UC Berkeley

UC Berkeley Previously Published Works

Title

Why high-risk, non-expected-utility-maximising gambles can be rational and beneficial: the case of HIV cure studies.

Permalink

<https://escholarship.org/uc/item/7j58c5vk>

Journal

Journal of medical ethics, 43(2)

ISSN

0306-6800

Author

Buchak, Lara

Publication Date

2017-02-01

DOI

10.1136/medethics-2015-103118

Peer reviewed

Penultimate draft. Please cite the published version (Journal of Medical Ethics, 2016)

Why high-risk, non-expected-utility-maximizing gambles can be rational and beneficial: The case of HIV cure studies. Lara Buchak, UC Berkeley.

1. INTRODUCTION

Some early phase clinical studies of candidate HIV cure and remission interventions appear to have adverse medical risk-benefit ratios for participants. Why, then, do people participate? And is it ethically permissible to allow them to participate? Recent work in decision theory sheds light on both of these questions, by casting doubt on the idea that rational individuals prefer choices that maximize expected utility, and therefore by casting doubt on the idea that researchers have an ethical obligation not to enroll participants in studies with high risk-benefit ratios. This work supports the view that researchers should instead defer to the considered preferences of the participants themselves. This essay briefly explains this recent work, and then explores its application to these two questions in more detail.

2. DECISION THEORY AND RISK-TAKING

Decision theory addresses the question of what it is rational to choose, given one's aims, in situations involving risk. (This type of rationality is sometimes called 'instrumental' or 'means-ends' rationality.) The orthodox view has long been that individuals should *maximize expected utility*—a weighted average of the utility values a gamble might yield, each value weighted by the probability with which the gamble yields it. So, for example, if an individual faces a coin-flip between \$0 and \$100, represented as {HEADS, \$0; TAILS, \$100}, the value of that gamble should be $p(\text{HEADS})u(\$0) + p(\text{TAILS})u(\$100)$, where p corresponds to the probability the individual assigns to the possible states, and u corresponds to the utility value the individual assigns to the possible consequences.

It is important to note three things about the utility and probability functions. First, they are *the individual's* probabilities and utilities, not independent or objective values. Second, utilities attach not just to monetary values, but to consequences described so as to include everything an individual cares about. For example, in the context of choices about participating in HIV cure studies, one potential consequence of participation might be that the individual's condition improves—but other potential consequences include that other people's well-being improves, that one makes a contribution to science, or that one symbolically takes a stand against the disease. Finally, utility and probability functions are typically reverse-engineered from the individual's preferences: the decision-theorist examines a subject's choices and finds values which make sense of these choices as expected-utility-maximizing. It turns out

that if we assume that an individual maximizes expected utility, and if we know enough of her preferences, then there will be a unique way to assign these values—this is guaranteed by “Representation Theorems” that connect preferences about gambles to the existence of a unique probability and unique (up to affine transformation) utility function relative to which the individual prefers gambles that maximize expected utility.[1-3] To make the discussion easier to follow, I will here assume a “psychological realist” conception of utility: I will assume that utility corresponds to an actual psychological property such as the desirability of a consequence..

Two familiar phenomena are risk-averse and risk-seeking preferences. A simple example of risk-aversion is a preference for a sure-thing \$50 rather than a coin-flip between \$0 and \$100. More generally, to be risk-averse is to prefer, of two gambles with the same mean monetary value, the one that is less spread out.[4] And to be risk-seeking is just to have the opposite preference: to prefer, of two gambles with the same mean monetary value, the one that is more spread out—e.g., to prefer the \$0/\$100 coin-flip to \$50. (To be risk-neutral is to be indifferent between any two gambles with the same mean monetary value.)

There are a few possible explanations for risk-averse and risk-seeking preferences. The first is the standard explanation that follows from expected utility maximization: an individual with risk-averse preferences has a utility function that diminishes marginally—as she gets more money, additional bits of money add less value ($u(\$50) - u(\$0) > u(\$100) - u(\$50)$). And vice versa for the risk-seeking individual: as she gets more money, additional bits of money add more value. Thus, the explanation for risk-averse and risk-seeking behavior is that the individual has a utility function with a particular shape. This is an explanation that locates the source of the behavior in the desirability of various consequences.

This explanation is somewhat intuitive, but there are a number of problems with it. First, it seems to conflate two very different phenomena: how much an individual values particular consequences and how the individual approaches risky gambles that contain these consequences.[5, 6] For example, there is at least a psychological difference between preferring the sure-thing \$50 because one needs exactly \$50 for a bus ticket home, and preferring the sure-thing \$50 because even though one values money linearly (each \$50 matters as much as the next), one does not like to take risks—one thinks that “a bird in the hand is worth two in the bush,” so to speak. Second, even if we do not care about this psychological difference, the explanation generalizes poorly: it rules out many typical cases of risk-averse or risk-seeking preferences.[5, 7-9] This is to say: there are sets of preferences that many people find intuitive that are *incompatible* with expected utility maximization.

A second possible explanation for risk-averse and risk-seeking preferences comes in two versions, one of which is compatible with expected utility maximization and one of which is not. This explanation locates the source of the behavior in the probabilities. In the version compatible with EU-

maximization, the individual assigns probabilities consistently, but erroneously. (This possibility still leaves EU-maximization unable to capture the aforementioned counterexamples.) In the version incompatible with EU-maximization, the individual assigns more or less probability to events that are bad for her: she is irrationally pessimistic or optimistic in believing that the world will go in the way that's worst or best for her.[10]

Because expected utility maximization does not capture many intuitive preferences, a number of alternatives have been proposed. Typically these alternatives hold that *rational* individuals maximize expected utility, but that most individuals deviate from rationality in predictable ways. *Risk-weighted expected utility (REU) maximization*, however, is an alternative explanation for risk-averse and risk-seeking preferences according to which individuals do not maximize expected utility but are nonetheless *rational*. [6] This explanation locates risk-averse and risk-seeking preferences not in the utilities or the probabilities, but in a different kind of attitude altogether. The thought behind REU-maximization is this. Merely determining how much an individual values consequences and how likely she thinks various states of the world are to obtain is not enough to answer the question of how she should value a gamble that has some probability of realizing any one of a number of various consequences—it does not answer the question of how to aggregate the utility values of the possible consequences to arrive at a single value for the gamble. Taking an average weighted by probabilities is *just one way* to aggregate, a way that corresponds to holding that the weight of a possible consequence in one's practical deliberation is just the probability of that consequence.

Contra EU-maximization, some individuals might be more concerned about what goes on in relatively worse or relatively better states. Some individuals—*risk-avoidant individuals*—might be more concerned with what happens in worse states than better states, and thus might hold that the value of a gamble is closer to its minimum value than the expected utility maximizer holds. Other individuals—*risk-inclined individuals*—might be more concerned with what happens in better states than worse states, and thus might hold that the value of a gamble is closer to its maximum value. Finally, some individuals—*globally-neutral individuals*—might be equally concerned with what happens in all (equiprobable) states, regardless of their relative rank, and thus will simply be expected utility maximizers.

The way to account for these differences is to hold that each individual has, in addition to a utility and probability function, a *risk-function* that measures *the weight of the top p-portion of consequences in her practical decision-making*. Formally, risk-avoidant individuals have convex risk-functions: as the best consequences are realized in less likely states, they care proportionally less about them. (For example, the top 0.5 consequences may garner 0.4 weight, and the top 0.1 consequences may garner only 0.01 weight.) Risk-inclined individuals have concave risk-functions: as the best consequences are

realized in less likely states, they care proportionally more about them. (For example, the top 0.5 consequences may garner 0.6 weight, and the top 0.1 consequences may garner 0.2 weight.) Globally neutral individuals have linear risk-functions (the top 0.5 consequences garner 0.5 weight and the top 0.1 consequences garner 0.1 weight). An upshot of this proposal that will be important in what follows is that *a risk-inclined individual might prefer, of two gambles, the one with the lower expected utility*, provided that gamble has a more favorable ‘spread’—e.g. if the lower-expected-utility gamble has an unlikely consequence with very high utility, whereas the higher-expected-utility gamble is more concentrated around a few middling consequences.

The idea behind REU-maximization is that there are actually three psychological components in preference-formation and decision-making: how much an individual values consequences (utilities), how likely an individual thinks various states of the world are to obtain (probabilities), and the extent to which an individual is willing to trade off value in the worst-case scenarios against value in the best-case scenarios (the risk function). There are two different ways to think about the risk function: as a measure of distributive justice among one’s ‘future possible selves’—how to trade off the interests of the best-off possible self against the interests of the worst-off possible self—and as a measure of how one trades off the virtue of prudence against the virtue of venturesomeness.

Informally, when making choices, some people might care more about the potential downsides than the potential benefits, and others might have the opposite pattern of concern. Some might generally weight worse-case scenarios most heavily than best-case scenarios, others might do the opposite, and others might weight them equally. According to REU-maximization, *any of these patterns of concern is rational*. Of course, there will be some constraints on the risk function: for example, it must be non-decreasing in probability (informally: one must prefer a better chance at a good consequence rather than a worse chance). And there may be some patterns of concern that we deem unreasonable because they are too extreme in one direction or the other, just as we might think that some ways of valuing consequences are unreasonable. But in general, the range of acceptable attitudes towards risk is much broader than the orthodox view suggests.

Furthermore, since REU theory has a Representation Theorem, then if we assume that individuals maximize risk-weighted expected utility, we will be able to determine each of these three functions uniquely (again, with utility unique up to positive affine transformation) by examining an individual’s preferences. Thus, we now have a third possible explanation for risk-averse and risk-seeking preferences—an individual places proportionately more weight on what happens in worst-case or best-case scenarios, respectively—and a way to determine which of these explanations (or what precise combination of them) is behind an individual’s preferences.

3. RISK-ATTITUDES IN HIV CURE STUDIES

This last point brings us to our first application to the debate surrounding high-risk clinical trials. Risk-weighted expected utility maximization gives us a way to untangle several possible reasons that a subject might be willing to undergo a high-risk trial—and to evaluate whether, if these are the subject's reasons, we ought to allow him to undergo the trial. On the one hand, he might be willing to undergo the trials because he attaches utility to much more than the benefits to himself: for example, he may assign a high utility to both the unlikely consequence in which the trial is successful and to the likely consequence in which the trial is unsuccessful, on the grounds that either way he is helping HIV research progress. Or he might assign high utility to a successful trial because he assigns a high utility to no longer needing to take a daily drug regimen. On the other hand, he might be willing to undergo these trials because he (erroneously) estimates the probability of a successful trial to be much higher than it is, perhaps because of the language used to frame the study[11]. Finally, he might be willing to undergo these trials because he places more weight on what happens in the “good case” rather than the “bad case,” even though the latter is more likely—in short, he might be willing to undergo these trials because he is risk-inclined.

There are a few ways we might determine what role an individual's utilities, beliefs, and risk-attitudes are each playing in his willingness to undergo a high-risk clinical trial. The most “operational” way to do this is to solicit his preferences in a wide range of gambles involving health states. We know from the REU Representation Theorem that certain sets of preferences will suffice to separately reveal his utility function, probability function, and risk-function.[6] Roughly, his probability function is determined by looking at which events he would rather bet on. His utility function is determined by which pairs of consequences constitute “equal tradeoffs” in gambles that put consequences in the same order. For example, if our subject is just barely willing to accept a given chance of two weeks of severe flu-like symptoms in exchange for a given chance of having to take 3 pills a day rather than 6 a day for the rest of his life, and if he is also just barely willing to risk that same chance of two weeks of severe flu-like symptoms in exchange for that same chance of having to take 1 pill a day rather than 2, then the utility difference between taking 3 and 6 pills is the same as the utility difference between taking 0 and 1 pills: 3 vs. 6 and 1 vs. 0 are “traded off” for the same risk.¹ (Such a choice may hypothetically arise if the subject doesn't know whether he has HIV, but knows that if he undergoes a particular treatment then he will reduce the number of pills he must take if he has it, but he will experience negative side effects if he undergoes the treatment but does not have it.) The subject's risk-function is determined by the *magnitude*

¹ More formally: if the subject is indifferent between {p, 2-week flu; 1-p, 3 pills} and {p, no flu and no pills; 1-p, 6 pills} and is also indifferent between {p, 2-week flu; 1-p, 1 pill} and {p, no flu and no pills; 1-p, 2 pills}, then $u(3 \text{ pills}) - u(6 \text{ pills}) = u(1 \text{ pill}) - u(2 \text{ pills})$. We are assuming the subject would rather endure two weeks of flu than take any number of pills for the rest of his life, an assumption we can confirm from his preferences.

(in utility terms) of the tradeoff required: the ratio of utility that can be lost in the worst event relative to that gained in the rest of the gamble. For example, if we discover (through the above method) that two subjects have the same utility function concerning pills per day and concerning flu-like symptoms, and that our first subject is just barely willing to accept the possibility of two weeks of flu-like symptoms in exchange for the possibility of reducing the number of pills he must take if he has HIV but our second subject is willing to accept the possibility of two months of flu-like symptoms in exchange for this, then our second subject is more risk-avoidant: he is willing to accept a worse best-case scenario than our first subject, in exchange for the same improvement in the worst-case scenario. Of course, this method requires the subject to know his preferences about a wide range of gambles, which may make it cumbersome.

A slightly less formal method involves assuming that a subject has some access to his utility function and beliefs about likelihoods, and eliciting these more directly. There are a number of standard ways of eliciting a subject's utility function concerning health states, many of which don't involve gamble preferences.²[12] We can query the subject directly about how likely he thinks these states would be if he were to undergo a research trial or not do so. Finally, we can elicit his risk attitudes, as above, by looking at the magnitude of the tradeoffs he is willing to accept.

The tradeoff-magnitude method of eliciting a subject's risk-attitudes uses hypothetical choice. We might instead elicit his risk-attitudes via self-reported behavior—the researcher asks which risky activities he has engaged in over the past year and whether he endorses having engaged in those activities; via other-reported behavior—the individual's personal care physician or psychiatrist or spouse assesses how much the subject seeks or avoids risk; or via observed behavior—the researcher invites the subject to perform simple choice tasks which involve options with various levels of riskiness, and observes the subject's choices.

A few caveats are in order, regardless of which of these methods are employed. First, in order to determine individuals' actual risk-attitudes with respect to the actual trials at hand, it is important that we determine their preferences over hypothetical gambles involving *health*, rather than, say, money. This is because people tend to have different risk-attitudes in different domains.[13] Second, we need to make sure that individuals' preferences are free from influences that are genuinely irrational. One such influence is *framing effects*: individuals tend to have different preferences among the very same gambles

² One popular method which cannot be used here is the "standard gamble" method, which assumes that individuals maximize expected utility. This method is similar in form to the tradeoff-magnitude method discussed above. The difference is that the standard gamble method assumes a particular risk-attitude (global neutrality, i.e. EU-maximization) and uses a subject's gamble preferences to elicit utilities; whereas the tradeoff-magnitude method assumes particular utilities (elicited using a different method) and uses a subject's gamble preferences to elicit risk-attitudes.

if these gambles are presented in different ways, for example if an individual's perception of the status quo is shifted so that gambles are perceived as losses rather than gains.[14] Thus, we must be careful to determine which frame the subject thinks is correct—which frame the subject takes to correspond to his true preferences. (For example, we can make him aware that a particular consequence can be thought of either as a loss relative to the life he was enjoying before he contracted HIV or a gain relative to his current state, and ask him which way of thinking he endorses.) Another type of framing effect we may need to be wary of is referring to the same consequence as “curing HIV” or as “lifetime remission of HIV”.

4. ALLOWING RESEARCH SUBJECTS TO TAKE RISKS

If, indeed, individuals are willing to undergo high-risk trials because they are risk-inclined, then this naturally raises the question of whether researchers ought to allow individuals to undergo these trials. And if risk-inclined preferences are indeed rational, then we have a potential reason to allow individuals who indeed choose these trials because they are risk-inclined to undergo them. If rational individuals maximize risk-weighted expected utility according to their own risk-attitudes, then maximizing expected utility is just one way of making decisions. Therefore, there is nothing particularly special about a trial's having positive expected utility (a ‘favorable’ risk-benefit ratio).³ Again, a risk-inclined individual might rather have some small chance at a benefit, despite a larger chance of a downside, than forgo this chance and remain in the current state, even if taking the chance has negative expected utility—and this preference is no more or less rational than a preference for gambles that maximize expected utility.

More generally, the rule *only allow people to sign up for trials with positive expected utility* cannot be justified on the basis that rational people only undergo trials with positive expected utility. There is nothing special about a trial's having positive expected utility, for an individual who does not aggregate value in this way. A rule that allows individuals only to undergo trials when it is rational for them to do so would instead say: *only allow people to sign up for trials with positive risk-weighted expected utility, relative to their own risk attitudes*. (We could add: as long as their risk-attitudes are reasonable.) Of course, everything should be done to minimize the risks to the subject and maximize the benefits, but this alternative rule has important upshots for which trials we should impose on individuals:

³ I note that codes of ethics do not typically mention a precise formula for quantifying acceptable risk-benefit ratios[15-17], although this does not imply that these judgments cannot be made.[18] Furthermore, it is helpful to precisify this standard to consider the ideal case in which utilities and probabilities are known precisely, to illuminate the messier real-world case. Informally, this paper is directed against the view that a risk-benefit ratio that is unfavorable according to a ‘hypothetical reasonable person’ is unfavorable to actual trial subjects. I assume that ‘favorable risk-benefit ratio’ is precisified by ‘EU-maximizing’, since this is the orthodox view, but even without that precisification the general points of this paper will still hold. The key point is that there is a wide range of acceptable attitudes towards risky gambles, and there is no single risk-benefit ratio that all reasonable persons require.

risk-avoidant individuals require a higher risk-benefit ratio than we previously thought, and risk-inclined individuals require a lower risk-benefit ratio than we previously thought.

This alternative rule is consistent with the way we treat utility: since there is no privileged way to value health-states—some might place more value on longevity and others on physical abilities—we think it is up to individuals to value their own health-states however they want, and allow them to make choices accordingly. Similarly, since there is no privileged way to determine how much risk to accept—some prefer a guarantee of a not-too-bad consequence and others prefer a small chance of a terrific consequence—we ought to allow individuals to choose the tradeoffs they are willing to make between how things go in worse scenarios and in better scenarios.

Keep in mind that the rule *defer to others' risk attitudes* does not prioritize individuals' right to choose *over* doing what is in their best interests. Instead, taking high-risk gambles *can be in an individual's best interests*, because what is in an individual's interests depends on what she values: just as it is up to individuals to decide what their interests (utilities) are, so too is it up to individuals to decide how to weigh risks. When choosing under risk—when we don't know exactly how our decisions will turn out—and when no one choice is better than others under all circumstances, there is no answer to the question of which choice is in an individual's interests, apart from answering the question of what his attitude towards risk is. Thus, taking an individual's risk-attitudes into account when allowing him to undergo gambles such as medical trials seems not only permissible, but unavoidable.

So far, we've been discussing how to determine if a subject should be enrolled in a trial. But trials are also assessed by a research ethics committee before we know which individuals will potentially participate, and so we must also say when a study should be approved prior to the enrollment of subjects. On the model I argue against, a trial must maximize expected utility for the participant—or have a favorable risk-benefit ratio according to a hypothetical reasonable person—in order for a study to be approved; and a subject may then decline to participate after he is made aware of the risks and benefits. However, if I am correct, then there is not a single, objective standard for which risk-benefit ratios a reasonable person would accept. Furthermore, given the wide range of acceptable risk-attitudes, the 'hypothetical reasonable person' standard is likely to license trials that more risk-avoidant individuals would not accept. Thus, even proponents of the reasonableness standard cannot skip carefully determining the preferences of the individual himself. Consequently, the practical recommendation is that RECs should not assess whether a *hypothetical* reasonable person would participate in the trial, or whether *any* reasonable person would, but whether *some* reasonable people would—and then that participants who are too risk-avoidant to license participation should be screened off. Clearly, this puts more pressure on the protocols for determining whether a subject genuinely and rationally wants to take

the relevant risk. Nonetheless, our current practices have two problems: they withhold, from subjects who are more risk-inclined than the hypothetical reasonable person, possibilities that are genuinely in their interests by their own (reasonable) lights; and, worse, they may allow (by not sufficiently attending to the step of eliciting individual preferences) subjects who are more risk-avoidant than the hypothetical reasonable person to undergo trials that are genuinely not in their interests.

5. OBJECTIONS

I close by considering four objections to this proposal. The first objection is that we do consider risk-inclined behavior to be irrational or unreasonable, even though we don't consider risk-avoidant behavior to be irrational. For example, it is common to deride the practice of playing the lottery, while we tend not to deride the practice of buying (even non-EU-maximizing) insurance. However, I claim that the explanation for why we think that lottery players are being irrational is not that risk-inclination is itself irrational, but rather three other reasons. First, many who play the lottery are irrationally optimistic—they overestimate the probability of their winning. Second, many of those who play the lottery are not risk-inclined in other contexts, and so we think they are irrational in the sense of being inconsistent about what they value. Finally, *repeatedly* playing the lottery is irrational for nearly everyone, even very risk-inclined people, because repeatedly taking any gamble has the effect of concentrating all the probability near the mean value (in essence, eliminating the risk and producing a gamble with a near-sure-thing result of its mean value, which for lotteries is negative).

The second objection is that this approach could create a conflict between the researcher's recommendations and the trials the subject is allowed to undergo. Imagine that a subject is inclined to undergo a trial that the researcher himself would not undergo, because the subject is more risk-inclined than the researcher. We may balk at allowing the subject to undergo the trial for two reasons. First, we might speculate that there is an asymmetry in when we should defer to an individual's risk-attitude instead of our own: I can choose options for another that are *less* risky than options I myself would take if he is less risk-inclined than I am, but I can't impose risks that I wouldn't take myself, even if the other person wishes it. In response, note that it's not really the 'riskiness' per se of the options that is relevant here: what is relevant is that depending on how things turn out, undergoing the trial might be better than not undergoing it or it might be worse. One cannot justify removing the possibility of the good consequence from someone else, for the sake of removing the possibility of the bad consequence, when the other person does not wish this. When we carry out choices for others, and we know their wishes, and we know that these wishes are well-informed, not irrational, and not for something objectively more harmful than some alternative, then we ought to defer to their preferences, including their preferences about which risks they find worthwhile.

Second, we may balk because although in general a decision-maker ought to make decisions based on the risk-attitude of the individual affected rather than his own risk-attitude, things are different when the decision-maker is more knowledgeable than the subject, as is the case when a physician can back up his negative advice with medical evidence. However, in response, notice that medical evidence is not direct evidence about what treatment is best for a subject; rather, it is evidence about the probabilities of various consequences under each treatment. In short, medical evidence is about what “gamble” a particular treatment amounts to, not directly about what treatments a rational individual should undergo. As long as the researcher can make the subject aware of which gamble he faces, the researcher should defer to the subject’s preferences.

The third objection trades on another conflict between how the subject thinks about a research program and how the researcher thinks about it. Although the consequence of each trial is unknown, the researcher can expect to make predictions about consequences in the aggregate. Let us assume that each of a researcher’s subjects is highly risk-inclined, and that the possibility of a cure or positive consequence from some treatment is low (and let us assume for the sake of argument that whether each subject is cured is uncorrelated). Each of the subjects is only living one life, but the researcher has many subjects. If the researcher implements the risky strategy, then he can be sure that his subjects will be worse off in the aggregate than if he implements a less risky strategy. Therefore, regardless of his risk-attitude, he will, because he is treating a large number of subjects, think that the less risky strategy is ethically better.

In response, it is important to note that when a researcher chooses what course of action to implement, this shouldn’t primarily be thought of as a choice about what he himself values. Rather, it is a choice about which overall society is best to realize. And it is true that if one group of people is each making choices that maximize expected utility, and another is not, then (given the law of large numbers) the first group will have a higher average utility. However, why should we think that *average* well-being is what matters? After all, it is not as if each individual in the society gets to experience the average well-being. What matters is how each individual person fares, and, again, there is no way to answer this question ahead of time without knowing each individual’s attitude towards risk. In the envisioned scenario, each subject holds that the risky strategy is better *for him*, because it better realizes his aims when it comes to risky gambles.

The final objection is that the alternative rule would disproportionately impose burdens on certain segments of the population. We know that some groups of people tend to be more risk-seeking than others: for example, men tend to be more risk-seeking than women.[19] Indeed, it may be the case that those who get HIV in the first place are disproportionately risk-inclined, so those with HIV will be more inclined to agree to the relevant trials than, e.g., those with cancer will be to agree to similarly risky trials. (Incidentally, this may be one point in favor of risk-attitudes as an explanation for why people are willing

to participate in HIV trials.) It is true that we are giving each participant what he rationally wants, but there may be a societal cost to imposing burdens disproportionately on some group, even if members of this group willingly accept the burdens.

In response, it is important to point out that in the case of HIV trials, we do not merely impose burdens on individuals. There is some chance of a downside but some chance of an upside. So we're also disproportionately imposing the (small chance of) advantages on some group. This group thinks these two facts more than balance out, other groups think the opposite. And, importantly, individuals do not have the preferences they have because of some background disadvantage. When we envision a world free from injustice, it is one in which no one is so poor that it becomes rational to, say, accept a very-high-interest loan so that they can feed their family this week—but it needn't be one in which no one is so risk-inclined that they care twice as much about what happens in the best-case scenario than about what happens in the worst, because there is nothing inherently bad for the individual about having this pattern of concern.

Still, while the initial versions of these last two objections can be answered, the general points behind them merit further discussion. What are the effects, both on a research program and society at large, of allowing research participation to be guided by individuals' own risk-attitudes? If subjects are selected without regard to social group, but if a particular social group is more likely to want to participate because of a characteristic more common to this group, then does this constitute unfair subject selection? While I have shown that the reasoning *only give EU-maximizing gambles to individuals, because non-EU-maximizing gambles are irrational or contrary to individuals' well-being* is flawed, there might be other reasons not to allow individuals to take (even individually rational) risks. It may be, for example, that the best way to think about research participation is not in terms of whether it is rational (i.e. whether it benefits the individual) at all, but rather in terms of some other criteria. With the flawed reasoning for accepting the standard rule out of the way, we can move on to more fruitful discussion of this rule and the alternatives.

REFERENCES

- 1 Ramsey F. Truth and Probability. In: Braithwaite R.B., ed., *The Foundations of Mathematics and other Logical Essays*. London: Kegan, Paul, Trench, Trubner & Co., New York: Harcourt, Brace and Company 1931:156-198. First published 1926.
- 2 Von Neumann J, Morgenstern O. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press 1944.
- 3 Savage L. *The Foundations of Statistics*. John Wiley and Sons 1954.
- 4 Rothschild M, Stiglitz J. Increasing Risk: I. A Definition. *J Econ Theory* 1970;2:225-43.

- 5 Hansson B. Risk-aversion as a Problem of Conjoint Measurement. In: Gärdenfors P, Sahlin N., eds. *Decision, Probability, and Utility*. Cambridge: Cambridge University Press, 1988:136-158.
- 6 Buchak L. *Risk and Rationality*. Oxford: Oxford University Press 2013.
- 7 Allais M. Fondements d'une Théorie Positive des Choix Comportant un Risque et Critique des Postulats et Axiomes de L'Ecole Americaine. *Econometrie* 1953:257-332.
- 8 Rabin M. Risk-aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica* 2000;68(5):1281-92.
- 9 Starmer C. Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. *J Econ Lit* 2000;38:332-82.
- 10 Quiggin J. A Theory of Anticipated Utility. *J Econ Behav Organ* 1982;3:323-43.
- 11 Dubé K, Henderson G, Margolis D. Framing Expectations in early HIV cure research. *Trends Microbiol* 2014;22(10):547-49.
- 12 Froberg D, Kane R. (1989), Methodology for Measuring Health-State Preferences—II: Scaling Methods. *J Clin Epidemiol* 1989;42(5):459-71.
- 13 Rottenstreich Y, Hsee C. Money, Kisses, and Electric Shocks: On the Affective Psychology of Risk. *Psychol Sci* 2001;12(3):185-90.
- 148 Tversky A, Kahneman D. The Framing of Decisions and the Psychology of Choice. *Science* 1981;211:453-8.
- 15 National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. Washington, DC: US Government Printing Office; 1979.
- 16 World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects. *JAMA* 2013;310(20):2191-94.
- 17 *Electronic Code of Federal Regulations* Title 45 Part 46.111. <http://www.ecfr.gov/> Retrieved 2016-04-21.
- 18 Emanuel, E, Wendler D, Grady C. What Makes Clinical Research Ethical? (Reprinted) *JAMA* 2000;283(20):2701-11.
- 19 Byrnes, J, Miller D, Schafer W. Gender Differences in Risk Taking: A Meta-Analysis. *Psychol Bull* 1999;125(3):367-83.