, if effective, they will live longer—it's another question entirely as to whether there is utility or disutility in this assertion).[13]

The alternative is to conduct the research over a much longer period of time, sacrificing a lot of potential well-being, which is a significant cost. But the benefit is that the subjects themselves will not be deceived (and, if living longer is a benefit, they won't benefit much, either).

[13] There are statistical methods to account for possible confounding, but they wouldn't work because they require that researches be aware of what the potential confounding variables are and have an idea of what effect they might have on an outcome. In the case of clinical trials on AAIs, they wouldn't have this information. There are also statistical methods to adjust for the multiple comparisons of different independent variables. But there are so many multiple comparisons that to meet contemporary conventions of statistical significance, the effect of an AAI would have to be huge, or researchers would need enormous sample sizes, both of which are unlikely.

doi:10.1007/s11948-017-9917-z

A final alternative that avoids the ethical difficulties detailed above and the horns of the dilemma is to conduct research on AAIs using a piecemeal approach. First, do a clinical trial on an AAI and its effect on gene expression, then do another one on its effect on congestion in the heart, then do others on its effect on the body's inflammatory response, and so on.

It's fine to do all of this research, but this approach will not yield information on the extent to which the AAI increases quality life span, because quality life span is dependent on so many other variables, and it's impossible to conduct clinical trials on all of them. Short of this, the only way to understand the full extent of a potential AAI is to conduct a clinical trial in which the outcome variable is not some specific biomarker or effect on blood sugar or whatever, but the extension of quality life span.

I take no stand on which of the alternatives is preferable. What I do take a stand on is that if researchers are going to pursue the development of AAIs, they will have to contend with the dilemma between the confounding variable of pre-intervention health and behavior and the confounding variable of post-intervention health and behavior. The only resolution of this dilemma requires manipulating the epistemic position of the subjects, which requires either refraining from providing informed consent to subjects, or conducting the research over multiple generations of subjects.

References

45 CFR 46, Available at: http://www.hhs.gov/ohrp/regulations-and-policy/regulations/45-cfr-46/ [Accessed September 1, 2016].

Anisimov, V.N. et al., 2008. Metformin slows down aging and extends life span of female SHR mice, 7(17), pp.2769–2773.

Carpenter, C.J., 2010. A Meta-Analysis of the Effectiveness of Health Belief Model Variables in Predicting Behavior. *Health Communication*, 25(8), pp.661–669. Available at: http://dx.doi.org/10.1080/10410236.2010.521906.

Janz, N.K. & Becker, M.H., 1984. The Health Belief Model: A Decade Later. *Health Education & Behavior* , 11 (1), pp.1–47. Available at: http://heb.sagepub.com/content/11/1/1.abstract.

Martin-Montalvo, A. et al., 2013. Metformin improves healthspan and lifespan in mice. *Nature communications*, 4, p.2192.

Metformin in Longevity Study - ClinicalTrials.gov. Available at: https://clinicaltrials.gov/ct2/show/NCT02432287 [Accessed September 1, 2016].

Rosenstock, I.M., Strecher, V.J. & Becker, M.H., 1988. Social Learning Theory and the Health Belief Model. *Health Education & Behavior* , 15 (2), pp.175–183.

The Belmont Report. Available at: http://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/ [Accessed September 1, 2016].

UMIN-CTR Clinical Trial. Available at: https://upload.umin.ac.jp/cgi-open-bin/ctr/ctr.cgi?function=brows&action=brows&recptno=R000024575&type=summary&language=E [Accessed September 1, 2016].

WMA Declaration of Helsinki, Ethical principles for medical research involving human subjects. Available at: http://www.wma.net/en/30publications/10policies/b3/ [Accessed September 1, 2016].

Yoshino, J. et al., 2016. Nicotinamide Mononucleotide, a Key NAD^+ Intermediate, Treats the Pathophysiology of Diet- and Age-Induced Diabetes in Mice. *Cell Metabolism*, 14(4), pp.528–536. Available at: http://dx.doi.org/10.1016/j.cmet.2011.08.014.