Appearing in *Risk Analysis,* 32(8): 1326–1332.

Doing Our Best: Optimization and the Management of Risk

Yakov Ben-Haim[1]

Abstract Tools and concepts of optimization are widespread in decision-making, design and planning. There is a moral imperative to 'do our best'. Optimization underlies theories in physics and biology, and economic theories often presume that economic agents are optimizers. We argue that, in decisions under uncertainty, what should be optimized is robustness rather than performance. We discuss the equity premium puzzle from financial economics, and explain that the puzzle can be resolved by using the strategy of satisficing rather than optimizing. We discuss design of critical technological infrastructure, showing that satisficing of performance requirements—rather than optimizing them—is a preferable design concept. We explore the need for disaster recovery capability and its methodological dilemma. The disparate domains—economics and engineering—illuminate different aspects of the challenge of uncertainty and of the significance of robust-satisficing.

Keywords: optimizing, satisficing, robustness, uncertainty, info-gap, financial investment, equity premium puzzle, technological infrastructure, moral hazard, disaster recovery

1 Introduction

We admire excellence in all areas of endeavor: art, sport, science, business, and risk management. The fastest runner wins the race and our admiration. The lowest-risk design—all else being equal—is preferred. 'Better', by definition, is 'more desirable' and—by the logic of preference—the best is most preferred. The logic of preference is so compelling that there is a moral imperative to do our best. Optimization also has deep roots in the physical and natural realms. The laws of physics can be derived from optimization principles.(1) Biological evolution is a process of selection of the better over the less good leading—all else being stable—to optimal morphologies. Mathematical economics was quick to adopt the imperative of optimization, which underlies modern theories of economic dynamics.

Decision makers often face severe uncertainties. This has profound implications for any attempt to optimize the outcome of their decisions. In this essay we first discuss a paradox from financial economics that belies the cardinality of performance-optimization by economic agents. We contrast performance-optimization with a strategy of robustly achieving critical goals. We then apply this concept to technological risk analysis, and consider the schematic design of a critical but risky infrastructure. Finally, we discuss the importance of disaster recovery as an integral part of risk management. Our goal is to understand the utility as well as the limitation of performance-optimization. The consideration of diverse domains—economics and engineering—facilitates the generality of that understanding.

2 Financial Economics and Optimization

Economic agents are optimizers according to standard economic theories. Rational behavior is maximization in the pursuit of self interest. Firms try to maximize profits, households try to maximize utility, governments try to maximize social welfare.(2,3) Those who fall short of

[0]\papers\essays\Do-Best-Econ-Opt\risk-opt04.tex 12.3.2012

[1]Yitzhak Moda'i Chair in Technology and Economics, Faculty of Mechanical Engineering, Technion—Israel Institute of Technology, Haifa 32000 Israel, yakov@technion.ac.il.

the optimum are weeded out by competition. Optimization has its own moral imperative: we should all do our best, and if we don't, then it's the other guy's turn at bat.

Is this true? Should we rely on economic models that adopt maximization as an axiom?

Biological evolution is a powerful metaphor for economics. Consider a squirrel nibbling acorns, and noticing a stand of fine oaks in the distance. There are probably better acorns there, but also other squirrels and predators. How long should the squirrel forage here before moving there? What strategy should guide the decision? The squirrel needs a critical amount of energy to survive the night. Maximizing caloric intake is *not* necessary. Maximizing the reliability of achieving the critical intake *is* necessary. What is maximized is not the substantive "good" (calories), but confidence in satisfying a critical requirement.(4)

Fifty years ago Herbert Simon (Nobel Prize in economics, 1978) advanced the idea that economic agents lack information, understanding, and the ability to process data. These deficiencies, which he called their "bounded rationality", force agents to look for solutions that are good enough, though not necessarily optimal. The optimum may exist but it cannot be known by the resource- and information-limited agent. "Satisficing" is what Simon called this strategy of settling for a solution that is good enough, as opposed to optimizing.(5)

But academic economists seem to take scarce notice of Simon's work.(2,3) Like Twain's quip about the weather, they all talk about it (either weather or satisficing) but they don't do a damn thing. Rationality, we learn, is optimization of profit or utility.

This very conventional attitude to rationality may be related to the long list of unresolved economic paradoxes. One example began with an article by Mehra and Prescott (the latter won the 2004 Nobel Prize in economics) entitled "The equity premium: A puzzle".(6) A decade later Kocherlakota published "The equity premium: It's still a puzzle".(7) There are many theoretical explanations of the equity premium puzzle (EPP),(8,9) but no consensus. In fact, not all economists agree that a long-term equity premium even exists.

What is the EPP, and what can we learn about optimization in economics, and beyond?

Stocks are riskier than US government bonds, so the average return to stocks should be higher. Otherwise who would look at stocks? This is sound common sense, and many economists have shown that the annual return to stocks is higher than to bonds, typically by 7%, sometimes by as much as 20% or as little as 0.3%. Does this "risk premium" for stocks make sense in terms of rational (read: maximizing) behavior? The puzzle (assuming the premium is real) is that standard asset pricing models can explain the size of the risk premium only by assuming that investors are much too averse to risk. The observed behavior of investors in other risky settings would suggest that they would be willing to accept a much lower equity premium for stocks.

There are, as noted, many insightful attempts to resolve the EPP, including that it is a statistical chimera. But what these explanations have usually not challenged is the assumption of optimization. Explanations of the EPP usually assume that investors try to maximize their returns rather than trying to achieve adequately large returns (e.g., larger than the competition).

Robustness to uncertainty is the key to understanding the role of satisficing in explaining the EPP as shown in detail elsewhere.(10, sec.11.5) The same is true of another financial conundrum, the home-bias paradox.(11)

Suppose we hear that a new start-up offers higher returns than anyone else, though many risks are involved. While such large returns would be great, the investor would be satisfied with lower returns. The investor has a fixed budget with which to construct an investment portfolio by choosing from among a large number of assets. Usually only one portfolio has maximum expected return. However, for any specified increment below this maximum, many different portfolios will have the same expected return at this specified value. This means that by relinquishing the goal of maximizing (even on average), the investor is able to choose among more alternative portfolios, all of which are adequate in terms of expected average return. This is an algebraic result with a simple geographical analogy. A mountain usually has only one summit, but it has an infinity of points at any fixed distance below the summit.

Accepting a sub-optimal but adequate investment is an example of what Simon called satisficing, and it also motivates the idea of robustness. The investor who satisfices (rather than maximizes) can choose the alternative that would yield the required return over the greatest range of uncertain future scenarios. That is, the investor foregoes some aspiration for profit in exchange for some robustness against unacceptably low returns. In other words, satisficing is more robust to uncertainty than optimizing. Hence this strategy is called robust-satisficing. If satisficing—rather than maximizing—is in some sense a better bet, then it will tend to persist under uncertain competition.

Now we can understand that equity premia should be lower than predicted by theories which assume that investors try to maximize their returns. Satisficers have to satisfy their requirements (or those of their clients). Satisficers tend to survive because they are more likely to meet critical requirements. And since satisficing entails a preference for less-than-maximal options, we should expect the market to lower the equity premium for risk.

There is a broader policy implication of the logic of satisficing. Going after critical requirements is usually a better bet for "survival", than going after what seems optimal. This is true of many decisions under uncertainty in forecasting,[12] engineering,[13] economics, public policy, homeland security and elsewhere.[10, sec.11.4] We must ask what outcomes are critical, not what are the best outcomes predicted by our uncertain models, even if those models are probabilistic. Critical requirements are usually more modest than the best anticipated outcome, so there will usually be many more ways to achieve them. We should choose the option that will lead to the required outcome most robustly. We now apply this strategy to the design of a critical and dangerous infrastructure technology.

3 Sea Walls and Tsunamis

On 11 March 2011 a magnitude 9.0 earthquake struck northeastern Japan, followed by a massive tsunami. Walls of water flooded the coastal region causing vast damage and tens of thousands of deaths. The nuclear reactor complex at Fukushima Daiichi was flooded and extensive damage was done to 3 of the 6 reactors. We will apply the reasoning of robust-satisficing to a schematic design analysis of the flood protection problem, motivated by the Fukushima incident. Our goal is to illustrate the motivation for robust satisficing—as distinct from risk-aware performance-optimizing—in light of the uncertainties inherent in the design.

Let's start with the **uncertainties**. Tsunami heights can be enormous. A 1792 eruption of mount Unzen in Japan is reported to have produced a 100m tsunami. A 1958 landslide in Alaska's Lituya Bay caused an earthquake leading to an initial wave 524m high. A 1963 landside above the Vajont Dam in Italy resulted in a 250m surge.[14] These heights are rare and extreme, but tsunami heights of several tens of meters are not uncommon. The Fukushima tsunami was 15m high, the tsunami following the 1993 Okushiri earthquake in Hokkaido, Japan, was 30m. While there clearly is an upper bound to physically possible tsunamis, the upper tail of the probability distribution is long and fat and highly uncertain.

Let's now consider risk-informed **performance requirements**. Performance requirements may originate in various ways: by legislation, by administrative fiat, by public debate and collective decision making, and so on. Furthermore, multiple constraints of various sorts may be imposed, such as cost and safety constraints, engendering trade offs. The process by which performance requirements are established and balanced is beyond the scope of this paper. Rather, we study the implications of two broad classes of requirements: optimizing and satisficing requirements, which we illustrate through example.

Roughly speaking, we want the flood defense system to assure that large damage to the nuclear reactor is very rare. The designer is well aware of risks and wants to control them. More precisely, in our example the designer specifies the largest acceptable probability of unacceptably large damage. For instance, the designer (or a public commission or the legislature) may require that the probability of a loss of coolant accident (LOCA) be no greater than one in a million per year.[2]

We must note two attributes of a design requirement such as this.

First, while it is perhaps extremely demanding, is not an optimization requirement, but rather a satisficing requirement. An optimizer would require *minimal* probability of LOCA. The satisficer seeks *acceptably small* probability. Note however that both optimizer and satisficer attempt to model and manage risk, though in different ways.

Second, it is necessary to know, or reliably estimate, the upper tail of the probability distribution of tsunami damage in order to operationalize this performance requirement. That probability distribution is highly uncertain, as we just explained.

These two attributes are linked: the uncertainty of the probability distribution motivates the satisficing strategy. Reliable minimization of the probability of a LOCA, while desirable, is not feasible since the probability distribution is imperfectly known. It is therefore a better strategy to seek an acceptable (and perhaps very demanding) outcome for as wide a range of probability distributions as possible.

Engineers sometimes use the language of performance optimization. Nonetheless, satisficing requirements such as the one we have discussed are quite common in engineering. US Dept. of Defense design specfications ("mil-specs") and many other design codes routinely specify required performance thresholds, rather than specifying that the performance must be as good as possible. We cannot avoid noting the contrast between this tradition in engineering design, and the axiom of performance optimization that underlies much economic theory.

There are two broad categories of **design options** in defending against tsunamis. *Sea walls* are intended to keep the water—or a significant fraction of it—from reaching vulnerable areas. The sea wall design is specified by the location, height and durability of the wall. *Flood gates and channels* are intended to re-direct the water away from vulnerable areas. The design is specified by the location of the gates and channels and by their water capacity (volume per time).

Before discussing the robust-satisficing design of a flood defense system we briefly mention some of the many concepts of **robustness.**

Wald[(16)] studied the problem of statistical hypothesis testing based on a random sample whose probability distribution is not known, but whose distribution is known to belong to a given class of distribution functions. Wald states that "in most of the applications not even the existence of ... an a priori probability distribution [on the class of distribution functions] ... can be postulated, and in those few cases where the existence of an a priori probability distribution ... may be assumed this distribution is usually unknown." (p.267). Wald developed a decision procedure that "minimizes the maximum ... of the risk function." (p.267).

Many engineering researchers, beginning in the 1960s, developed estimation and control algorithms for linear dynamic systems based on sets of inputs. Schweppe[(17)] for instance develops inference and decision rules based on assuming that the uncertain phenomenon can be quantified in such a way as to be bounded by an ellipsoid, with no probability function involved. These robust estimation and control methodologies have flourished in engineering and beyond.

Hansen and Sargent have pioneered the introduction of robustness tools in economics. In their recent book[(18)] they quantify model misspecification by taking "a given approximating

[2]This numerical value is just for illustration. The probabilistic risk analysis of nuclear reactors is complex and multi-variate. For instance, a report prepared for the US Nuclear Regulatory Commission in 2005[(15)] cites mean expert estimates of the frequency of LOCAs in boiling water reactors ranging from 6.4×10^{-4} per year for very small pipe breaks to 6.4×10^{-9} per year for very large breaks.

model and surrounding it with a set of unknown possible data generating processes, one unknown element of which is the true process .... Our decision maker confronts model misspecification by seeking a decision rule that will work well across a set of models for which" the relative entropy is bounded. "The decision maker wants a single decision rule that is reliable for *all* [emphasis in original] models ... in the set" (p.11). They explain that "'Reliable' means good enough, but not necessary optimal, for each member of a set of models." (footnote 21, p.11). They then "maximize [an] intertemporal objective over decision rules when a hypothetical malevolent nature minimizes that same objective .... That is, we use a max-min decision rule." (p.12).

Many other concepts of robustness are to be found, including P-boxes for managing uncertainty in probability distributions in many applications,[19−22] imprecise probabilities,[23,24] info-gap theory[10] and others.

This essay does not attempt to compare or contrast these methodologies (some comparison is found elsewhere[25]), all of which have been usefully applied. The info-gap concept of robustness will be implicit in our discussion,[10] though similar conclusions would be reached by posing the analysis within other robustness frameworks.

We now sketch out the **robust design** of the flood defense system. In our hypothetical and schematic design problem we imagine that the designer identifies a suite of design alternatives. Borrowing the language of portfolio theory, we say that a fixed budget must be allocated among design options so that the total effect satisfies the design requirement. For instance, heightening the seawall reduces the available budget for floodgates, but together they must achieve the required probability of failure.

The designer has access to models of earthquakes, tsunamis, coastal hydrodynamics, nuclear reactor cooling, and so on. The models include probability distributions as well as deterministic models. These are the best available models, but they still contain imperfections, errors, and uncertainties due to the vast complexity of the processes. Next year's models will be better, but design decisions must be made today.

The strategy of performance optimization would proceed as follows. Use the best available models to predict the outcomes of each budgetarily feasible design. The models may be probabilistic, and predicted outcomes may be probabilities of LOCA or other events, or other risk-weighted performance functions with means, variances or quantile terms. The best-model performance-optimizer adopts the design whose predicted risk-informed outcome is best.

A performance-optimizer would advocate using the best risk-informed model to identify the design with the best predicted performance, even though the models are uncertain. Our analysis of the equity premium puzzle suggests that this is not a reliable strategy, as we now explain.

Suppose the designer has identified two budget-feasible designs, A and B, both of which satisfy the performance requirement (e.g., probability of LOCA less than one in a million). Moreover, the best available models predict that the probability of a LOCA for design A is less than for any other budget-feasible design.

The performance-optimizer would choose design A, and this might indeed be the design of choice. However, we must first ask the robustness question: how wrong can the models be (including the probabilistic models whose tails may be poorly known), and the performance of the design is still acceptable? We would like to know the true probability of LOCA, and the best-model estimate is only an approximation to the truth.

The central point is that a design has two attributes: its best-model predicted behavior, and the sensitivity of its predicted behavior against error in the best models. Its predicted behavior is evaluated from one type of calculation (applying the models to the design in question), and its sensitivity to error is evaluated by a different calculation (a robustness analysis of one sort or another).

Predicted behavior, and robustness-to-model-error of predicted behavior, are distinct and independent attributes of a design. It is not true that good predicted outcome implies large robustness to model error, even when the model is risk-based. Model predictions of the best risk-informed outcome may be very sensitive to uncertainty in the sense that small changes in estimates or assumptions may cause much worse outcomes. The idea of robustness quantifies this intuition. The robustness is a "rate of change" attribute of the design: how "fast" does the performance change as the models change? The predicted outcome, in contrast, is evaluated from a specific realization of the models (which may be probabilistic and risk-informed).

A graphical analogy will illustrate the difference between performance optimization and robustness optimization. A straight line has two independent attributes: a slope (rate of change) and an intercept (specific value). An infinity of lines have the same intercept and different slopes, and a different infinity of lines have the same slope and different intercepts. In the same way, the predicted behavior and the robustness-to-model-error of predicted behavior, are distinct and independent attributes of a design. The predicted best design is not necessarily the most robust to uncertainty, even when the model itself involves risk-based terms.

Let's return to the decision between designs A and B, where A is the best-models' predicted optimal design. If both designs are predicted to satisfy the risk-based performance requirement, and if A's predicted behavior is also more robust to model-error than B, then both the optimizer and the satisficer will choose A over B. Optimizer and satisficer agree in this case, but for different reasons. The optimizer chooses A over B because it is the putative optimum. The satisficer chooses A over B because A is more robust than B for satisfying the design requirement.

It is clear that the optimizer and the satisficer will not always agree. Even though A is predicted (based on the best models) to out-perform B, it can and does happen that B satisfies the performance requirement over a larger range of models (surrounding the best available models) than A. Both A and B are predicted to be satisfactory, and B is predicted to be satisfactory over a larger range of model-deviations from the best available models. The satisficer argues that—since the best models are wrong in unknown ways—it is a better bet to choose the more robust design.

This conclusion—that robust-satisficing may or may not agree with best-model performance-optimization—is reminiscent of the resolution of the equity premium puzzle that we discussed earlier. Practical decision makers—squirrels, financial investors, or engineers—face the consequences of their decisions. When those consequences are fatal in some sense, less successful decision makers (or their decision strategies) tend to be removed from the game. Survival depends on robustness to ignorance in achieving critical goals.

The decision-maker who maximizes robustness-to-uncertainty in seeking critical outcomes is indeed optimizing something: robustness. Both the performance-optimizer and the robust-satisficer employ the concepts and tools of optimization in attempting to manage risk. But there is a conceptual difference between their strategies, and their decisions may disagree. Both strategies employ the best available models. The performance-optimizer chooses the design whose best-model risk-informed predicted behavior is best. The robust-satisficer picks the design whose behavior is satisfactory for the widest range of deviation from the best models.

4 Disaster Recovery

The distinction between robust-satisficing and performance-optimizing has implications for our attitudes towards disaster recovery.

The failure of critical technologies can be disastrous. The crash of a civilian airliner can cause hundreds of deaths. The meltdown of a nuclear reactor can release highly toxic isotopes. Failure of flood protection systems can result in vast death and damage. Society therefore insists that critical technologies be designed, operated and maintained to extremely high levels of reliability. We benefit from technology, but we also insist that the designers and operators "do their best" to protect us from their dangers.

Industries and government agencies who provide critical technologies almost invariably act in good faith for a range of reasons. Morality dictates responsible behavior, liability legislation establishes sanctions for irresponsible behavior, and economic or political self-interest makes continuous safe operation desirable.

The language of performance-optimization—not only doing our best, but also achieving the best—may tend to undermine the successful management of technological danger. A probability of severe failure of one in a million per event is exceedingly—and very reassuringly—small. When we honestly believe that we have designed and implemented a technology to have vanishingly small probability of catastrophe, we can honestly ignore the need for disaster recovery.

Let's contrast this with an ethos that is consistent with robust-satisficing. We now acknowledge that our predictions are uncertain, perhaps highly uncertain on some specific points. We attempt to achieve very demanding outcomes—for instance vanishingly small probabilities of catastrophe—but we recognize that our ability to reliably calculate such small probabilities is compromised by the deficiency of our knowledge and understanding. We robustify ourselves against those deficiencies by choosing a design that would be acceptable over a wide range of deviations from our current best understanding. Not only does "vanishingly small probability of failure" still entail the possibility of failure, but our predictions of that probability may err.

Acknowledging the need for disaster recovery capability (DRC) is awkward and uncomfortable for designers and advocates of a technology. We would much rather believe that DRC is not needed, that we have in fact made catastrophe negligible. But let's not conflate good-faith attempts to deal with complex uncertainties, with guaranteed outcomes based on full knowledge. Our best models are in part wrong, so we robustify against the designer's bounded rationality. But robustness cannot guarantee success. The design and implementation of DRC is a necessary part of the design of any critical technology, and is consistent with the concept of robust-satisficing.

One final point: moral hazard and its dilemma. The design of any critical technology entails two distinct and essential elements: failure prevention and disaster recovery. What economists call a 'moral hazard' exists since the failure prevention team might rely on the disaster-recovery team, and vice versa. Each team might, at least implicitly, depend on the capabilities of the other team, and thereby relinquish some of its own responsibility. Institutional provisions are needed to manage this conflict.

The alleviation of this moral hazard entails a dilemma. Considerations of failure prevention and disaster recovery must be combined in the design process. The design teams must be aware of each other and even collaborate in the design process because a single coherent system must emerge. But we don't want either team to relinquish any responsibility. On the one hand we want the failure prevention team to work as though there is no disaster recovery, and the disaster recovery team should presume that failures will occur. On the other hand, we want these teams to collaborate on the design.

The presence of this moral hazard and its dilemma does not obviate the need for both elements of the design. It highlights the special challenge of high-risk critical technologies: design so failure cannot occur, and prepare to respond to the unanticipated.

Acknowledgement

The author is indebted to Mark Burgman for valuable comments.

5 References

1. Lanczos C. The Variational Principles of Mechanics. 2nd edition, Dover Publications, New York, 1970.
2. Mas-Colell A, Whinston MD, Green JR. Microeconomic Theory. Oxford U. Press, 1995, p.6.
3. Sen A. Rational Behaviour. in Eatwell, John and Murray Milgate and Peter Newman, eds., New Palgrave: A Dictionary of Economics. 1998.
4. Carmel Y, Ben-Haim Y. Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice? American Naturalist. 2005; 166:633–641.
5. Simon HA. Models of Bounded Rationality. Vol. 3, Empirically Grounded Economic Reason. MIT Press, 1997.
6. Mehra R, Prescott EC. The equity premium: A puzzle. Journal of Monetary Economics. 1985; 15:145–161.
7. Kocherlakota NR. The equity premium: It's still a puzzle. Journal of Economic Literature. 1996; 34:42–71.
8. Cochrane JH. Asset Pricing. Princeton University Press, Princeton and Oxford, 2001.
9. Fama EF, French KR. The equity premium. The Journal of Finance. 2002; LVII(2):637–659.
10. Ben-Haim Y. Info-Gap Decision Theory: Decisions Under Severe Uncertainty. 2nd edition, Academic Press, London, 2006.
11. Ben-Haim Y, Jeske K. Home-bias in financial markets: Robust satisficing with info-gaps. Federal Reserve Bank of Atlanta, Working Paper Series, 2003; 2003-35.
12. Ben-Haim Y. Info-gap forecasting and the advantage of sub-optimal models. European Journal of Operational Research. 2009; 197:203–213.
13. Ben-Haim Y. Robust satisficing and the probability of survival. ICVRAM 2011: 1st International Conference on Vulnerability and Risk Assessment and Management, April 11–13, 2011, University of Maryland, College Park, 2011; pp.963–970.
14. Megatsunami. http://en.wikipedia.org/wiki/Megatsunami
15. Tregoning R, Abramson L, Scott P. Estimating Loss-of-Coolant Accident (LOCA) Frequencies Through the Elicitation Process. US Nuclear Regulatory Commission, NUREG-1829, p.xix, 2005.
16. Wald A. Statistical decision functions which minimize the maximum risk. Annals of Mathematics. 1945; 46(2):265–280.
17. Schweppe FC. Uncertain Dynamic Systems. Prentice-Hall, Englewood Cliffs. 1973.
18. Hansen LP, Sargent TJ. Robustness. Princeton University Press, Princeton and Oxford, 2008.
19. Ferson S. RAMAS Risk Calc 4.0 Software: Risk Assessment with Uncertain Numbers. Boca Raton: Lewis Publishers, 2002.
20. Karanki DR, Kushwaha HS, Verma AK, Ajit S. Uncertainty analysis based on probability bounds (p-box) approach in probabilistic safety assessment. Risk Analysis. 2009; 29:662–75.
21. Sander P, Bergbäck B, Öberg T. Uncertain numbers and uncertainty in the selection of input distributions—Consequences for a probabilistic risk assessment of contaminated land. Risk Analysis. 2006; 26:1363–1375.
22. Ferson S, Burgman M. Correlations, dependency bounds and extinction risks. Biological Conservation. 1995; 73:101–105.

23. Miranda E. A survey of the theory of coherent lower previsions. International Journal of Approximate Reasoning. 2008; 48:628–658.
24. Walley P. Measures of uncertainty in expert systems. Artificial Intelligence. 1996; 83:1–58.
25. Ben-Haim Y, Dacso CC, Carrasco J, Rajan N. Heterogeneous Uncertainties in Cholesterol Management. International Journal of Approximate Reasoning. 2009; 50:1046–1065.