

## The Risk-Escalation Model: A Principled Design Strategy for Early-Phase Trials

Spencer Phillips Hey and Jonathan Kimmelman

### Abstract

Should first-in-human trials be designed to maximize the prospect of therapeutic benefit for volunteers, prioritize avoidance of unintended harms, or aim for some happy medium between the two? Perennial controversies surrounding initiation and design of early-phase trials hinge on how this question is resolved. In this paper, we build on the premise that the task of early-phase testing is to optimize various components of a potential therapy so that later, confirmatory trials have the maximal probability of informing drug development and clinical care. We then explore three strategies that investigators might use to manage trial risks while optimizing a therapy, using cell therapy for Amyotrophic Lateral Sclerosis (ALS) as an example. We argue that an iterative application of maximin strategies over successive cohorts and trials, which we call the “risk-escalation model,” establishes a moral principle that should guide decision-making in early-phase trials.

### INTRODUCTION

Should first-in-human trials be designed to maximize the prospect of therapeutic benefit for volunteers, prioritize avoidance of unintended harms, or aim for some happy medium

#### Contributors

*Joseph J. Fins*, MD, MACP, is the E. William Davis, Jr., MD Professor of Medical Ethics and Chief of the Division of Medical Ethics at Weill Cornell Medical College where he also serves as Professor of Medicine. He is the immediate past-president of the American Society for Bioethics and Humanities. His forthcoming book, *Rights Come to Mind: Brain Injury, Ethics and the Struggle for Consciousness* will be published by Cambridge University Press.

*Daniel P. Sulmasy*, MD, PhD, is the Kilbride-Clinton Professor of Medicine and Ethics in the Department of Medicine and Divinity School at the University of Chicago, where he serves as Associate Director of the MacLean Center for Clinical Medical Ethics and Director of the Program on Medicine and Religion. He is a member of the Presidential Commission for the Study of Bioethical Issues. His latest book is the second edition of *Methods in Medical Ethics*.

*Rebecca Dresser*, JD, is the Daniel Noyes Kirby Professor of Law and Professor of Ethics in Medicine at Washington University in St. Louis. Since 1983, she has taught medical and law students about legal and ethical issues in end-of-life care, biomedical research, genetics, assisted reproduction, and related topics. From 2002 to 2009, she was a member of the President’s Council on Bioethics.

*Spencer Phillips Hey*, PhD, is a postdoctoral fellow in Biomedical Ethics at McGill University and a member of the Studies in Translation, Ethics, and Medicine (STREAM) research group. He received his PhD in philosophy from the Rotman Institute of Philosophy at the University of Western Ontario.

*Jonathan Kimmelman*, PhD, is Associate Professor in Biomedical Ethics, Experimental Medicine, and Social Studies of Medicine at McGill University, and directs the Studies in Translation, Ethics, and Medicine (STREAM) research group.

*Franklin G. Miller*, PhD, is a member of the senior faculty in the Department in Bioethics, National Institutes of Health (NIH), and Special Expert, National Institute of Mental Health Intramural Research Program. His principal current research interest is examination of ethical issues in clinical research, including study design, informed consent, and the ways in which clinical research differs from medical care.

*Alberto Giubilini*, PhD, is a research fellow at the Centre for Applied Philosophy and Public Ethics (CAPPE) at Charles Sturt University in Canberra, Australia. He holds a PhD in Philosophy from the University of Milan. His areas of interest include bioethics, moral philosophy, and moral psychology. He has published on a variety of topics in bioethics and applied ethics, including abortion, euthanasia, ethics of the family, informed consent, and organ donation.

between the two? Some of the most recurrent and hotly debated controversies surrounding initiation and design of early-phase trials hinge on how this question is resolved.

For example, in disease domains as diverse as spinal cord injury (Wirth, Lebkowski and Lebacqz 2011), neurodegenerative disease (Holden 2009), rheumatology (Sugarman and Bingham 2008), and genetic disorders (Kimmelman 2007), many commentators have debated whether initial trials should enroll patients who might benefit from trial enrollment (because they have recent disease onset) or patients who are less likely to be harmed (because they have advanced disease). Debates over starting doses, pace of enrollment and dose escalation, choice of delivery method, and preclinical evidence have revolved around similar questions (Tibbitts et al. 2010; Dresser 2009; van der Worp et al. 2010).

In what follows, we build on the premise that the task of early-phase testing is to optimize various components of a potential therapy so that subsequent confirmatory trials have the maximal probability of informing later drug development and clinical care. We then explore three strategies that investigators might use to manage trial risks while optimizing a therapy, using cell therapy for Amyotrophic Lateral Sclerosis (ALS) as an example. (1) One approach would use an iterative “maximax” strategy, whereby investigators aim to maximize gains by designing aggressive trials as they search for optimal components of a therapy. (2) A second approach would use an iterative “maximin” strategy, where investigators would design studies to minimize harms from the worst possible outcomes by designing the least risky trials. (3) A third strategy would blend elements of the two approaches.

We argue that the iterative application of maximin strategies over successive cohorts and trials, which we call the “risk-escalation model,” establishes a moral principle that should guide decision-making in early-phase trials. We ground our claim in appeals to patient–subject welfare and the social goals of clinical research. We close by addressing implications, exceptions, and limitations for our analysis.

## THE SOCIAL MISSION OF EARLY-PHASE TRIALS

Successfully translating a therapy (by which we mean any intervention, including biologics, devices, vaccines, and/or procedures) entails that researchers learn how to intervene in a pathophysiological process. This requires two types of discovery: first, discovery of an agent that has activity in a disease process; and second, identification of specific conditions that effectuate the clinical utility of this agent. The latter almost always involves discovering an appropriate dose, timing or schedule for administration, diagnostic procedures for identifying patients who are candidates for the drug, and strategies for managing side effects.

Consider, for example, the discovery process for cell therapies in ALS. ALS is a neurodegenerative disorder associated with progressive and generalized loss of motor neuron function. It typically starts in one anatomical location, and radiates to others, causing progressive and fatal paralysis. Presently contemplated cell-therapy strategies aim at interrupting the propagation of degenerative processes by creating a neuroprotective “firewall” in the spinal cord.

The uncertainties associated with cell therapy are myriad, beginning with the composition and safety of cells, surgical methods for delivery, dose, whether immunosuppressive drugs should be co-administered, and whether anti-inflammatory drugs should be co-delivered. Although rodent in vivo ALS models are available, their value in predicting therapeutic activity is limited. Indeed, the only licensed drug for treating ALS did not show activity in mouse models (Scott et al. 2008; Boulis et al. 2012). The task of first-in-human and early-phase clinical trials of ALS therapies is therefore to determine the appropriate preparations, doses, surgical methods, immunosuppressive regimes, and co-interventions to carry forward into confirmatory trials.

More generally, this process of determining the necessary and sufficient components of an effective and approximately optimal intervention—what has elsewhere been described as an “intervention ensemble” (Kimmelman 2012)—is the central goal of early-phase testing. Completion of this discovery step is crucial for three reasons. First, if the elements of an effective intervention ensemble are not clarified before later-phase testing (i.e., Phase 3 trials) or regulatory approval, patient–subjects enrolled in late-phase trials or receiving clinical care can be overdosed, injured from delivery, or deprived of co-interventions that would maximize the clinical utility of the therapy.

Second, by discovering the necessary and sufficient components of an intervention ensemble, early-phase studies create the moral and epistemic conditions for randomizing patients to experimental therapy in controlled confirmatory trials (that is, they establish necessary conditions for clinical equipoise).

Third, optimization of ensemble parameters is necessary for designing confirmatory trials that are informative. Confirmatory trials are far more likely to produce “positive” findings if researchers know and apply the necessary and sufficient components of an active intervention ensemble. “Negative” confirmatory trials are far more likely to be interpretable and useful to the research community if scientists can exclude the possibility that confirmatory tests used a wrong dose, delivered agent to the wrong compartment, intervened at a wrong stage of disease, etc.

## ENSEMBLE SPACE

The process of optimizing a therapeutic ensemble can be thought of as exploring a multidimensional landscape (Piantadosi 2005), which we call “ensemble space.” Each dimension of this space corresponds to some aspect of the intervention or its administration, such as dose of drug, dose of co-interventions, timing of intervention with respect to disease progression, location of delivery, etc. Many of these dimensions scale with increased risk: higher drug and co-intervention doses or delivery closer to an anatomically sensitive region involves greater risk. Administration of a drug to patients who are medically stable is also generally riskier, since persons with medically stable disease bear greater opportunity costs by exposing themselves to untried drugs than patients with advanced and refractory disease.

A complete ensemble space for ALS cell therapy would thus include numerous dimensions: cell-therapy dose, level of a particular receptor or marker on cells, location along spine for

delivery, dose of immunosuppressive drugs, time in disease process to intervene, etc. The current uncertainties surrounding these dimensions lead to many controversies about how to design initial trials. For example, should initial trials be run in patients with advanced disease, or should they be run in patients who are newly diagnosed? The former are less likely to respond, because degeneration may have already occurred. But they also have less to lose from safety events, given the inexorable disease course. Should cells be delivered to the lumbar region of the spine, or into the cervical region? If successful, the latter would preserve diaphragm (essential for ALS survival) and arm function (the loss of which many patients identify as an important source of morbidity). However, untried surgical techniques in this region could cause fatal injury to nerves that control diaphragm function, whereas the worst-case scenario for a lumbar region delivery would be paralysis of the legs.

In figure 1, we illustrate two 3-dimensional versions of this ensemble space—one representing a landscape where dimensions can be explored independently (1a) and another representing a landscape where they cannot be explored independently (1b). The  $x$ -axis represents dose, the  $y$ -axis represents delivery technique (which we take to be an aggregate of variables, such as number of injections, anatomical location, etc.), and the landscape along the  $z$ -axis represents a signal that correlates with benefit/ risk ratio (for example, a surrogate for clinical response combined with hard safety endpoints) with each dose–technique ensemble. For ALS, such a “signal” might be represented by an alteration in the slope of decline on the Amyotrophic Lateral Sclerosis Functional Rating Scale (ALSFRS), which is often used as a surrogate for prolonged survival. Negative signals (correlating with unfavorable benefit/risk ratios) are represented on the landscape by the flat, dark gray regions; positive signals (correlating with favorable ratios) are represented by peaks of lighter gray; and neutral signals are white.

At the point where first-in-human trials are contemplated, researchers know very little about the contours of the landscape. However, there is one region that is well understood and recognized by the relevant expert community as representing a low-risk and subtherapeutic area of the landscape. This is depicted as the white region near the origin in our figures. We call this region “the base.” For ALS cell-based interventions, the base would correspond to a small number of low-dose injections in the lumbar region of the spine.

Although the rest of the landscape is unknown at the outset of testing, researchers do have a bank of preclinical evidence and perhaps experience with related interventions such that they can make educated estimates about some of the ensemble dimensions and their optimal values. By definition, however, these estimates are uncertain. In figure 1, we represent  $P$  as an intervention ensemble that, at the launch of clinical testing, is projected to produce the best signal. This is in contrast to the “true” best intervention ensemble, which is located at the peak. The question under consideration here—strategies for designing and unfolding early-phase trials—is therefore a question about the most morally appropriate way of searching through the ensemble space for this peak.

But before discussing particular optimization strategies, a few more things need to be said about the representation: First, although figure 1 depicts the peak as positioned at lower values on the dose dimension than the estimated ensemble at  $P$ , this is an arbitrary choice

made only for illustrative purposes. Researchers might also underestimate the dose needed to achieve therapeutic signal, which would place the peak in a riskier region of the landscape. Second, the peak in landscape figure 1b is steep and narrow. This reflects an ensemble that has a narrow therapeutic index (i.e., the margin between therapeutic and toxic exposure is small). Many intervention ensembles have this property, including cytotoxic cancer drugs (which are often active only at levels just under the limits of tolerability). Nevertheless, some intervention domains will have wide therapeutic indices, in which case their intervention ensemble landscapes resemble rolling hills or large plateaus (more similar to figure 1a). In our example of ALS cell-based interventions, we are assuming a narrow therapeutic index and significant risks throughout much of the space. Such landscapes place greater pressure on ethical decision-making.

## THE INNOVATIVE CARE MODEL

One strategy for discovering the regions of optimal benefit/risk signal would be to begin clinical testing by enrolling a cohort of patients at doses and using surgical techniques that are believed, on the basis of preclinical evidence, to be the most likely region to observe a positive signal ( $P$  in figure 1). This is what decision scientists call a maximax strategy—it aims to maximize benefit by testing an ensemble that is estimated to have the best possible outcome (i.e., maximize the maximum gain). In ALS, this would entail that initial trials intervene on the cervical region of the spine, using the exact cell line and dose that seemed most promising in animal models, in populations of high-functioning patients. Once outcomes are observed, the research team would recalculate their estimate for the peak and enroll a new cohort of patients at that position. We call this iterative strategy for finding optima the “innovative care model.” Figure 2 illustrates best- and worst-case scenarios for this strategy.

*In the best-case scenario* (figure 2a), investigators are correct in their projections and the patients in the initial cohort experience dramatic, positive responses. Efficacy trials can then immediately proceed using that ensemble. This is also beneficial for other patients, since they will be able to access a new intervention sooner due to the expediency of clinical translation. Further, a dramatic success against a morbid disease like ALS would attract other researchers to explore other cell-based strategies, thus promoting medical advance. Finally, this strategy might reduce resources needed to validate a therapy, since it would be easier to recruit patient-subjects to trials involving a prospect of therapeutic benefit.

However, the innovative care strategy has three significant drawbacks. First is the risk of an untoward medical event. When there are large uncertainties about the properties of an intervention, small errors in the intervention ensemble are to be expected. For ensemble spaces with narrow therapeutic indices, small errors can translate into greater risk of serious adverse events. *In a worst-case scenario*, a drug-related death occurs, and the study provides no read-out on therapeutic signal (figure 2b). Since very little useful information has been gained about the contours of the landscape—beyond, for example, unacceptable toxicity with that particular ensemble—the researchers remain highly uncertain about where in the ensemble space to sample next or even whether to sample a new region at all.

The second type of risk is harm to the research effort. Unexpected and major harms can trigger crises of confidence among stakeholders engaged in a research program, as has occurred following deaths in other research areas—e.g., gene transfer in 1999 (Watson 2009) or CD28 superagonists in 2006 (*Nature Biotechnology* 2006). Depending on how events unfold, such crises of confidence can revolve around particular drugs, drug classes, institutions, companies, regulators, or particular research teams. For instance, physicians might begin to have doubts about referring their patients to such trials or to an institution; trainees opt for other research areas; sponsors withdraw funding. Such harms to a research effort can impede development of a useful therapy.

Third, *even in a best-case scenario*, where the first test produces a significant response, the innovative care strategy returns little information about the contours of the ensemble landscape. As described above, innovative care successes and failures can be informative for particular points in the ensemble space, but because these cohorts or trials are not anchored to a known region in the landscape, they may not provide later investigators or physicians with sufficient information about how to modify the ensemble. For example, in addition to risk, many dimensions of operation also scale with cost. A crucial task in medicine is to determine the safest—and preferably least costly—way of achieving a given therapeutic outcome. As a consequence of not clarifying the boundaries surrounding therapeutic optima, an innovative care strategy can leave researchers and downstream caregivers uncertain about whether there are still safer ways of achieving therapeutic outcomes. We will return to this problem below.

## THE RISK-ESCALATION MODEL

A more cautious approach begins by enrolling a cohort just outside the base region (e.g., subtherapeutic doses, refractory patients, delivery to the least sensitive anatomical target). For ALS, this would mean starting with subtherapeutic doses at the lumbar region of the spine in populations with advanced disease. Once the landscape in this region is confirmed, researchers then update their beliefs about the surrounding landscape and enroll the next cohort of patients just outside the now-revised known region, in the direction of the preclinical estimate. The process is then iterated until the optimum is discovered. This “risk-escalation” strategy is based on using a maximin decision rule at each iteration. Maximin aims to maximize benefit of uncertain decisions in the event of the worst outcomes (i.e., maximize the minimum benefit) by minimizing the possibility of unintended harm. Figure 3 illustrates best- and worst-case scenarios for the risk-escalation strategy.

In the *best-case scenario* for risk escalation (figure 3a), the peak is discovered after a short series of cohorts and perhaps a few trials, wherein patient–subjects were exposed to relatively little risk through the entire optimization process. While this is obviously less efficient than the best case for the innovative care approach, it could still accrue some of the benefits from an expedient translation.

In the *worst-case scenario* for risk escalation (figure 3b), the peak of the benefit signal is far out into the high-risk end of the landscape. Since this strategy searches cautiously outward from the origin, it will still eventually reach the peak, but it may require exposing more

cohorts of patients to inactive intervention ensembles in the process. This means greater resources are expended in a translation effort, which can also have consequences for sustained collaboration (for example, biotechnology companies might view timelines too lengthy to warrant product development).

Nevertheless, a risk-escalation strategy is attractive for all the reasons that the innovative care strategy is not. First, *catastrophic losses are avoided* when maximin decision rules are used, thus avoiding serious harm and also avoiding debacles that threaten a withdrawal of confidence in a research program. Second, initiating trials in the base ensures that *the study outcome*, whether net positive or negative, *is at least minimally informative* about the contours of the landscape. Finally, we noted that a crucial task in clinical translation is to determine the safest way of achieving a therapeutic outcome. Beginning in a safe region of the landscape and escalating risk provide assurance that *the lower boundary of acceptable benefit/risk will always be defined*.

## INTERMEDIATE STRATEGIES

Maximax decision rules may, perhaps, be used in innovative clinical settings or realms like surgical innovation, where the boundary between research and care is porous. However, we suspect that pure maximax approaches are rare in translation of novel drugs or biologics. Instead, aggressive translation strategies are likely to blend elements of innovative care and risk escalation, wherein initial trials use maximax decision rules along some dimensions, but maximin decision rules along others. For ALS, this could mean starting with low-doses at the cervical region of the spine with populations of mixed high-functioning and advanced disease patients. Figure 4 illustrates the best- and worst-case scenarios for this blended strategy.

In the *best case* (figure 4a), the peak will be fortuitously located at the high-risk end of one dimension, but the low-risk end of the other. This allows the intermediate strategy to discover the peak in one (or very few trials) and is advantageous for all the same reasons as the innovative care model. However, even if the initial trial is not exactly near the peak, the intermediate strategy has an advantage over the innovative care model in that the risks of untoward events (and the accompanying crises) are not as great.

In the *worst-case scenario* (figure 4b), the peak is far out into the high-risk end of the landscape. Just as we saw with the risk-escalation strategy, it will take many cohorts and trials to discover it. This has all of the drawbacks of the risk-escalation strategy—i.e., increased time and cost of translation. It also has some of the drawbacks of the innovative care strategy, since each of these trials is riskier than it would have been with risk escalation.

Finally, in terms of knowledge gain, again the consequences are a mixture of risk-escalation and innovative care strategies: Lower boundaries of ensemble dimensions are likely to be clarified for the dimensions governed by maximin decision rules. Lower boundaries will not be established for dimensions explored using the maximax rule.

## EVALUATING THE OPTIONS

What kinds of principles and evidence should guide the selection of one of these strategies over the others? It is tempting to appeal to probabilities to resolve this question. That is, if we could say something meaningful about the probability of best- or worst-case scenarios, this might help decide whether a given rule is too conservative or not conservative enough. However, early-phase testing of novel interventions are by definition subject to ignorance (Djulgovic 2007): the very aim of this phase of clinical development is to transform a state of ignorance into knowledge of risk by mapping relevant sections of the ensemble space (Kimmelman 2012).

Nevertheless, we can say that best-case scenarios for innovative care strategies are probably rare: few interventions show dramatic, imatinib-like clinical utility in first-in-human studies; indeed, few interventions show clinical utility across whole translation efforts. Worst-case scenarios for innovative care approaches leading to deaths and crises of confidence are also rare, but Jesse Gelsinger-like debacles (where the death of a study volunteer led to a retrenchment of investment in gene transfer research) are impossible to rule out. Reasonable people are likely to disagree about the probabilities of each.

One principled way of resolving this uncertainty would appeal to patient autonomy by allowing cohorts of patients and clinical investigators to select or opt into particular exploration strategies. As patients with advanced and refractory illnesses are often willing to take large risks (Weinfurt 2007), and clinical investigators often begin trials convinced of the utility of study interventions, this would lead to many translation efforts employing innovative care and intermediate strategies. However, we offer two interrelated reasons why there should be a presumption in favor of the risk-escalation model.

The first derives from the fact that clinical research—as opposed to clinical care—is not merely a private transaction, and hence cannot be governed solely or even primarily by the principles of autonomous choice. The mission of clinical research is to supply health care and public health systems with the evidence needed to address priority health needs. Any testing strategy that threatens to compromise this social mission is in contradiction with the very purpose of research. One of us has argued elsewhere that early-phase studies serve this social mission in two key ways: first, by establishing all necessary and sufficient components for a therapeutically useful intervention ensemble; and second, by establishing lower boundaries along relevant dimensions for therapeutically active intervention ensembles (Kimmelman 2012). Controversy surrounding trials for oxygen saturation targets—where the medical community belatedly discovered that lower oxygen saturation targets were associated with significantly greater mortality in extremely premature infants—illustrate some of the losses that occur when boundaries that scale with risk are not defined before an intervention is clinically translated. They also testify to the social and ethical challenges of collecting this information systematically once a strategy has been translated (Carlo, Bell, and Walsh 2013). The risk-escalation strategy is the only one that provides assurance that all relevant lower boundaries will be defined before translation.<sup>1</sup>



Second, a risk-escalation strategy—though costly in the short term—better achieves the above social goals by configuring research communities in ways that are more likely to sustain drug development efforts. Clinical development is a prolonged process, and requires sustained collaboration of many different stakeholders, all of whom enter the collaboration pursuing different sets of ends (London, Kimmelman, and Emborg 2010). Investigators pursue studies primarily for professional rewards; companies finance studies to enable commercial opportunities; and patients often enter studies seeking care options or to discharge aspirations. Catastrophic events leading to crises of confidence threaten such sustained collaborations. For example, academic medical centers are less likely to invest in research programs that are viewed as presenting both legal as well as reputational liability.

Perhaps more importantly, choice of a strategy generates incentive structures that interact with the level of risk for subjects and the research enterprise. The present system of drug development disproportionately rewards drug companies, investigators, and other actors that are the first to translate a new strategy. Such an incentive structure will tend to penalize researchers who adopt slower exploration strategies—especially where gains are partly realized by collectivities like the research enterprise. Ceding choice of exploration strategies to private actors would likely amplify the risk of catastrophic events, since research teams are more likely to adopt innovative care exploration strategies in a bid to be the first to translate. Establishing risk escalation as a moral principle that is superimposed over patient–investigator transactions would help solve a collective action problem by diminishing the prospect of catastrophic events that undermine collaborative efforts.

## IMPLICATIONS, EXCEPTIONS, AND LIMITATIONS

We have argued that a risk-escalation strategy is the most appropriate way of managing risk and uncertainty while exploring intervention ensembles in early phases of drug development. While our account provides a novel theoretical argument for this position, we are hardly the first to advocate a risk-escalation strategy for early-phase trials. Similar approaches are generally used in cancer (where dimensions that are escalated are dose and scheduling). The first trial of cell therapy for ALS also was described as using a “risk escalation” design (Boulis et al. 2012).<sup>2</sup> Commentators designing trials of the highly immunogenic cancer strategies, like chimeric antigen receptor therapies, have recommended an approach that is similar to risk escalation (Junghans 2010). Indeed, what we have described as the “risk-escalation model” may represent a more formalized description of what one sociologist called the “indigenous morality” of contemporary medical research (Halpern 2006).

We can nevertheless envision several circumstances where intermediate or innovative care models would be preferable in early-phase drug trials. One condition would be where

<sup>1</sup>Of course, these can always be clarified in subsequent risk deescalating trials. However, deescalating studies are more difficult to conduct, because patients and physicians facing dire medical circumstances are more reluctant to give up a sure benefit for a shot at a possible benefit that is safer.

<sup>2</sup>This design consisted of six cohorts of escalating risk. In the first two, nonambulatory patients are given (a1) unilateral injections to the lumbar region; (a2) bilateral injections to the lumbar region. In the next two cohorts, ambulatory patients are given (b) unilateral lumbar injections; (c) bilateral lumbar injections. In the final two cohorts, ambulatory patients receive cervical injections, (d) unilaterally; (e) bilaterally.

dimensions are known to scale with risk, but where investigators are unable to get a “read” on the relationship between dimension values and therapeutic signal. Consider the use of immunosuppressive regimes in recent ALS cell-therapy trials. Rather than escalating immunosuppressive treatments, first-in-human trials started with a very aggressive immunosuppressive strategy. This was motivated by the lack of viable, in vivo markers of immune rejection.<sup>3</sup> Conditions like these would favor intermediate exploration strategies where all dimensions are explored using maximin decision rules save those dimensions that cannot support a read-out of potential clinical utility.

A second condition where risk escalation may be inappropriate is in the context of public health crises. The prospect of immediate, catastrophic losses that are expected in the absence of treatment strategy both diminishes the importance of defining lower bounds of treatment ensemble dimensions and increases the social costs associated with slow exploratory strategies. When these unusual conditions hold, innovative care strategies are more appropriate.

A third condition where risk escalation may be inappropriate is in drug development for ultra-rare disorders. Here, difficulties populating trials and attracting resources for them might necessitate intermediate strategies that promote discovery of optima using fewer research subjects and trials. In such circumstances, there may be fewer social and patient costs associated with not defining lower boundaries of all ensemble dimensions.

Fourth, there are circumstances where all three strategies merge. Some medical domains involve interventions that are well understood and safe. This might occur for a me-too drug with a wide therapeutic index, or in pediatric testing of drugs that have been licensed for adult disorders. In such circumstances, well-understood, safe regions of landscapes may extend into regions that are very close to projected peaks. When this occurs, innovative care and risk-escalation models would lead to identical trial design.

We close by noting an important question that our analysis leaves unresolved: how base regions of the ensemble landscape are demarcated from unknown regions such that investigators can establish starting conditions for clinical testing. Above, we suggested that this demarcation should be established by expert communities, rather than beliefs of individual researchers—much like clinical equipoise. Yet the epistemic norms governing this demarcation are going to be different in areas of high uncertainty than for areas where evidence is relatively mature. How base regions are demarcated and expanded as a trial progresses will dictate the rules governing risk escalation within protocols as well as across them.

## CONCLUSION

Early-phase trials are conducted at the point of greatest uncertainty about the effects and appropriate application of new drugs. The charge of such trials is to identify optimal methods of applying a new drug, and to define the safest possible way of achieving

---

<sup>3</sup>Nicholas M. Boulis, personal communications with the authors, June 17, 2013.

therapeutic effects. However, the costs and burdens associated with doing so can be considerable for human subjects and the research enterprise.

In most circumstances involving novel interventions, the risk-escalation strategy achieves these objectives while minimizing risk to volunteers and to the research enterprise. As a moral rule, it also institutionalizes incentive structures that control risk and foster the kinds of sustained collaborations that are necessary for clinical translation. The strategies proposed above should inform the planning, funding, design, ethical review, and reporting of early-phase research programs.


## Acknowledgments

This work was funded by a Canadian Institutes of Health Research operating grant (EOG102823). Jonathan Kimmelman wishes to dedicate this manuscript to the memory of Harold Kimmelman. At home, he lived by maximin; in his art, by maximax. And in both, he found the optimum.

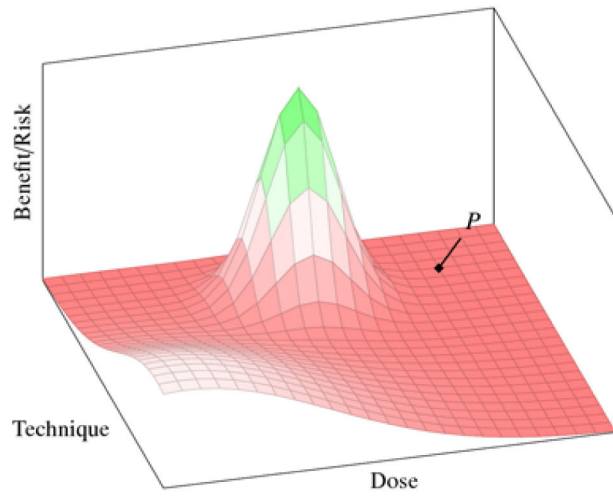
## References

- Boulis, Nicholas M., Federici, Thais, Glass, Jonathan D., Simon Lunn, J., Sakowski, Stacey A., Feldman, Eva L. Translational Stem Cell Therapy for Amyotrophic Lateral Sclerosis. *Nature Reviews Neurology*. 2012; 8:172–76.
- Carlo, Waldemar A., Bell, Edward F., Walsh, Michele C. Oxygen-Saturation Targets in Extremely Preterm Infants. *New England Journal of Medicine*. 2013; 368(20):1949–50. [PubMed: 23593943]
- Djulgovic, Benjamin. Articulating and Responding to Uncertainties in Clinical Research. *Journal of Medicine and Philosophy*. 2007; 32(2):79–98. [PubMed: 17454416]
- Dresser, Rebecca. First-in-Human Trial Participants: Not a Vulnerable Population, but Vulnerable Nonetheless. *Journal of Medical Ethics*. 2009; 37(1):38–50.
- Nature Biotechnology*. 2006; 24:368. Consequences Editorial.
- Halpern, Sydney A. *Lesser Harms: The Morality of Risk in Medical Research*. Chicago: University of Chicago Press; 2006.
- Holden, Constance. Fetal Cells Again? *Science*. 2009; 326(5951):358–59. [PubMed: 19833936]
- Junghans, Richard P. Strategy Escalation: An Emerging Paradigm for Safe Clinical Development of T Cell Gene Therapies. *Journal of Translational Medicine*. 2010; 8:55. [PubMed: 20537174]
- Kimmelman, Jonathan. Stable Ethics: Enrolling Non-Treatment Refractory Volunteers in Novel Gene Transfer Trials. *Molecular Therapy*. 2007; 15(11):1904–06. [PubMed: 17948045]
- Kimmelman, Jonathan. A Theoretical Framework for Early Human Studies: Uncertainty, Intervention Ensembles, and Boundaries. *Trials*. 2012; 13:173. [PubMed: 22999017]
- London, Alex John, Kimmelman, Jonathan, Emborg, Marina Elena. Beyond Access vs. Protection in Trials of Innovative Therapies. *Science*. 2010; 328(5980):829–30. [PubMed: 20466907]
- Piantadosi, Steven. *Clinical Trials: A Methodologic Perspective*. 2. Hoboken, NJ: Wiley Interscience; 2005.
- Scott, Sean, Kranz, Janice E., Cole, Jeff, et al. Design, Power, and Interpretation of Studies in the Standard Murine Model of ALS. *Amyotrophic Lateral Sclerosis*. 2008; 9:4–15.
- Sugarman, Jeremy, Bingham, Clifton O. Ethical Issues in Rheumatology Clinical Trials. *Nature Reviews Rheumatology*. 2008; 4:356–63. [PubMed: 18542108]
- Tibbitts, Jay, Cavagnaro, Joy A., Haller, Christine A., Marafino, Ben, Andrews, Paul A., Sullivan, John T. Practical Approaches to Dose Selection for First-in-Human Clinical Trials with Novel Biopharmaceuticals. *Regulatory Toxicology and Pharmacology*. 2010; 58(2):243–51. [PubMed: 20558225]
- van der Worp, Bart H., Howells, David W., Sena, Emily S., et al. Can Animal Models of Disease Reliably Inform Human Studies? *PLoS Medicine*. 2010; 7(3):e1000245. [PubMed: 20361020]

- Watson, James M. A History Lesson for Stem Cells. *Science*. 2009; 324(5928):727–28. [PubMed: 19423804]
- Weinfurt, Kevin P. Value of High-Cost Cancer Care: A Behavioral Science Perspective. *Journal of Clinical Oncology*. 2007; 25(2):223–27. [PubMed: 17210944]
- Wirth, Edward, Lebkowski, Jane S., Lebacqz, Karen. Response to Frederic Bretzner et al. ‘Target Populations for First-in-Human Embryonic Stem Cell Research in Spinal Cord Injury. *Cell Stem Cell*. 2011; 8(5):476–78. [PubMed: 21549322]




(a) Independent dimensions

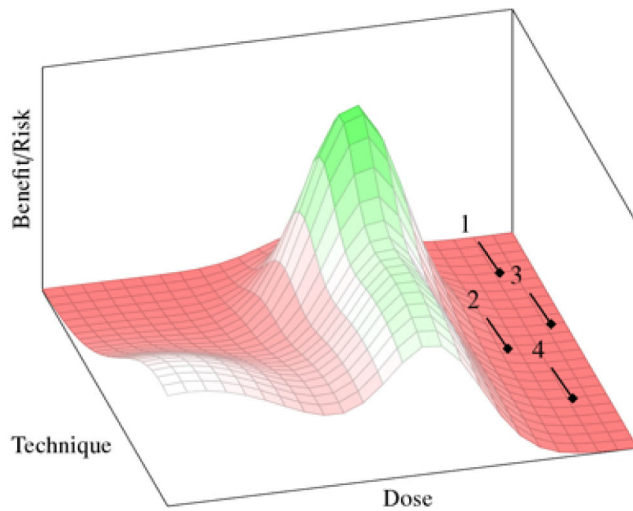


(b) Dependent dimensions

**Figure 1.** Benefit/risk landscape of ensemble space. The x- and y-axes (dose and technique, respectively) scale with increased risk. The z-axis represents benefit signal: light gray peaks are positive signal; flat, dark gray regions are negative signal; white is a neutral outcome (little harm, little benefit signal). The point *P* represents the estimated optimal therapeutic ensemble based on preclinical evidence. In landscape (a), there is a ridge of benefit signal for a specific dose value that extends across much of the technique dimension. This would permit investigators to explore the dimensions independently—first identifying optimal dose and then exploring for optimal technique. In landscape (b), since the narrow peak of benefit signal is surrounded by negative signal, the dimensions cannot be explored independently.




(a) Best-case maximax




(b) Worst-case maximax

**Figure 2.**

Best- and worst-case outcomes for the innovative care strategy. (a) The peak of benefit signal corresponds almost exactly to the estimates from preclinical data. Only one trial is needed to identify a promising ensemble, which can then be advanced into later phase trials. (b) The peak does not correspond to the preclinical estimates, and the deaths or toxicities in each negative trial do not provide informative evidence about how to adjust ensemble parameters. Each subsequent trial is then an arbitrary and risky guess in the neighborhood of the original preclinical estimate.




(a) Best-case maximin




(b) Worst-case maximin

**Figure 3.** Best- and worst-case outcomes for the risk escalation strategy. (a) The peak of benefit signal is discovered after few trials, exposing patients to minimal risks. (b) The peak is deeper into the high risk end of the landscape and many trials, of increasing risk, are needed to discover it.



(a) Best-case intermediate



(b) Worst-case intermediate

**Figure 4.** Best- and worst-case outcomes for the intermediate strategy. (a) The true peak of benefit signal lies at the high-risk end of one axis and the low-risk end of another. It is discovered after only single trial. (b) The peak is deeper into the high-risk end of the landscape and many high-risk trials are needed to discover it.