A Decision Analysis Framework for the U.S. Nuclear Fuel Cycle

by

Lara Marie Pierpoint

S.M. Technology and Policy and Nuclear Science and Engineering
Massachusetts Institute of Technology, 2008

B.S. Physics
University of California, Los Angeles, 2004

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy
at the
Massachusetts Institute of Technology
September 2011

Signature of Author……………………………………………………………………………………
Engineering Systems Division
August 12, 2011

Certified by………………………………………………………………………………………
Dr. Richard K. Lester
Head of the Department, and Professor of Nuclear Science and Engineering
Director, MIT Industrial Performance Center
Committee Member

Certified by…………………………………………………………………………….............…
Dr. Ernest J. Moniz
Cecil and Ida Green Professor of Physics and Engineering Systems
Director, MIT Energy Initiative
Committee Chair

Certified by…………………………………………………………………………….............…
Dr. Mort D. Webster
Professor, Engineering Systems
Committee Member

Accepted by…………………………………………………………………………….........…
Olivier L. de Weck, PhD
Professor, Aeronautics and Astronautics and Engineering Systems
Chair, Engineering Systems Division Education Committee

A Decision Analysis Framework for the U.S. Nuclear Fuel Cycle

by

Lara Pierpoint

Submitted to the Engineering Systems Division
on September 5, 2011 in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy in Engineering Systems

ABSTRACT

If we are willing to pay a premium, we may be able to mitigate some of the long-lasting impacts of nuclear waste. Deciding how to navigate this tradeoff, between cost and waste, is a central challenge for stewards of nuclear power. It is made more difficult by uncertainties that characterize the global future of nuclear electricity generation.

The recent increase in concern about climate change has prompted U.S. policymakers to back strategies favorable toward nuclear power, so much so that some experts see a "nuclear renaissance" on the horizon. Whether such a renaissance will come to pass, involving the construction of a vast new fleet of nuclear plants, is unclear – especially in light of the March 2011 nuclear accident at the Fukushima Daiichi reactors in Japan. Even more unclear is what should be done with the commercial U.S. nuclear waste, given an array of technical options and a large amount of uncertainty about how much waste will ultimately need to be managed.

This study introduces a framework for analysis of strategies to evolve the nuclear fuel cycle which may be helpful in analyzing decision problems for similarly complex, long-lived technical infrastructure systems. The framework consists of a system dynamics simulation coupled with a decision analysis model. The system dynamics code is developed specifically for this study to be simple, fast-running, and also to echo the results of many previous nuclear fuel cycle simulations in demonstrating how various technical options impact important parameters (like uranium consumed, waste generated, etc.). Code results are benchmarked to more complex fuel cycle simulations for the parameters relevant to the decision space. The decision analysis model takes information from the simulation and makes it useful to policymakers, by allowing the explicit analysis of desirable decision pathways under uncertainty, and also considering tradeoffs among system goals.

The framework is applied to three nuclear systems, the light-water reactor (LWR) once-through fuel cycle, which represents the status quo, an advanced, traditional, plutonium-fed self sustaining fast reactor fuel cycle, and a fast reactor fuel cycle for which initial fast reactor cores are composed of enriched uranium rather than recycled LWR fuel. Fast reactors are highly likely to cost more than LWRs, but they can produce electricity from some of the elements that most plague the long-term management of a nuclear waste repository. A value function compares how these options fare under different scenarios, incorporating system-wide costs and the system waste burden as the two attributes in the function.

The primary result is that the best strategy, under a strong preference for eliminating LWR spent nuclear fuel waste, consists of building a few traditional fast reactors now, and then building a full fleet more rapidly later in the century. This allows both for a significant amount of waste mitigation compared to an all-LWR fuel cycle, and for the costs associated with the more expensive fast reactor technology to be incurred primarily later in the century. On the other hand, if cost is the main consideration, the framework advises moving forward with the once-through LWR fuel cycle and avoiding fast reactors altogether, or at least until later in the century.

These results are examined from a traditional decision analysis perspective, and then from one that departs somewhat from the assumption of a fully powerful decision maker. In reality, a government decision maker can only offer incentives to industry in order to induce a strategy change. Changing the decision model to reflect this reality causes the framework to more strongly advise moving forward with traditional fast reactors. This occurs because any single attempt at offering incentives to industry might be unsuccessful, and thus prevent a waste-concerned government from achieving any significant mitigation.

The most important contribution of the methodology is its ability to illuminate which parameters represent strong drivers of system decisions. Preferences across competing attributes are always important: in general, if decision maker preferences for reducing cost vs. waste were to shift significantly, the framework would show a change in the desirable decision strategy. Decision results are not very sensitive, on the other hand, to the rate of nuclear power growth or to the cost of fast reactor technology.

A second contribution comes from the initial foray into studying a more complex decision maker perspective, and shows how a different view can complement results using the traditional decision analysis assumption of an "ideal" decision maker. Ultimately, the system dynamics/decision analysis framework presented here helps identify desirable pathways for complex system evolution, identifies factors that bear strongly on decisions and which are deserving of more study, and begins to show how strategy implementation can be considered within the framework in order to further improve decision-making.

Thesis Supervisor: Ernest J. Moniz
Cecil and Ida Green Professor of Physics and Engineering Systems

Committee Member: Richard K. Lester
Japan Steel Industry Professor of Nuclear Science and Engineering

Committee Member: Mort D. Webster
Assistant Professor of Engineering Systems

Acknowledgements

My greatest, most heartfelt thanks goes to Professor Ernie Moniz, my advisor, for leading me and mentoring me though the work that has resulted in this thesis. Despite his unbelievably busy schedule at the helm of the MIT Energy Initiative, he gave me very generous amounts of his time, challenged me in every meeting, and (almost obnoxiously) remembered nearly everything I ever gave him to read or told him about my work. I deeply appreciate his support, and am grateful to have him as a role model scientist and policymaker.

Professor Mort Webster worked with me through the deepest weeds of this thesis. Without his guidance and encouragement, I could never have completed this project. I am forever grateful for his willingness to share his war stories and his understanding of life and work as an MIT graduate student. Professor Richard Lester allowed me to trick him into a position on my committee, and I am deeply indebted to him for the valuable and grounding comments he gave me at every committee meeting I held.

Professor George Apostolakis served on my committee for nearly three years before being called to service as an NRC Commissioner. I thank him for the sound introduction to decision analysis I received from him via classes as well as from numerous meetings about value trees in his office.

I also thank Professors Joe Sussman, Richard deNeufville, John Parsons, Charles Forsberg, Mujid Kazimi, Jessica Trancik, Eugene Schwageraus, and Pavel Hezjlar for answering endless streams of my questions and supporting me in my thesis work. Their generosity in giving their time to a student who is not “theirs” is deeply appreciated.

My fellow students in ESD, NSE, and the MIT Energy Club taught me so many important things, on my own topic and on their subjects of expertise. They provided me with relevant academic papers, constructive feedback, great discussions, moral support, and, most importantly, comic relief and the best social life a grad student could ask for. I will cherish piles of great memories, relish our collective grousing, and plan to keep connections with you all.

I thank my yoga family, and chiefly Caroline Harvey, for keeping me whole and teaching me how to breathe. Lynne Begier, Josh Summers, Vanessa Morris, Peter Crowley, and Ana Forrest are all responsible for my health and clarity of mind during this thesis process.

My family at home has been endlessly supportive. Mom, Annie, and David, and many others cheered for me, hugged me, reminded me to relax and offered that rare combination of faith in my success along with unconditional love. You all have made me who I am and this work what it is. I also offer great appreciation to my new family, Maurice, Kerry, Kaily, and Stirling Weir. I love you all not only for providing the best vacation atmospheres ever, but also for the words of encouragement you so readily give.

Finally, I deeply, wholeheartedly thank the love of my life, Tristan Weir. Tristan has believed in me, without fail, every nanosecond of every day I have spent at MIT. I thank him for cheering me on, for attempting to teach me how to say no to commitments, for making me laugh, and for taking all of my crazy in stride. The way you love me has always been exactly what I've needed to thrive here. Getting a PhD will feel great, but you are by far the best thing I found at MIT. I can't wait to be your wife.

Contents

Figures

Tables

Table 3-1: FANTSY Core Simulation Input Variables ... 38
Table 3-2: FANTSY Outputs ... 41
Table 4-1: System Nuclear Fuel Cycle Costs ... 48
Table 5-1: Sample weights for simple tree ... 61
Table 5-2: Standard weighting scheme for partial build analysis ... 73
Table 5-3: Growth rates for three-period tree ... 87
Table 5-4: Remaining second period decisions after paring under sophisticated calculation ... 102
Table 6-1: Value function structures tested for sensitivity of results ... 113

List of Acronyms

AFCI: Advanced Fuel Cycle Initiative

CAFCA: Code for Advanced Fuel Cycle Assessment

COSI: COmelini-SIccard fuel cycle simulation code

CLIOS: Complex, Large-scale, Interconnected, Open, Sociotechnical (System)

DANESS: Dynamic Analysis of Nuclear Energy System Strategies

EUFR: Enriched-Uranium fed Fast Reactor

FANTSY: Flexible Advanced Nuclear Technology Simulation by Year

FCRD: Fuel Cycle Research and Development (Program)

FR: Fast Reactor

LWR: Light-Water Reactor

SNF: Spent Nuclear Fuel

TFR: Traditional (plutonium-fed) Fast Reactor

TRU: TRans-Uranic elements (typically of spent fuel)

VISION: VerIfiable fuel cycle SImulation Of Nuclear fuel cycle dynamics

Chapter 1: Introduction

Complex, high-technology systems are becoming more prevalent in our lives. One category of these could be considered "consumer technologies," which provide a service directly to a device owner. These objects have short life spans, rapid version updates, and low per-unit costs. Cell phones and some internet applications fall into this category. Other high-tech systems, however, enable these consumer technologies. Such enabling systems are large, complex, and generally involve huge investments in long-lasting infrastructure. Many electricity generation and large transportation systems are typical of this group. Managing and evolving these types of complex technologies, that last far beyond political and industrial cycles, presents a unique challenge.

For a certain class of big infrastructure systems, the notion that we can develop a huge suite of technology options to solve a problem and then "let the market decide between them" is not realistic. The investment required to accomplish this is simply too great, and often the returns on investment depend on the continuity of government policy. The more expensive, complex, and long-lasting the system, the more important it becomes to make careful design decisions ahead of any investment or implementation. Making these decisions, however, is challenging because large enabling systems are required to meet an array of sometimes-conflicting objectives. Usually these objectives include low costs, low environmental impacts, high degrees of safety and reliability, and others.

Modeling and simulation tools are increasingly aiding decision makers in defining complex system designs and selecting options. Simulation and decision analysis tools together can help establish a tractable strategy for system evolution, by helping to illuminate the tradeoffs between systems and their objectives. Complex technologies that would benefit from this approach include the U.S. electrical grid (onto which renewable energy sources are unlikely to prosper unless there is significant centralized planning involving storage and backup systems, in addition to cost reductions), and the nuclear fuel cycle (which will never advance beyond its once-through, throw-away status quo unless a careful path is charted).

All of these systems must evolve over the course of decades, and they face numerous uncertainties during that time. It is not yet clear what electricity demand will be, how many nuclear plants will be desirable, or how quickly new generations of technologies will become available. These systems and others like them also suffer, in varying amounts, from murky definitions of *who*, exactly, is responsible for changing them. For the U.S. nuclear fuel cycle and the electricity grid, complex relationships between private and public decision makers govern their futures.

This thesis provides a framework intended to help groups of decision makers compare and select options for evolving complex, long-lasting infrastructures serving multiple objectives. The framework comprises a decision analysis model linked with a system dynamics simulation, allowing for the explicit enumeration of options and uncertainties set over a temporal axis. Establishing a value function, as required by the decision analysis model, highlights the drivers for making system changes, and forces decision makers to consider their relative importance. The framework is applied here to the nuclear fuel cycle, because the nuclear fuel cycle strongly exemplifies the characteristics of a long-lived, technologically complex system with disparate groups of decision makers responsible for its evolution. In the U.S. and elsewhere, discussions focus on how to manage the conflicting objectives inherent in fuel cycle options, including reductions in CO_2, safety, waste management, and others. This thesis concentrates solely on the U.S. perspective.

.

1.1 Motivation: The nuclear fuel cycle as a complex infrastructure

Nuclear power may be poised for a dramatic renaissance in the United States. Chief among the drivers for new nuclear are concerns about the environmental impacts of anthropogenic carbon dioxide, and a desire to improve the sustainability of our energy systems. Though the nuclear meltdowns at Fukushima, Japan are having a negative impact on global public opinion, it is unclear whether the Japanese accident will significantly stall planned expansions of nuclear energy in the U.S. Nuclear remains the only widely available, commercialized technology that can provide baseload electricity while producing negligible amounts of CO_2.

Despite its advantages, nuclear power confronts many challenges besides public opinion in the U.S. electricity market. Costs are paramount among these, because the capital costs for a single nuclear plant range in the several billions of dollars and thus require utilities to bet their entire balance sheets on one project.(J. Deutch & et al., 2009) The difficulty in financing nuclear projects is compounded by the roughly 30-year U.S. hiatus in nuclear power plant builds, causing considerable uncertainty for investors about costs and regulatory risks.

A second major issue with nuclear power involves the unique characteristics of the waste. Spent nuclear fuel (SNF) is highly radioactive, dangerous, and long-lived. It also, however, contains significant potential to produce more electricity from various elements generated in or remaining in the fuel. These valuable components are stripped out and re-fabricated for use again in reactors to produce electricity. This reprocessing and reuse of SNF has the potential not only to increase energy extraction from the material, but also can in principle reduce some of the long-lived radiotoxicity and volume of nuclear waste. Whether or not SNF is reprocessed, high-radiation wastes will require disposal. Scientists have long agreed that the best method of disposal consists of underground burial.(Hess, 1957) Yet establishing a program for nuclear waste disposal has proved difficult in the U.S.

The 1987 amendment to the 1982 Nuclear Waste Policy Act designated Yucca Mountain, Nevada as the site for a nuclear waste repository. Since then, over $13 billion has been spent characterizing and designing the repository location.(Yucca mountain cost estimate rises to $96 billion.2008) The Department of Energy used this information to submit a license application for construction and operation of the repository to the Nuclear Regulatory Commission (NRC) in June 2008. In March of 2010, however, the Obama administration moved to withdraw the application from the NRC, plunging the future of Yucca Mountain and of the pathway to nuclear waste management into doubt. At the same time, President Obama appointed the Blue Ribbon Commission on America's Nuclear Future to chart a course for SNF management in the U.S. The Commission will make recommendations, but both public and private actors will still be responsible for making decisions at many points over time in order to determine the fate of the fuel cycle.

This thesis is directed toward developing and applying a methodology for decision making with regard to the U.S. nuclear fuel cycle. The intent is to help provide tools to navigate complex decisions between different pathways for fuel cycle development and investment. An example presented in this thesis concerns the choice between employing advanced recycling or maintaining the existing and well-understood once-through fuel cycle based on light water reactors (LWRs). Recycling nuclear fuel may make sense, but the utility of any fuel cycle will depend heavily on the trajectory for nuclear power growth (i.e. whether a nuclear renaissance indeed takes place), the cost of advanced reactors and recycling technologies, and proliferation considerations, among other things. Each of these parameters is highly uncertain, so that the risks of investing in infrastructure that will last 50 years or more are especially acute. The high degree of uncertainty present for this system calls for a management approach considering flexibility and the preservation of options.(Saleh, Hastings, & Newman, 2001)

1.2 Focusing the Problem: the Cost-Waste Tradeoff

Although many different challenges confront nuclear power, two are often cited as the major barriers to a true nuclear renaissance. These two problems are the cost of the technology, and the nuclear waste it produces. While challenges associated with nonproliferation, safety, and public perception are important, the issues of waste and cost seem to be especially central to arguments for and against various configurations of the nuclear fuel cycle.

Federal policy currently requires the U.S. government to take title to commercial spent nuclear fuel and dispose of it. But, as described above, this has proved extremely difficult given large amounts of political and scientific opposition. The lack of a solution is damaging: several states have enacted bans on building new nuclear plants until a waste disposal program is operating.

Reprocessing spent LWR fuel and reusing it in advanced reactors could, however, be one mitigating solution. Such a fuel cycle would sequester waste into reactor cores and put it to work producing electricity, and could potentially reduce the long-term waste burden posed by traditional spent LWR fuel. This might allow for building a smaller final repository, and prevent the need for additional repositories beyond the first one. Less waste would be more palatable politically, potentially enabling a true life-cycle nuclear waste solution and incentivizing greater participation from industry in building and operating reactors as part of the energy mix.

This advanced fuel cycle, however, will almost certainly be more expensive than the traditional once-through fuel cycle.(Kazimi, Moniz, Forsberg, & et al., 2011) Given that even LWRs operating in once-through mode are expensive (see above), increasing the cost of nuclear power could simply ensure that it is never used. A true tradeoff exists between options for waste mitigation and options for keeping the cost of nuclear as low as possible.

One caveat is in order: the extent to which LWR waste can be truly "mitigated" by an advanced fuel recycling is still the subject of debate. For example, although plutonium recycle decreases the volume of high-level waste destined for disposal, volume may be an irrelevant metric in a repository for which the loading factor is determined by the waste's heat output. Fast reactor recycle may lower the radiotoxicity, volume, and heat output of high-level waste compared to the once-through fuel cycle, but the extent of the benefit must be confirmed. Analysis is needed on a life-cycle basis, in conjunction with the relevant (and yet unknown) final disposal site characteristics. This thesis makes the assumption that an advanced fuel cycle will provide a tangible waste benefit. Indeed, if it is shown that the waste benefits of advanced recycling are negligible, considering such cycles is pointless at this point in time (under the plausible assumption that fresh fuel in a once-through fuel cycle is not a significant constraint).

Deciding which fuel cycle to implement in the face of the cost-waste tradeoff is difficult. One could claim that the decision should depend on how much new nuclear we build; if we build lots of LWRs, and potentially will need vast amounts of repository space, we should look into reprocessing. But this argument makes an assumption about how desirable it is to mitigate the waste burden. Even if we build very few new nuclear plants, deploying a small number of advanced reactors and associated recycling facilities could be an attractive way to reduce the waste we have.

Further complicating the decision, a large number of entities have a vested interest in the structure of the nuclear fuel cycle and also have different perspectives on which aspects of the problem matter. The government agencies responsible for waste management would like a politically expedient and socially acceptable way to handle and dispose of spent nuclear fuel. Industry actors, including the major reactor vendors and nuclear utilities, are also concerned about waste, but are more immediately interested in the cost of producing electricity.

The problem is complex, fraught with different actors of varying sizes, roles, and decision making power. These groups have decisions to make at many different strategic levels and across long timescales. Focusing on the cost-waste tradeoff with simplified representations of the multiple decision makers helps narrow the analysis, and sheds light on some of the better and worse pathways for evolving the U.S. nuclear fuel cycle.

1.3 Research Objectives

Over the last three decades, credible institutions have made extensive and deep system studies of advanced nuclear fuel cycle deployment. These studies can be broadly classified in one of three categories: in-depth design and analysis of individual advanced systems (e.g. of an advanced fast reactor or an aqueous reprocessing plant), static or equilibrium analysis of system interactions (e.g. on how much fuel would regularly be transferred in a fast reactor-reprocessing plant system), and dynamic analyses (e.g. on how quickly fast reactors can penetrate the nuclear market). Each type of study provides information relevant to decision makers at some level. This thesis aims to synthesize that information and put it into a framework relevant for policymakers at the level of the federal government, by creating a lens for viewing the results of dynamic fuel cycle analyses.

The primary research objective is to create a decision analysis framework for understanding and comparing various advanced fuel cycle options. The value function used to evaluate the utility of various nuclear power scenarios will be flexible and expandable, allowing decision makers to ultimately decide the criteria by which technologies will be chosen. The decision analysis methodology further allows for explicit enumeration and evaluation of uncertainties that may confront nuclear fuel cycle development. This increases the usefulness of the model, and adds a new dimension to the field of fuel cycle analysis. Various flexible patterns of development can be analyzed within the framework to evaluate them as potential responses to the uncertainties.

Among the limitations of decision analysis is the traditional methodological assumption, described more fully in section 2.1, that the system is governed by a single decision maker.(Keeney, 1982) This assumption is extremely restricting for many systems of interest, and especially for the nuclear fuel cycle. Nuclear reactor vendors, utilities, Congress, regulators, and local governments will all make decisions shaping fuel cycle evolution. The framework presented here is flexible enough to begin considering these interactions between private and public decision-makers, and one example is presented wherein industry actions are modeled from the perspective of the government.

Research questions addressed in this thesis include the following:

- Given an initial, simple value function, what are some of the preferred pathways for U.S. nuclear fuel cycle development?

- Does considering a broader range of options (including those that can flexibly respond to uncertainty) change the suite of desired pathways?
- What are the effects of recognizing some the limitations of government as a decision-maker that interacts with the nuclear power industry (in effect acknowledging that the system has no all-powerful decision-maker)?

1.4 Methodological Overview

The primary focus of this work is a decision analysis model intended to both assist policymakers and expand the discussion of options and uncertainties that deserve the attention of fuel cycle studies. The full methodology underlying the decision tree and attendant calculations is described in Chapter 4.

Providing solid inputs to the decision tree, in the form of fuel cycle scenario parameters, is essential to conducting a valid analysis. To this end, a new fuel cycle simulation is created and linked to adjustable sets of decision trees. The fuel cycle simulation is designed to be very simple, fast-running, and inclusive of only the variables most central to the analysis problem at hand. A full description of the MATLAB-based system dynamics fuel cycle model is presented in Chapter 3.

The decision tree and fuel cycle models, plus their linkage through a set of Microsoft Excel spreadsheet calculations, form the methodological basis of this work (see Figure 1.1; methodological tools are listed in orange boxes, and software packages are represented in grey boxes).

Figure 1-1: Outline of methodology: two linked models

The inputs to the overall process include “traditional” fuel cycle simulation inputs. They are the advanced fuel cycle technology types in question, the times at which they are introduced, and the assumptions for nuclear power growth. These parameters are varied and arranged into scenarios, corresponding to scenarios of interest that are modeled as paths through the decision tree. Each branch of the decision tree model in Figure 1-1, ranging from left to right, represents a scenario with certain set values for each of the model parameters. The structure of the scenarios in the context of the decision tree is discussed more fully in Chapter 4.

The outputs of the entire process include groups of scenarios that prove desirable under various conditions. Rather than assume the preferences of the decision makers, results are reported for a full range of “preference values,” any one of which in principle could represent the true desires of the decision making body. Other methodologies address the elicitation of these preference values in the context of decision analysis; such elicitation is beyond the scope of this work (and instead is replaced by analyzing the full range of preference values that could be chosen). Because preference elicitations are controversial, especially for groups of decision makers (and likely even more so for disparate groups public and private institutions), reporting desirable decisions over a range of preference values and demonstrating each decision's relative robustness was deemed more appropriate for this work. Further investigations should explore whether strategy "solutions" are more useful to decision makers when calculated for elicited preferences or when reported for a preference range.

1.5 Thesis Organization

This document is organized in eight chapters and two appendices. Chapter 2 presents a focused literature review of the decision analysis methodologies and fuel cycle analyses relevant to this work. Chapter 3 describes the system dynamics fuel cycle code, the Flexible Advanced Nuclear Technology Simulation by Year (FANTSY), developed to provide inputs to decision analysis models in this thesis. Chapter 4 explains the decision analysis methodology which forms the basis for the contributions of this work. Chapter 5 describes a series of studies performed using the decision analysis framework, each surrounding an important question about options for evolution of the nuclear fuel cycle. Chapter 6 presents seven sensitivity studies, which demonstrate the robustness of the results presented in Chapter 5. Chapter 7 discusses the impact of relaxing the assumption that a single decision maker has full control over the system, and Chapter 8 presents conclusions and recommendations for further work. The appendices provide the MATLAB script for the FANTSY fuel cycle analysis code, and a comparison of FANTSY results to a benchmark study of fuel cycle simulations.

Chapter 2: Selected Background on Decision Analysis and Nuclear Fuel Cycle Studies

The primary analytical foundations of this thesis include areas of decision analysis and simulation and modeling studies of the nuclear fuel cycle. Both subjects are exceptionally broad, so this chapter does not seek to provide a complete review of either one. Rather, the aspects of each most relevant to this work are highlighted, and the final section reviews the small body of literature on combining system dynamics methods with decision analysis.

2.1 The "Canonical Method" of Decision Analysis

Decision analysis is often divided into three "schools" of study: normative, descriptive, and prescriptive; an overview of the decision analysis field is described in (Tang, 2006). Normative theory describes how humans *should* make decisions if they are following ideal rules of logic and rationality.(D. Bell, H. Raiffa, A. Tversky, 1988) Descriptive theorists analyze how decisions are actually made. Prescriptive theory, the school in which this thesis is grounded, is concerned with improving the process and results of decision-making.

Within the prescriptive school of decision analysis, one major theme is of interest for this work, providing the foundation for a contribution to the field. This is the "canonical method" of decision analysis, consisting of a set of steps which a decision maker (or analyst) should follow in order to identify a recommended choice. Several versions of this method have been presented since the birth of decision analysis, all of them more or less following a pattern of problem definition, followed by the generation of options, and then option examination.

John Dewey created one of the first models for phased decision making steps in 1910.(Hansson, 2005) His model included five stages: problem identification, problem definition, generation of alternatives, evaluation of alternatives, and further observation leading to acceptance or rejection of the alternative. Herbert Simon proposed modifications to this same basic outline in 1960(Hansson, 2005), and then others followed by proposing sequential and circular/iterative versions of the process. Ralph Keeney, considered to be one of the "fathers" of modern decision analysis, proposed an iterative set of steps that added one for explicitly quantifying the values of decision makers, and another that prescribes performing a sensitivity analysis.(Keeney, 1982) Among others famous for their versions of this canonical set of decision analysis steps are (Clemen, 1996), and (Howard, 1988), who is often credited with coining the term “decision analysis.” Figure 2-1 shows one version of the canonical model, after the models described by Clemen and Keeney.

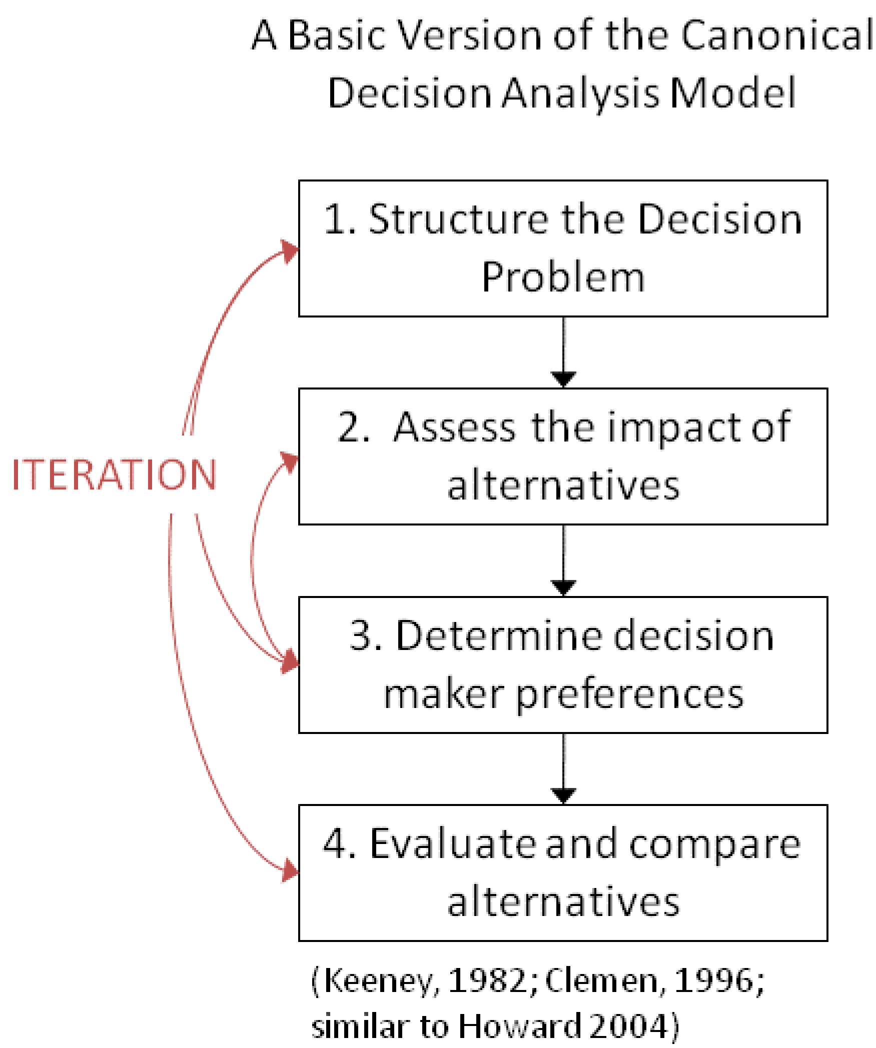

Figure 2-1: One general description of a canonical decision analysis model, with proposed iterative steps

The vast majority of initial literature on decision analysis recognized a common problem: the canonical model and the tools of decision analysis assumed a single, all-powerful decision maker.(Keeney, 1982) Keeney and others realized that few decisions are made by such ubiquitous beings, but argued that decision analysis could nonetheless provide a useful perspective on many problems. Subsequent research has tackled this problem in many ways, including by incorporating game theory for cooperative or competitive players (who are making decisions as they play).

A study by a National Research Council committee suggested a different avenue for handling multiple decision makers, producing a report in 1996 that recommended decision analysts incorporate stakeholder deliberation while moving through the steps of a canonical decision analysis model.(Committee on Risk Characterization, National Research Council, 1996) The benefits of deliberation include improved communication among decision makers, such that discrepancies in viewpoints are more easily discovered and consensus can be more easily reached.(Weil, R. and Apostolakis, G., 2001) The framework presented in this thesis is intended to be used along with a deliberative process, but testing it with actual (or acting) human decision makers was beyond the scope of work and is left to further study.

It is important to note that performing this analysis with the full participation of several stakeholders would not completely solve the multi-decision maker problem. For the nuclear fuel cycle and similar systems, the problem is not simply that there are multiple people involved with decision making; there are multiple *groups* responsible for the relevant decisions, each composed of other groups with varying aims, and it is not always clear which groups possess decision making authority. This study offers initial steps toward a solution, by suggesting a modification to the most recent versions of the canonical decision analysis process.

Variations of the canonical procedure can now be found outside traditionally-defined boundaries of the decision analysis field. One example of a canonical model version comes from Joseph Sussman and his students in the MIT Engineering Systems Division, known as the "Complex, large-scale, interconnected, open, sociotechnical" (CLIOS) process.(Sussman, Dodder, McConnell, Mostashari, & Sgouridis, 2007) The three major stages of the CLIOS process consist of: (1) system *representation*, (2) *design, evaluation, and selection* of alternatives, and (3) *implementation* of the selected alternatives. CLIOS is a more thorough expression of the canonical decision analysis model (see Figure 2-2). Another example, designed specifically in the context of nuclear waste cleanup, is called the "KONVERGENCE" Framework (for which the goal is to find where Knowledge, Values, and Resources converge).(Piet, Gibson, Joe, Kerr, & Nitschke, 2003) Like CLIOS, the KONVERGENCE model includes explicit consideration of implementation of options.

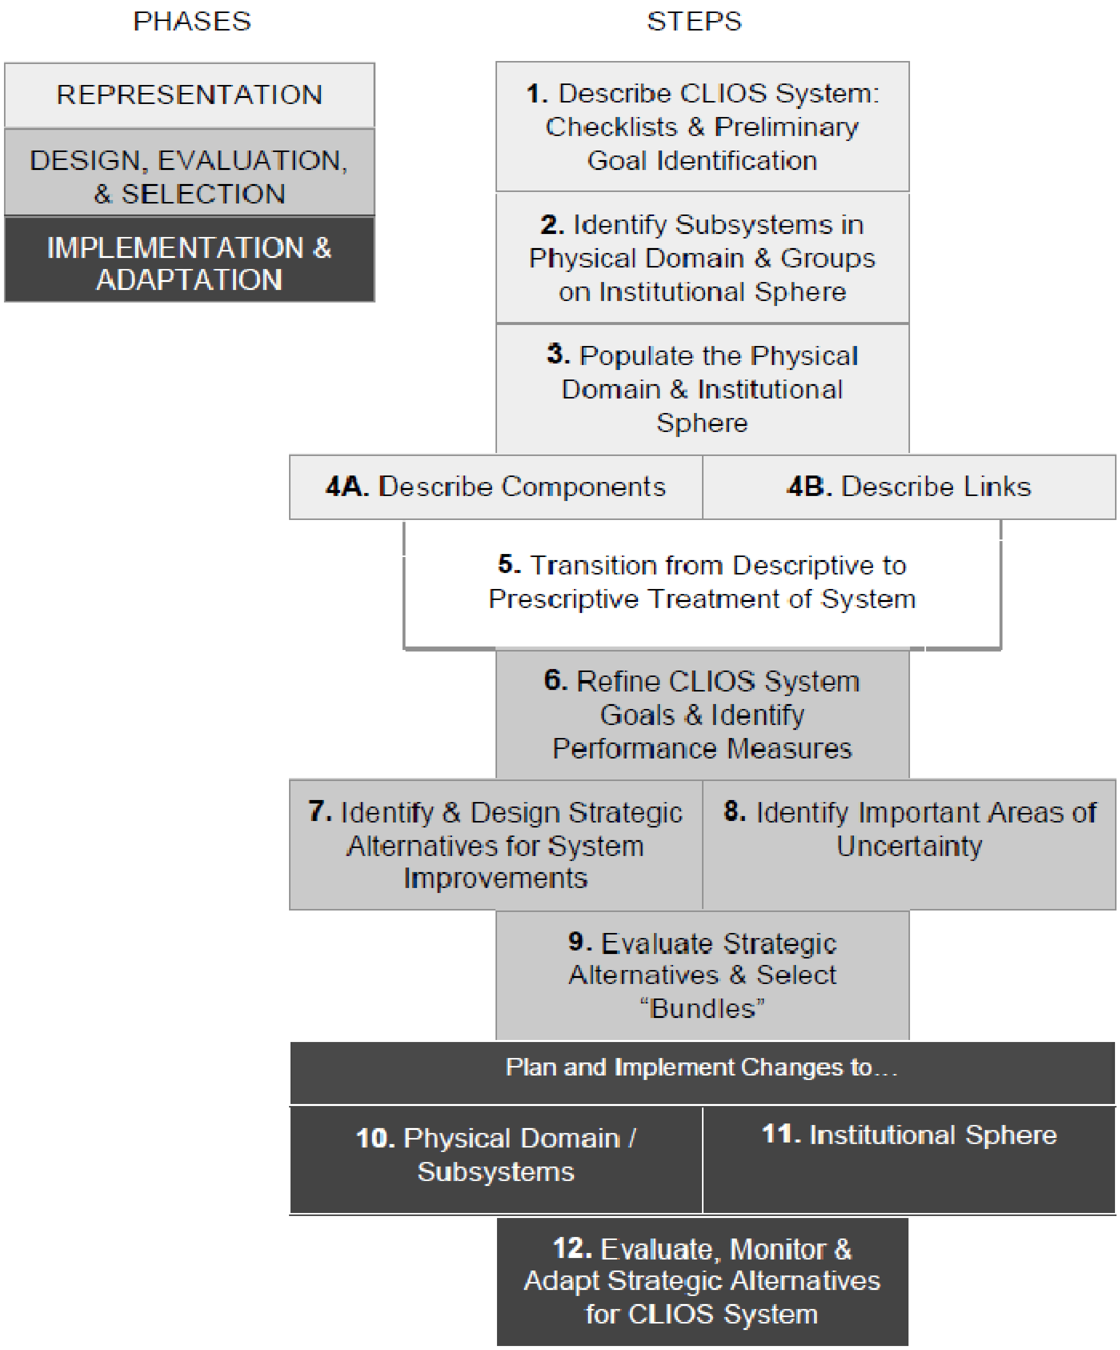

Figure 2-2: Stages of the CLIOS processes (Sussman, 2007)

These and other models offer comprehensive decision support for ranges of complex systems, without necessarily invoking decision analysis. Each recognizes the importance of iteration between steps, the need for a strongly creative process to generate options, and the critical importance of implementation phase planning. The framework established in this thesis could be aligned with any one of these system analysis and evolution procedures.

A new contribution to this class of processes, demonstrated here, is the utility of performing the set of procedural steps *twice*. The stages of problem definition, alternative generation, and analysis of options are first performed under the assumption that a single, all-powerful decision maker can alter the system in order to maximize public welfare. In some sense, this perspective is aligned with that of the federal government, which acts to benefit the common good. But an examination at the *implementation planning* phase quickly illuminates the fact that the government will not have the ability to enact the desired options as they have been defined. A second iteration is performed, where the alternatives are adjusted to reflect the interaction between government regulations and industry responses in the nuclear sector. The information from the first iteration about "best" options can then be compared with those that are more realistic, and a better option set ultimately chosen.

This thesis lays the groundwork for how such a "dual-iteration" procedure might work; further study will be needed to determine its effectiveness for other types of problems and whether continued iteration with more complex decision maker definitions could illuminate even better sets of desired alternatives.

2.2 Nuclear Fuel Cycle Studies

The U.S. Department of Energy (DOE), along with many other organizations, has been conducting research on the nuclear fuel cycle for over 40 years. In 2001, the DOE began studying fuel cycle systems under the aegis of the Generation IV initiative. In 2003, the system studies program was expanded as part of the Advanced Fuel Cycle Initiative (AFCI). Work continues under the AFCI, now called the Fuel Cycle Research and Development (FCRD) Program (after a brief period under an umbrella initiative called the Global Nuclear Energy Partnership). This work includes annual reports and other documents produced by researchers at Idaho National Lab and Argonne National Lab, and by FCRD-supported university research. A small selection of U.S. fuel cycle reports, many (but not all) produced with full or partial support of the AFCI/FCRD, is described in section 2.2.2 below.

2.2.1 Why Study the Nuclear Fuel Cycle: AFCI/FCRD Goals

Policymakers and technologists have produced several iterations of the goal set to which advanced nuclear fuel cycle technologies should aspire. Each time these goals/values are restated, the theme is the same: the nuclear fuel cycle should be sustainable, imposing as small an impact as possible on resources and nuclear proliferation. To varying extents, every fuel cycle study invokes goals along these lines, and qualitatively or quantitatively discusses the impacts of advanced technologies on the goals.

The current DOE program, FCRD, states explicitly that its objective is to conduct research and help develop sustainable fuel cycles. It states goals for fuel cycle technologies through a definition: "Sustainable fuel cycle options are those that improve uranium resource utilization, maximize energy generation, minimize waste generation, improve safety, and limit proliferation risk."(U.S. Department of Energy, 2011) Though the program managers did not explicitly name low cost as a goal, they do indicate later in the document that the fuel cycle must be "acceptable to the American public." As such, the final option must adhere to a cost goal that was defined in the AFCI mission, to "ensure that advanced fuel cycle technologies cause no significant decrease in the economic competitiveness of nuclear electricity."(Dixon et al., 2008) Indeed, given the current competitive environment in which all nuclear and non-nuclear energy technologies are judged, a fuel cycle that fails to meet this objective is sure never to be implemented.

2.2.2 Nuclear Fuel Cycle Reports

Among the factors that make fuel cycle decision making difficult is the sheer volume of information available about technological options. Reports on the subject range from the purely qualitative to those based entirely on quantitative models. On the qualitative end, industry has suggested paths forward for fuel cycle evolution by combining statements from experts.(*Nuclear energy for the future: Required research and development capabilities - an industry perspective*2008) Similarly, recommendations to government by advisory committees also tend to come in the form of the gathered wisdom of fuel cycle experts, who themselves are familiar with a wide range of studies (e.g. (Martin, Ahearne, & and the Nuclear Energy Advisory Committee, 2008)). Hybrid qualitative/quantitative studies abound as well. One example is the MIT Future of the Nuclear Fuel Cycle report, which includes fuel cycle system modeling for scenario analysis and a separate model for uranium consumption, alongside qualitative sections assessing the pros and cons of various fuel cycle options.(Kazimi, Moniz, Forsberg, & et al., 2011)

Highly quantitative modeling and simulation fuel cycle studies tend to fall into three broad categories: individual facility analyses (such as in-depth modeling of an advanced reactor), static or equilibrium system analyses (usually discussions of *system-wide* impacts of the fuel cycle, including linkages between reprocessing plants and advanced reactors), and dynamic analyses (fleet-level system analyses that examine the development of the fuel cycle system over time). Each type of study adds crucial information to the field of fuel cycle analysis.

Among the most important single-facility analyses are those that examine advanced reactors. Reactor studies often include calculations performed using core neutronic and thermal-hydraulic models, in order to simulate the performance of reactor fuel and the potential response of the reactor system to design-basis accidents. Examples include numerous MIT theses (see e.g. (Pope, Driscoll, & Hejzlar, 2004)). One of the most cited and influential studies in the area of fuel cycle analysis is Edward Hoffman's design of a sodium-cooled fast reactor (SFR).(Hoffman, Yang, & Hill, 2008) Hoffman performed an analysis of core operation, and in the process defined the isotopic compositions needed to sustain reactor operation for a range of fast reactor conversion ratios. His isotopic vector "recipes" are used in this thesis as well as Idaho National Lab's fuel cycle simulation called VISION.

The in-depth studies of individual facilities are vital for any advancement of nuclear technology, but they do not provide all of the information needed to make decisions about fuel cycle evolution. Advanced reactors in fact exist within a complex system of reprocessing plants, light water reactors, fabrication plants, and disposal systems, and these all need to be technologically and economically harmonized in order to close or significantly alter the fuel cycle.

Many system studies focus on collections of multiple facility types. Wigeland and Bauer, for example, consider the repository impact of thermal recycling by looking at the interaction of light-water reactors (LWRs) and aqueous reprocessing facilities.(Wigeland & Bauer, 2004) They are able to quantify the benefit of limited thermal recycle on repository loading in terms of a mass factor increase that can be loaded into Yucca Mountain. They do not make any analysis of fleet characteristics or nuclear power growth, but instead concentrate on a representative LWR and reprocessing facility. In this sense, the analysis can be said to be “static” – it does not address growth or change over time.

Studies that estimate the costs of advanced recycling also often take static or equilibrium viewpoints. Parsons and De Roo perform a calculation of advanced fuel cycle costs assuming that LWRs and fast reactors are in equilibrium with one another; this means they assume that the numbers of LWRs and fast reactors are attuned to one another such that there is no (or very little) excess fuel passing between them, and also that the isotopic composition of multi-recycled fuel changes very little with each pass through a fast reactor.(De Roo & Parsons, 2009) This approach enables them to calculate a levelized cost of electricity for the advanced fuel cycle system that can be compared to the levelized cost of nuclear power today.

The third category consists of dynamic fuel cycle studies. Many of these studies are relatively recent, because complex fuel cycle simulations have become gradually more accessible as the speed of computing has increased. The use of dynamic fuel cycle analysis ranges from single-growth rate, single-technology scenarios intended to bolster qualitative arguments, to extensive use of complex codes that include sensitivity analyses.

The codes which are the workhorses of these analyses similarly range from the very simple to the very complex. Among the most complex and highly-developed fuel cycle codes are VISION (Idaho National Laboratory), DANESS (Argonne National Laboratory), and COSI (CEA – the French Atomic Energy Commission). Of the three, COSI may be the most sophisticated, because it includes linked neutronics models to calculate fuel compositions as it runs. VISION and DANESS are both written on system dynamics platforms (PowerSim and iThink, respectively), and include millions of input options and extensive output data.

Researchers at Idaho National Lab used VISION in 2008 to perform one of the most extensive fuel cycle analyses to date, called the Dynamic Systems Analysis Report for Nuclear Fuel Recycle (DSARR).(Dixon et al., 2008) DSARR’s primary objective is a comparison of three fuel cycle systems: (1) once-through, (2) “1-tier” (LWR spent fuel is sent to immediate fast reactor recycle of all transuranics), and (3) “2-tier” (LWRs recycle spent fuel in a limited manner, followed by a later introduction of fast reactors). System scenarios are run assuming a 1.75% growth rate in nuclear energy demand, and assuming single introduction dates for each advanced technology. The technology options are then compared in terms of their uranium usage and waste impacts (which is divided into dimensions of dose, heat, and radiotoxicity). The costs of each system are separately estimated using static assumptions, to provide a further metric of comparison. Proliferation concerns are not addressed in the report.

The notions of uncertainty and flexibility are addressed qualitatively in DSARR. The authors suggest three ways in which fuel cycle systems could be flexible, including (1) burning Pu-MOX in LWRs if fuel separations have begun but fast reactors are not ready for deployment, (2) fast reactors could use excess weapons material or enriched uranium if separations are not operational in time, and (3) demonstration plants could be converted upward to commercial-scale plants. Each of these possibilities could prove to be useful pathways for fuel cycle system evolution, but they all are more like solutions to potential problems than they are options that can be exercised for a range of uncertain futures. Furthermore, they do not address uncertainties in nuclear power growth, costs of advanced technologies, or the possibility that other events could dramatically alter the desired suite of technologies at any given time.

The authors do, however, quantitatively assess the possibility that advanced technologies might be desired (or only available) later in the future. A sensitivity analysis on the date of fast reactor introduction shows that late deployment of FRs actually slightly *increases* the number of fast reactors that are ultimately built, because the build rate of FRs is so sensitive to the availability of used LWR fuel. The effect, however, is minimal. These results track with some of the scenarios studied in this thesis, but they are not discussed further in DSARR.

DSARR is extremely successful in accomplishing its objective of describing some canonical technology options for fuel cycle evolution, and comparing them to the stated goals of the AFCI. The authors have made an enormous contribution to the information required by policymakers to move development of the fuel cycle forward. The picture painted by DSARR and many other dynamic fuel cycle studies, however, is incomplete. In order to make good decisions about the future of nuclear power technologies, policymakers will need to better understand the impacts of various uncertainties and the tradeoffs between fuel cycle impacts. Nearly all studies like DSARR address a range of metrics (e.g. resource sustainability in the form of uranium consumption, heat output of the waste), but do not offer a quantitative assessment of how the objectives of sustainability, cost, safety, and security trade off among one another.

The framework outlined and demonstrated in this thesis forces explicit consideration of sources of uncertainty, the level of flexibility with which various systems address those uncertainties, and the tradeoffs between system values. The technology options evaluated here are few, because the intent was to demonstrate the method more than to provide a comprehensive fuel cycle option analysis. The hope is that as fuel cycle information becomes richer, new data and important options can be added into the framework in order to provide policymakers with a holistic analysis of fuel cycle evolution pathways. Ultimately, this framework can help identify the most (or least) promising ways to balance competing goals and move fuel cycle development forward.

2.3 Hybrid System Dynamics–Decision Analysis Studies

System dynamics and decision analysis are two tools that were developed independently to aid system managers. Each has strengths and weaknesses. Decision analysis offers the ability to quantify the benefits of real options in system evolution, so that managers can take advantage of flexibility and understand when to exercise options. Within a decision analytic framework, however, accounting for feedbacks, endogenous uncertainties, and nonlinearities is exceedingly difficult. System dynamics models make short work of handling nonlinearities and feedbacks, and can to some extent incorporate simple decision rules and flexibilities with use of “if-then” statements, but consideration of multiple-decision, multiple-uncertainty problems in the search for an optimum decision rule is intractable.

Several researchers have taken initial steps toward combining system dynamics and decision analysis, in order to reap the benefits of both. Hovmand and Ford applied a hybrid SD-DA approach to a social question, on what the best methods are for preventing domestic abuse while simultaneously maintaining low rates of victim arrest. Their approach included developing a system dynamics model, detailing the impacts of a domestic violence mandatory arrest policy. They then constructed decision sequences, involving options for advocate and community training, in case primary victim arrests increased after implementation of the mandatory arrest policy. These were carefully added into the decision model as “if-then” statements, and the model was run to determine which decision rule produced the most desirable results. They noted that coding the decision rules into the model was a difficult task, and that the best learning from the process came from thinking about how to structure new decision strategies.(Hovmand & Ford, 2009)

Nathaniel Osgood showed that he could drastically reduce the calculation time with a different method: he aligned decision tree branches with system dynamics runs and then performed a backward-induction calculation on the tree to determine the best decision rule. In this way, decision rules did not need to be included in the system dynamics model.(Osgood, 2005) He demonstrates his method with a very simple case, involving a single decision, single decision criterion, and one uncertainty. The method used in this thesis is based on Osgood's approach, but is extended to incorporate multiple decisions, uncertainties, and decision criteria.

Burcu Tan combined real options valuation approaches with system dynamics in a series of papers. He uses a system dynamics model to generate cash flow distributions for each time period of a decision tree, and then performs backward induction on the tree to calculate the best course(s) of action and to discern the value of the options modeled. He uses NPV as the single decision attribute, which is often appropriate for a firm's decision to invest in a project. He also, interestingly, models all uncertainties within the system dynamics model rather than specifying them in the decision tree. This has the advantage of keeping the tree simple, but obscures the direct impacts the uncertainties are having on the results.(Tan, Anderson, Dyer, & Parker, 2010)

None of these hybrid approaches so far have tackled national-level multi-decade strategic planning problems, nor have they considered more than one objective to optimize. This thesis addresses those gaps, applying decision analytic approaches to the relatively well-understood system dynamics characterization of the nuclear fuel cycle, and incorporating multiple objectives and explicit analyses of uncertainty in order to inform decision maker discussion.

Chapter 3: The Flexible Advanced Nuclear Technology Simulation by Year (FANTSY)

Chapter 2 described several of the most important fuel cycle analysis codes. An effort was made in the course of this work to use an available, well-tested fuel cycle simulation in order to obtain information about different scenario options. Unfortunately, neither of the two accessible options was viable, mostly for logistical reasons. VISION in its third version contained some non-working parts, which were resolved too late to serve as the basis for this thesis; an MIT simulation called the Code for Advanced Fuel Cycle Assessment (CAFCA) is too limited in the options it offers, and as of writing had no active stewards or users of the code. A new fuel cycle code was created to be maximally flexible, to run quickly, and to interface easily with Microsoft Excel and TreeAge software programs.

The new fuel cycle code is written in MATLAB script and is called the Flexible Advanced Nuclear Technology Simulation by Year (FANTSY). FANTSY consists of two primary files: a basic 100-year fuel cycle simulation script, and a wrapper that runs the fuel cycle simulation multiple times with different inputs for each run. Each time a new study is done with a new decision tree, the two files are adjusted to incorporate items of interest. One version of each file (complete code) is presented in Appendix A.

3.1 Code Description

The wrapper code is designed to run the primary FANTSY simulation file one time for each branch of a decision tree. Each branch is associated with a number of inputs, including the growth rate for nuclear power at each epoch, the type (or types) of nuclear plants to build and the relative amounts. FANTSY also contains options to define a start date for centralized storage, and/or to specify the loss percentages for spent fuel recycling; neither of these variables is required for the simulation to run. See Table 3.1 for a complete description of the required input variables, where LWR = light water reactor, TFR = traditional (spent TRU-fed) fast reactor, and EUFR = enriched-uranium fed fast reactor (see section 4.2 for a more complete description of the reactor decision options).

Table 3-1: FANTSY Core Simulation Input Variables

Input Variable	Var Type	Allowed Values	Description
Nuclear Growth	Scalar	Positive reals	% growth per year of nuclear demand
Reactor Type(s)	String	LWR, TFR, EUFR	Reactors to build
Reactor Fraction	Scalar	[0, 1]	Fraction of FRs to build of those allowed by available amounts of fuel feedstock

The wrapper code iterates through each possible combination of inputs by assigning appropriate values for growth, reactor types, and reactor amounts to each time period between decisions. Once the wrapper specifies a complete set of inputs, it calls the core 100-year simulation function.

The FANTSY core simulation code shares several characteristics with CAFCA, but brings inputs to outputs in a linear fashion. The simulation file begins with initialization, taking the inputs defined above and establishing a 100-year trajectory for nuclear power demand. The initialization also includes designation of the current U.S. reactor fleet and the setting of "fuel recipes." In practice, the isotopic composition of the fuel changes with each pass through a fast reactor. FANTSY, however, follows CAFCA and VISION by setting the isotopic composition of fresh and spent fuels and keeping these constant throughout the simulation; in this case, fuel recipes were taken from VISION (which in turn credits Hoffman for the specification of fast reactor isotopic ratios). Like CAFCA, FANTSY groups the constituents of the fuel into categories (fission products and actinides); CAFCA includes an option for separate tracking of plutonium, but this was deemed unnecessary for FANTSY at this time given that no Pu-separation fuel cycles were evaluated for this work. VISION tracks all isotopes individually and performs decay calculations at various steps throughout the cycle. This level of detail was also deemed unnecessary for this work, owing to the breadth of questions under consideration, but further studies could incorporate the sophistication of VISION.

Following initialization, FANTSY commences a set of commands performed for a single year and loops these commands 100 times. The first step is to create a forecast for reactors needed five years hence. Like CAFCA and VISION, FANTSY assumes “perfect vision,” or that reactor builders have accurate knowledge of what nuclear power demand will be in five years (the time it will take to build the reactors started in the current year). The next step is to fuel currently-operating LWRs with 4.2% enriched uranium fuel, in the process calculating uranium and SWU usage. Enriched uranium-fed fast reactors are fueled next if any are starting up in the current year.

The next step is construction of reprocessing plants. The benchmarking exercise performed to calibrate CAFCA, VISION, DANESS, and COSI (see Appendix B) determined that the amount of available reprocessing capacity has a significant impact on the speed with which traditional fast reactors can be built.(Guerin & et al., 2008) Many of the codes specify limits on how fast reprocessing plants can be built (e.g. that it makes no sense to have 100 plants built per year starting as early as 2020). FANTSY follows these codes, specifying a generous limit on reprocessing plant builds that allows 500MT/year of LWR SNF reprocessing capacity adds until 2050 when this increases to 1000MT/year, and 50MT/year reprocessing capacity for FR spent fuel increasing to 150MT/year in 2065. More stringent limits on additions of reprocessing capacity have the effect of “smoothing” the curve of fast reactor builds, but have little impact on the results in this study (see Chapter 6.5 for the sensitivity analysis on these parameters).

After reprocessing plants are constructed, spent fuel from LWRs and from FRs is reprocessed (if the scenario allows it). The amount of fuel reprocessed is equal to either the total cooled spent fuel in inventory or the maximum reprocessing capacity; whichever is lower. In this way, reprocessing is modeled as a “push” process in the parlance of VISION: as soon as spent fuel is generated, the utility gets rid of it as fast as possible by pushing it to any available reprocessing plants. In practice, this is likely not how the fuel cycle system would operate, because it could lead to situations where separated TRU is stockpiled at reprocessing plants. Nonproliferation experts might prefer a “pull” process where TRU is only separated if it is needed by the FR fleet.

The distinction, however, is not particularly important to this very simplistic model at this stage of analysis (though further work on nonproliferation aspects of scenario comparisons might warrant a second look at the assumption). Modeling the process as “push” makes the yearly fuel cycle calculations, based primarily on TRU availability, much easier. It is also a conservative choice, because the stockpile of TRU from year to year thus represents a “worst-case scenario” for amount of separated TRU. None of the decisions, however, depend on this particular parameter (they depend only on TRU that is designated as waste and intended for disposal), so decision outcomes remain unaffected.

After reprocessing occurs, existing fast reactors receive their fuel from the now-separated TRU stockpile. This is mostly a straightforward calculation. There are scenarios, however, where the amount of separated TRU in the stockpile is not sufficient to meet the fuel demands of operating reactors. CAFCA suffers from the same idiosyncrasy, which is kept in both codes in order to maintain simplicity of the reactor ordering algorithm. A lack of fuel happens most often for EUFR scenarios, when reprocessing capacity builds are not keeping pace with the needs of FR fuel (reactor builds can outpace reprocessing capacity expansion because EUFR startup is not constrained by TRU availability). For most scenarios, a maximum of about 10% of the fuel is “missing.” Were these situations to occur in the actual fuel cycle system, the “missing” fuel could be replaced by enriched uranium, or perhaps by TRU or SNF purchased or taken for free from a foreign country under a nonproliferation-based fuel take-back agreement. There would be a cost for this fuel, although it would likely be very small (especially if the U.S. were taking title to another country’s SNF). Another alternative to keep FRs fueled would be to increase the TRU yield of the FR fleet by increasing the conversion ratio, so that more TRU would be produced per kg of FR SNF reprocessed. A final option involves lifting restrictions on the pace of reprocessing plant builds. Whatever action is taken to ensure reliable refueling for FRs, the cost penalty for missing fuel at a system-wide level is likely to be relatively small and is not likely to impact fuel cycle decisions. The only major impact of fuel availability is in the prediction of how many FRs can be built, and because the missing fuel never “builds up,” the pace of reactor builds still tracks with those of other fuel cycle codes.

The final major step in the yearly fuel cycle simulation involves the building of reactors. Fast reactors are built first, depending on fuel availability (for TFRs) and on the ratio of “allowed” fast reactors that will be built. The number of TFRs allowed is determined by taking the current year’s TRU inventory, multiplying it by 60 (in effect assuming that the TRU available that year will be the same over the lifetime of a reactor built) and dividing this number by the lifetime TRU needs for a single TFR. This equation is somewhat arbitrary, but it is the same rule employed by CAFCA and it leads to the same general build patterns observed in VISION and other codes. LWRs are then built to fulfill the remaining nuclear energy demand.

The remaining "clean-up" steps include first the discharging of spent LWR and FR fuel, placing it in a cooling pool where it will remain for five years before becoming available for reprocessing. Reactors and fuel in cooling pools are then aged for one year before beginning the next year-cycle.

The outputs of the fuel cycle simulation are many. In principle, any of the arrays generated by the core fuel cycle code is accessible within MATLAB and could be considered an output of the simulation. Seventeen important arrays are passed back to the wrapper code, and written into Excel for analysis by decision trees. These outputs are described in Table 3-2; all take the range of positive real numbers.

Table 3-2: FANTSY Outputs

Output	Var Type	Units	Description
NaturalUraniumConsumed	# Array	MT	Nat U used to fuel LWRs/EUFRs by yr
SWU	# Array	kgSWU	SWU required to fuel LWRs/EUFRs by yr
LWRfuelfabricated	# Array	MT	Uranium oxide fuel fabricated by yr
FRfuelfabricated	# Array	MT	FR U/TRU fuel fabricated by year
ReactorsConstructed	# Array	#	LWRs constructed by year
TotalFRsConstructed	# Array	#	FRs (EUFR or FR) constructed by year
ReactorsOperating	# Array	#	LWRs operating by year
FRsOperating	# Array	#	FRs operating by year
SNFreprocessed	# Array	MT	LWR SNF reprocessed by year
FRreprocessed	# Array	MT	FR spent fuel reprocessed by year
ReactorsDecommissioned	# Array	#	LWRs decommissioned by year
FRsDecommissioned	# Array	#	FRs decommissioned by year
SNFcooled	# Array	MT	SNF leaving cooling pool by year
FRcooled	# Array	MT	FR SNF leaving cooling pool by year
SNFstock(100)	Scalar	MT	SNF in cooling + dry storage at 100 years
Sum(FP)	Scalar	MT	Total fission products separated over 100 yrs

Sum(TRUlosses)	Scalar	MT	Total TRU lost to waste over 100 yrs

3.2 FANTSY Benchmark Results

In order to confirm that the results produced by FANTSY are commensurate with those of other fuel cycle codes, various scenarios were simulated and compared to those described in the MIT Center for Advanced Nuclear Energy Systems benchmark study.(Guerin & et al., 2008) A full report of this benchmarking exercise is provided in Appendix B. Overall, FANTSY tracks well with the scenario results of all four codes (and especially well with CAFCA, which provided the primary modeling example).

The only major point of departure tends to center around whether restrictions on the pace of fast reactor builds are implemented. For most of the work in this thesis, the number of fast reactors that can be built in a given year is unrestricted. At particularly high nuclear demand growth this does mean that as many as 90 fast reactors are built in a single year. Though such high numbers are unlikely, especially in the near term, they are consistent in shape with boom-and-bust cycles that could actually exist due to delays within the system. Adding in restrictions on fast reactor builds produces curves that are more consistent with results from the other four codes, but these restrictions in a sense are artificial. Ultimately, whether restrictions are implemented or not, FANTSY winds up building about the same number of reactors as the other fuel cycle codes as long as fast reactor builds are not overly constricted for the entire simulation. Furthermore, the differences are only stark when extremely high growth rates are assumed (see Figure 3-1 and Figure 3-2 for an example). A full sensitivity analysis of FR build restrictions on optimal decisions is presented in section 6.5.

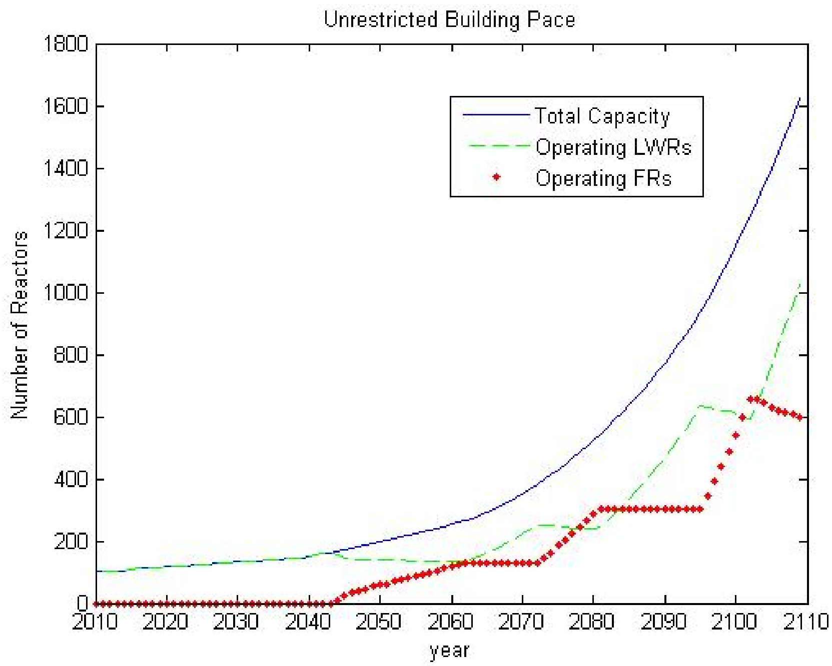

Figure 3-1: FANTSY reactor fleet results when FR build case is unrestricted and growth is high

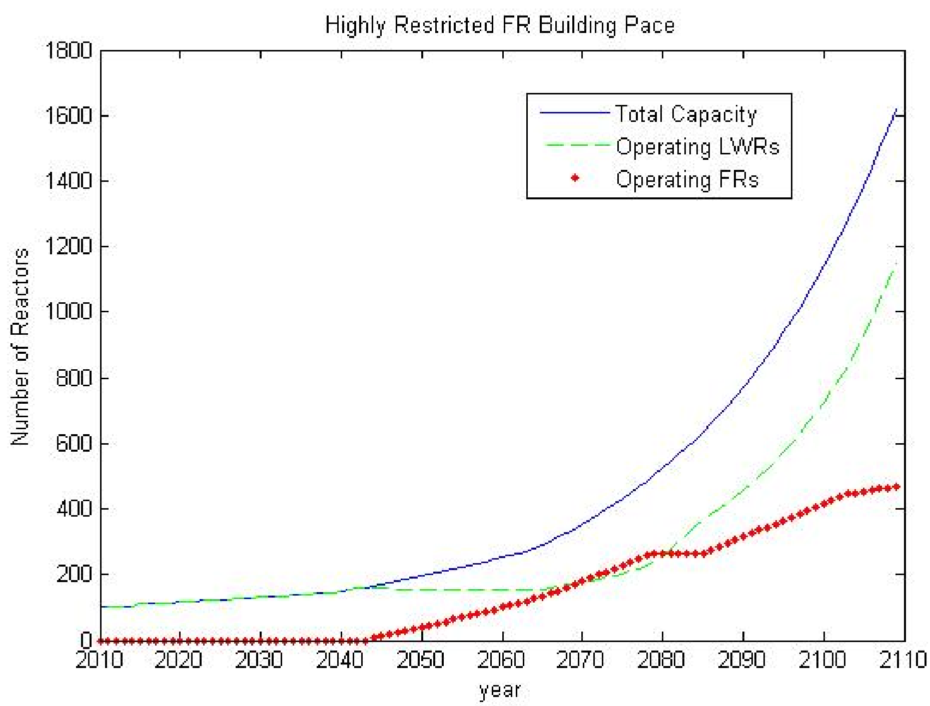

Figure 3-2: FANTSY reactor fleet results when growth is high and the FR building pace is highly restricted (beginning at 2 FRs per year, expanding to 10 FRs per year by the end of the simulation)

Chapter 4: Methodology for a Decision Analytic Model of the U.S. Nuclear Fuel Cycle

The decision tree model of the U.S. nuclear fuel cycle represents the heart of this thesis work. Each tree includes three component types: values, decisions, and uncertainties, which are used to define and compare scenarios for fuel cycle evolution. The structure adds two major types of knowledge to traditional fuel cycle analysis. First, the decision tree allows for clear construction and analysis of flexible scenarios, where decisions to deploy advanced recycle technologies can be re-evaluated later in the century. The second major contribution is explicit designation of metrics by which to compare fuel cycle strategies. Though these metrics can never perfectly capture the goals and preferences of any set of decision makers, they can provide a common language by which to discuss the impact options have on fuel cycle objectives.

Each tree component is discussed in turn. Throughout this thesis, common conventions for designating decisions, uncertainties, and terminal nodes are used. Figure 4-1 shows the basic decision tree structure: blue squares represent decision nodes, green circles represent uncertainty (chance) nodes, for which probabilities represent the likelihood that the scenario will take on a particular value for the uncertain parameter, and red triangles represent terminal nodes. The branches run from left to right, and in this case, each branch represents a separate scenario (which is also a separate run of the basic fuel cycle simulation model). There is, in effect, an invisible time axis along the bottom of the tree. Decisions on the left are made before gaining any knowledge about uncertainties that appear on the right. Decisions represented to the right of uncertainty nodes are made with full knowledge of the previous uncertainty outcome.

The tree begins with a question to be answered by comparison of the branches, and ends with a variable intended to capture the value of each branch (designated in this thesis as "ScenarioValue"). TreeAge software calculates the value of each branch and then rolls back the tree to determine which choice is best at the initial decision node. TreeAge also allows extensive sensitivity analysis of this initial choice on a wide set of parameters represented in the tree.

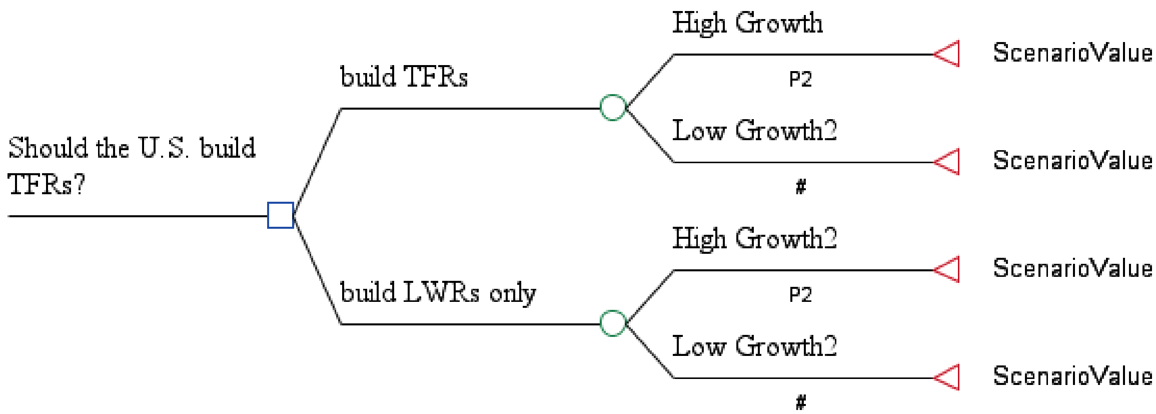

Figure 4-1: Simple illustrative decision tree

4.1 The Value Function

The goals laid out by the DOE Fuel Cycle Research and Development (FCRD) program and other organizations (see section 2.2.1) provide guidance on how to compare and choose fuel cycle evolution pathways. They are not, however, specific or quantitative enough to serve as true metrics for comparison. For example, one goal of the FCRD program involves the minimization of waste generation. But what qualifies as waste? Is used nuclear fuel a waste, or a resource? Should fuel cycles minimize waste mass, volume, heat, or radiotoxicity?

Because no definitive resource specifies the metrics by which fuel cycles should be judged, fuel cycle metrics have been defined differently across a range of studies. Ideally, for a comprehensive and policy-focused study like the one conducted here, an analyst would gather fuel cycle experts and stakeholders and conduct an elicitation procedure to determine important fuel cycle goals, overarching values, and metrics by which fuel cycle impacts can be measured. An extensive elicitation process is beyond the scope of this thesis. Instead, several plausible metrics are constructed and evaluated for a range of scenarios in order to determine the sensitivity of decision results to the value function.

A common metric used to evaluate options in complex engineering and finance projects is Net Present Value (NPV).(de Neufville, 2010) This number adequately captures the costs of making investments, and is incorporated here as a “net present cost” metric. NPV, however, does not fully account for the range of fuel cycle system objectives defined above. Indeed, as Bob Budnitz noted at the March 31, 2011 CANES conference on Nuclear Power in 2050, if the U.S.

decides to move to an advanced fuel cycle system, it will not be because that system is cheaper than the one currently in place. This implies that there is some value in closing or advancing the fuel cycle (tied to waste management, security, and/or sustainability) unrelated to cost.

The value function for evaluating fuel cycle scenarios must therefore account for multiple objectives, which can be done using a multi-attribute utility function. Ideally, each fuel cycle goal would be translated into a metric, which could then be compared to others in an “apples-to-apples” fashion through a utility function assessed across the range of values the metric takes. These utility values would then be traded off using preference weighting. Again, the full elicitation required with human decision makers or actors is beyond the scope of this work, and the process is simplified.

One example of a “full” set of values and metrics, that could be similar to one that fuel cycle policy makers would construct, is presented in Figure 4-2. This is by no means the best or most comprehensive value tree imaginable; for example, one objective that has been left out includes the desire for security of fuel supply. Rather, it has been designed to incorporate the values evinced in the AFCI and FCRD program goals, in addition to some other potentially important objectives (like minimizing the use of land). A multi-attribute utility function that contained all of the metrics and goals in Figure 4-2 would be large and somewhat unruly. Nevertheless, working with such a function might ultimately be necessary in order to understand all of the important tradeoffs between fuel cycle options and values.

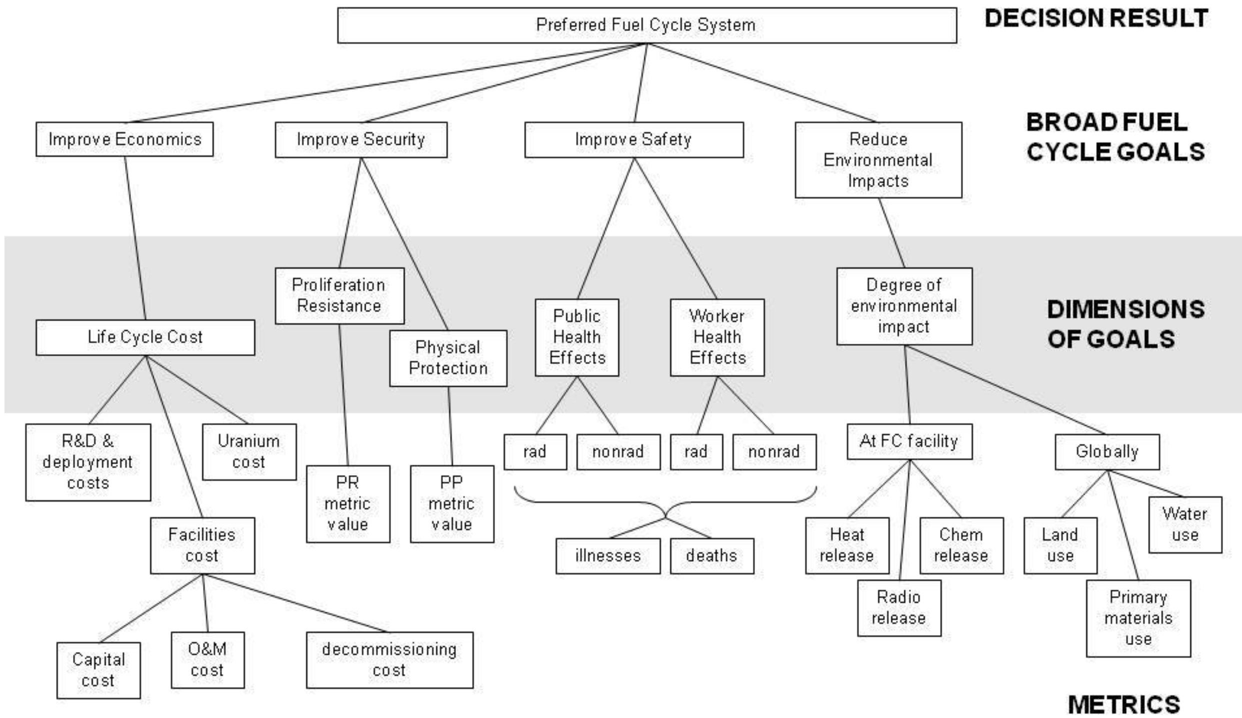

Figure 4-2: Example value tree for nuclear fuel cycle decision-making

For this study, a much simpler set of values and attendant value functions are employed. The methodology is more easily demonstrated with a small number of metrics, but the framework is intended to be flexible enough to handle other or more values and value function types according to the wishes of particular decision makers.

The two metrics chosen for study throughout the remaining chapters relate to system costs and amounts of waste generated. Cost and waste were chosen because they come up often in arguments for and against advanced fuel cycles (proponents of closing the fuel cycle argue that recycling reduces the waste management burden, while those opposed to recycling argue that it is too costly). The two metrics trade off: paying more for an advanced fuel cycle may allow the U.S. to reap a benefit in terms of waste management.

One major challenge in the studies that follow is a lack of complete knowledge about what the costs and waste management impacts *are* for the technologies under consideration. This problem is not unique to the nuclear fuel cycle: it applies to any system that, like the fuel cycle, requires heavy investment along a particular pathway before all economic and technological uncertainties are resolved. In this thesis, the cost uncertainty is dealt with explicitly and probabilistically (see section 4.2). This is the recommended way to treat uncertainties that have a strong impact on scenario outcomes. For waste, one aspect of the uncertainty is treated explicitly (one sensitivity analysis investigates the impact of separations efficiency on the set of desirable decisions; see section 6.2). Several other unknowns related to waste management are treated using best estimates. These include, for example, the assumption (explained in Chapter 3) of how much TRU is produced during each FR cycle. Significantly more research is needed on recycling systems and reactors in order to reduce the uncertainties in this realm, and projects sponsored by the FCRD will help in the coming years. For example, more detailed neutronics/reactor studies would clarify the extent to which fast reactor recycling results in actinide destruction.

In the following sections, the components of the value function are explained more fully. The sections discuss the cost metric, the waste metric, the disutility functions used to compare cost to waste, and the weights used to explore tradeoffs between the metrics.

4.1.1 The Cost Metric

The cost metric is calculated as the total system cost of the nuclear fuel cycle, including the reactor capital cost. Outputs (number of reactors constructed, SNF reprocessed, etc.) are written directly from the FANTSY code to an Excel spreadsheet, and a cost model immediately calculates the per-year costs. These yearly costs are then discounted back to the present year and added together to get the total system cost of a given scenario. Table 4-1 shows the cost assumptions used in the model. Note that the waste fee and other waste-related costs are not considered in the model; this is done in order to make the cost and waste metrics as independent as possible. The financial cost associated with waste is accounted for by capturing the extent to which waste is "disliked."

Table 4-1: System Nuclear Fuel Cycle Costs

costs		amount	unit	source
natural uranium		80000	$/tonneHM	(De Roo & Parsons, 2009)
conversion		10000	$/tonneHM	(De Roo & Parsons, 2009)
SWU		160	$/SWU	(De Roo & Parsons, 2009)
fabrication		250000	$/tonneHM	(De Roo & Parsons, 2009)
capital		4645	$/KWe	CAFCA

fixed O&M		56000000	$/Gwe/year	(J. Deutch & et al., 2009)
variable O&M		0.00042	$/kWh	(J. Deutch & et al., 2009)(Bunn, Holdren, Fetter, & van der Zwaan, 2005)(Bunn, Holdren, Fetter, & van der Zwaan, 2005)(Bunn, Holdren, Fetter, & van der Zwaan, 2005)(Bunn, Holdren, Fetter, & van der Zwaan, 2005)(Bunn, Holdren, Fetter, & van der Zwaan, 2005)(Bunn et al. 209-230)
decommissioning		700000000	$/plant	(De Roo & Parsons, 2009)
incremental cap		40000000	$/GW/yr	(De Roo & Parsons, 2009)
waste fee		0.001	$/kWh	IGNORED
dry cask storage**		200000	$/tHM	(Bunn et al., 2005)
FR Capital premium		20%	%	(De Roo & Parsons, 2009)
FR O&M premium		20%	%	(De Roo & Parsons, 2009)
UO2 reproc		1600000	$/MT	(De Roo & Parsons, 2009)
FR reproc		3200000	$/MT	(De Roo & Parsons, 2009)
FR fuel fab		2400000	$/MTHM	(De Roo & Parsons, 2009)
FR fuel managed storage		20%	on 200/kg	guess
FR decommissioning premium		20%	%	guess
annually compounded discount		7%		(De Roo & Parsons, 2009)

In addition to the per-year nuclear system costs, a "cost liability" term is added for each scenario. This liability accounts for the commitment of nuclear power plants to continue operating for their full 60-year lifetime, past the rather artificial end of the 100-year simulation. The cost liability is calculated by determining the total number of reactor-years remaining for the nuclear fleet in place at 2110, and multiplying this number by the rough O&M, fuel, and SNF management costs each plant type incurs each year. This value is discounted to 120 years from the present (2120), and added to the total decommissioning costs of the reactors discounted to 150 years. The values and discount times are somewhat arbitrary, but are intended to give a very rough estimate of the financial commitment in place at the end of the simulation. Discounting over the long time horizon of the simulation means that adding in this cost liability has only a very small effect, but it is included for completeness.

An alternative cost metric defined as the per-kWh levelized cost of electricity was considered in the early stages of this study. Though it has some advantages, including its ease in communicating the differences between scenarios as cents per kWh (vs. tens of billions of dollars in system costs), the levelized cost was ultimately not pursued as a metric. One main reason is that the levelized cost in effect would change rapidly throughout each individual simulation, owing to the change in the ratio of advanced to traditional reactors. Creating an economically sound and useful metric proved extremely difficult, and initial attempts in this direction did not produce results significantly different from those obtained with the total cost metric. An exploration of an analogous definition of the waste liability is explored in section 6.5.

4.1.2 The Waste Metric

The waste metric has three components: spent (or used) nuclear fuel (SNF), transuranic high level waste (TRU), and fission products (FP). FANTSY tracks the mass of each of these produced throughout the simulation, and they are combined to form the waste metric. The amount of SNF for a scenario is defined as the amount in dry storage (after 5 years cooling) present at the end of the 100-year simulation. For TRU and FP, the amount of each is accumulated throughout the simulation due to separations activity (losses producing TRU to waste, and FP explicitly separated as a waste product). For the systems considered here, the FP stream is in fact commingled with the TRU waste. The total amounts produced over 100 years are reported as the TRU/FP metrics. The apparatus to track them both separately remains throughout the simulation, but in fact they are one stream whose management characteristics are largely dominated by the presence of TRU.

A common criticism leveled at waste metrics based on mass (or volume) is that the waste management burden, or the difficulty of nuclear waste disposal, rarely depends at all on the mass of the waste. Instead, repository loading is generally constrained by the heat output of the waste and sometimes on waste radiotoxicity. This is an important point that would render a purely mass-based waste metric useless if significantly different reactor technologies were compared. For the studies presented here, however, only one basic fast reactor technology is compared to LWRs. This means that the particular makeup of the SNF is constant across scenarios (equal to that of a once-through LWR) and the makeup of separated TRU and FP is also constant (equal to that for a self-sustaining sodium-cooled fast reactor). In turn, this means that the SNF, TRU, and FP terms expressed in terms of heat or radiotoxicity would simply be a constant multiplied by the mass, and mass is thus a sufficient metric.

Combining the three waste terms is done in the traditional multi-attribute way: weighting reflects the relative difficulty in managing each (note that the weights therefore take the place of a heat/radiotoxicity constant that differentiates the three waste types). The combined waste metric for a scenario i is expressed as:

$$WasteMetric_i = wSNF * SNF_i + wTRU * TRU_i + wFP * FP_i$$

where wX is the weight for variable X, and X is the mass of the waste type in question (actually a utility representation of the mass, explained in section 4.1.3), and wSNF + wTRU + wFP = 1.

As a check, two other versions of the waste metric were defined and tested for most of the study vignettes presented in Chapter 5. These included a repository-based version of the metric, for which the value of waste depended on the number of repositories needed to contain the SNF (assuming a 100,000 legal MTHM limit) plus the number of boreholes required to contain the TRU and FP (assuming 6.5 MTHM of TRU per borehole, with corresponding fission products added in because they are not heat-limiting). A third version of the metric involved multiplying each waste type by the amount of heat it produced at a given time after discharge. As discussed above, this amounted to multiplying already-weighted terms by constants. Replacing the mass metric with these did not significantly impact decision results for any of the questions analyzed; see the sensitivity analysis of section 6.1.

4.1.3 Utility Functions

One major challenge with multiple attributes is finding a way to compare them consistently. Utility theory offers one approach to this problem. By eliciting a decision maker's utility for one unit of cost vs. one unit of waste, the two metrics are easily compared. This elicitation process, however, is involved and requires time and committed participants. Instead of eliciting preferences, several potential versions of utility functions are defined here and analyzed to see how they impact the sets of desirable options. Because greater cost and larger amounts of waste are undesirable, the utility functions defined below are actually *disutility* functions, representing, as they get larger, greater and greater disutility.

One utility function is applied to the cost metric. The function maps cost numbers onto the range [0,1], by assigning the highest cost scenario in a given tree to a utility of 1, the lowest cost to a utility of zero, and the other costs to corresponding values in between. The utility function for the cost of scenario i is thus:

$$U_c(cost_i) = \frac{cost_i - cost_{min}}{cost_{max} - cost_{min}}$$

where $cost_i$ is the total system cost for scenario i, $cost_{min}$ is the minimum of all scenario costs, and $cost_{max}$ is the maximum of all scenario costs.

Two different types of utility functions are associated with the waste metric. The first mimics the cost utility function above; waste values are mapped onto [0,1] and

$$U_w(waste_i) = \frac{waste_i - waste_{min}}{waste_{max} - waste_{min}}$$

The utility function is applied to each of the waste types (SNF, TRU, FP) individually before they are weighted.

The second type of utility function is a diminishing returns function. The first function assumes linearity of decision makers' disutility, because a single unit more of waste is equally undesirable whether it is added on to a small or large stockpile. This need not be the case for nuclear waste. For example, having one MTHM more than 100,000 MTHM, if a repository limit is set at 100,000 MTHM, would be much more problematic than one MTHM more than 95,000 MTHM. As long as the next MTHM can be put in an already-open repository, it is not very threatening. But if another unit of SNF means opening a new repository, there is a significant disutility associated with that unit. Indeed, this could imply a step function, where each step increases the disutility once a repository is filled. A step function was explored, but it did not provide qualitatively different results from a smooth diminishing returns function. The disutility function shape is demonstrated in Figure 4-3 for the waste masses after they are multiplied by a heat value.

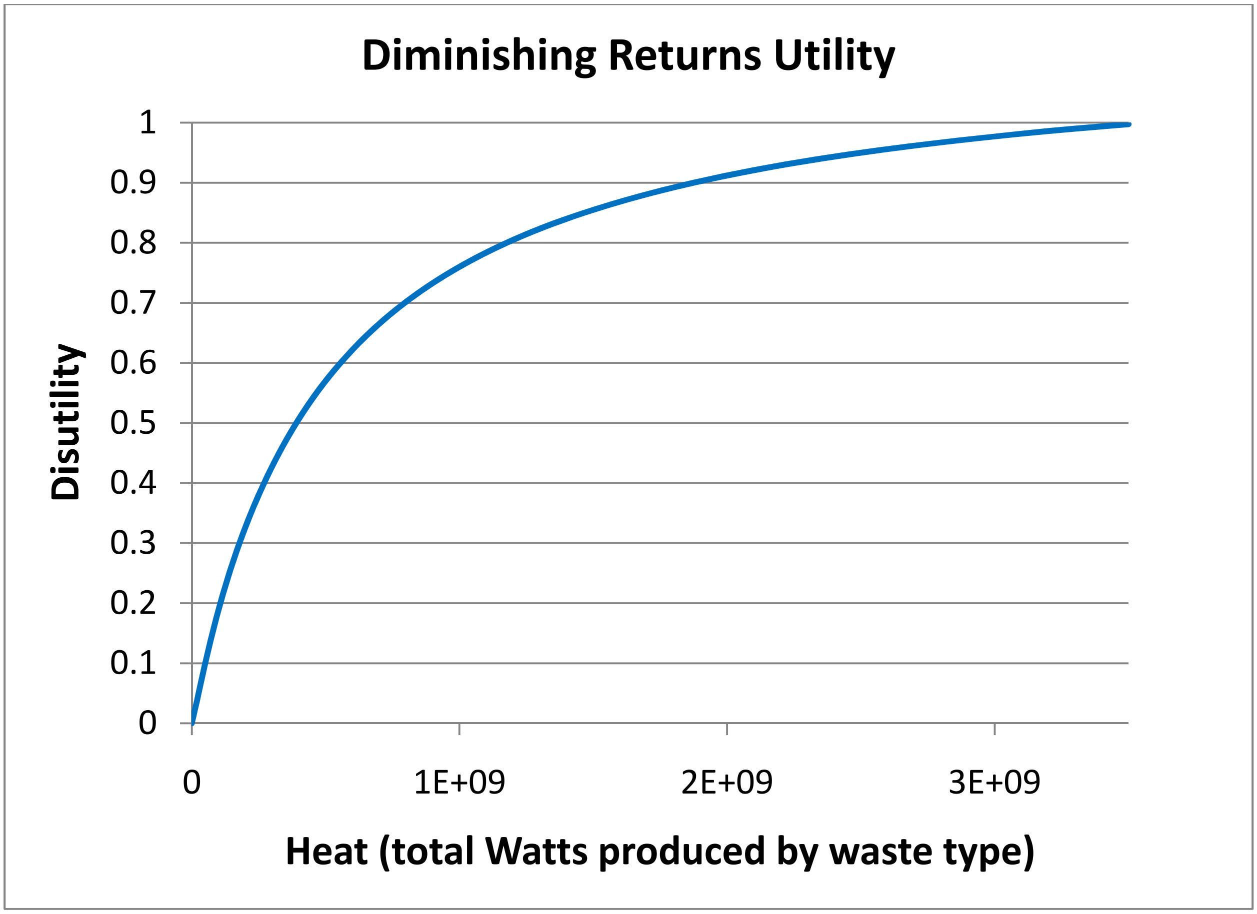

Figure 4-3: Diminishing returns disutility function for waste

The equation used for the diminishing returns disutility function in Figure 4-3 (which was also used to test the heat-multiplied waste metric) is:

$$U_w(waste_i) = \left(1 - \frac{1}{\frac{waste_i}{5E08} + 1}\right) * 1.14$$

where each waste type utility is calculated individually. Note that for some values of waste heat input to the function, the utility could be greater than one; these values were not observed in the course of this analysis, so the function is appropriate for the relevant ranges.

It is important to note here that the waste metrics and utility functions assume a material benefit to decreasing the amount (either in terms of mass or heat) of spent LWR fuel. As mentioned at the beginning of the chapter, even the most well-understood reprocessing and transmutation systems would benefit from more research in order to quantify these benefits. The extent of the benefit, furthermore, depends heavily on the characteristics of the waste repository, and the repository location for U.S. commercial waste is currently unknown. If it was determined that the benefit of reducing waste is negligible, the decision would be easy: the inexpensive and understood once-through fuel cycle would handily dominate a more expensive advanced cycle, unless attributes other than waste and cost were considered. This decision space could very well be the best representation of attitudes. If it is not, however, the complexity of the solution is increased and requires more study: this thesis therefore focuses on the realm where a serious waste benefit to an advanced fuel cycle trades off with the lower cost of the once through cycle.

4.1.4 Attribute Weights

The final piece of the value function is the weighting used to trade off the primary attributes. The total, final value for scenario i is thus calculated as:

$$ScenarioValue_i = wC * U_c(cost_i) + wW * [wSNF * U_w(SNF_i) + wTRU * U_w(TRU_i) + wFP * U_w(FP_i)]$$

where wC + wW = 1, and the utility functions and metrics are defined as above.

In order to get a definitive “answer” about which scenario is the best, an analyst would need to fix all weights in the equation. As with the function structure, the analyst would obtain these weighting values through a decision maker elicitation process. For this work, results are instead reported for the entire range of values for wC and wW. wC is varied from zero to one, and wW is calculated correspondingly. Similarly, several values of the waste weighting are explored (i.e. the relative magnitudes of wSNF, wTRU, and wFP are varied, under the condition that they always sum to one). This allows the reader to see how robust certain decision pathways are to a range of preferences.

For problems of this type, with many different disparate groups of decision makers, it may never be feasible to do a comprehensive elicitation process to obtain a definitive set of weights (let alone a definitive form for the value function). The results and following discussion show that valuable information is gleaned by showing answers across the whole weighting spectrum. With this information, decision makers may be able to choose options that are robust to a range of values, making the best choice given that stakeholder preferences are highly varied and are likely to change over time.

The possibility exists that decision results will be highly sensitive not only to the weighting scheme, but also to the structure of the value function (including the forms of the metrics and utility functions). For these studies, it turned out that the results were far more sensitive to the weights than to the metric structures or utility function forms. A sensitivity analysis is presented in section 6.3, but in reality, every study in Chapter 5 was checked with all four forms of the value function. None were significantly sensitive to the value function structure. The results in Chapter 5 are thus presented using the first functional form only: the metrics are dollars and mass of waste, and utilities are composed of a linear mapping of the metrics onto the range [0,1].

4.2 Decisions

The purpose of the decision tree is to determine which options, or initial decision paths, are desirable. It is thus critical to appropriately frame available decisions and represent them in the tree structure, so that they can then be compared according to the value functions defined in section 4.1.

As discussed in Chapter 2, this thesis will initially follow other fuel cycle reports by framing decisions from the perspective of a "benevolent dictator," who has full control over implementing any desired technology in the system. This decision analysis assumption will be relaxed later (see Chapter 7). The question asked by this benevolent dictator is "What fuel cycle should the U.S. employ?" This requires an analysis of technological options, including the reactor type(s) in use by the nuclear fleet and the components of the fuel to reprocess. At the highest level, one could ask whether to maintain the status quo (the once-through fuel cycle) or move to a closed fuel cycle. There are many other layers of choices that must ultimately be answered in order to employ an advanced fuel cycle, such as: what type of fuel (oxide or metal) should be used, what type of reprocessing (pyroprocessing, aqueous, or other) should be implemented, etc. Some of the broader classes of options are delineated in Figure 4-4, which is a hierarchy of decisions rather than a decision tree.

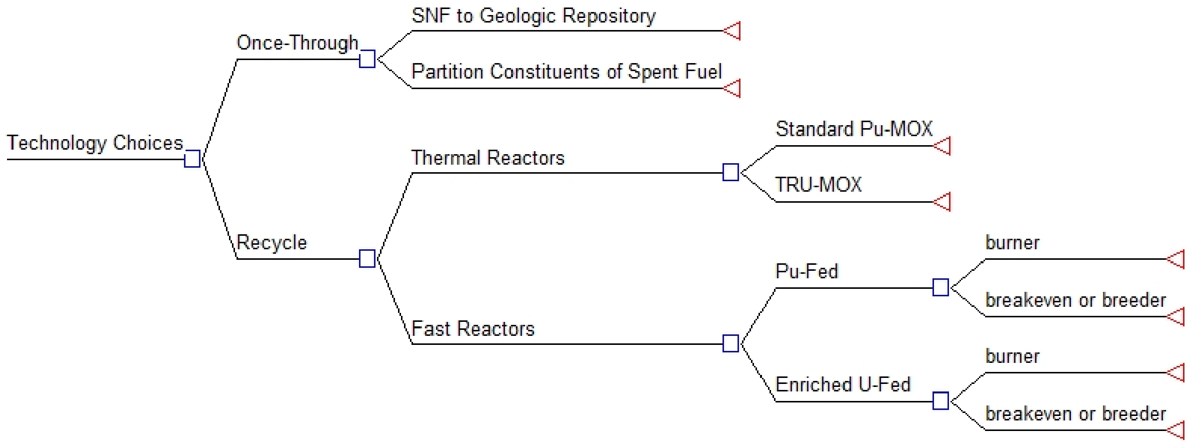

Figure 4-4: Technology options for an advanced nuclear fuel cycle

As mentioned above, the highest-level decision involves whether to close the fuel cycle. It is impossible, however, to understand which of the two cycles is more desirable according to the value function without specifying a little more about what each system would entail. The next level of options considers the type of waste disposal (for once-through) or the reactor class in which recycling will occur. At the next level, the type of fissile material used in the fuel cycle is important to calculate the waste and cost metrics. One must also consider the reactor conversion ratio; though it may be flexible for a given reactor type, the conversion ratio has a substantial effect on fuel cycle outcomes.

Choices between the options in Figure 4-4 will strongly affect waste management outcomes, costs, uranium usage, and security outcomes. Because these outcomes are directly tied to fuel cycle goals (see section 2.2.1) and to the value function, this level of decision is appropriate. By contrast, decisions about fuel type etc. can be considered second-order for the purposes of this analysis, which seeks to advise high-level policy decisions. This may not remain true, however, as innovation proceeds in the future; for example, if a new separations technology were introduced that virtually eliminated waste losses and substantially reduced recycling costs, analysts might need to consider the new technology at this high level since it would impact the important fuel cycle outcomes.

Ultimately, decision makers will want to analyze and understand all of the technological options defined above. This thesis, which seeks to demonstrate the methodology appropriate for making decisions on these types of complex systems, focuses on a small subset of the options.

The options compared in the scenarios below include the status quo (LWRs operated in once-through mode); traditional, self-sustaining fast reactors using spent LWR TRU as their initial batch of fissile material; and the same type of fast reactor started with enriched uranium rather than LWR TRU. The fast reactor model used is a scaled-up version of Hoffman's sodium-cooled fast reactor (see Chapter 2 for a description of Hoffman's study).

Fast reactors were chosen over thermal recycling reactors because the fast reactor fuel cycle involves greater system changes. Recycling in thermal reactors does not entail significantly different reactor costs from the status quo, and does not provide actinide reduction for the sake of waste management to the extent that fast reactors can.(Piet, Matthern, Jacobson, Laws, & Cadwallader, 2006) The self-sustaining fast reactor was chosen because it allows for significant introduction of fast reactors (traditional fast reactors with low conversion ratios are constrained by the availability of LWR TRU) and it has a better chance than low-conversion reactors of meeting uranium resource sustainability targets in the distant future. High conversion ratio reactors may eventually be needed in order to address the resource issue, but in the nearer term, self-sustaining reactors are far more enticing for their benefits in terms of minimized recycle throughput and more options for reactor design.(Kazimi et al., 2011)

4.3 Uncertainties

Modeling uncertainties adds important richness to the scenarios considered. Decisions for complex systems like the nuclear fuel cycle are never made with complete knowledge of the future. A huge number of uncertainties may be important to the nuclear fuel cycle system, including the level of nuclear power demand, the cost of advanced fuel cycle technologies, the timing and availability of advanced technologies, public opinion of nuclear power and of waste management, the existence or not of a carbon price, further nuclear catastrophes, etc. Some of these uncertainties can be grouped. For example, if a carbon price is enacted, that will increase demand for nuclear power; so both of these uncertainties can be modeled by varying a single proxy parameter for nuclear demand. For this work, nuclear power growth and advanced technology cost are the two core uncertainties investigated. Follow-on studies should consider more sources of uncertainty.

Nuclear power growth is modeled in FANTSY as a percent increase in demand per year. For most of the studies described in Chapter 5, “low” growth is considered to be an increase of 0.05% per year, while “high” growth includes 2.5% growth in period 1 (2040-2065) and 4% growth in period 2 (2065-2110). For most scenarios, growth is tacked at 1.2% per year for the period in which no decisions are made (2010-2040); this number was chosen as a median between the MIT Future of the Nuclear Fuel Cycle projections of 1-1.5% for that period.(Kazimi et al., 2011) The high and low growth values are arbitrarily chosen to cover a big, bounding range. It is possible that growth will be lower than 0.05% per year (see section 5.1), but cases where there is zero or declining growth in nuclear power are not interesting to study because fast reactors would not be needed. 4% is in fact an absurdly high growth number, and continued nuclear power growth on that trajectory would lead to nuclear supplying 100% of U.S. electricity (and its growth rate would eventually outstrip that of total electricity demand). This is impossible in practice, but again allows for investigation of a bounding scenario. More exotic versions of growth uncertainty are examined in section 5.8.

Cost uncertainty modeling focuses solely on reactor costs. Because reactor costs are by far the highest of any fuel cycle system component (De Roo & Parsons, 2009), adjusting them is a sufficient way of understanding the differences in fuel cycle costs. Fast reactor capital costs are expressed as a premium over LWR costs: low-cost FRs have a premium of 5%, whereas high-cost FRs have a premium of 55%. As with nuclear power demand, these values were chosen to span a very wide range in order to clearly illustrate differences in desirable decision paths. Though LWR costs are also uncertain, especially at this time in nuclear history, their cost uncertainty is not modeled. One tree analysis was completed with uncertain LWR costs, but this confirmed that what matters for this formulation of the problem is in fact the *relative* cost of the two technological options rather than the absolute cost.

4.4 Linking the Fuel Cycle and Decision Analysis Models

Choices on which decisions and uncertainties to model are made first by the analyst, and then implemented in a decision tree structure. As described in Chapter 1, each branch through the decision tree represents a single scenario. For each scenario, uncertainty and decision values are adopted at each node (i.e. at each time period). These scenario groupings are then fed into the fuel cycle model as inputs, one scenario grouping per 100-year run of the model.

The outputs from the fuel cycle model consist of waste and construction data for each 100-year simulation. These results are fed into an excel model, which calculates the scenario value for each run, and tabulates the set of tree scenario values for the decision analysis software. The use of decision trees both for setting up the problem and for calculating decision results, as well as interaction between the trees and the fuel cycle model, is depicted in Figure 4-5.

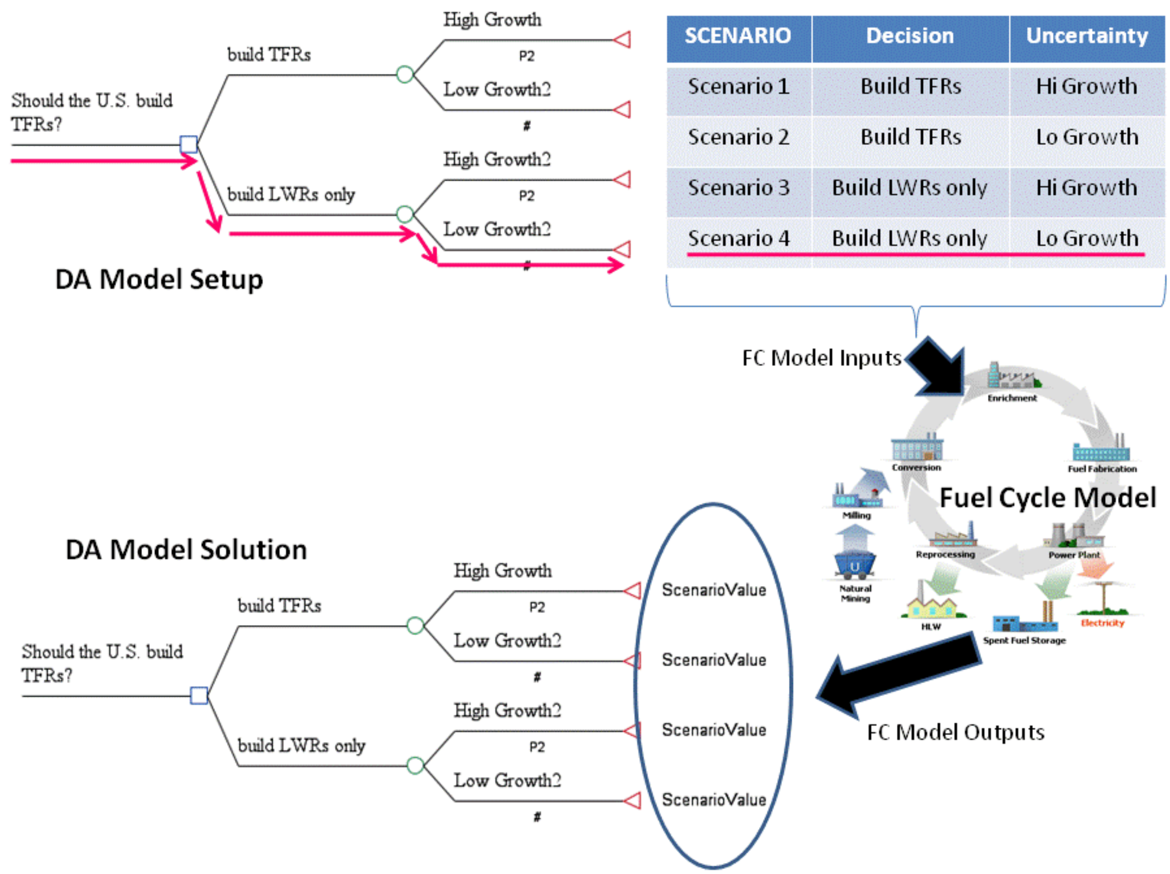

Figure 4-5: Linking of fuel cycle simulation with decision tree

Chapter 5: Studies on Desirable Fuel Cycle Evolution Alternatives under Uncertainty

The decision tree framework described in Chapter 4 is designed to be highly flexible, in order to address a range of systems, options, and value structures. It has the capacity to answer a range of questions about the possibilities for system evolution. In this chapter, ten “vignettes” will demonstrate the flexibility and utility of the methodology by exploring possibilities for improving the U.S. nuclear fuel cycle. Each mini-study fundamentally improves on current system perspectives, by altering the set of strategic alternatives or adding uncertainty explicitly into fuel cycle analysis. These more sophisticated alternatives recognize a need for flexibility, and so consider options as phased and changing deployments of advanced fuel cycle systems for which impacts depend on uncertain outcomes.

5.1 One Period Analysis

The simplest version of the decision tree includes a single period with a single decision and one source of uncertainty. This is already more sophisticated than most fuel cycle scenarios analyzed, because it recognizes that we may not be correct about nuclear power growth after we make a decision about evolving the nuclear fuel cycle. Figure 5-1 shows the decision tree representing the simple case.

For the first seven studies in this chapter, the initial decision node always comes at 2040, reflecting the assumption that we will not be able to begin operating an advanced fuel cycle until then. As explained in Chapter 4, nuclear power grows at 1.2% from 2010-2040. At that point, the decision maker implements either as many TFRs or as many EUFRs as possible (for TFR scenarios, the remaining nuclear energy demand is met by LWRs), or chooses to continue building LWRs only. This choice is made without any knowledge of the eventual nuclear power growth rate. Immediately after the choice is made, nuclear power takes on a growth rate of 2.5% per year (high growth) or 0.5% per year (low growth). P1 represents the probability that growth will be high.

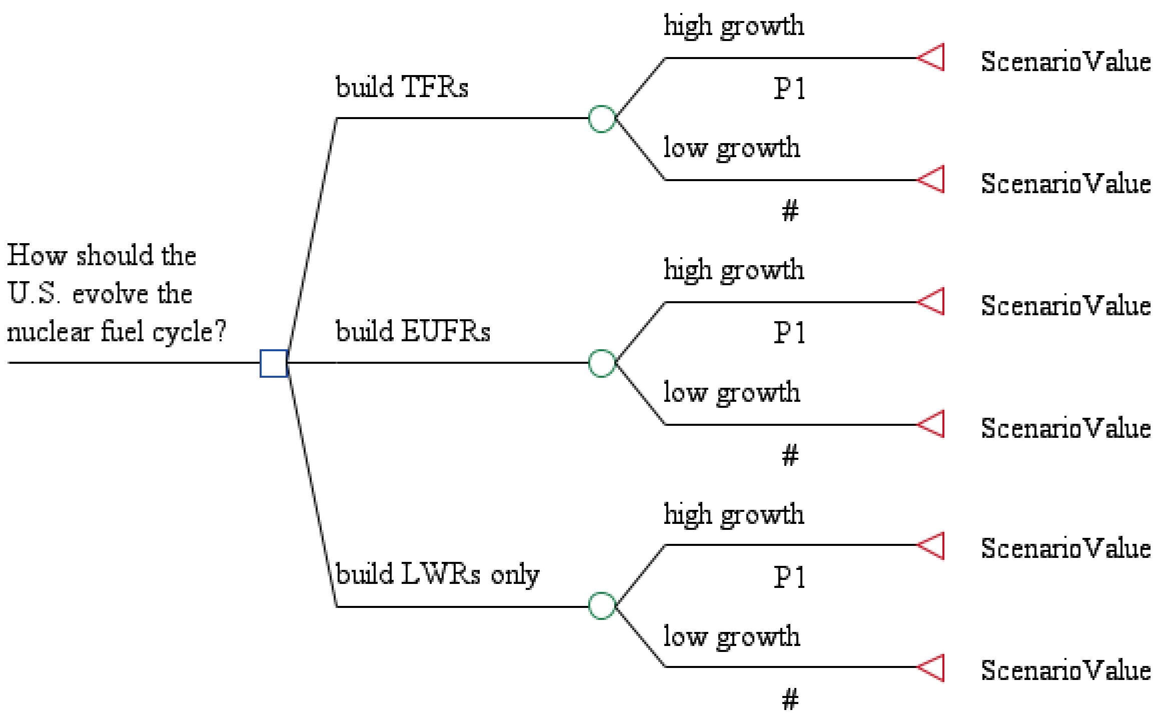

Figure 5-1: Simple one-period tree

Once values are chosen for the parameter P1 and for the weights that comprise the value function, the tree can be "rolled back" and the best decision calculated, given uncertain outcomes and preferences. Table 5-1 shows one possible weighting scheme (recall the requirement that wC + wW = 1 and wFP + wTRU + wSNF = 1). This particular weighting scheme reflects a high dislike for waste, and especially for SNF. FP is considered easier to manage than TRU.

Table 5-1: Sample weights for simple tree

Weight	Value
wC	0.12
wW	0.88
wFP	0.05
wTRU	0.15
wSNF	0.80

As mentioned in Chapter 4, the assumption holds here that an advanced fuel cycle will provide a material waste benefit. If reprocessing will not provide a waste benefit, or if the public or decision makers do not care about the waste burden, there is no point in considering a cycle other than the relatively inexpensive, better-known once through fuel cycle. The interesting decision space lies where employing advanced recycling is desirable from the perspective of waste management, and where reducing the waste burden is important to society. The following discussions concentrate on this regime where there is a difficult decision to be made between two options.

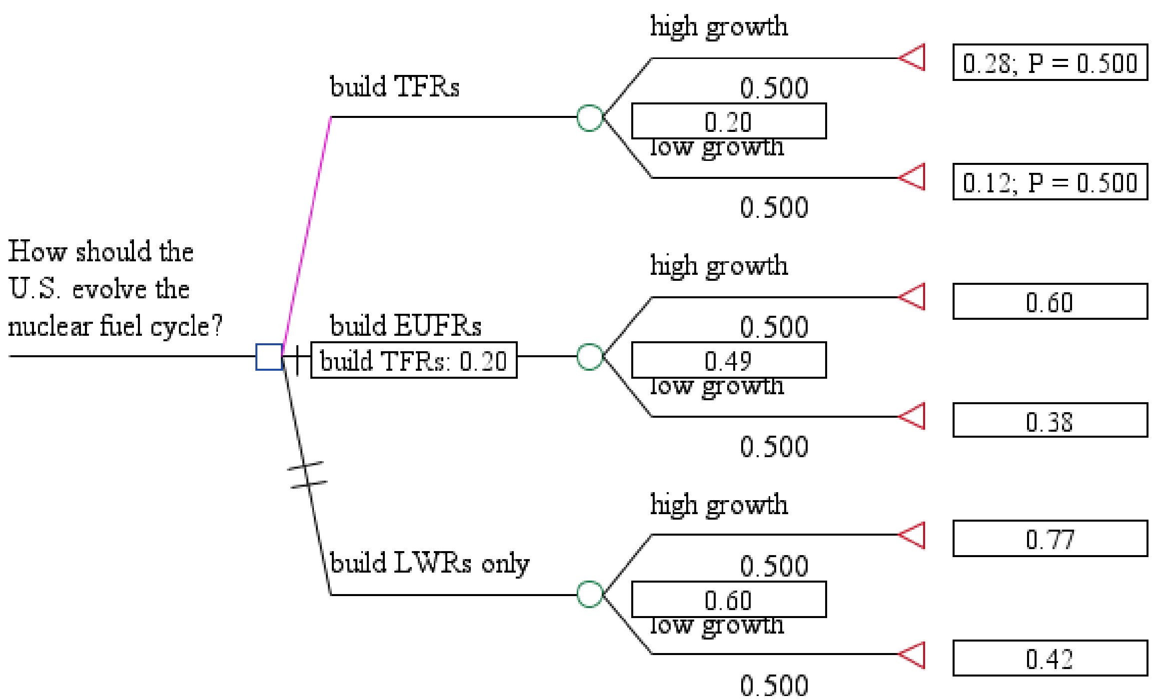

Figure 5-2: Simple tree rollback

Figure 5-2 shows the tree rolled back for the weights in Table 5-1. The probability of high nuclear growth has been set to 0.5. TreeAge software calculates the scenario value for each branch according to the value function, and displays it to the right of each terminal node. The best path is the one with the lowest value (because the value function calculates disutilities rather than utilities). Choosing to “build TFRs” entails a 50-50 chance of getting a value of 0.28 or 0.12. The expected value, 0.2, is the lowest of any branch, so the best choice is to build TFRs. The choice determined “best” can change if we make different assumptions about the probability of growth or the weights for the elements in the value function.

Figure 5-3 shows the result of a sensitivity analysis simultaneously over P1, the probability of high growth, and the cost weight. The color on the graph corresponds to the optimal decision (see legend) given the particular combination of cost weight and probability of high growth.

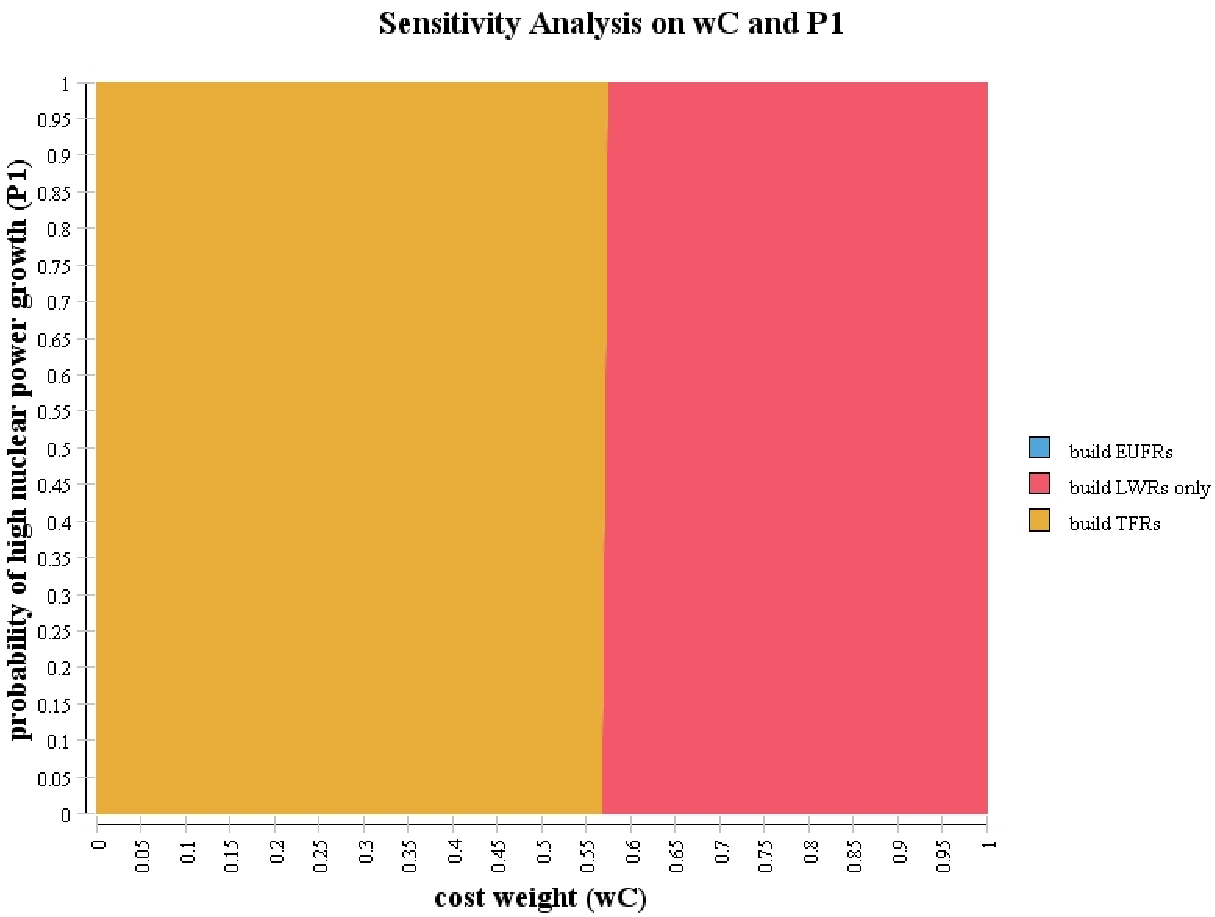

Figure 5-3: Sensitivity analysis for simple tree over wC and P1

Several conclusions can be drawn from the graph. The first is that only two decisions appear: TFRs are optimal (as expected) if the cost weight is low and waste weight is high, whereas with weights reversed, LWRs are optimal. EUFRs never appear because they are effectively “middled-out.” EUFR scenarios wind up costing slightly less than TFR scenarios for these growth values, owing to the early need for TFRs to reprocess lots of SNF, but EUFR scenarios fare significantly worse than TFR scenarios on SNF buildup. The second observation is that the decision is clearly more sensitive to the cost weight than to the probability of nuclear power growth: for any value of growth, there is a point at which the cost weight could change a decision makers’ mind, but the same is not true for a given value of the cost weight.

Increasing wTRU to 0.5 has the expected effect of increasing the space where LWRs are the favored alternative (see Figure 5-4). The increased TRU weight indicates that the decision maker would prefer handling SNF to TRU, so LWR scenarios (where no TRU is separated) are more optimal. These results follow intuition, and provide a baseline for the studies that follow.

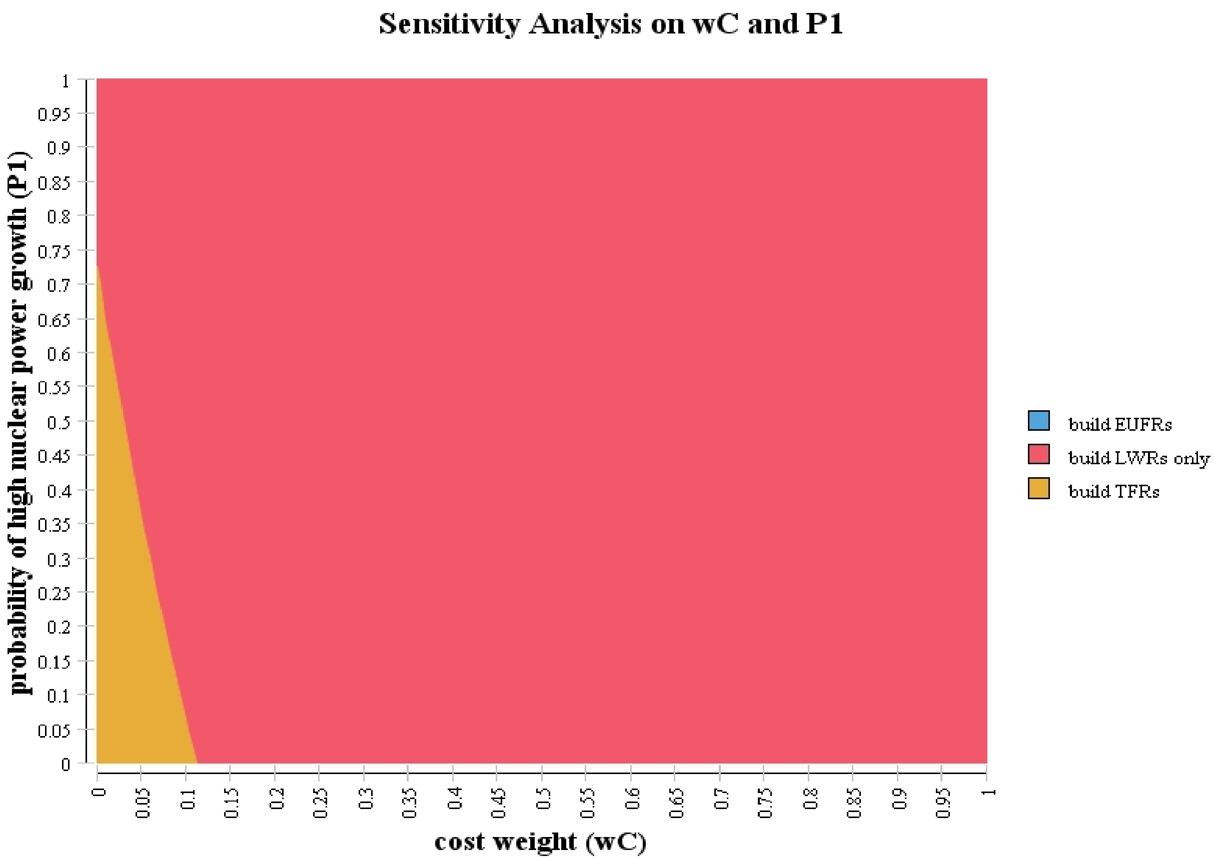

Figure 5-4: Sensitivity on cost weight and growth probability with high wTRU

5.1 Key Takeaways from the One-Period Analysis

The one-period analysis provides a baseline for the analyses of more complex, changeable decisions that follow. The analysis shows that traditional fast reactors are desirable if mitigating LWR SNF is preferable both to decreasing costs or avoiding stocks to separated TRU. We also learn that the decision is more sensitive to preference values than to the probability of nuclear power growth.

5.2 Two Period Analysis

The one-period representation of the fuel cycle decision is highly unrealistic. Decision makers in fact will have opportunities later in the century to re-evaluate the evolution of the fuel cycle, and to change course depending on the outcomes of various uncertainties. A two-period decision tree represents a more realistic option structure that includes the possibility of changing course.

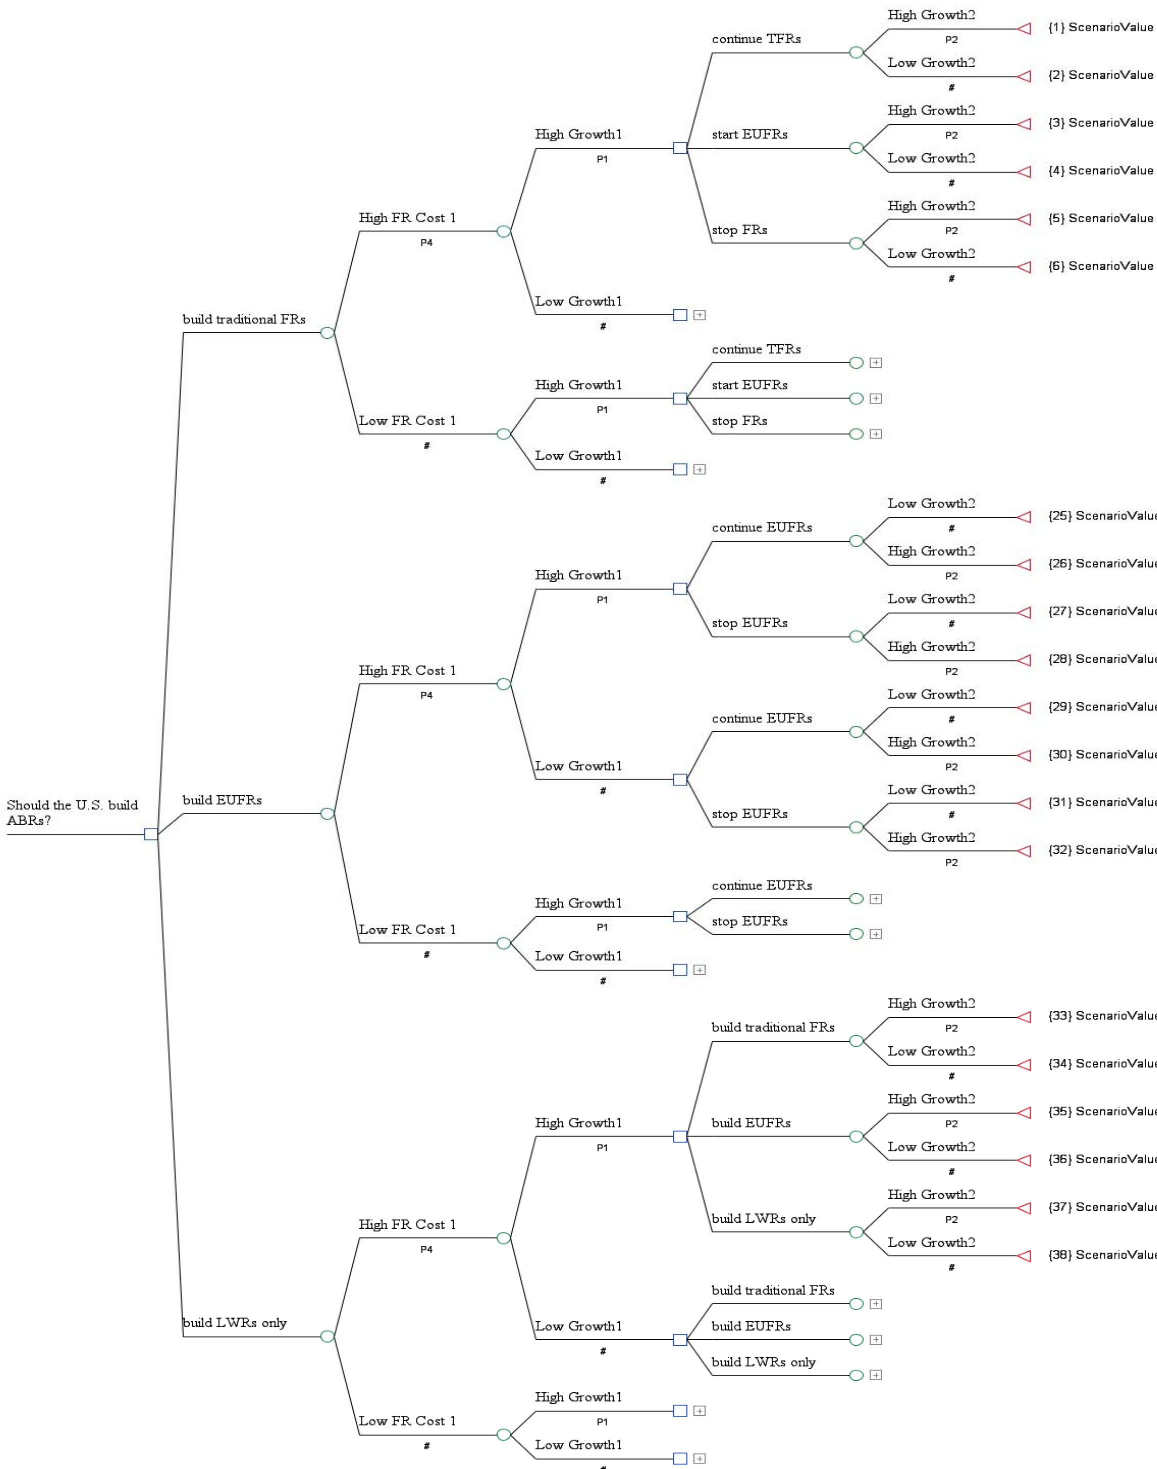

Figure 5-5: Two period decision tree with cost uncertainty

Figure 5-5 shows a two-period decision tree, now with cost uncertainty added; note that some decision paths are truncated to save space. As before, the decision maker first chooses between TFRs, EUFRs, and LWRs at 2040 (after 30 years of LWR growth at 1.2% per year). Immediately after choosing, the decision maker finds out if fast reactors are expensive (55% more expensive than LWRs) or cheap (only 5% more expensive than LWRs). The probability for high cost is represented by P4. Growth is then high or low as before. Then, in 2065, the decision maker has the opportunity to re-evaluate his decision. If he chose TFRs in the first period, he can continue them, stop them, or switch to EUFRs. If EUFRs are chosen first, the options are to continue or stop (in principle, switching from EUFRs to TFRs is also a possibility, but this was eliminated because it would obviate one of the main arguments for EUFRs: avoidance of LWR SNF reprocessing infrastructure). A choice to build LWRs only in the first period means that the same options are repeated in 2065. Growth can then take on a high or low value once again in the second period.

Results for this tree, commensurate with the parameter values in section 5.1, are shown in Figure 5-6. As before, wTRU = 0.15 and wFP = 0.05, meaning that wSNF = 0.80. P4, the probability of high FR cost, is set to 0.5. The primary difference from Figure 5-3 is obvious: considering two periods means that the "build TFR" decision is considerably less attractive overall. As expected (and as before), when growth is expected to be higher, TFRs become more attractive because there is more waste to mitigate. Also as expected, increasing the probability of high FR cost means that TFRs are less favorable.

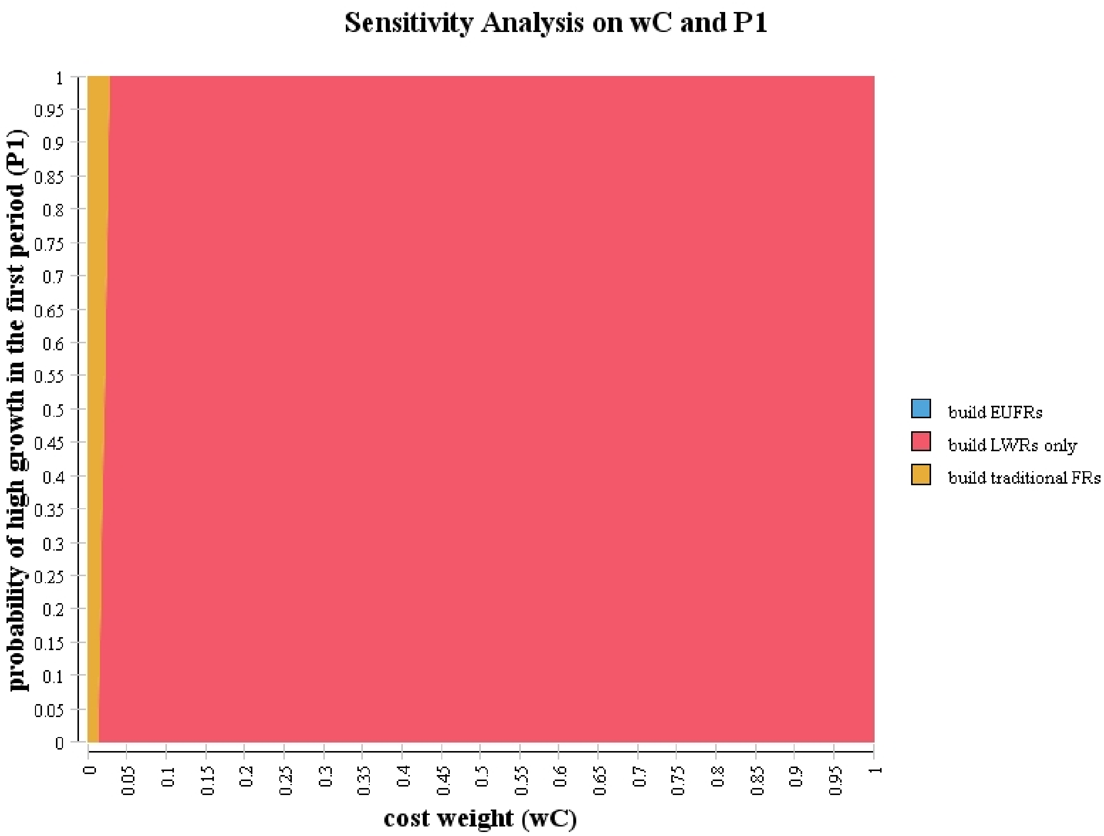

Figure 5-6: Basic results for two-period, two-uncertainty tree

Explaining why TFRs are less desirable in general for a two-period tree requires examination of the second period decision. Figure 5-7 shows the desirable decisions for the second period, given the first period decision. If, during the first period, your weighting and estimation of probabilities means that you decide to build TFRs, in 2065 you will face the decision space shown in the left-hand graph of Figure 5-7. For most weight values, your optimal decision is to continue building TFRs (and consider that in order to arrive at that decision set, your cost weight was very low in period 1).

If, on the other hand, the decision maker chooses to build LWRs in the first period, the optimal choice for most weights is to begin building TFRs in the second period (see right-hand graph of Figure 5-7). The pathway most robust to a wide range of weights is thus to build LWRs in the first period and TFRs in the second.

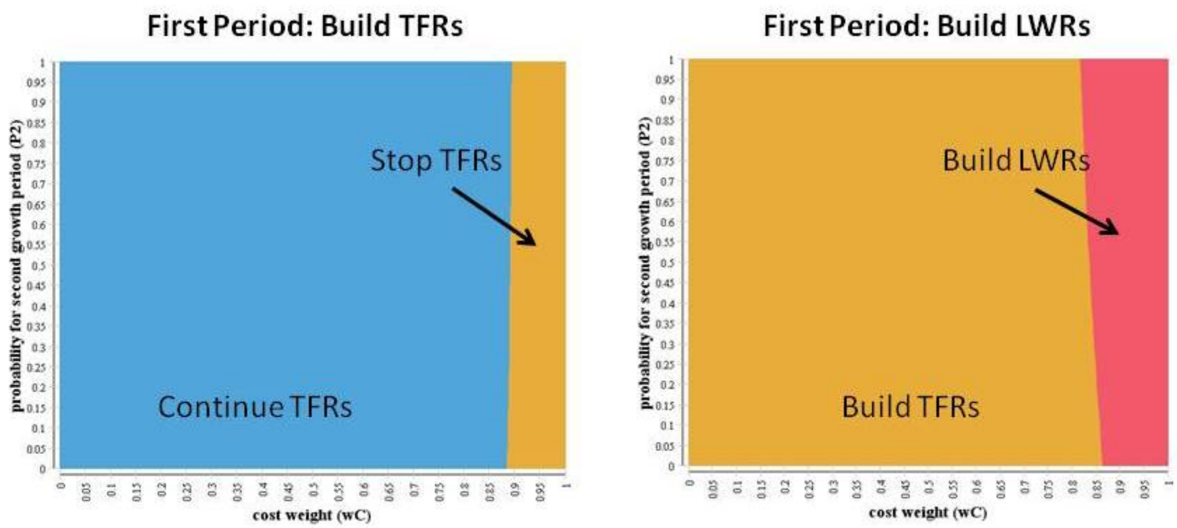

Figure 5-7: Second period decisions for the two-period tree

The reason for the desirability of the "build LWR then TFR" path is that waiting until period 2 and *then* building TFRs entails the lowest system costs, while still providing a waste benefit (under this value function structure). Figure 5-8 shows the operating reactor profile for the two most desirable pathways. The blue line shows the decision to build TFRs starting in 2040, and then continuing in 2065. The green dotted line shows the operating reactor profile when TFR builds begin in 2065. For both scenarios, roughly the same number of fast reactors is built by the end of the century. The reason for this is that starting earlier does not mean you get ahead: the pace of TFR builds is severely restricted by the availability of SNF from LWRs. Starting early *does* mean, however, that the system costs are higher, because more FRs are built earlier when they are discounted less heavily.

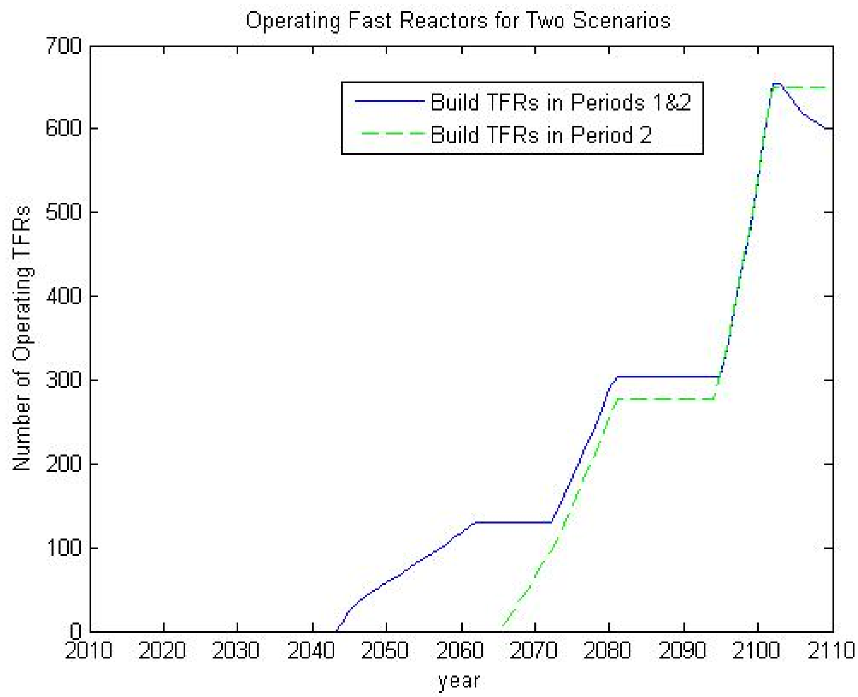

Figure 5-8: Operating FRs for two desirable fuel cycle pathways

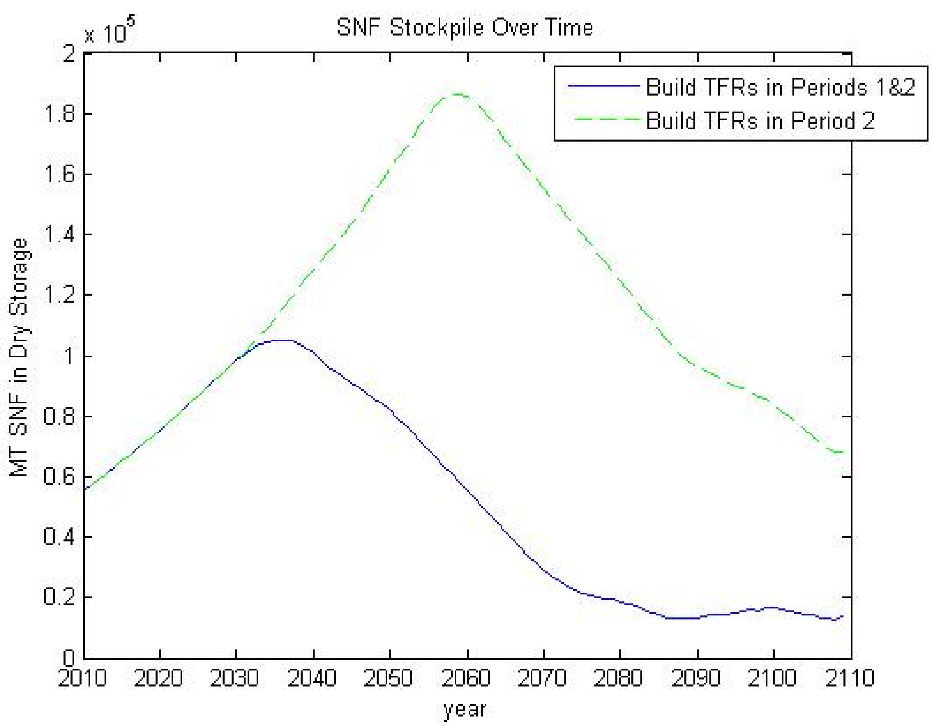

Figure 5-9: SNF stockpile for two desirable fuel cycle pathways

Figure 5-9 shows the SNF stockpile for the same two scenarios. One clear advantage of starting FR builds earlier in the century is that the peak amount of SNF in the stockpile is about half that of the peak if FR builds start later. Ultimately, however, the waste benefit is substantial compared to a continuing increase under LWRs, regardless of which TFR path is chosen. Note that the graphs here represent the highest-growth scenarios possible; for lower-growth scenarios, the SNF stockpile for each decision path is nearly the same by the end of the century. The conclusion is that the best decision, given this formulation, is to build LWRs only for a significant period of time, and to plan to build traditional fast reactors later in the century.

5.2 Key Takeaways from the Two Period Analysis

Adding a second period to the decision tree substantially decreases the desirability of building traditional fast reactors in the first period. This happens because waiting until later in the century to build the reactors entails less cost (because the later builds are discounted more), while a significant waste benefit compared to LWR scenarios can still be realized. Building LWRs in the first period thus appears to be the best decision for almost all preference weightings.

5.3 Building the FR Fleet Gradually

Another odd assumption many fuel cycle analyses make is that building fast reactors is an “all-or-nothing” proposition. In reality, the nuclear industry will need time to become comfortable with a completely new technology, and utilities are unlikely to build fast reactors immediately and as fast as possible once they become available. Indeed, there is an advantage to building a small fleet of FRs early: we can observe how well (and at what cost) they operate, and let those data inform our decisions about whether to more extensively deploy FRs later. Adding in the option to build a fraction of the “allowable” FRs (given availability of SNF) paints a different picture again of the desirability of TFRs.

The tree including options for partial build of TFRs is shown in Figure 5-10. Everything is the same as for the two-period tree above, except now there are two more options for fast reactors at the first decision node. Decision makers can choose to build as many TFRs as are allowed by SNF availability, as before, or can choose to build TFRs at a rate that is 25% the allowed amount throughout the first decision period. Similarly, decision makers can choose to build EUFRs at 25% their “allowed” rate, but recall that EUFRs are not restricted by the availability of LWR SNF. The “25% EUFR” decision thus corresponds to building EUFRs at

25% of electricity demand, with the remaining demand met by LWRs. Partial build options are similarly available in the second period. In the second period, the decision maker can choose to build 50% of the allowed FR amount, or can build 100% as before.

The decision results for the first period with partial build options are shown in Figure 5-11. Each line represents the demarcation between the spaces where FRs vs. LWRs are more desirable for the first period course of action. The line through the green triangles indicates the same basic result presented in the previous section (compare to Figure 5-6). Waste weights are the same (Table 5-2) as for the previous two sections.

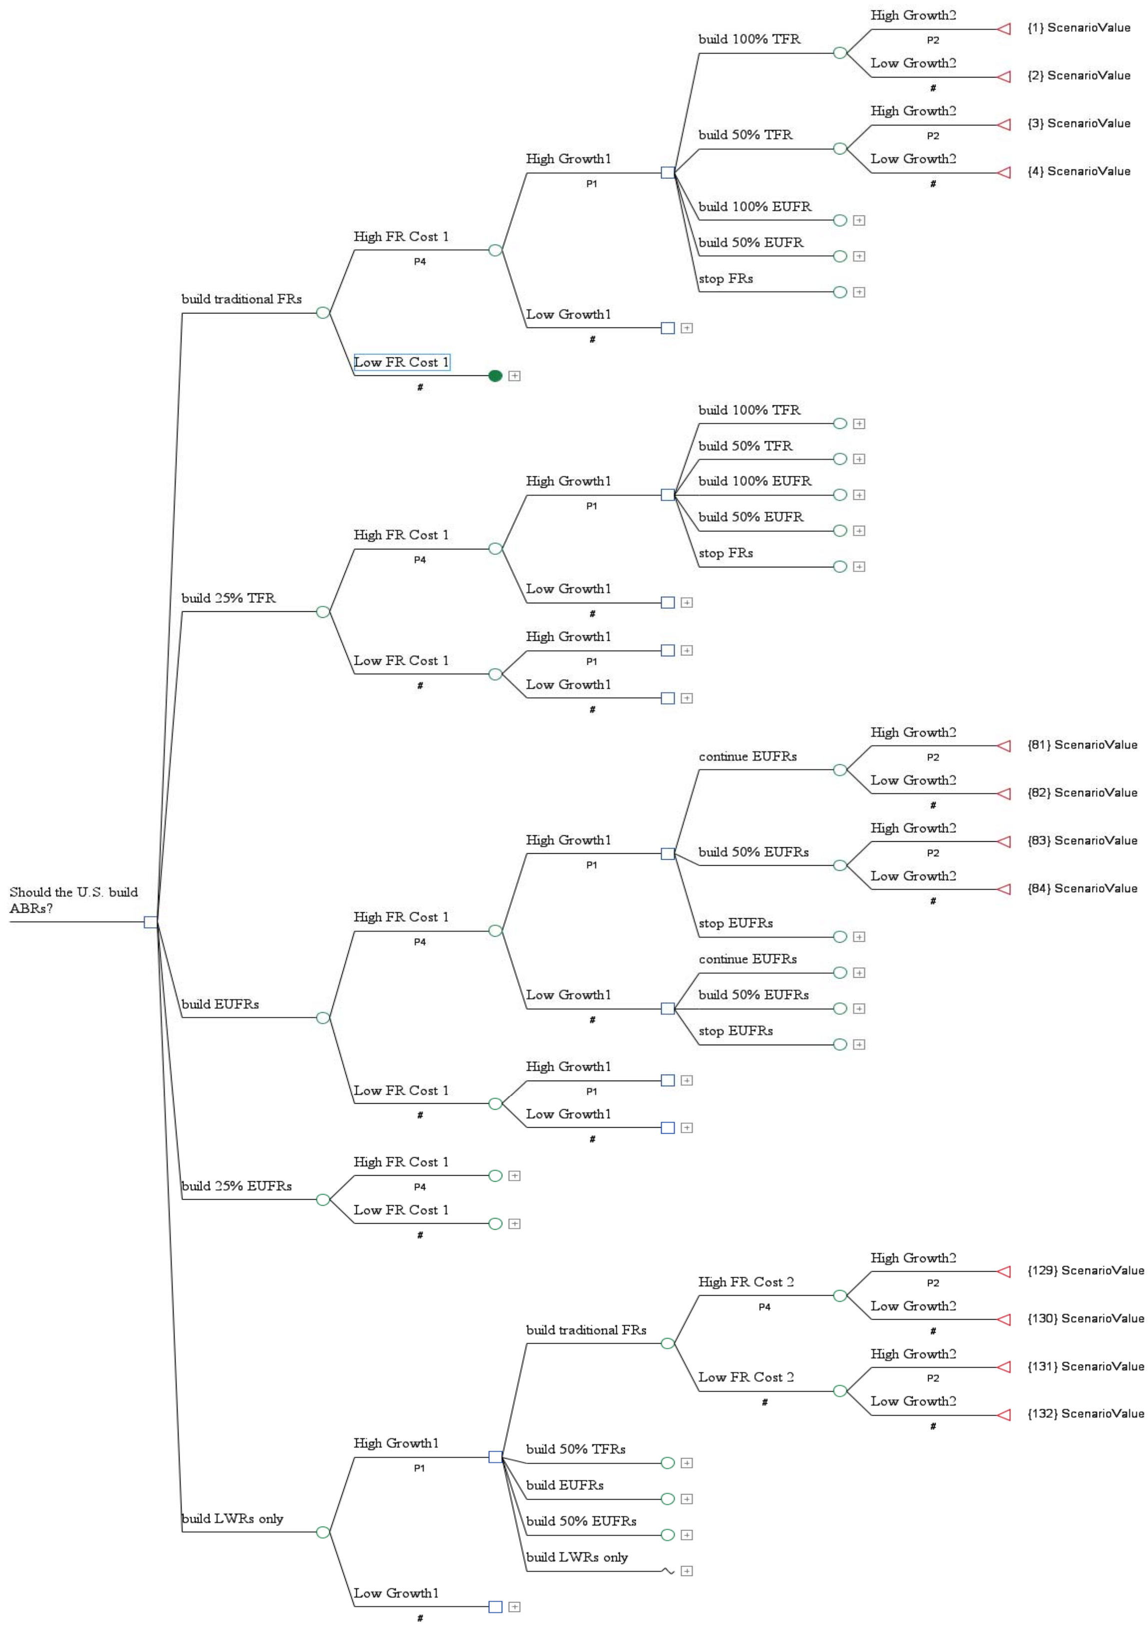

Figure 5-10: Two-period tree with partial build options

Table 5-2: Standard weighting scheme for partial build analysis

Weight	Value
wC	0.12
wW	0.88
wFP	0.05
wTRU	0.15
wSNF	0.80

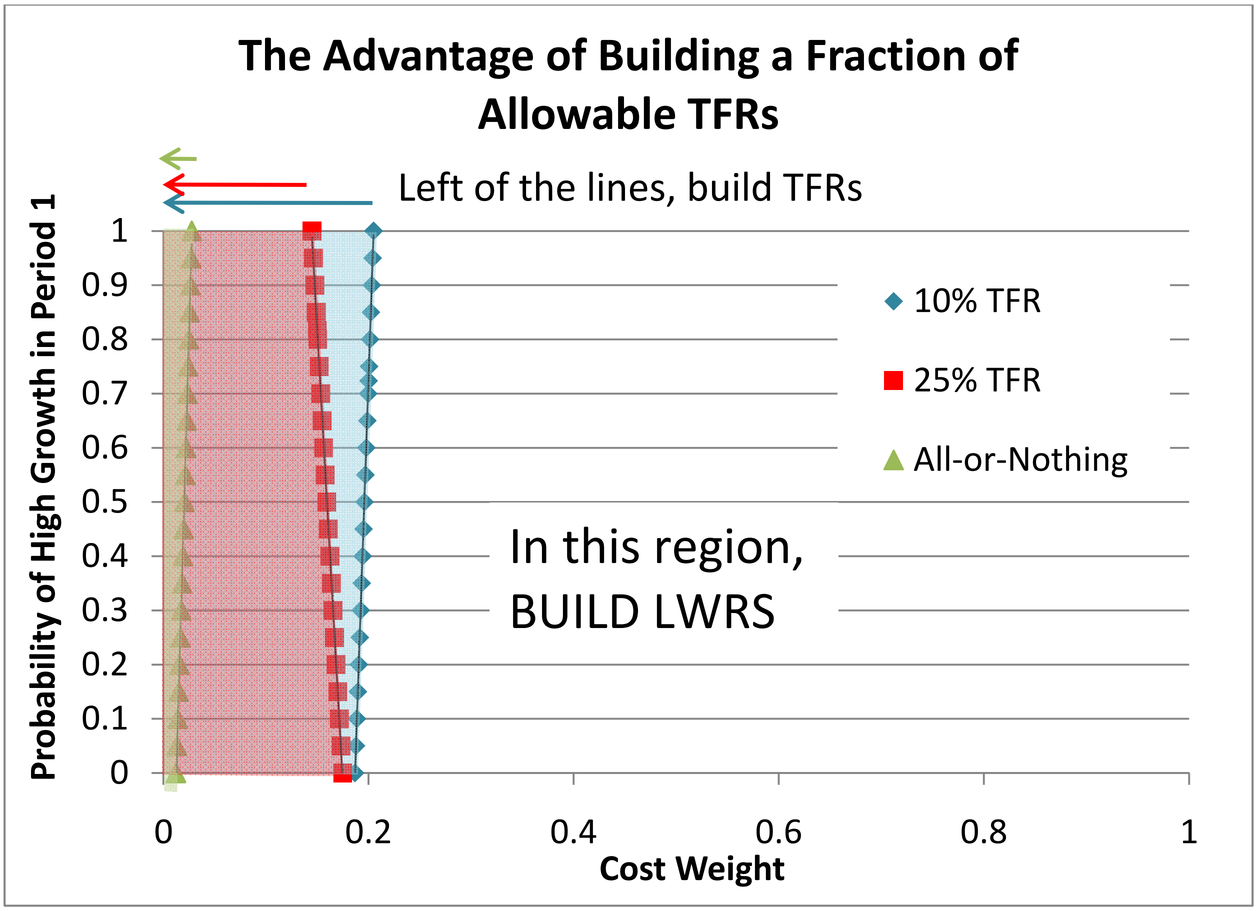

Figure 5-11: Desirable first-period decisions when building fractions of FRs is allowed

To the left of the line through the red squares, building TFRs in the first period at 25% of the allowable amount is desirable. Similarly, to the left of the line through the blue diamonds, building TFRs at 10% of the allowable amount is the desirable course of action. Decreasing the amount of FRs built in the first period actually makes the choice to begin builds a better one for a wider range of cost weights.

As mentioned above, this makes intuitive sense because we would learn about TFRs by building a smaller number of them and would be able to make future decisions based on real deployment data. This effect, however, is not modeled in this tree structure. In fact, re-ordering the tree branches to model a version of path dependency (where we only learn about the cost of fast reactors if we build some first) has virtually no impact on the decision result. Rather, the effects seen in Figure 5-11 are due to the waste management benefit that comes from building a few FRs early in the simulation, for relatively less cost than building a large number of FRs.

A comparison between the partial and full build TFR scenarios is presented in Figure 5-12. The figure clearly shows that building 10% or 25% of the allowed TFRs in the first period, then switching to full deployment of as many TFRs as possible (red and green lines), entails a waste management benefit close to that of starting with 100% TFRs immediately in 2040 (solid blue line). The waste penalty for starting with 10% or 25%, vs. starting immediately with 100% TFRs, is a little less than 20,000 MTHM of SNF. Compare this to a much larger penalty associated with waiting until 2065 to start building TFRs at 100% (dotted blue line: a nearly 60,000 MTHM difference from starting 100% TFRs in 2040 by the end of the century). The cost, moreover, for these small-percent scenarios is significantly less than for building 100% TFRs at 2040. For all three 2040 TFR scenarios, close to the same number of fast reactors is built by the end of the century. Building 100% TFRs early, however, means that more reactors are built sooner, when costs are not as steeply discounted.

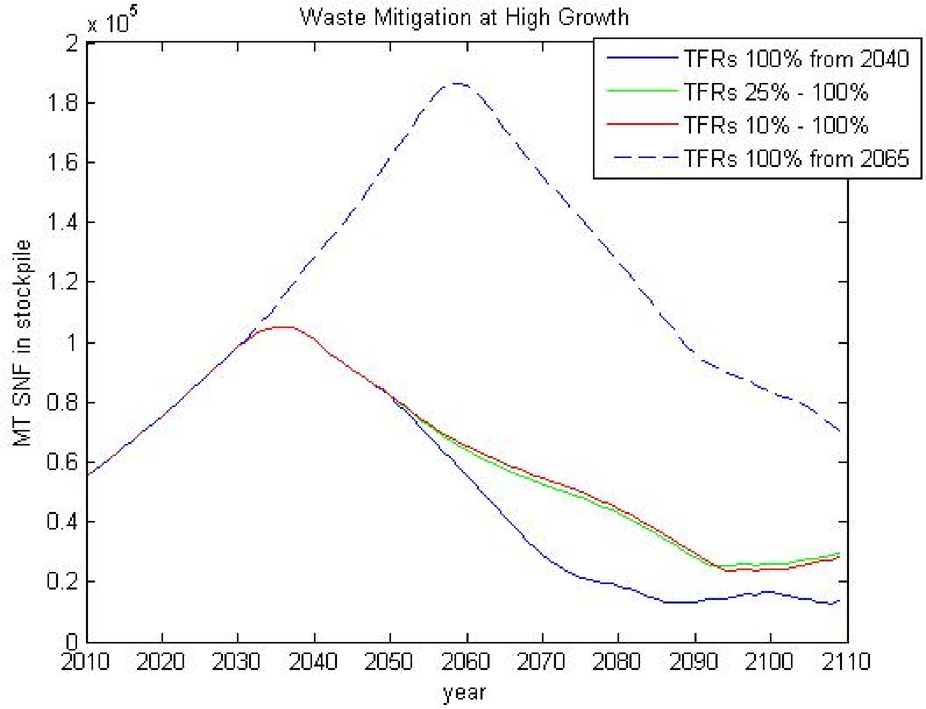

Figure 5-12: Waste benefit of building a few reactors early

Growth effects are responsible for the difference in slope (seen in Figure 5-11) between the results for 10% vs. the 25% build profile. At low growth, the differences between the two percentage build choices are almost imperceptible (and hence the two demarcation lines between optimal decisions meet on the x-axis). The 25% pathway entails building very slightly more TFRs early on, but because the percentages are taken out of a small growth number, the difference between 10% and 25% of possible builds are small. Almost exactly the same number of fast reactors is operating at the end of the century in both scenarios if growth is low. If growth is high, however, the 25% TFR scenario is more costly because it does require building more FRs early on. The end-of-century waste benefit, however, is not larger for the 25% scenario, as can be seen in Figure 5-12. Therefore, at high growth, the 25% scenario is less desirable (and the TFR-desirable space is smaller) than the 10% build scenario.

The ultimate conclusion from this study is that even *before* considering the value of information learned by building a small fleet, building some smaller number of fast reactors early on is desirable. The final decision will depend heavily on the decision makers' cost weights, and building LWRs is still more robust to a wider range of weightings, but decreasing the amount of FRs built in the first "wave" makes an early FR pathway much more attractive.

5.3 Key Takeaways from the analysis on Building the FR Fleet Gradually

Building a small number of traditional fast reactors in the first period (starting in 2040) is much more desirable than immediately building as many FRs as possible given the availability of spent nuclear fuel. The additional desirability comes because there is a waste benefit to starting with some fast reactors, but relatively little cost penalty; the gains from the "value of information" in terms of lessons learned from a small fast reactor fleet would be an additional reason to build a few early on.

5.4 Very Low Nuclear Power Growth or "Nuclear Catastrophe"

The scenarios above all assume the possibility for very high nuclear power growth. Even low growth is modeled as an increase in nuclear power, because if we build no new reactors, we have no need to consider advanced fuel cycle systems (other than possible partitioning of existing waste in the context of the once-through cycle, which is not an option evaluated in this thesis). A natural question decision makers have is, "what should we do if nuclear electricity demand grows very little or stagnates?" Another natural question to ask is what happens if a large catastrophe were to strike, such that all nuclear reactors and processing facilities were suddenly shut down.[1] Each of these questions is addressed in this section.

5.4.1 Very low growth and nuclear stagnation

Intuition tells us that building a fast reactor infrastructure is unlikely to be worthwhile if we will not be generating significant amounts of waste. In order to see if the analytical framework here agrees with intuition, new values for growth are inserted into the "gradual build", five-option tree of section 5.3, with 10% TFRs as the slow build option (see Figure 5-10). Now, instead of low growth at 0.5% per year and high growth at 2.5% or 4% per year, low is 0% and high is 0.1% for 2040-2065 and 0.01% for the second period (2065-2110) These growth values are demonstrated in tree form in Figure 5-13.

[1] Prior to the March 2011 accident at the Fukushima Daiichi reactors in Japan, this case was considered to be of small probability. The response of many major nuclear countries has been moderate; most are calling for safety reviews, but neither France nor the U.S. has abandoned plans to build more plants. Italy, Switzerland, and Japan have all announced that they will forego any new builds, but will continue operating the plants they have in place (though the Swiss government is considering a phase-out by 2034). Only Germany has so far definitively called for a before end-of-life shutdown of all its operating nuclear power plants.(Hewitt, 2011)

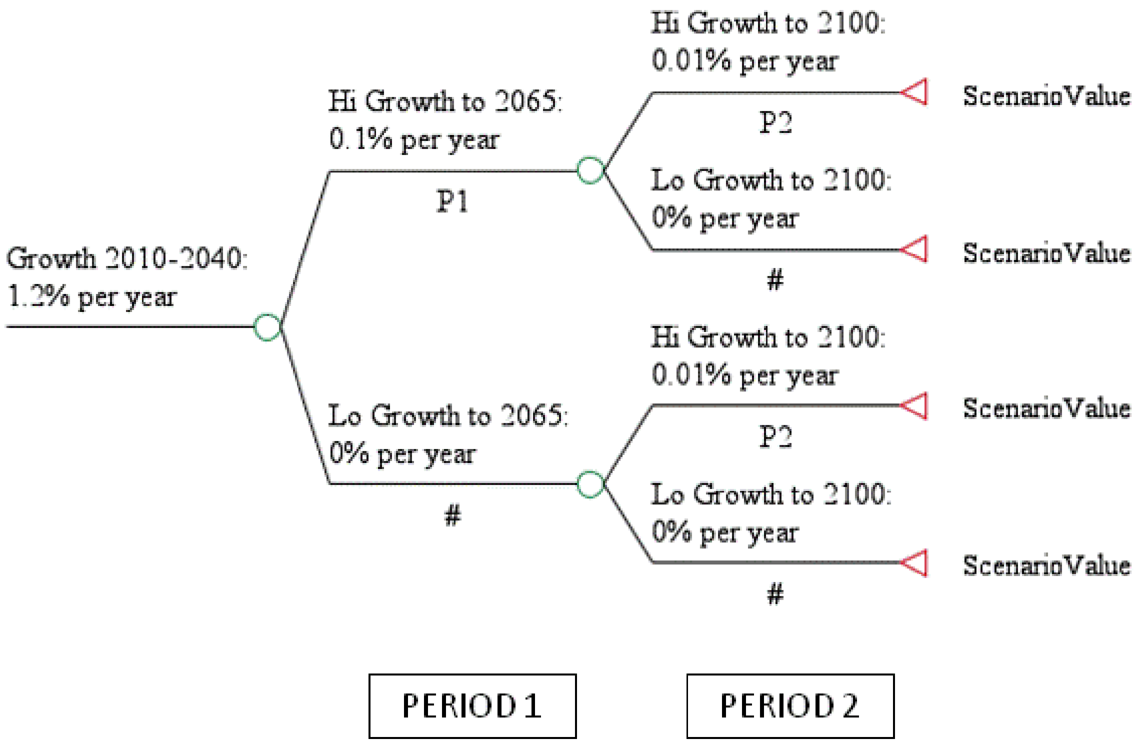

Figure 5-13: Tree of uncertainties in nuclear demand for low growth scenarios

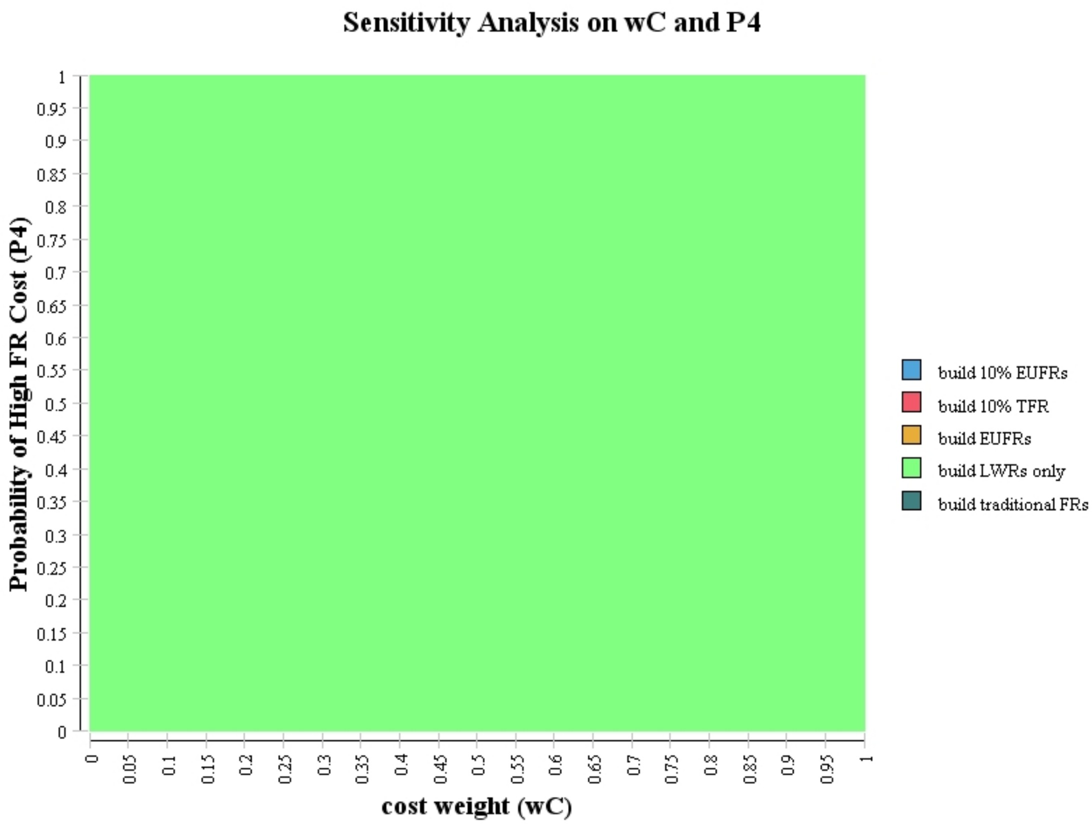

Figure 5-14: Desirable decisions under conditions of extremely low growth

Results for desirable first-period decisions are shown in Figure 5-14. Clearly, the only desirable option for conditions of low growth is to build LWRs at the first decision node. The results are shown for sensitivity to cost weight and to probability for high FR cost, but the same graph appears for the sensitivity of cost weight to growth in period 2 (P2 vs. wC). Note that these results are valid for a particular set of weights over the different waste types; in this case, as before, wFP = 0.05, wTRU = 0.15, and wSNF = 0.8, indicating an extreme dislike for handling SNF vs. other waste products.

Results for the second period show a marked and somewhat surprising affinity for building TFRs at 50% of their allowed pace (see Figure 5-15). The results are presented as a sensitivity between the cost of FRs and the cost weight, but the results are nearly identical for the relationship between wC and the probability of high growth in the second period (although with an even more vertical slope between the regions, owing to the very slight difference in growth values). Even at very low growth, a fair number of traditional fast reactors would be built later in the century both to replace decommissioning LWRs and to build the few more that would be needed to fill demand. Because costs are much more heavily discounted in the second period, fast reactors are attractive even at high cost weights because they come with a waste advantage.

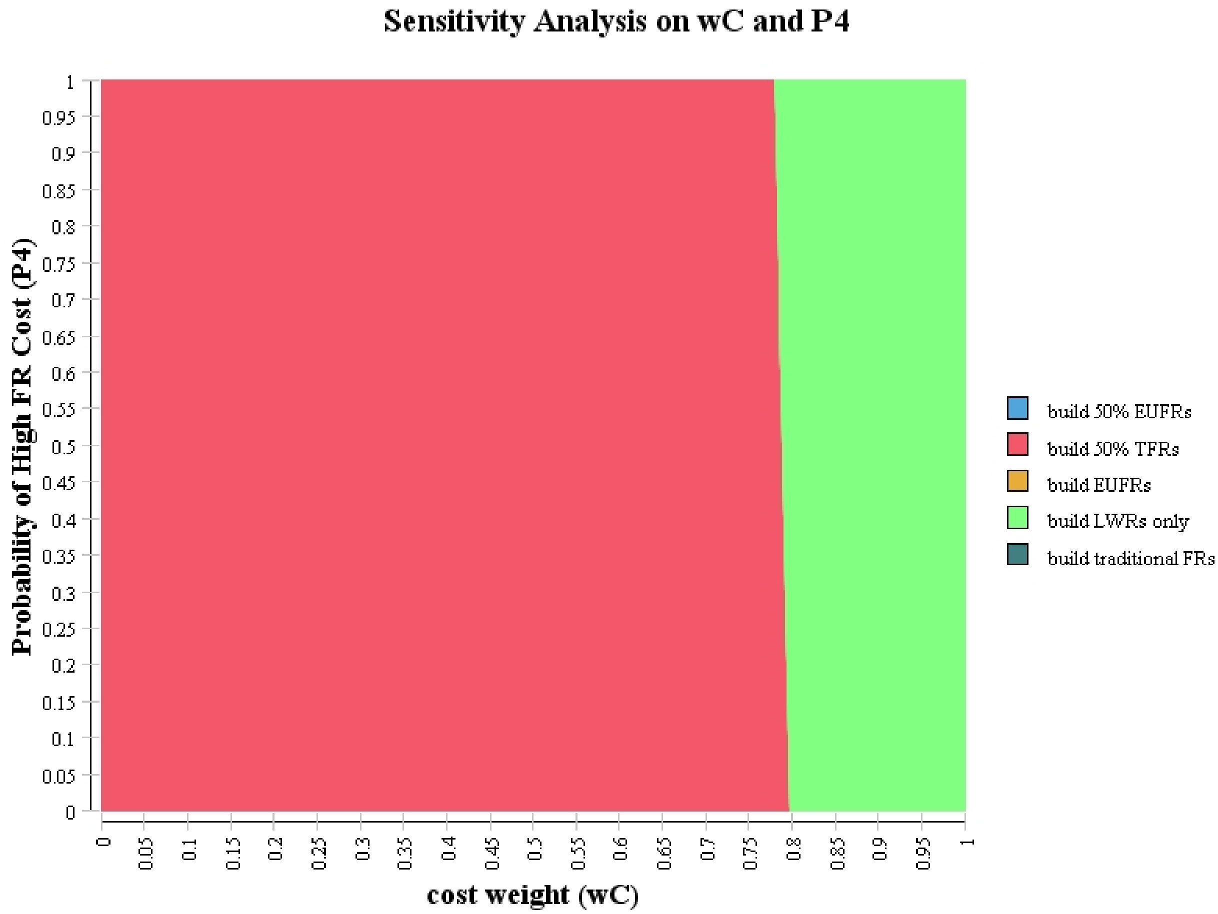

Figure 5-15: Desirable options for the second period, very low nuclear power demand

To some nuclear fuel cycle experts, the desirability of building an entire infrastructure of fast reactor systems later in the century in order to support only 100-odd reactors does not track with intuition. Indeed, it may be that the analysis framework over-emphasizes the value of waste management under these low growth conditions. In practice, however, the overall decision analysis results should not be the sole decision-making device (they would be supplemented by deliberation in which points of this nature would rise to the surface), and furthermore, the second period decisions recommended by the framework are not automatic. Rather, a decision would be taken for the first period under consideration of the second period pathways, but later, analysts should conduct a similar study with updated information (rather than simply relying on the optimal second period decision as decided several decades ago). The information from this low-growth exploration thus tells us that first-period decisions track well with intuition, but that for these analyses we should understand that the framework includes (for these waste weights) a measure of bias toward using FRs for waste management purposes.

Changing the relative weighting of wTRU and wSNF has an impact on the results of the low-growth analysis. Figure 5-16 shows the sensitivity of the decision result to the TRU weight and cost weight. The fission product weight is held constant at 0.05, so the SNF weight changes opposite to the TRU weight, as wSNF = 0.95-wTRU. The graph shows, as expected, that at very high cost weights and/or very high TRU weights (indicating that management of a separated TRU stream is undesirable), LWRs are the preferred option. At this end of the spectrum, having lots of SNF is more desirable than having any separated TRU. If, for example, repository space were plentiful, and keeping TRU mixed with FP and other spent fuel products were desirable from a proliferation perspective, the weight for minimizing SNF would be very low. The once-through fuel cycle would be strongly favored in this case. For many values of the TRU weight, building 50% TFRs in the second period is desirable, because the decision maker would be very keen to get rid of SNF. If the cost weight *and* TRU weight are both very low, the decision maker might even want to build 100% TFRs (this is essentially a case where SNF must be minimized at all costs).

The first period decision, on the other hand, is not very sensitive to the TRU vs. SNF weighting. Because FRs are more expensive in the first period, and because waste can be mitigated later, the best choice is to build LWRs for nearly every value of cost and TRU weights.

A scenario in which built reactors continue to operate but no new builds occur (even to replace decommissioned reactors) is examined in the three-period tree of section 5.5.

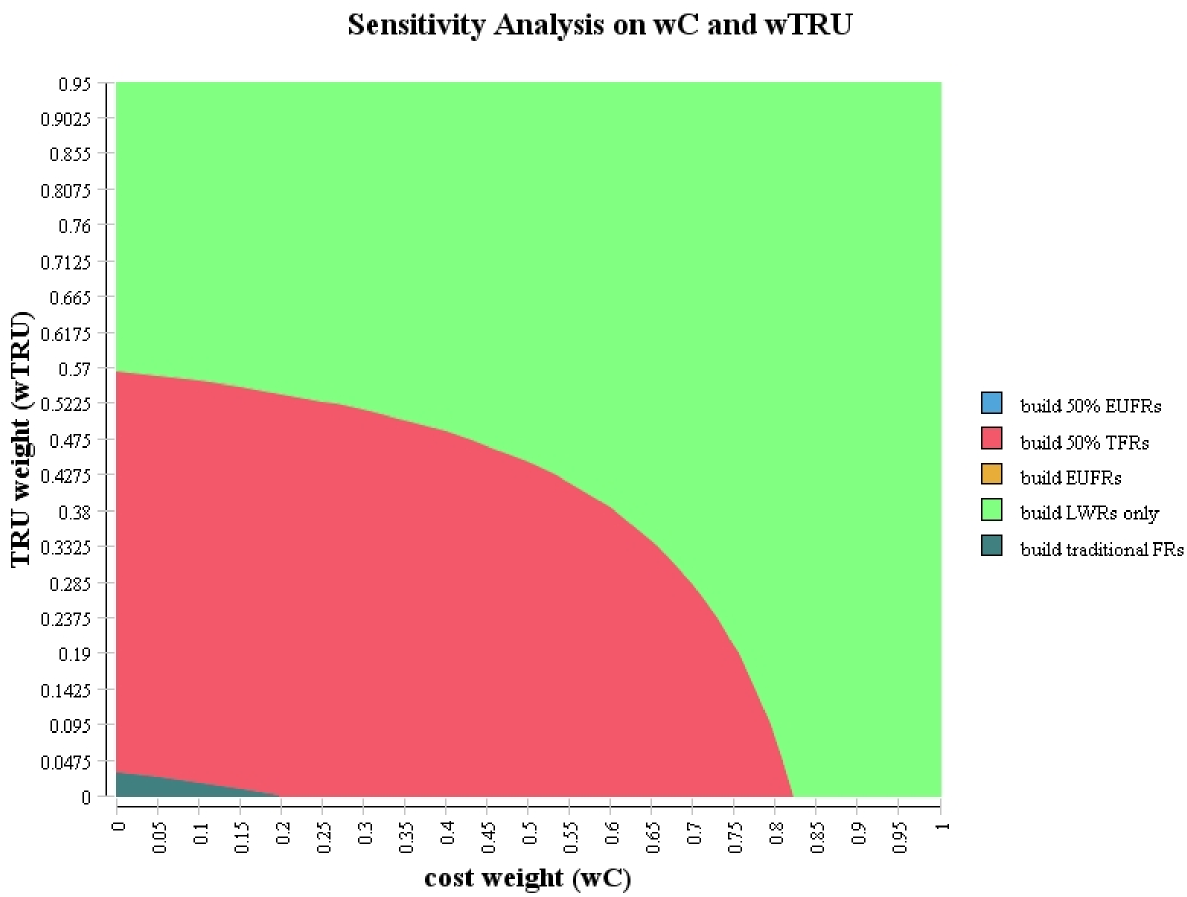

Figure 5-16: Sensitivity of low growth results to TRU/SNF weighting

5.4.2 Nuclear Catastrophe

Despite the severity of the nuclear accidents at Chernobyl, Three Mile Island, and Fukushima, society has never required a complete and immediate shutdown of all nuclear power plants either in the affected countries or elsewhere (though the situation in Japan is still evolving, and immediate shutdown of all Japanese plants is a future possibility). Advancements in nuclear safety continue to minimize the risks associated with nuclear disasters, so that a sudden abandonment of operating nuclear power due to a catastrophe remains fairly unlikely. Nevertheless, one can imagine a very small-probability event with economic and public health ramifications so big that all U.S. nuclear power plants are ordered to shut down, after which the whole nuclear system would be dismantled.

The decision tree is the same one used in section 5.3, with five options (including 10% TFRs and EUFRS) and two periods. This time, however, “Low Growth” in the second period is modeled as a shutdown scenario rather than as 0.5% growth. P2 represents the probability of attaining high growth in the second period (still modeled as 4% per year), so that (1-P2) represents the probability of full nuclear shutdown.

For the shutdown scenario, the decision is made and all nuclear reactors cease operating in 2065. The reactors are then decommissioned over the next 1-2 decades (though no reactors operate during that time, the time of decommissioning determines when decommissioning costs are incurred). All reprocessing, enrichment, and fabrication plants cease operation immediately, and all reactors discharge their full cores into cooling pools in the same year the decision is made (2065). After cooling for five years, LWR fuel that was discharged in 2065 enters the SNF stock. The FR fuel discharged from FR cores requires a new management route for this scenario (for the previous scenarios where all plants continue operation, we assume that only losses from FR fuel recycling need to be managed, because the TRU from decommissioned reactors would continue to be recycled). The discharged FR fuel would contain both TRU and fission products, and separation would not be possible because all recycling facilities would be shut down. To account for this fuel, the TRU component (14.1% of the fuel mass) is added to the TRU waste bucket. Fission products remain with that TRU, so nothing is added to the FP bucket. In essence, this assumes that the FR core fuel would be cooled for a period of time until the TRU characteristics alone dominate the waste behavior and the fission products are no longer relevant.

Note that for such an extreme scenario, if all reactors were shut down at exactly the same time, they would all be at different points in their fueling/outage cycles. Much of the fuel coming from the reactor cores would thus not have achieved full burnup, and this would impact the fission product and TRU ratios in the fuel. This effect is ignored, because the decision outcomes depend on relative values between the different fuel cycles, and the effect of the lesser TRU/FP fractions is likely to similarly impact both once-through and closed fuel cycles.

The results for a potential shutdown scenario are shown in Figure 5-17. This is starkly different from the basic 10% TFR result (see Figure 5-11). Before, the 10% TFR decision was desirable for low cost weights, and building 100% TFRs was never desirable. Now, however, building 10% TFRs is never desirable, and at low cost weights, the decision maker should build TFRs at 100%. This makes intuitive sense: if there is a possibility that all nuclear will shut down entirely in the second period, there is no reason to build a small number of TFRs for some added cost with little relative waste benefit (especially because these reactors will have to discharge their cores). If, however, the cost weight is low, the analysis tells the decision maker to build TFRs quickly in order to reap as much waste benefit as possible before a possible shutdown. The large area for which the 100% TFR decision is attractive goes counter to what we learned in sections 5.2 and 5.3: waiting is no longer quite as advantageous if the possibility exists that the nuclear system will cease functioning.

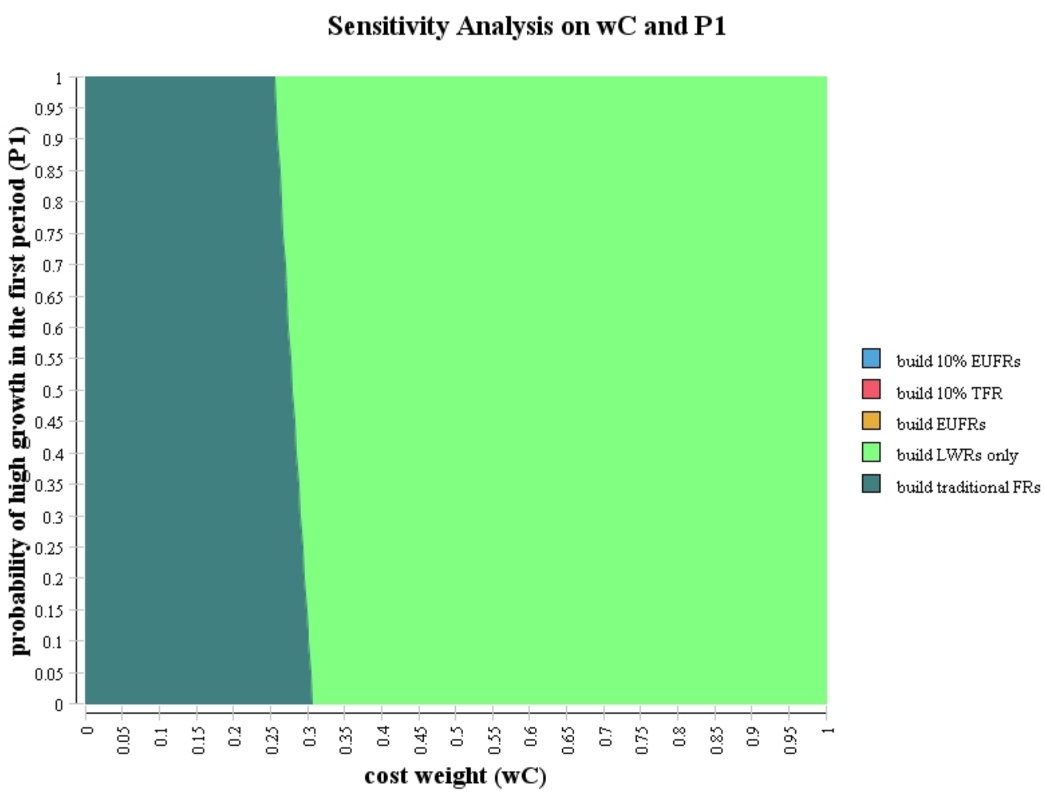

Figure 5-17: Decision sensitivity to cost weight and period 1 growth for shutdown scenario

If we are fairly certain, however, that growth will be high in the second period (and thus that we will not face a nuclear shutdown), the first period decision to build 10% TFRs reappears as a desirable option. Figure 5-18 shows how the first period decision depends on growth in the second period. For many values of the probability for high growth (recall that the probability for shutdown is 1-P2), we would not want to bother with 10% TFRs, but at high P2, we face a set of desirable options similar to those that rose to the top in section 5.3.

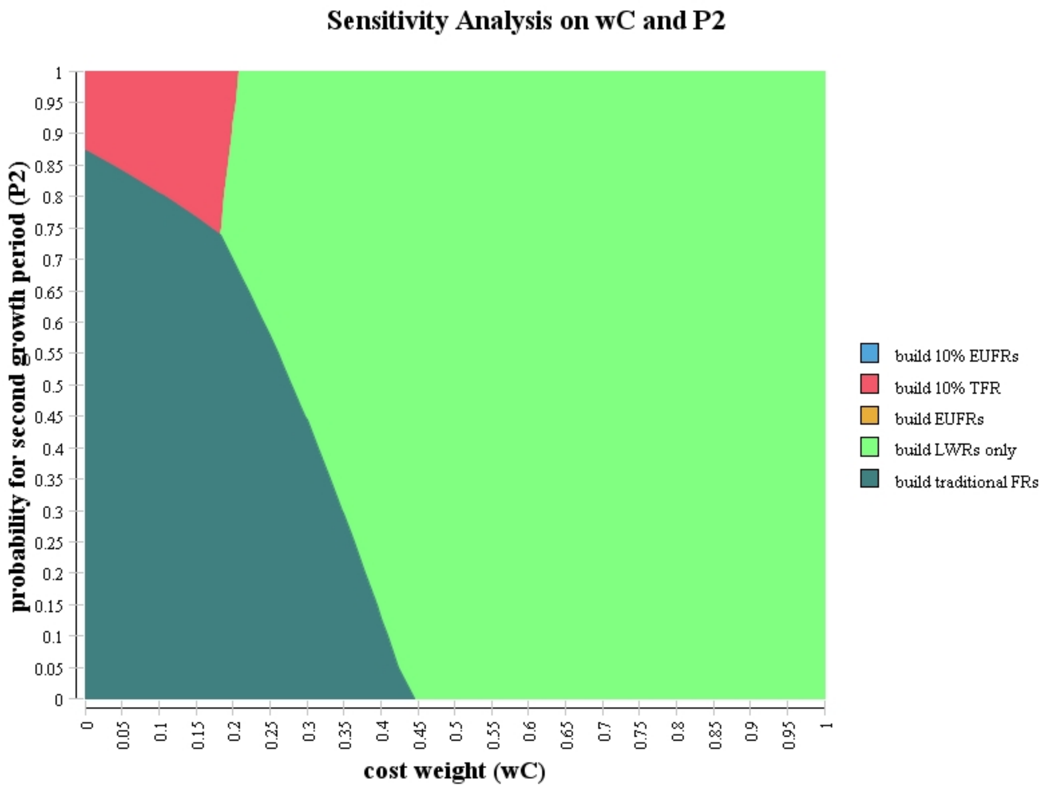

Figure 5-18: Sensitivity between second period growth and cost weight for nuclear shutdown scenario

In section 5.3, we learned that in general, a higher probability of low growth means TFRs become less desirable in the first period (but more desirable in the second period). Here, on the other hand, if the possibility exists with some decently high probability that nothing will be built in the second period, and indeed that U.S. nuclear infrastructure will stop operating, decision makers are encouraged to build fast reactors quickly in the first period (provided their cost weight is low enough). Most likely, the probability of such a catastrophe scenario is low. The decision result space in section 5.3 is thus probably more relevant to decision makers. The next section explores a less-dire "zero build" scenario where new power plants are never built, but a plant built before a moratorium continues to operate for its full 60-year lifetime. The probability of this type of scenario is much higher than the probability of full nuclear shutdown.

5.4 Key Takeaways from Low Growth and Nuclear Catastrophe Analyses

If growth is likely to be very low throughout the century, building LWRs in the first period is the best option for all possible values of growth and preference weights. For the same conditions in the second period, however, the framework advises decision makers to build traditional fast reactors at 50% of their allowed pace. If there is a probability of “nuclear catastrophe” (i.e. a sudden shutdown of all nuclear facilities) in the second period, the framework advises a rapid buildup of traditional fast reactors in the first period if the waste weight is high, in order to achieve *some* waste mitigation before the catastrophe occurs. The emphasis on TFRs in the catastrophe scenario and in the second period during low growth is evidence of a bias in the framework toward waste mitigation (since intuitively, for example, we would be unwilling to develop fast reactors if we knew nuclear growth would remain extremely low).

5.5 Three Periods

The three-period analysis evaluates the impact of dividing the problem into three stages, rather than just two. In principle, this is a more realistic model: policymakers are likely to have more than two opportunities to make changes to the nuclear fuel cycle within the century. If, however, we consider similar ranges for costs and growth trajectories, we would expect the insights from a three-period model to be similar to a two-period model, with some greater fidelity in the results but similar trends. The three-period model may still be useful: it allows more freedom to explore exotic growth patterns, including full stops and starts to nuclear infrastructure at the low end of nuclear growth assumptions.

A “bumpy” pattern of nuclear power growth is assumed for this analysis. This allows an exploration of the desirable decisions under conditions where we begin building a new technology, then stop and start again. Figure 5-19 gives a general impression of the tree layout. The same options presented in section 5.3 are preserved for the first node, including building 100% or 10% of the possible fast reactors; but now, the first decision is made at 2025. At the second decision node (2050), the decision maker has a choice between 100%, 50% or 0% of the same advanced technology. If LWRs are chosen in the first period, the choices in the second period are 100% TFR, 50% TFR, 100% EUFR, or 50% EUFR, or LWRs again. The third period (a decision made at 2075) repeats the decision options of the second period.

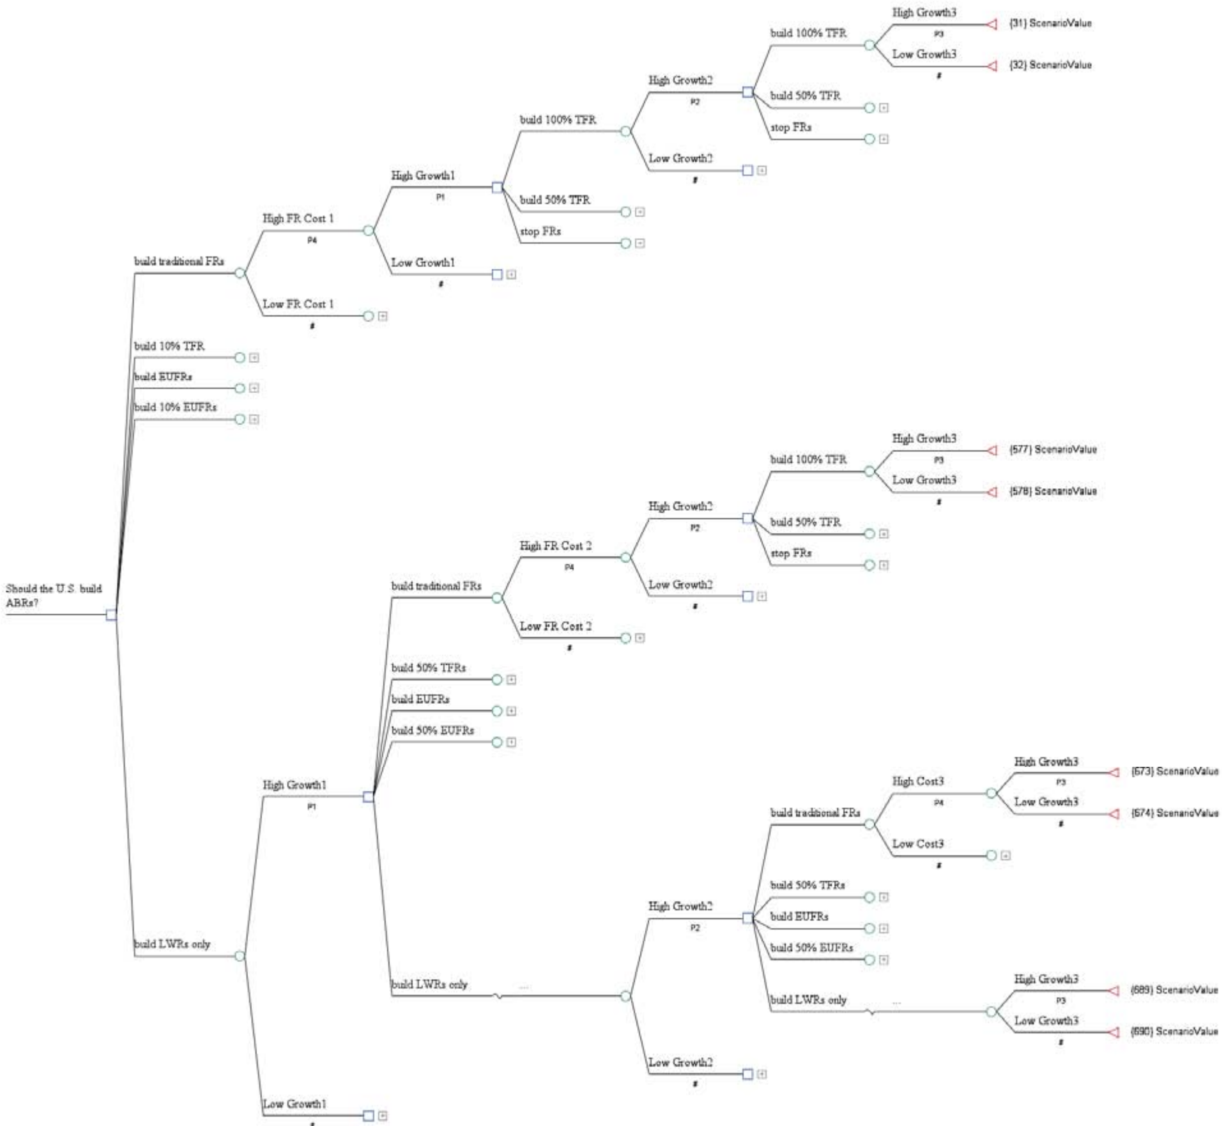

Figure 5-19: Three-period tree

At each node, one of two growth rates materializes. The possibilities for each uncertainty node are listed in Table 5-3. A growth rate of 0% per year means that reactors are built when needed to replace those that are decommissioned. The designation “no builds” means that zero reactors are built (and no LWR fuel is reprocessed), even if reactors are decommissioned; electricity demand is met by non-nuclear sources. Any built reactor continues to operate until it has served its full 60 years, regardless of the growth rate.

Table 5-3: Growth rates for three-period tree

	Period 1 Growth Rates	Period 2 Growth Rates	Period 3 Growth Rates
High Growth	1.75% per year	0% per year	1.75% per year
Low Growth	0% per year	No builds	No builds

The decision results are similar to those for the low-growth study presented in section 5.4. Figure 5-20 shows the desirable options given different possibilities for the cost weight and for the probability of high growth in the first period. The desirability of traditional fast reactors actually *increases* compared to the 2-period, positive growth scenarios. This occurs because the possibility exists that zero reactors will be built in later years, so to guarantee a waste benefit, fast reactors should be built sooner. This effect is similar to the shutdown effect observed in the preceding section, which pushed 100% TFRs into the range of desirability.

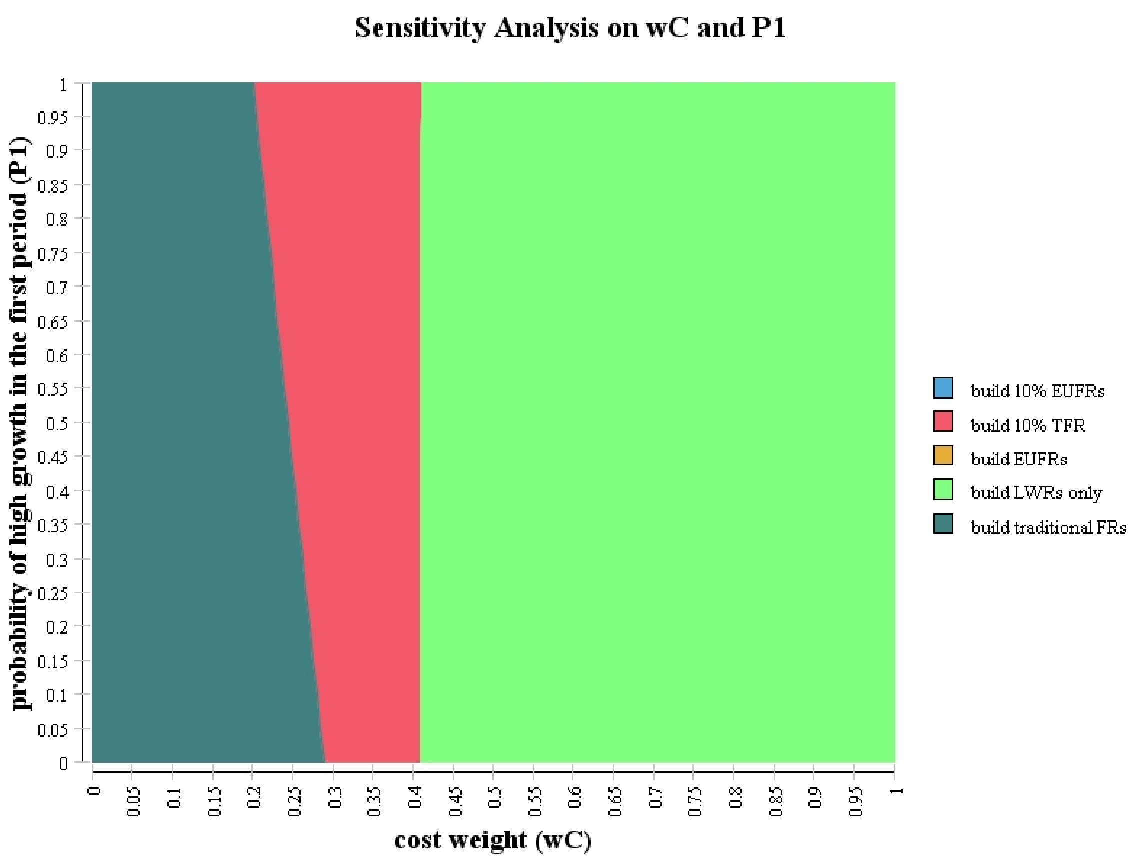

Figure 5-20: Desirable first period decisions for the three-period tree

Desirable decisions for the second period are shown in Figure 5-21, for the scenario in which 10% TFRs are build in the first period, and then high FR cost and high growth rates materialize. At low P2, the probability of requiring any nuclear reactors to be built in the second period is extremely small. In this regime of P2/wC space, decision makers should be indifferent to the choice between 50% and 100% TFRs (because in fact, zero reactors will be built). TreeAge is not able to represent indifference on the 2-way sensitivity graphs, so instead presents a strange sharp shape between the blue and pink regions at low cost weights.

The curve separating 50% TFRs and stopping FRs is a real result. It stems from the assumptions made for these scenarios: "stopping FRs" entails stopping all reprocessing of FR spent fuel. For these simulations, the stockpile of TRU is enough to keep built FRs operating. On the other hand, if demand is low or zero but the "desire" still exists to build FRs, reprocessing of spent FR fuel continues. The result is that putting serious brakes on the entire FR infrastructure (both reprocessing plants and reactors) becomes more desirable if growth is very probably going to be nil. Changing the assumption, so that reprocessing of spent FR fuel continues regardless of whether new FRs will be built, would make the decision maker indifferent between the stop/build choices.

The results for third period decisions look similar to second period decisions in their dependencies on cost weight and probability of growth.

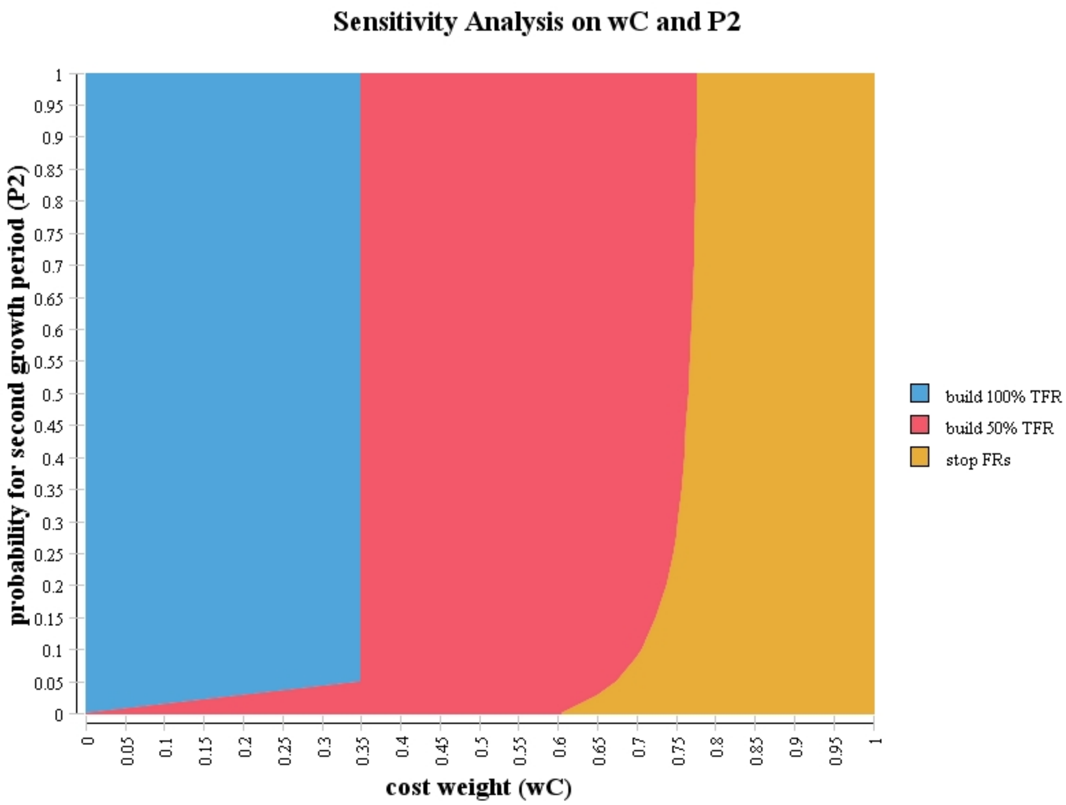

Figure 5-21: Desirable second period decisions for the three-period tree, if 10% TFRs are built in the first period

The three-period results do not produce wildly different results or offer new insights compared to the two-period studies. The stopping-starting build profile assumed in this section, however, would actually wreak havoc on the nuclear industry, causing great uncertainty for the nuclear construction workforce and for the institutions that finance nuclear builds. But the value function, dependent only on the scenario costs and heavily biased toward waste considerations, is not able to capture those effects. This is a limitation of the present model and value function structure. Further examination of the impacts of stopping/starting will be necessary before any significant investments are made in a new direction for the nuclear fuel cycle.

5.5 Key Takeaways from the Three Period Analysis

Adding a third period to the decision tree allows a more exotic growth profile to be examined, including starts and stops to the nuclear industry. The decision results under these conditions, however, are similar to the results from low-growth and catastrophe scenarios examined in the preceding section. Decisions are still not especially sensitive to the probability of high vs. low growth, but the possibility that no reactors will be built in later periods makes building fast reactors more desirable early on, in order to achieve at least some waste benefit.

5.6 The Impact of Uranium Cost

Concerns about uranium availability were a primary motivation for France to build its reprocessing systems, (World Nuclear Association, 2011b) and were instrumental in Japan's decision to close its fuel cycle. Fast reactors also address uranium availability challenges, because they can use LWR SNF as a feedstock and/or be operated in a "breeder" mode where fissionable U-238 is transmuted into fissile Pu-239. Though this investigation does not consider breeder reactors, it does explore self-sustaining, plutonium and enriched uranium-fed reactors with the potential to extend uranium resources. In fact, for some growth scenarios, enriched-uranium fed fast reactors can save even more uranium than traditional fast reactors, because EUFRs can take over the fleet quickly and obviate the need to continue building LWRs in order to have feedstock for starting up fast reactors.(Kazimi et al., 2011) A natural question is whether increasing uranium prices will increase the desirability of building these FRs.

In order to answer this question, the tree depicted in Figure 5-22 is analyzed. To simplify the results, only the 3-option tree with "build at 100%" choices is examined. At low nuclear power growth, uranium resources are not as likely to constrain LWR expansion. The low-growth nodes are thus pared, and high growth scenarios are retained in order to highlight the maximum possible benefits FRs can offer in terms of fuel resource extension. The tree reflects zero uncertainty in nuclear power demand.

Full tree evaluations are performed with different uranium costs. In this way, the examination of uranium cost uncertainty is a deterministic scenario analysis. The simple scenario analysis is easier to model than a stochastic representation of uranium cost, so it was conducted first to determine if more complex uranium cost modeling is necessary.

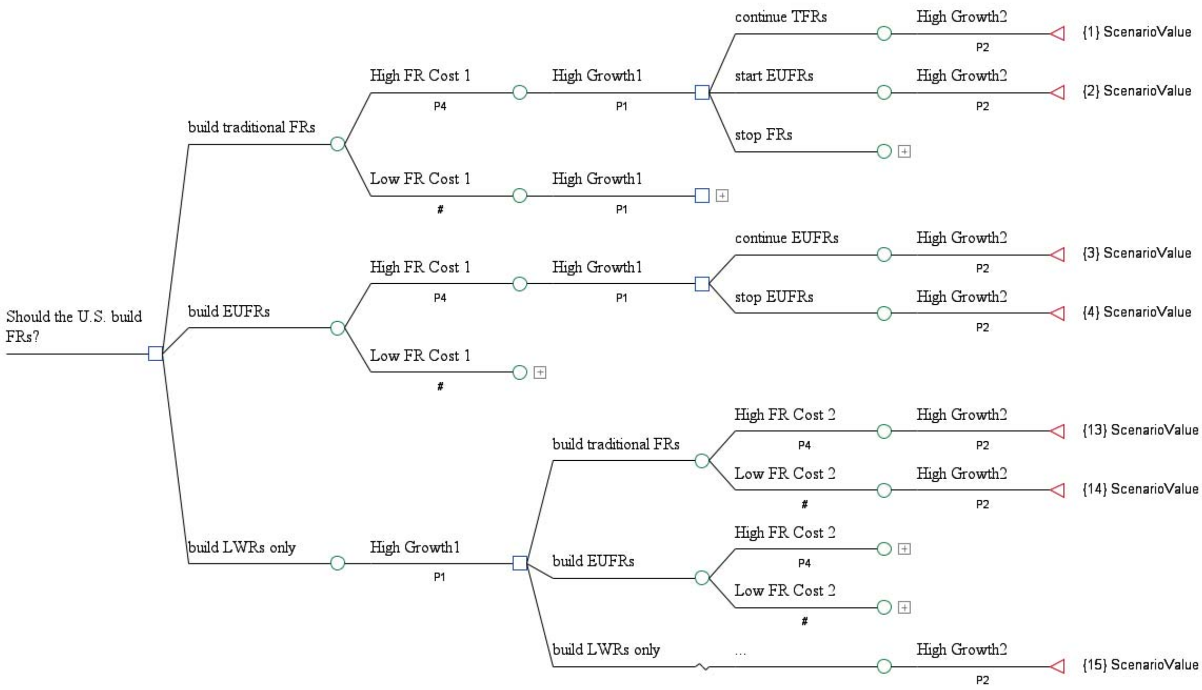

Figure 5-22: Uranium price impacts tree

The results in Figure 5-23 for the first period decision show that uranium costs do have an effect, but that a significant change in the desirability of fast reactors requires a very high uranium cost. $80/kg represents the approximate current uranium contract price (for delivered U_3O_8 in 2010), and is the assumption used so far in this study as well as by (De Roo & Parsons, 2009). The green region thus reflects the area where TFRs are desirable, given this uranium cost and the promise of high nuclear power growth. The next uranium cost threshold discussed in the NEA/OECD “Red Book” includes recoverable resources above $130/kg, (A Joint Report by the OECD Nuclear Energy Agency and the International Atomic Energy Agency OECD, 2010); at $400/kg, extracting uranium from seawater (a gigantic resource) becomes economically viable.(Kazimi et al., 2011) Yet TFRs only take over significant portions of the decision space at far higher uranium costs, and costs above $1000/kg are not likely to ever be attainable (as other forms of energy would begin to substitute for uranium-fueled nuclear power). From Figure 5-23, we conclude that swings within the band of likely near-term uranium costs will not have a significant effect on the decision results. The reason for this is that uranium comprises only a small portion of the total fuel cycle costs.

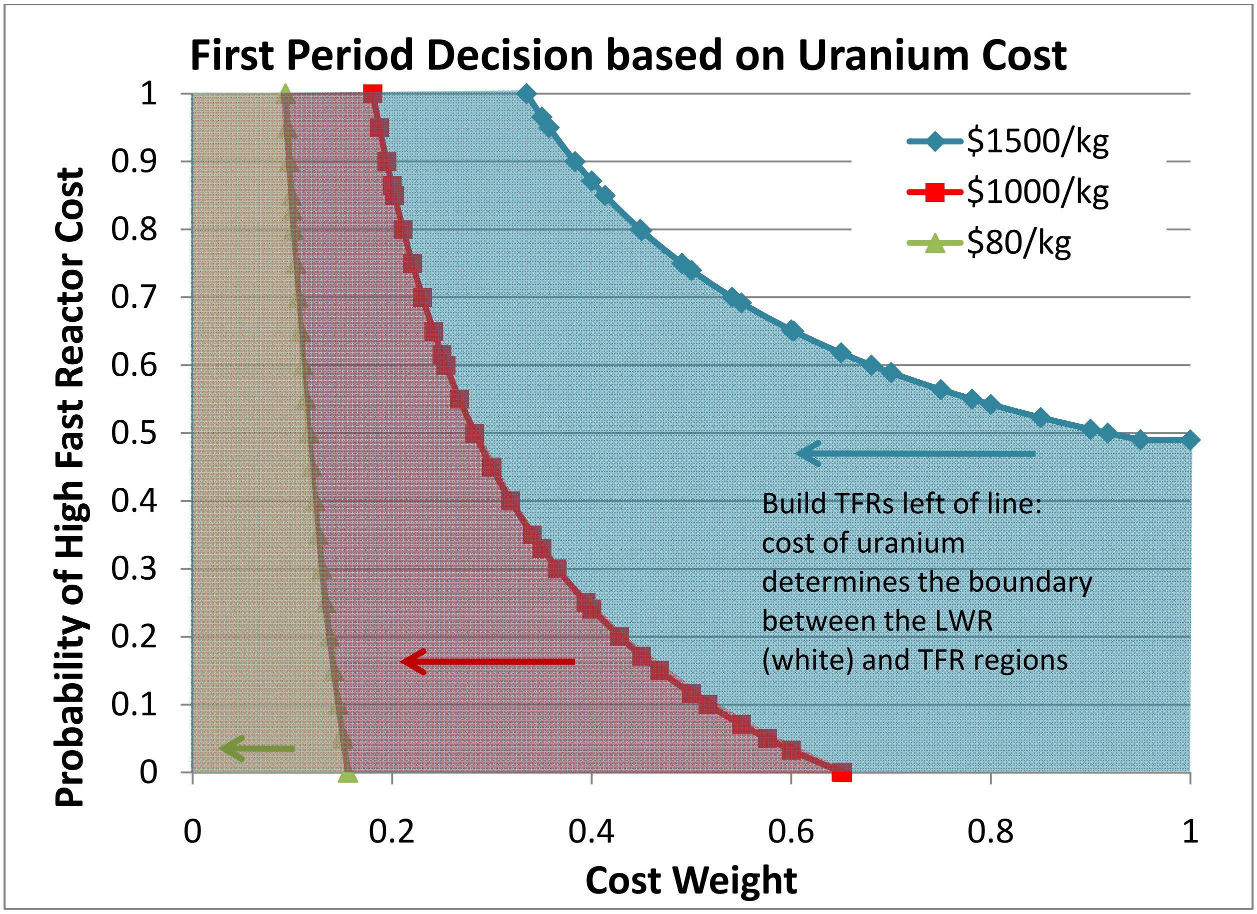

Figure 5-23: Effect of uranium price on desirability of traditional fast reactors

It is notable, however, that enriched uranium fast reactors do not appear as a desirable option. One might think they would, given that at high nuclear power growth, EUFRs eventually consume less uranium than TFRs (see Figure 5-24). But EUFRs are not desirable because for high growth scenarios, far more EUFRs are built than TFRs. The EUFR scenarios are thus much more expensive even though they use less uranium. Note, however, that the analysis did not consider building EUFRs alongside TFRs (future work could examine this option, and see if an "ideal" mix exists under a uranium constraint). These conclusions are valid for a value function that depends only on cost and waste management. If decision makers wish to consider security of fuel supply as an additional metric, EUFRs and U-saving TFR scenarios will become more attractive.

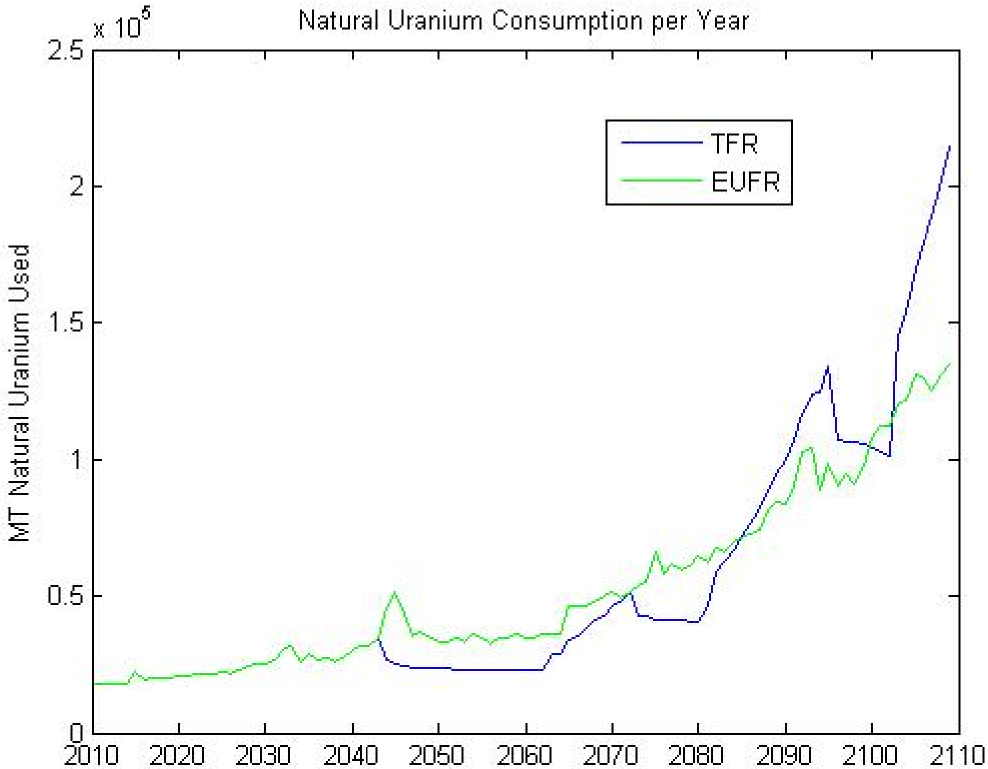

Figure 5-24: Natural uranium consumption by TFRs vs. EUFRs at very high growth

5.6 Key Takeaways from the analysis of Uranium Cost Impact
An increase in uranium cost increases the desirability of building traditional fast reactors, because fast reactors do not depend on uranium for feedstock. The increases in the decision space are not extremely significant, however, because uranium is such a small portion of the cost of the reactor system.

5.7 The Impact of Uncertainty in Reprocessing Efficiency

The separation of the constituents of used nuclear fuel, whether done by pyroprocessing or an aqueous process, will never be 100% efficient. Nuclear chemical engineers aim to develop a commercial-scale process that is as effective as possible, separating out high fractions of fission product elements for disposal while maintaining very high levels of TRU in the fuel stream. It is not yet clear what level of efficiency is attainable, but the DOE stated a goal in 2005 to “remove more than 99.5% of TRU from waste destined for geologic disposal.”(Shropshire et al., 2009)

The assumption used so far is that the separations efficiency for TRU is 99%. This means that when LWR or FR SNF is reprocessed, 99% of the TRU is recovered for fuel fabrication, while 1% is lost to a waste stream that includes the TRU and the fission products. This analysis also assumes that 0% of the fission products remain with the TRU fuel feed; in principle, this will not be the case. Both of these assumptions are the same as those made by VISION.

Changing the loss fraction may have an impact on the desirable fuel cycle decisions. If we are only able to recover 90% of the TRU from spent fuel, and the rest goes to disposal, closing the fuel cycle should become less attractive because we are doing less to reduce the most noxious waste burden. Similarly, if we can separate 99.9% of the TRU, closed scenarios should become more desirable. Figure 5-25 shows the simple three-decision two-period decision tree with the uncertainty in reprocessing losses modeled. We make a decision about what type of reactor/fuel cycle to employ, and then find out (if we build fast reactors and begin fuel recycling) what the maximum separations efficiency will be for the remainder of the simulation.

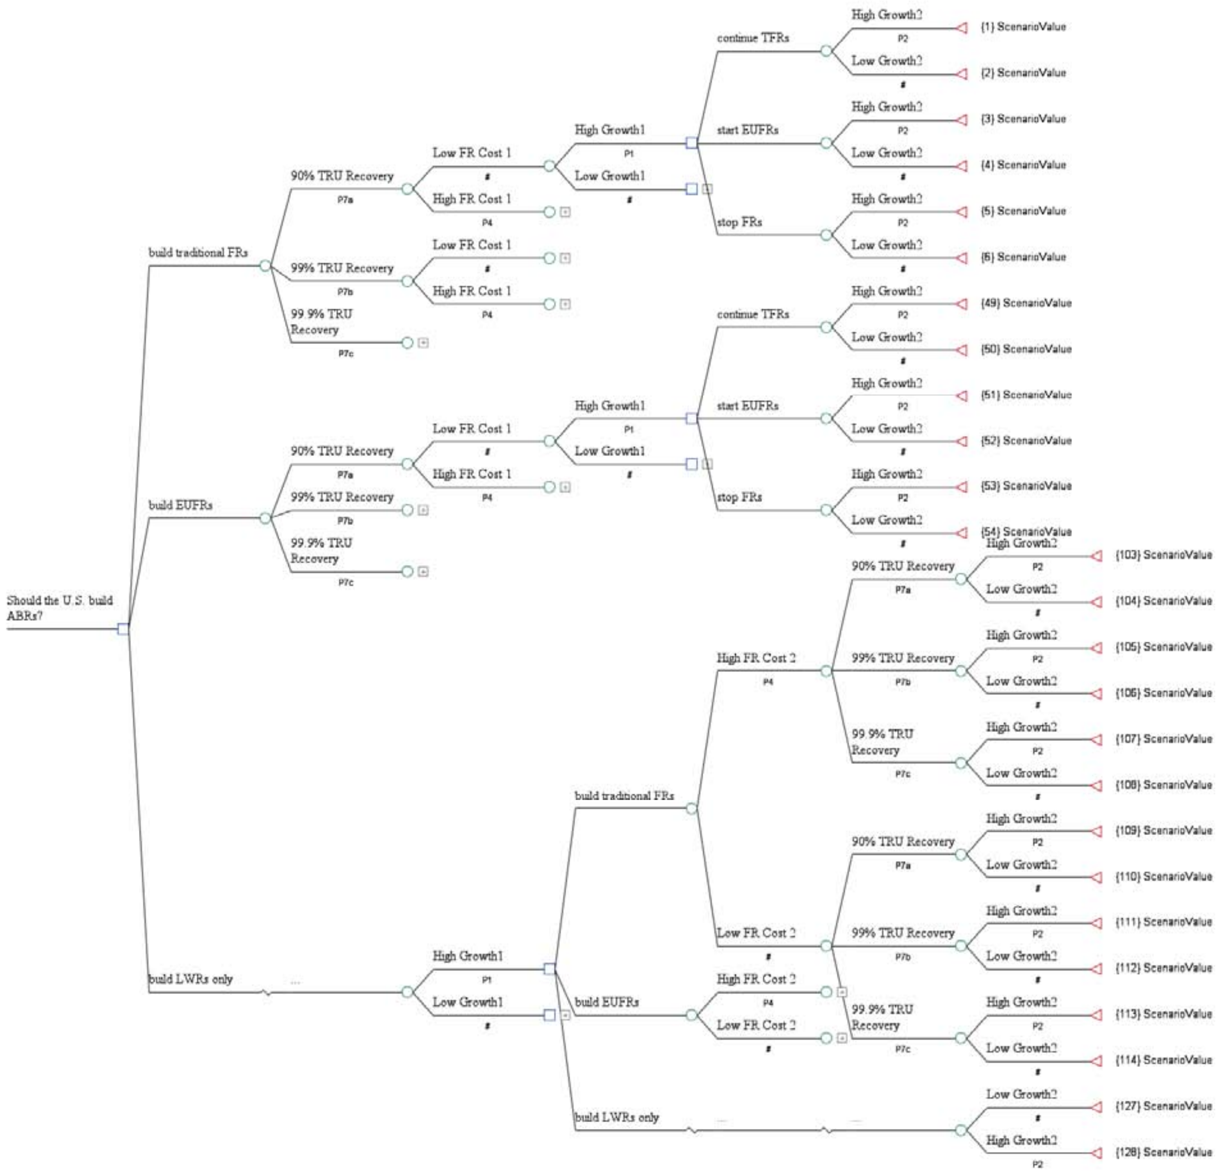

Figure 5-25: Decision tree with reprocessing loss uncertainty

The desirable decisions across a range of cost weights and growth probabilities are shown in Figure 5-26. For this graph, the probability of attaining each of the three possible loss fractions is set to 33%. The space over which TFRs are attractive (yellow) is larger for this formulation than for the 99% fraction analyzed in section 5.2. The yellow space grows because there are now more options for waste outcomes, and the 99.9% recovery rate entails so little TRU waste that the possibility it would happen pushes closed cycles to a slightly more desirable position.

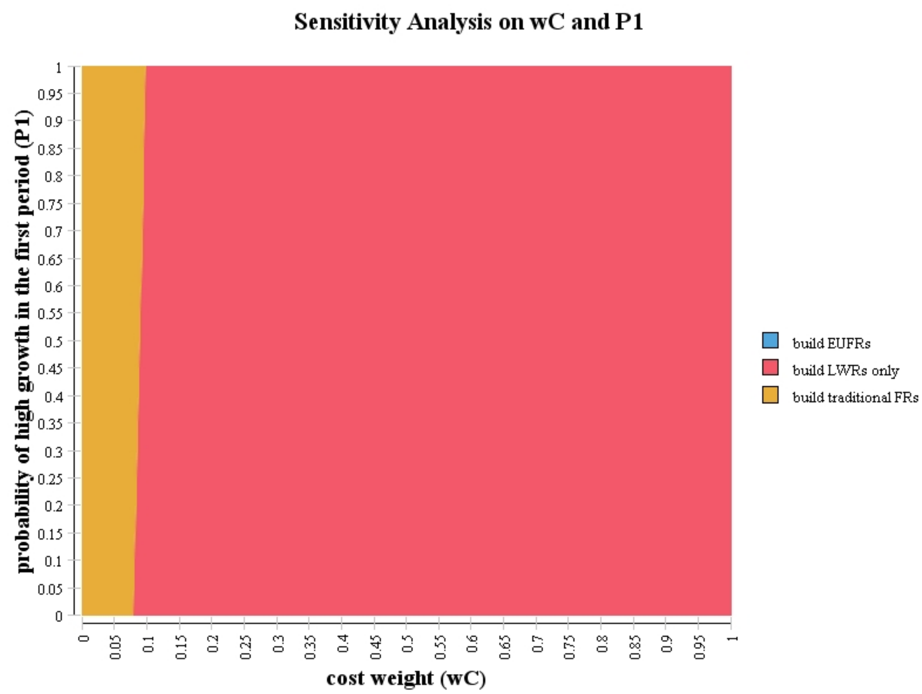

Figure 5-26: Desirable decisions given an even probability between three reprocessing loss fractions

Figure 5-27 and Figure 5-28 illustrate the effect of different probabilities for the various loss fractions. The nearly-straight line between the TFR and LWR regions shows that the decision is far more sensitive to the cost weight than to the loss probability. A rollback of the tree indicates that for TFRs, the expected values for the different loss branches are nearly identical. When the loss fraction is low, the amount of TRU going to waste is less, so we have a waste benefit. There is also a cost penalty, however, because now we must build more TFRs in order to handle the extra TRU in the fuel stream. The additional TFR cost counter-balances the benefit we get from separating more waste (though both changes are quite small: the differences are on the order of a few 10,000ths of utility points for both waste and cost). The fact that different loss fractions come with both costs and benefits, and the fact that the changes are so tiny over the utility scale, means that the achievable reprocessing loss fraction is not a very important parameter from the perspective of the decision maker.

This decision analysis, however, does not consider the impact of e.g. fission product impurities in the fuel, which are tied to the loss fractions and may have a significant impact on reactor performance. Furthermore, the particular concentrations of radionuclides in the product and waste streams may affect worker safety and the cost of protection measures at a commercial reprocessing plant. The System Analysis Campaign at Idaho National Lab continues to evaluate these and other losses-related questions, and may ultimately conclude that a certain loss fraction is highly desirable. A complete understanding of the losses issue will be necessary for an entity that builds a commercial plant. It may not be as important, however, for the highest-level decision makers to account for achievable loss levels in making broad strategic decisions about closing the fuel cycle. Depending on how significantly they impact safety and security attributes of fuel cycle systems, loss levels may not need to be considered at all when deciding the broad U.S. nuclear fuel cycle strategy.

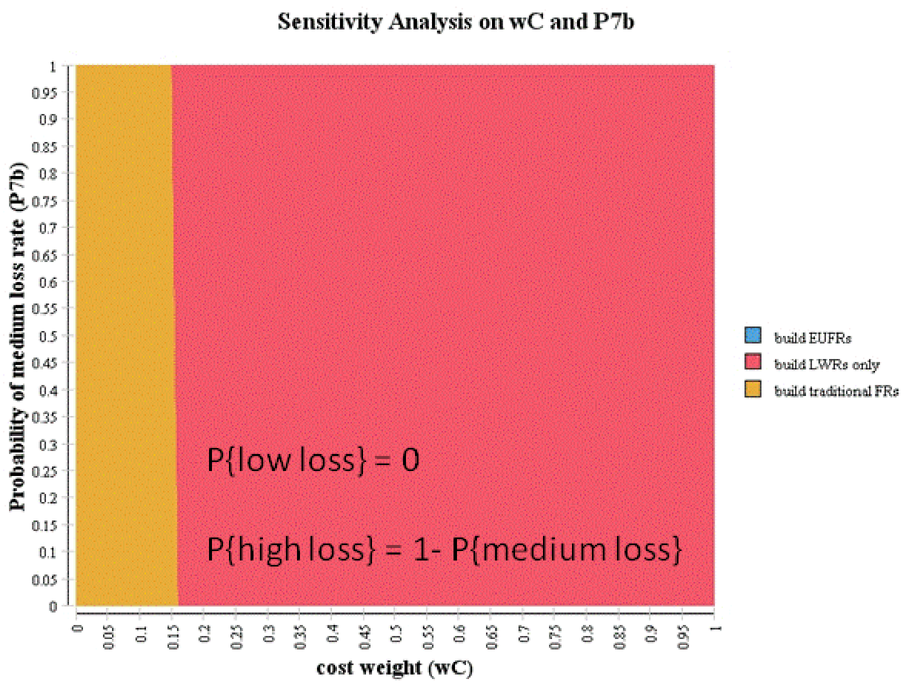

Figure 5-27: Desirability of TFRs when loss fraction is either low or medium

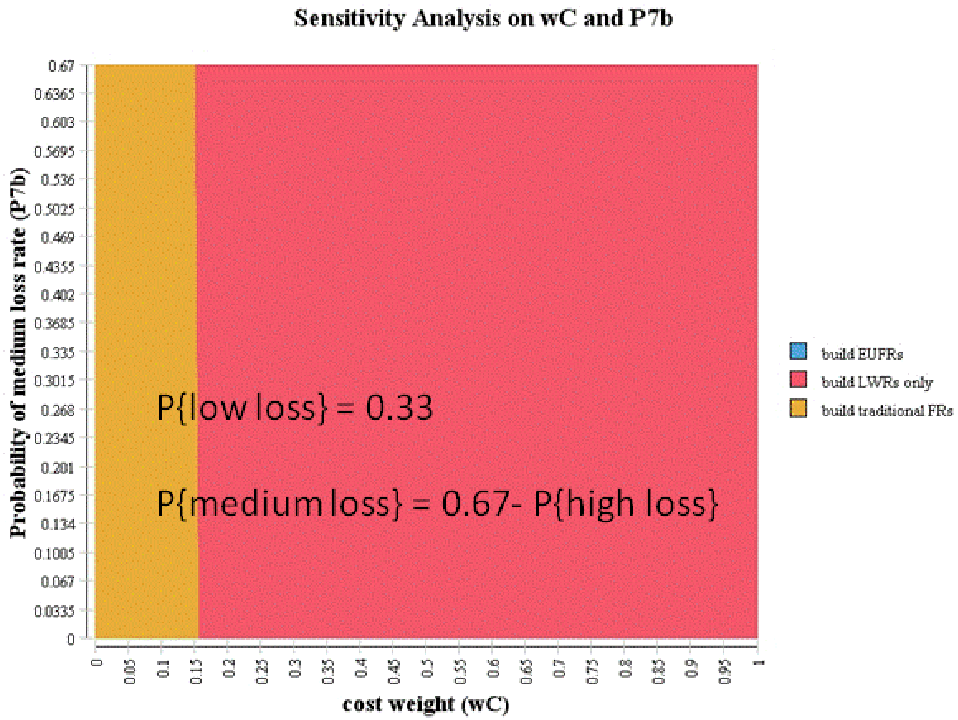

Figure 5-28: Desirability of TFRs when low loss fraction is set at a stable value

5.7 Key Takeaways from the Reprocessing Efficiency Analysis

The efficiency ultimately obtainable for nuclear fuel reprocessing does not have a significant impact on the choice between traditional self-sustaining, enriched-uranium startup, and light water reactors. In fact, improving the separation efficiency for TRU has the impact of increasing system costs because more traditional fast reactors will be required to handle the additional TRU. This analysis, however, does not account for the fuel performance benefits of higher efficiencies.

5.8 The Impact of Changing Preferences Over Time

The desires of decision makers and of society are not likely to be the same over the course of an entire century. Preferences will be shaped by the performance of nuclear power, by the availability of other energy technologies, and by broad attitudes on environmental protection. The results presented in the previous five sections are all shown for the full range of the primary preference tradeoff, and thus show which decisions are robust to large swings in opinion.

The calculations underlying the previous graphs, however, assume constant preferences throughout time. That is, each calculation of the optimal first-period decision assumes that the optimal second-period decision will be made given the same weight value. A natural question is whether we will get different decision results if we re-do the calculations, allowing the weights to change between the first and second period; in some sense, this might represent a more realistic decision scenario.

One major challenge with relaxing the assumption of constant preferences is that a standard method for performing a mixed-preferences calculation does not exist. For this study, two calculation regimes are proposed and compared, modeled after the time-dependent decision analyses of Hammond, Grout, and many others.(Grout, 1982; Hammond, 1976) A “myopic” or “naïve” calculation assumes that the decision maker does not consider the second period (or any periods thereafter); he or she simply makes a decision for the period at hand, and then makes remaining decisions after reaching later nodes. A “sophisticated” calculation starts at the last decision period, and calculates the best decision contingent on each possible first period decision being made. The tree is pared so that only the best second period decisions remain, and the first period decision is then made with the remaining information. The terms “myopic” and “sophisticated” do not, in fact, imply that one calculation mode is better than the other. In fact, (Grout, 1982) determines that for all but the simplest decision trees, myopic calculations can Pareto dominate the best sophisticated results.

The first myopic evaluation will determine the best first-period decision. Figure 5-29 shows the decision tree corresponding to the calculation. The fuel cycle simulation is run from 2010 to 2065, and the decision is made at 2040 based only on knowledge of what has come before and estimates of what will occur between 2040 and 2065 (rather than before, where decisions at the 2040 node considered the entire century). This amounts to a one-period decision under uncertainty. Cost and wastes for each scenario are tabulated at 2065.

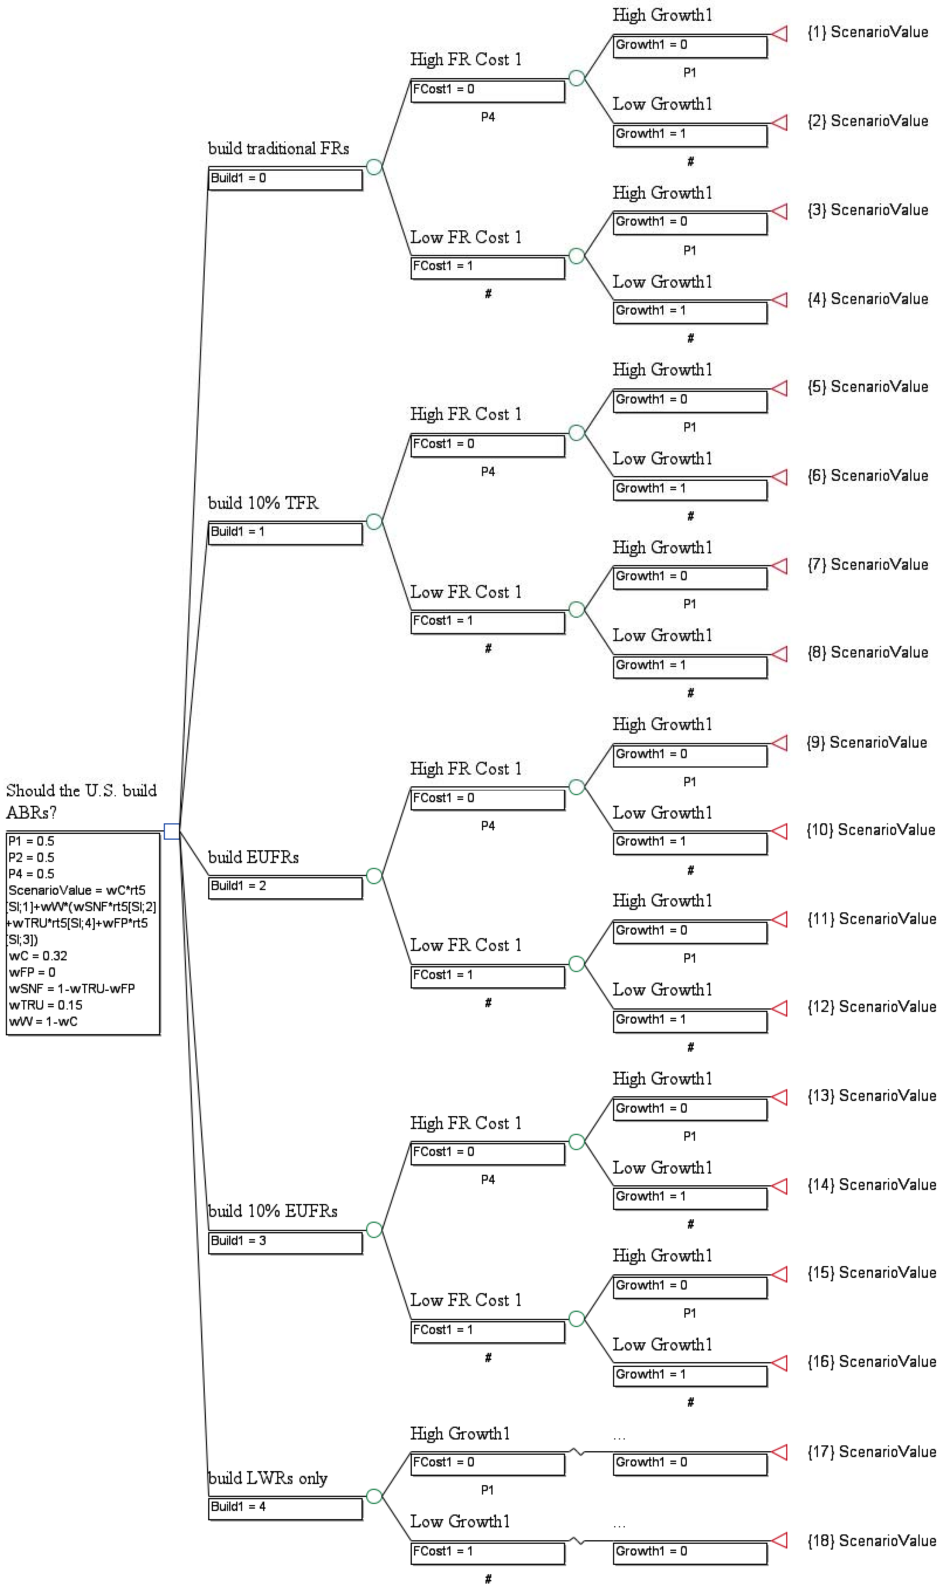

Figure 5-29: Decision tree for first period myopic calculation

The waste weights are set through the entire investigation to wSNF = 0.8, wTRU = 0.15, and wFP = 0.05. The cost weight changes: for the first-period calculation described above and depicted in Figure 5-29, the cost weight has a limited range between 0.5 and 1 (indicating a period in which society prefers reducing cost to minimizing waste). In the second period, the cost weight will range between 0 and 0.5 (indicating that in later decades, minimizing waste will become more important).

The second-period decision is calculated using the entire tree, but assumes that the only possible first-period decisions are those identified by the previous calculation. The results for the two possible first period decisions and the subsequent second period decisions, using the myopic view, are presented in Figure 5-30.

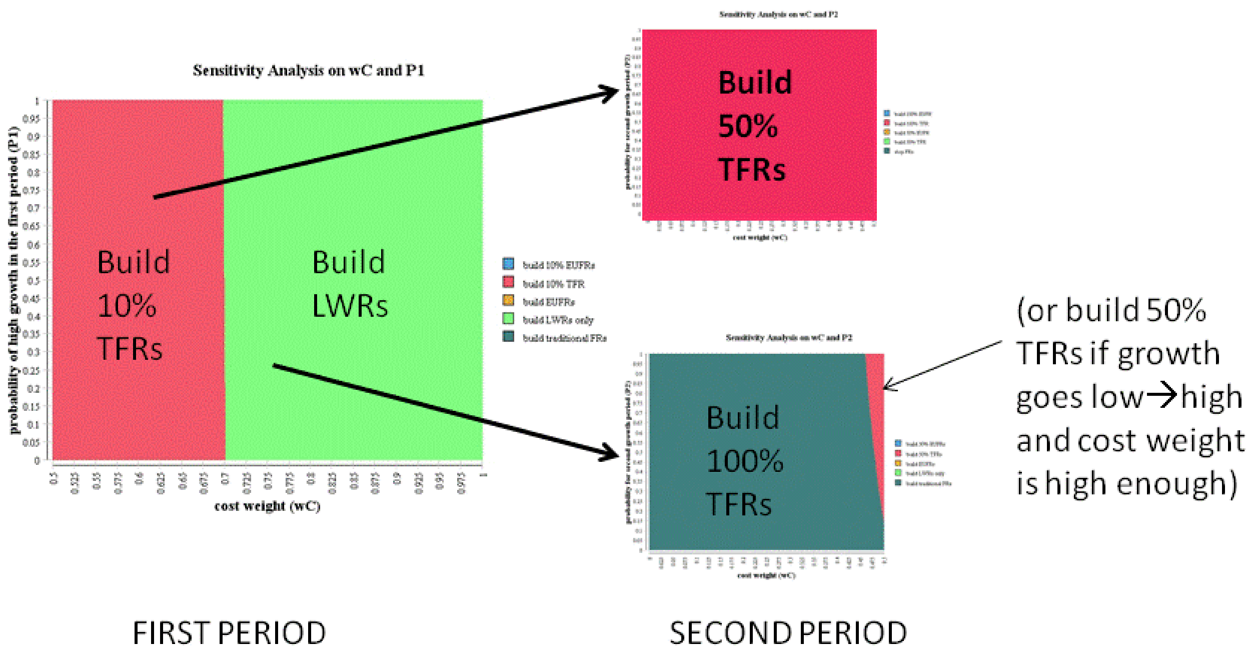

Figure 5-30: Results for myopic calculation of decisions with changing preferences

The myopic calculation suggests the same general suite of desirable decisions as the traditional analysis of the five-option tree with 10% FR builds. Again, low cost weight means that 10% TFRs are desirable in the first period, and that LWRs are preferred otherwise. Building LWRs first means that 100% TFRs are preferred in the second period, but a slower TFR build is preferable if 10% TFRs are chosen earlier in the century.

The primary difference between these and previous results is that 10% TFRs are preferred in the first period up to a much higher cost weight. Recall that in the first period and on the left graph, wC is constrained to [0.5, 1], so decision makers would need a cost weight above 0.7 to want to build LWRs only (vs. above about 0.2 in section 5.3). The reason for the difference is that for the myopic calculation, the decision maker does not need to consider any impacts (i.e. costs or wastes) after 2065. The decision maker also acts as though there is not going to be an opportunity to change course later, so to get any waste benefit, the results emphasize an early need to build TFRs.

The sophisticated decision calculation focuses on the second period first. The calculation begins with the traditional five-option tree (see Figure 5-10), including both periods and all uncertainties with fuel cycle simulation results tabulated at the end of the century. The weight sets are the same as in the myopic calculation, with waste weights constant throughout the evaluation, and high cost weights for the first period but low cost weights for the second period. The best decisions for the second period are identified using wC = [0, 0.5], for each possible decision that could be made in the first period.

The results are similar to those for the myopic second period, with 50% TFRs being desirable if TFRs are built early on, 100% EUFRs are desirable if EUFRs are built early on, and 100% TFRs are desirable if LWRs are the first-period decision. Any pathway that does NOT appear on a sensitivity graph (where wC is varied over the appropriate range) is pared. If a pathway does not appear on any sensitivity graph, this means it will never be desirable, regardless of the possible values of the uncertainty outcomes and decision variables. The remaining second period possibilities after paring are listed in Table 5-4: for each possible first period decision (listed on the left), the corresponding attainable second period decision(s) are identified.

Table 5-4: Remaining second period decisions after paring under sophisticated calculation

First Period Decision	Second Period Decisions
100% TFR	100% TFR, 50% TFR
10% TFR	50% TFR
100% EUFR	100% EUFR
10% EUFR	100% EUFR

LWR	100% TFR, 50% TFR

The next step is to solve for the best first period decision, given this pared tree (the tree is shown in Figure 5-31). The pared tree will be solved as a traditional two-period tree, but with a cost weight range restricted to [0.5, 1]. The first period results, which ultimately are of most import to the decision maker, are shown in Figure 5-32.

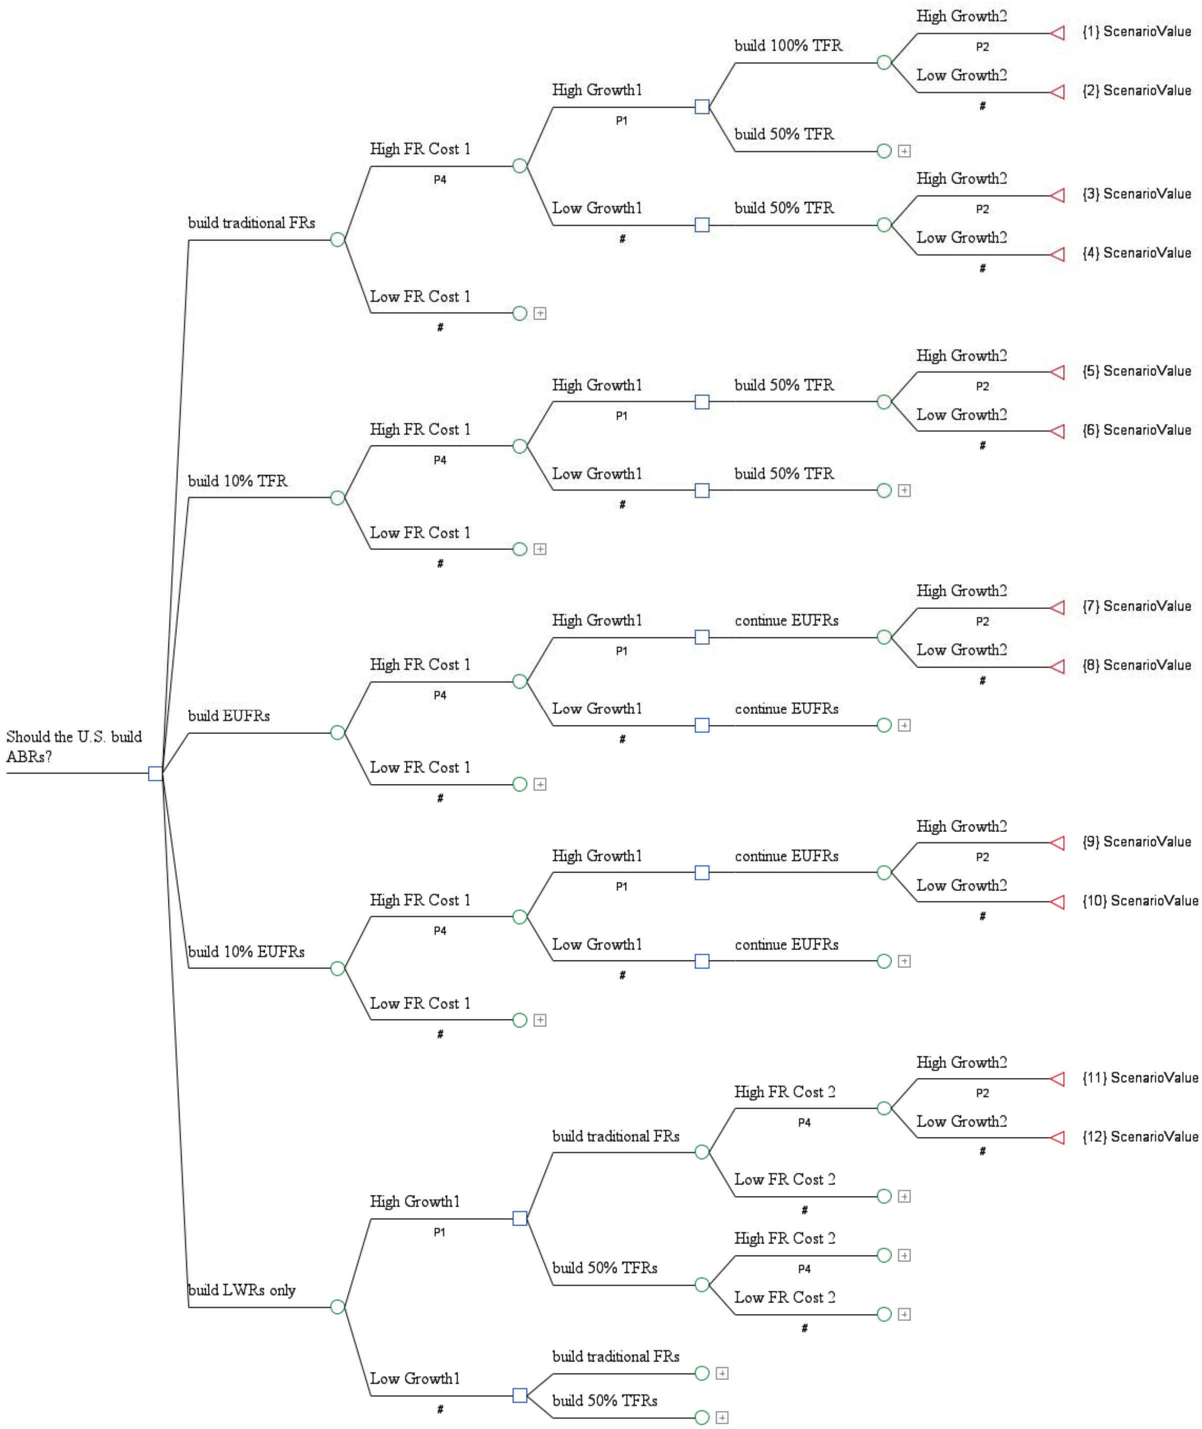

Figure 5-31: Pared decision tree for calculation of period one decision in sophisticated mode

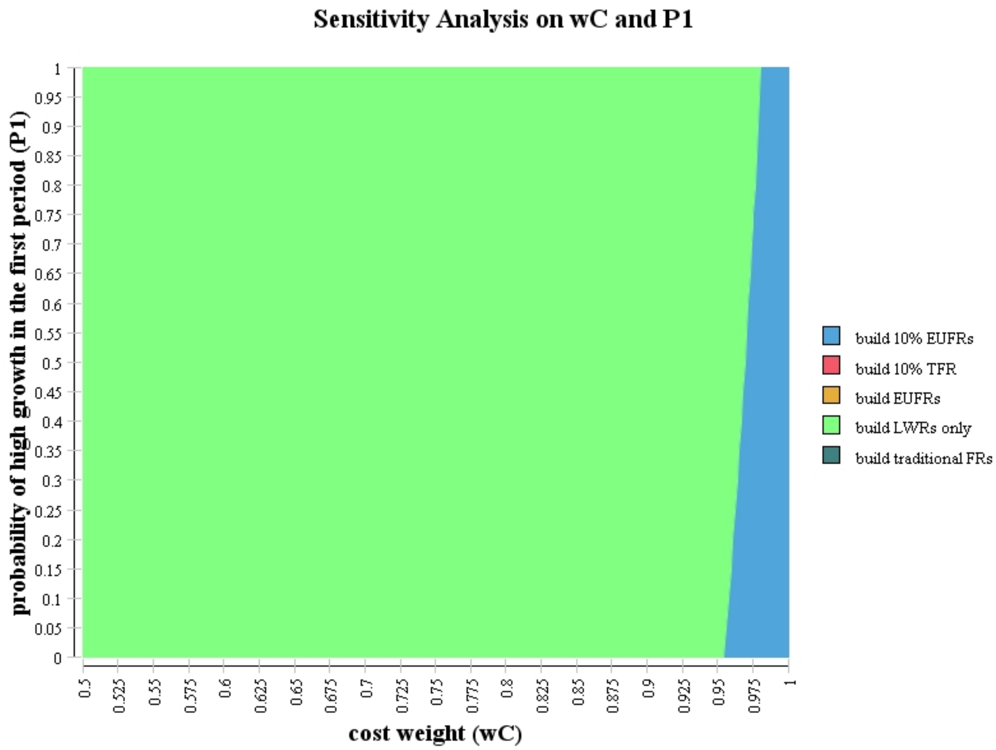

Figure 5-32: Decision results for the first period, sophisticated calculation

The results for the sophisticated calculation are quite surprising: at a very high cost weight, EUFRs are the most desirable course of action. EUFRs appear because, in fact, they are the cheapest of all remaining options in this formulation (and at wC near 1, cost is all that matters). Building TFRs first is always cheaper for the two-period formulation before the discount rate is applied, but TFRs require extensive reprocessing of LWR SNF early in the simulation. With discounting (at 7% or at 3%), the TFR calculation comes out just slightly more expensive as a result of this early and heavy cost. Note, however, that this is only true at low growth, when the number of EUFRs built is similar to the number of TFRs (otherwise one builds many more EUFRs than TFRs and the EUFR scenarios are more expensive – but the effect of full LWR takeover at high growth is only visible in the sensitivity P2 vs. wC because the second period is much longer).

EUFRs also manage to edge out LWRs at high cost weights: recall that we pared the branch “build LWRs” in the second period, so at high cost weight, building LWRs only through the whole century is no longer an option. We MUST either build 100% TFRs or 50% TFRs if we begin with LWRs only, and the EUFR-first option is solidly cheaper than LWRs → 50% TFRs (now because of the combination of effects that SNF reprocessing is costly, and building LWRs first allows a rapid buildup of TFRs because so much SNF is available).

We could conclude from the sophisticated calculation that if we are absolutely sure we will highly favor waste mitigation in the future, it may make sense to build EUFRs now. But this conclusion is not particularly useful, because it is impossible to be sure about our later preferences (especially for such a volatile issue as nuclear waste management).

The fact that the results from the sophisticated analysis are at odds with those of the myopic or unchanging preference calculations means we likely have a scenario where myopic solutions Pareto-dominate sophisticated ones. Neither the myopic nor sophisticated calculations are very realistic; it does not make sense to choose a pathway without regard for future decision opportunities, or to artificially constrain our options in this time period according to what we believe we will want later on. It thus seems likely that the regular formulation provides the best information for a range of decision preferences, even though the calculation is done assuming that those preferences stay constant. If the desirable decisions depended more heavily on the attribute weights (i.e. if a tiny change in the cost weight always produced a change in the desirable decision path), the effect of changing preferences might have to be explored more fully.

The analysis was repeated for the balance between the waste weights. The cost weight (and waste weight) varied from 0 to 1 for both periods, while different values of wSNF and wTRU were tested (wFP remained constant at 0.05). The results from this study were similar to the above: nothing differed substantially from the traditional formulation that was not an artifact of the extra constraints on the tree.

5.8 Key Takeaways from the Changing Preferences Analysis

Two methods of calculating best decisions under changing preferences are performed, one "myopic" or forward-looking, and one "sophisticated," starting at the end of the simulation time. The "sophisticated" calculation involves constraining the available decisions in the first period in a way that is probably not realistic. The "myopic" calculation produces results somewhat similar to those seen in section 5.3, with a larger area of optimality for fast reactors. The result shows that the time period over which the decision is considered has an impact on the general desirability of fast reactors, and the timeframe should thus be chosen carefully.

5.9 The Impact of Learning

So far, analyses of fuel cycle decisions have assumed that LWR and FR costs are constant over the entire simulation. While the methodology does model uncertainty in the relative FR cost, once this uncertainty is resolved, there are no further changes. We know the assumption of perfectly constant costs (for any type of nuclear reactor over time) to be inaccurate. Costs may change for any system in the nuclear fuel cycle, but for purposes of this discussion, capital costs are the primary focus because they are the strongest driver of nuclear electricity cost.

Capital costs for either reactor type could increase over time, depending on macroeconomic factors like commodity costs, or regulatory delays. At least one study, for example, has concluded that nuclear reactors costs actually *increased* in real terms over the main buildup of the French fleet, because of escalation in complexities of reactor design and construction processes.(Grubler, 2009) The U.S. fleet, with far less standardization than that present in the French fleet, has also surely seen a real escalation in nuclear costs.

It is possible, however, that for the large number of builds envisioned in many of these growth scenarios, the capital costs of plants will come down due to learning effects, and indeed other studies have found positive learning rates for the nuclear industry. The University of Chicago performed an extensive economic analysis of nuclear power, focusing mostly on LWR technology; among their investigations is a survey of literature on nuclear learning rates.(The University of Chicago, 2004) They conclude that a learning rate of 3-10% cost improvement with each doubling of units built is a fair estimate, and they provide conditions under which various points along that range are appropriate.

The actual learning rate for LWRs, and for less economically-understood FRs, is a subject of intense debate. For purposes of this study, estimates within the 3-10% range are applied to LWR and FR costs, in order to determine the basic effect on learning decision analysis results. A 3% learning rate is assumed for LWRs (intended to reflect the likelihood of near-term slow growth), while a 5% learning rate is assumed for FRs (because FRs are built later in the century and generally at a faster, more continuous pace, allowing for higher worker retention and better construction learning). The initial cost premiums of FRs are increased, because the 55% and 5% cost premiums were based on an assumption of an “nth of a kind” plant. Now, the first

FR will cost 100% more than a LWR at the high range, and 25% at the low range, but the learning curve will bring those costs down as more FRs are built.

Learning is implemented in FANTSY by tracking the number of each reactor type built, and calculating the cost adjustment factor (CAF) for year y by:

$$CAF_y = (1 - d)^{\ln(n)/\ln(2)}$$

where d is the learning rate, and n is the number of reactors built up to the beginning of year y (equation adapted from (The University of Chicago, 2004)). The CAF for year y is then multiplied by the initial capital cost of the reactor, and that cost is used for all reactors built in year y. For the particular learning rates assumed and at the highest-growth scenario, fast reactor costs reach 60% of their initial cost by the end of the simulation. This represents the most extensive cost reduction due to learning of all pathways in the analysis.

The new cost model is applied to the five-option decision tree that includes the option of building FRs or EUFRs at 10% of the allowed rate (tree depicted in Figure 5-10). The results for decision desirability across the range of cost weights and period 1 growth probabilities are shown in Figure 5-33. The dark purple region shows the growth and cost weight conditions under which building 10% TFRs was desirable before, given constant reactor costs (the assumptions of section 5.3). When learning is included as described above, the 10% TFR space expands out to the blue area.

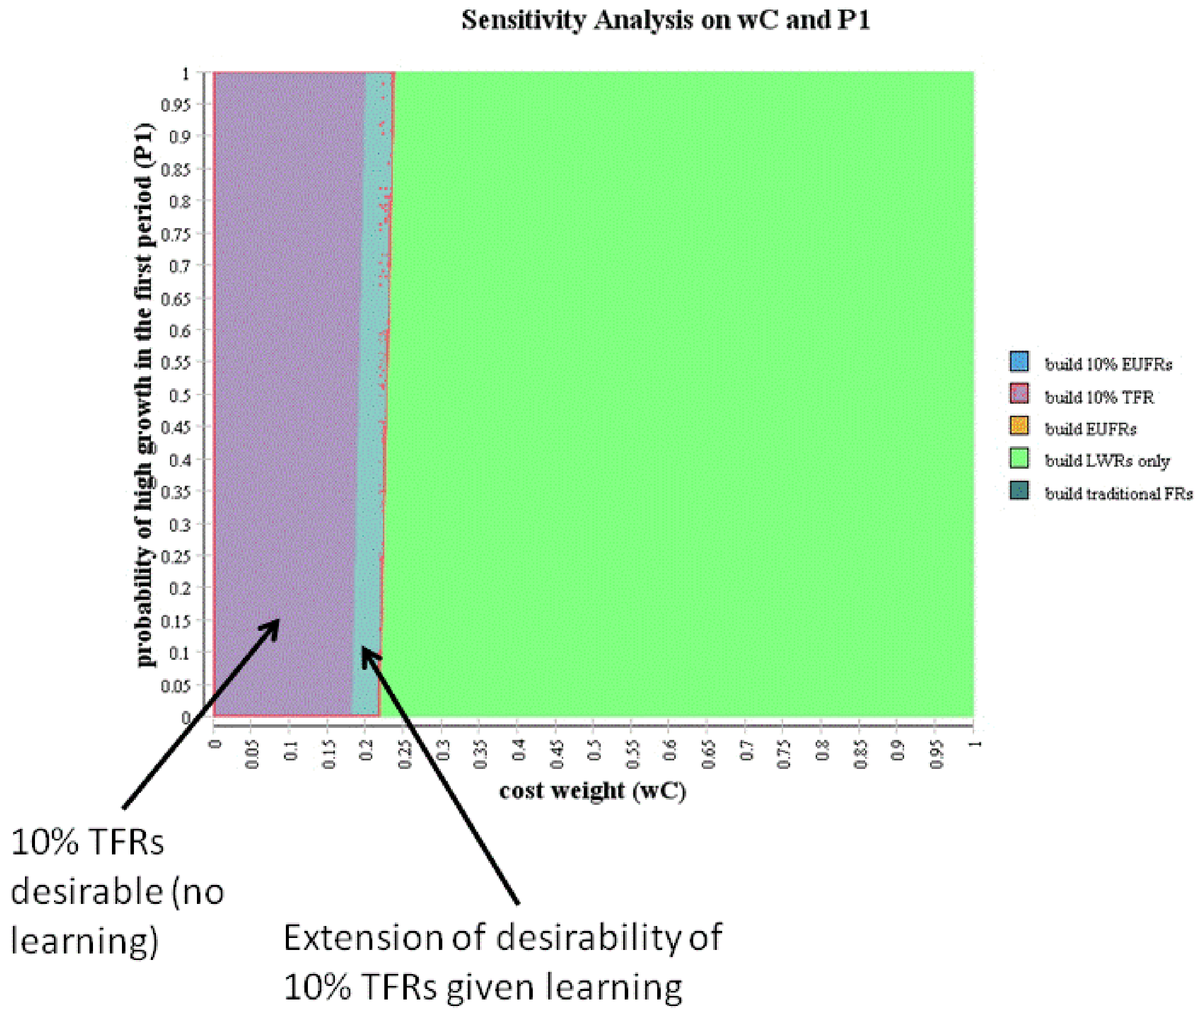

Figure 5-33: Increased desirability of building 10% TFRs given learning

We would expect TFRs to become more desirable in general if the learning curve brings the cost of fast reactors closer to the cost of LWRs. In fact, for this particular set of assumptions, the cost gap between FRs and LWRs is *greater* than it was in section 5.3, and for most decisions throughout the tree, FRs become less desirable than they were in previous studies. For the second period decisions that follow the best first period options, the decrease in FR desirability is small but evident.

Figure 5-34 shows the impact of learning on the first period decision for a value function incorporating the relative heat output of each type of waste (a more full explanation of the differences between value functions is presented in section 6.1). A comparison of the results for learning vs. no learning elucidates a seeming contradiction: the space favoring 100% TFR builds shrinks when learning is considered, but the space favoring 10% TFRs grows. As mentioned above, it makes sense that the space favorable to TFRs shrinks, because the particular assumptions about learning mean that FRs are especially expensive relative to LWRs throughout the simulation.

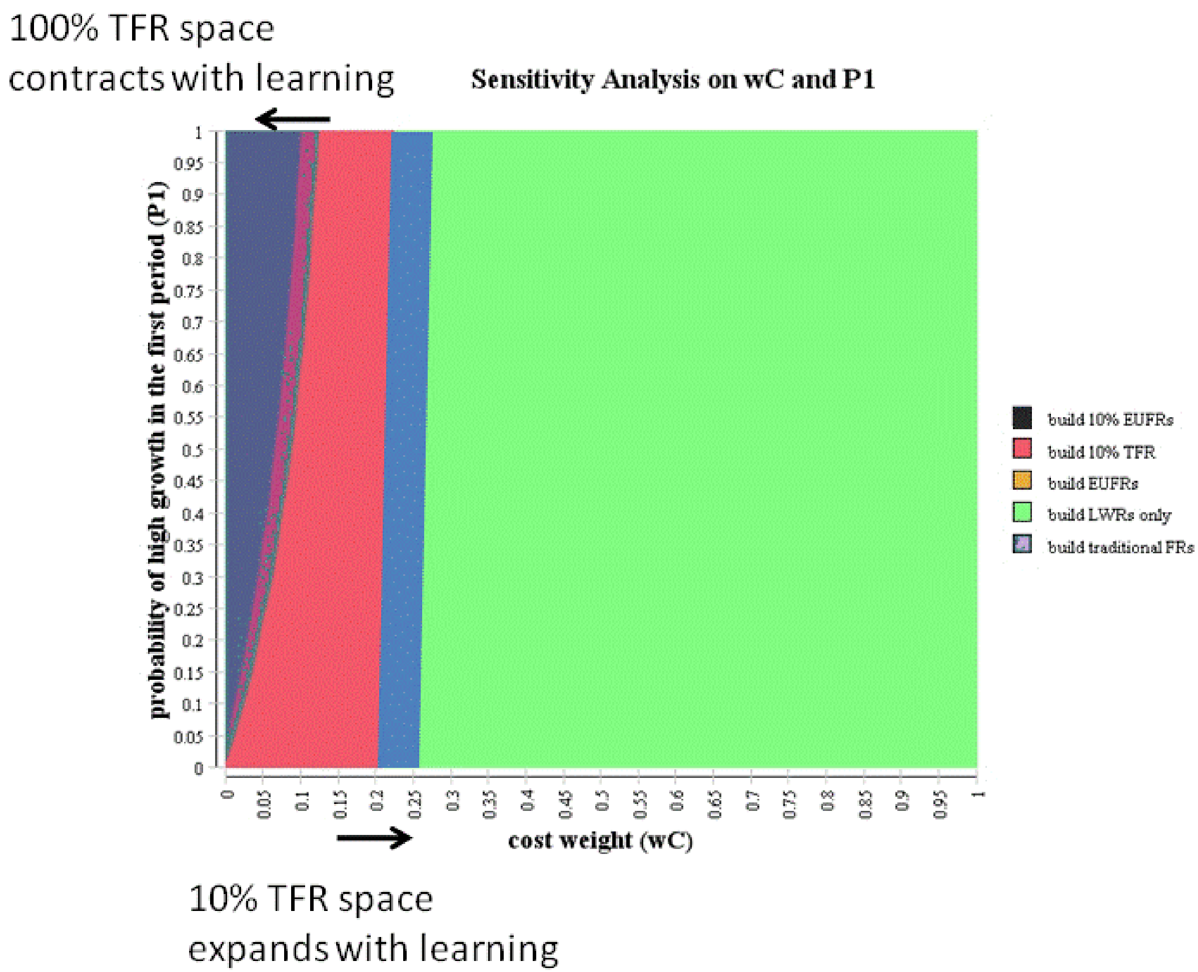

Figure 5-34: Effect of learning on the first-period decision, for heat metric value function

In fact, the 10% TFR decision responds differently to learning because of the particular pattern of FR construction in that scenario. Figure 5-35 shows the FR construction pattern for the three decision options represented in Figure 5-34 (all at high growth only). The three sets of data represent three different first period decisions, modeled along with the optimum second period decision. For both the LWR and 100% TFR scenarios, the optimal second period decision is to build TFRs at 100%. The graph shows that “clumped” construction results, whereby many reactors are built in short bursts. For the 10% TFR decision, the optimal second period choice is to build 50% TFRs, and the result includes slightly more total TFRs built than for the other two scenarios in addition to a much smoother build pattern.

The learning factor is calculated each year. This means that when FRs are built smoothly across many years, more reactors are built at lower cost. Rather than build 15 reactors in a single year at one point on the learning curve, those 15 reactors are stretched across several years and the cost decreases each year. Ultimately, this means that the 10% TFR scenario comes out with a slight cost advantage relative to its position when learning is not considered.

It is of note, however, that inclusion of a learning cost model overall has only a small impact. Even with dramatic assumptions about initial costs, learning rates, and growth rates, learning by doing has a tiny effect on the optimal decision space. These results do confirm, however, that sustained, smooth building is preferable to a “bunched” build profile in order to maximize the effects of learning. Further exploration of the relationship between the building curve shape and decision results is presented in section 6.5.

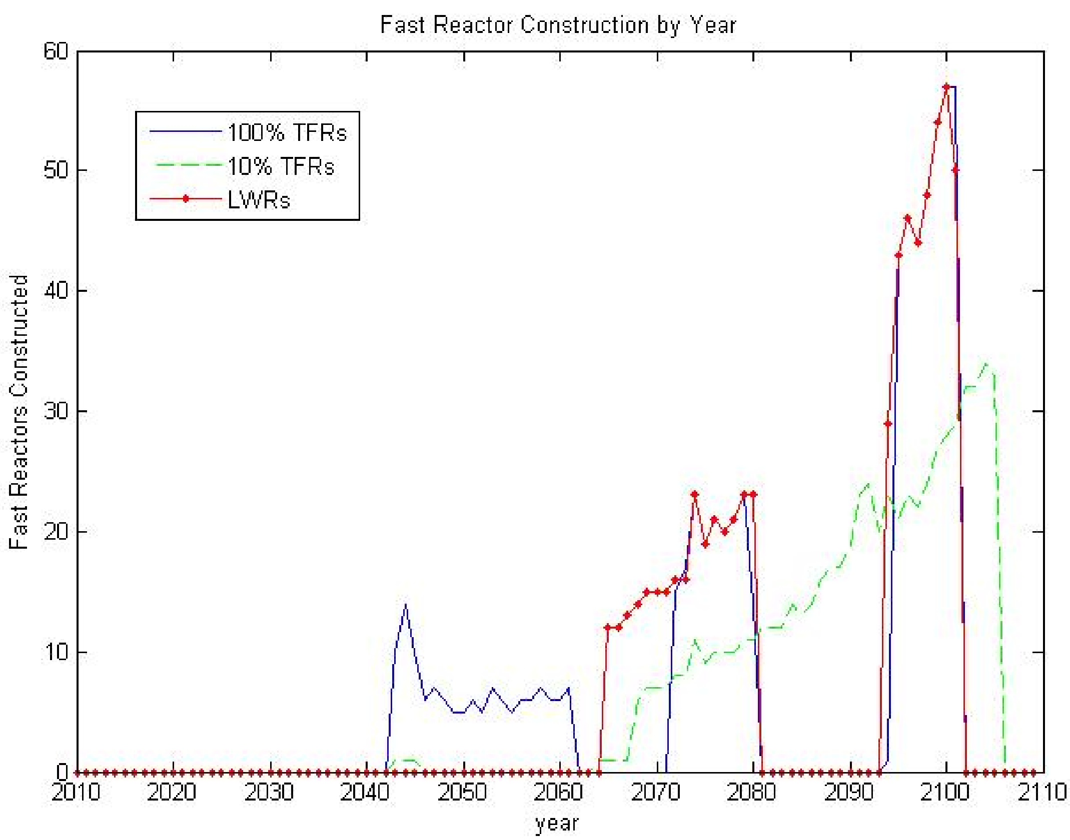

Figure 5-35: FR construction patterns for three scenarios

5.9 Key Takeaways from Learning Curve Analysis

Assuming that nuclear reactor costs will decrease over time due to learning effects has only a small effect on the decision space. The effect is to increase the overall desirability of building traditional fast reactors at 10% of the allowed pace and increasing the pace later, because this particular configuration allows for the smoothest building curve. Given that learning is likely to disseminate relatively slowly (under long construction times), spreading reactor builds out across more years allows better leverage of learning cost reductions.

Chapter 6: Sensitivity Analyses

In order to ensure that the results above are robust to some of the assumptions made, five sensitivity analyses are performed. Overall, the results are encouraging and demonstrate that the conclusions of Chapter 5 are qualitatively sound. Sensitivity analyses have not been performed for every possible contingency or parameter, however, and before using the results for decision making, decision makers should confirm that the most important assumptions have been addressed.

6.1 Sensitivity to Metric and Utility Function Structure

One criticism often leveled at decision analysis studies is that the mathematical construction of the metrics and utility functions may not accurately reflect desires of decision makers. For this reason, four versions of the waste metric and utility function structure were compared for each analysis presented in Chapter 5. Obviously, four value functions do not capture the full range of plausible structures, but they give a sense for how the results respond to variations of the assumed equations.

A full description of the metrics and utility functions is provided in Chapter 4 and summarized below in Table 6-1. Note that only the waste metrics and waste utilities change; the cost metric for each value function is total dollars, and the cost utility function is always a linear mapping of those dollar amounts.

Table 6-1: Value function structures tested for sensitivity of results

Value Function Label	Waste Metric	Utility Function
Value Function 1	Mass (kg)	Linear
Value Function 2	Repositories/Boreholes (No.)	Linear
Value Function 5	Repositories/Boreholes (No.)	Diminishing Returns
Value Function H1	Heat (W)	Diminishing Returns

Originally, twelve value functions were considered, each representing different combinations of the metrics and utility functions (including, at first, a diminishing returns utility function for cost). This explains the seeming arbitrariness of the function labels in Table 6-1: the functions listed were those retained throughout all analyses, while others were discarded because they consistently produced results nearly identical to at least one of the value functions retained. Three of the remaining four structures seem to give a good sense of the range of decision responses to different functions; VF2 is kept despite giving consistently similar results to VF1, in order to provide a check.

Figure 6-1 shows a sample of one-period decision results for each of the four utility functions. These are calculated using a one-period tree for which five options are available in the first period, corresponding to the same options available in section 5.3 (including 100% TFRs, 10% TFRs, 100% EUFRs, 10% EUFRs, or LWRs). For illustrative purposes, only the relationship between P1 and wC is graphed (but other parameters were tested and gave similar qualitative results in terms of the differences between value functions).

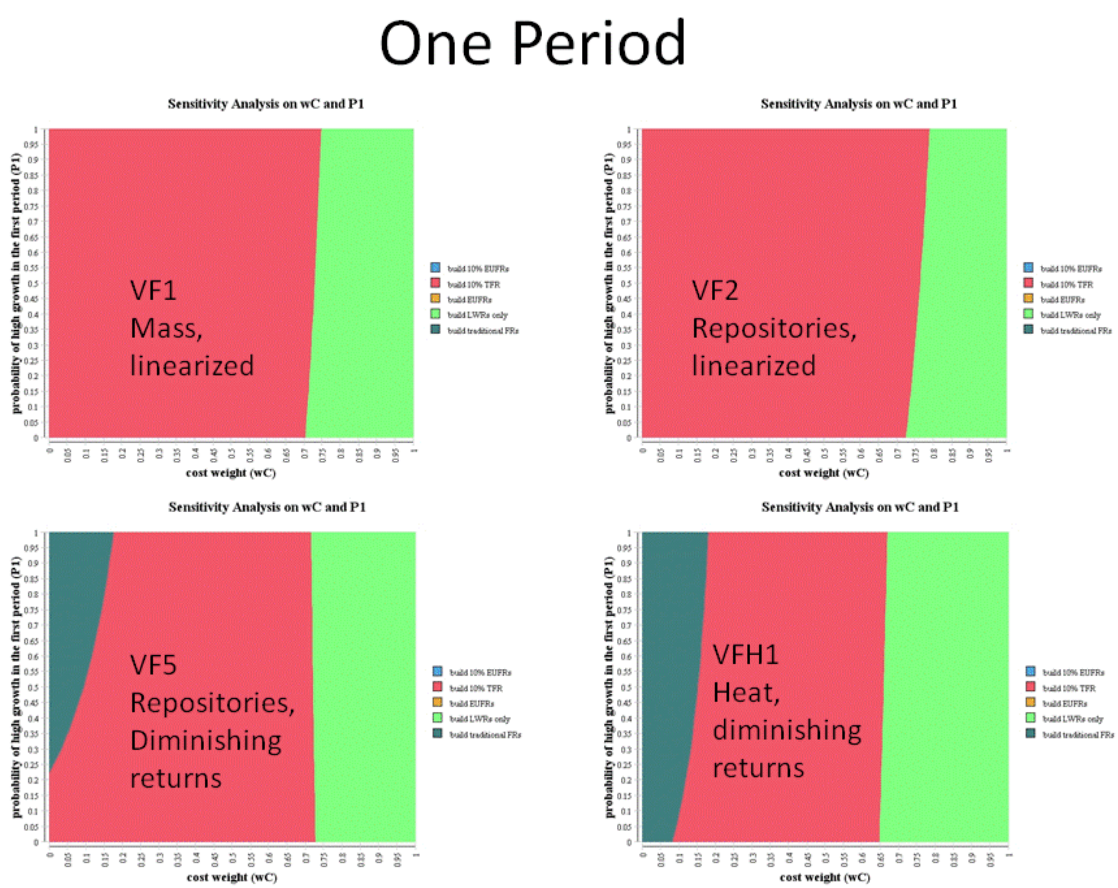

Figure 6-1: Comparison of utility functions for the one-period tree

The one-period results show some variation between value functions, especially between linear and diminishing returns utility functions. The first two value functions, VF1 and VF2, give essentially identical results. This makes sense because the only difference between the two is a conversion of the waste metric unit: for VF2, instead of considering the mass of each waste type, the number of repositories needed to contain the SNF (based on a 100,000 MTHM repository capacity) and number of boreholes needed (based on 6.5 MT TRU per borehole) is compared. The only material difference between the functions is that for VF2, the fission products are assumed to be commingled with the TRU, and because TRU is the limiting factor for borehole loading, the effect of fission products is completely and explicitly removed. Already with VF1, however, the fission products had effectively no bearing on the results because wFP is generally kept low.

VF5 uses the same repository/borehole metric as VF2, but uses a diminishing returns function to express the utility. The decision results are similar at the higher ranges of the cost weight, but there is a reversal of the nearly-vertical slope between the decision regions in VF5 compared to VF1 and VF2. This reversal happens because the diminishing returns function increases the waste disutility of all but one scenario: building LWRs at high growth. For the linearized disutility functions, the LWR/high growth scenario had an SNF utility of 1 because it produced the most of that type of waste. For the diminishing returns function, however, the disutility function effectively asymptotes to 1 in that range, so the LWR/high growth disutility is less than one. At high growth, LWRs therefore become more desirable compared to TFRs than they were in the first two utility functions.

The second major difference between utility functions is the appearance of the 100% TFR decision for VF5 and heat-based VFH1. 100% TFRs are able to appear on the bottom graphs, beating out 10% TFRs for low cost weights. This happens because the diminishing returns function increases the 10% TFR disutility compared to a linear function, while 100% TFR disutilities remain low regardless of which utility function is used. Figure 6-2 illustrates this: the clusters of points at the lower left represent the linear (diamond) and diminishing return (square) disutilities for the 100% TFR situation. A diminishing returns function only increases the 100% TFR disutility by a small or negligible amount. By contrast, the 10% TFR points in the middle show that there can be a larger increase in disutility when a diminishing returns utility function is applied in place of a linear one.

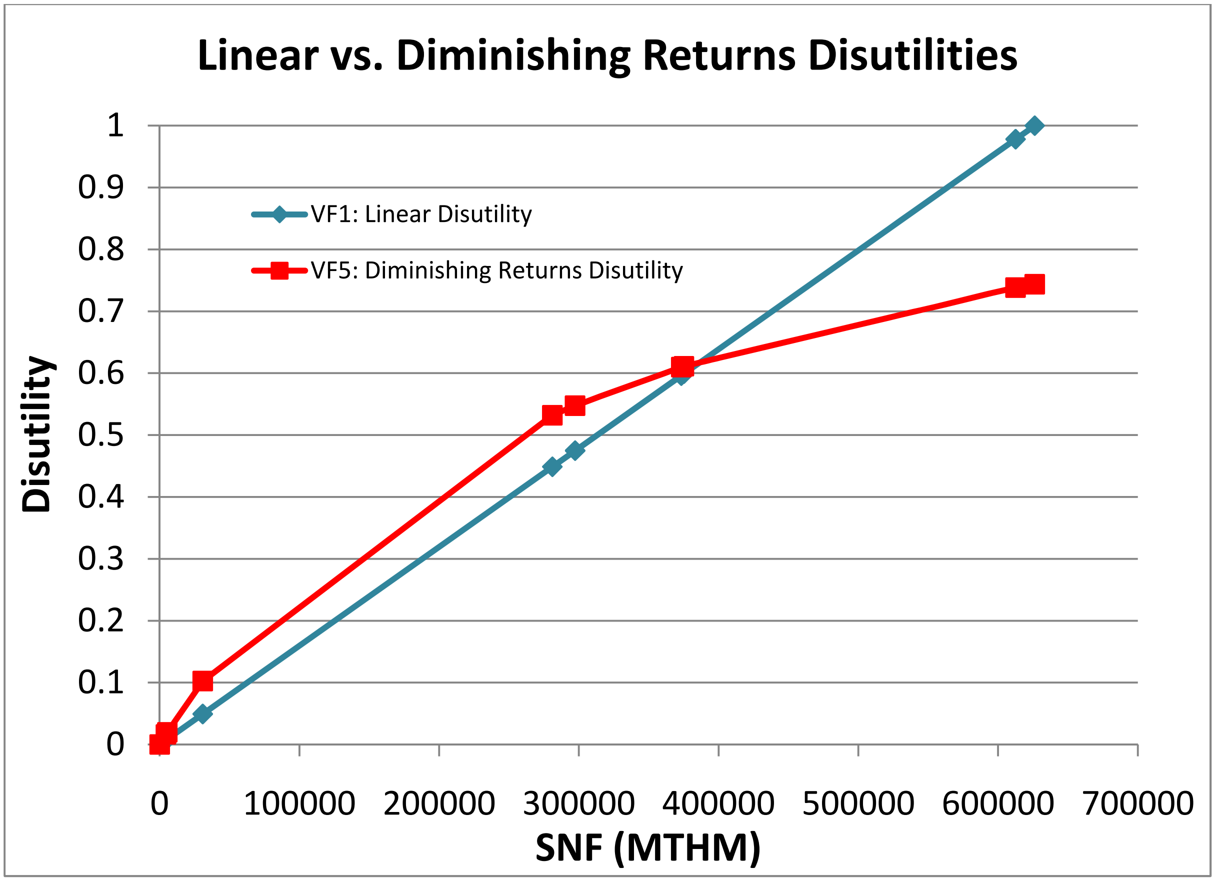

Figure 6-2: SNF disutility for linear and diminishing returns functions -- 100% TFR points are clustered at the lower left, 10% TFRs in the middle, and LWRs in the upper right

Finally, the last function is slightly different in that it has the most space where 100% TFRs are desirable: this is due to the particulars of the heat metric. The waste metric for VFH1 is obtained by calculating the heat output of each waste type, at an assumed emplacement time of 100 years for SNF and FP and 5 years for TRU. The SNF always has a much greater heat rate for the scenarios considered, so the effect is that the SNF value within the function receives extra weighting. As a result, 100% TFRs are highly desirable at very low cost weight.

The utility function test is next applied to the five-option, two-period version of the decision tree. Compared to the one-period sensitivity analysis, the primary difference is the expected shrinking of the space where 10% or 100% TFRs are desirable. Figure 6-3 shows the cost weight vs. first-period growth probability graphs for the standard two-period tree discussed in section 5.3, and, as explained in that section, the reduction of TFR desirability in the first period happens because TFRs can be built later in the simulation. These graphs, therefore, confirm that regardless of the utility function, the qualitative conclusion that waiting to build TFRs still holds.

Figure 6-3: Comparison of utility functions for the two- period tree

As with the one-period tree, the first two utility functions produce nearly identical results. The only departure from the one-period analysis is that the slopes between all decisions are positive for each version of the utility function, indicating that for this tree, high growth is always more favorable toward fast reactors than low growth. This happens because at the boundary between LWRs and 10% TFRs, the best second-period decision is to build 100% TFRs. For the one period scenario, the decision “build LWRs” meant locking in LWRs-only for the entire simulation. Now, the LWR decision (at a medium cost weight) still allows for 100% TFRs later on, and this pathway to 100% TFRs is optimal. This means that these LWR scenarios include an SNF reduction such that they fall into a mid-range category along the x-axis of Figure

6-2, and thus increase their disutility substantially when a diminishing returns function is applied. This increase in disutility is commensurate with the increase for building 10% TFRs first; LWRs gain no new advantage by applying a disutility function.

Figure 6-4 shows three-period results according to various utility functions. The set of graphs confirms the results of section 5.8, with qualitatively similar patterns of decision desirability across the various functions. As in Chapter 5, the three period results exhibit greater desirability for TFRs at low growth, in a departure from the two-period results. This happens because the three-period analysis considered possible stagnation and very low growth of nuclear power, and a probability for no builds in later periods (equivalent to low P1) means TFRs are more desirable now, in order to ensure at least some are built to provide waste benefits.

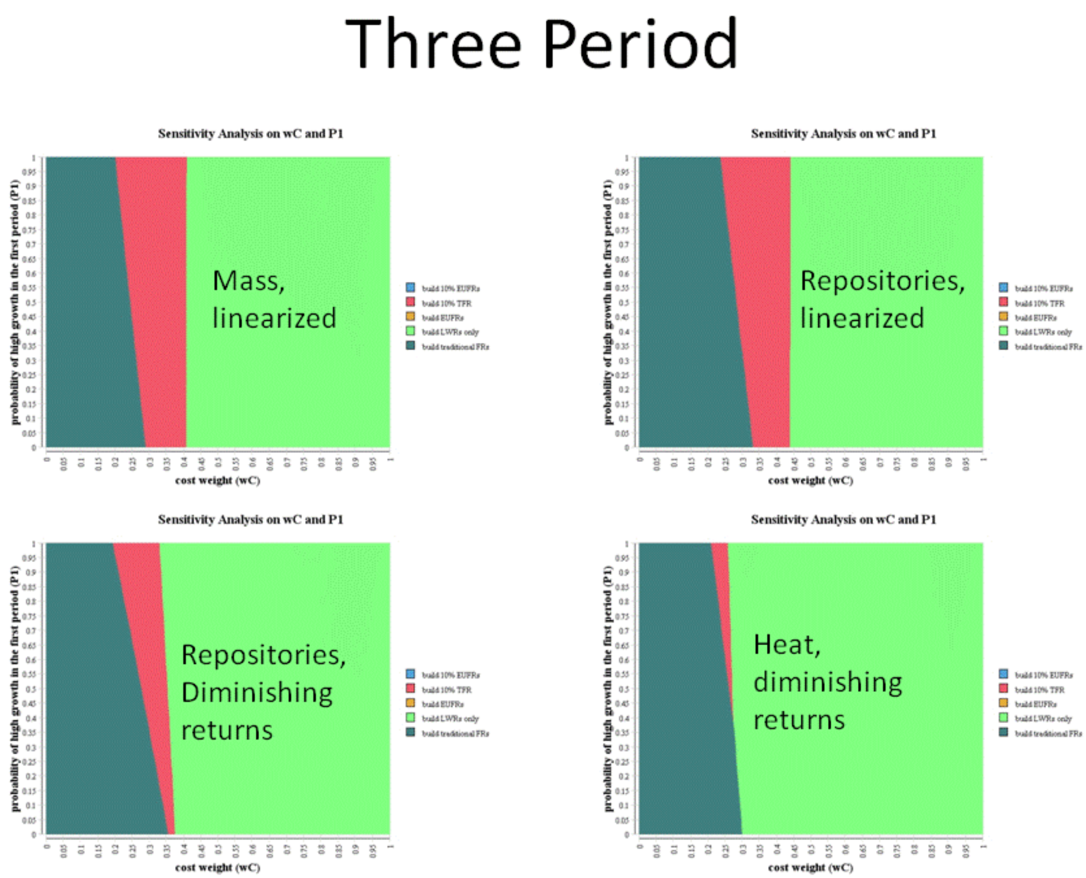

Figure 6-4: Comparison of utility functions for the three-period tree

The most important message from the sensitivity analysis of the value function is that changes in the decision space due to variations in value function structure are small. The general pattern of desirable decisions, and thus the conclusions in Chapter 5, hold true regardless of the function applied. Of the small differences, the most significant arise when a new shape of utility function is applied to the waste metric. In general, the diminishing returns function places greater desirability on scenarios that can bring the SNF stockpile close to zero; if highly exact decision results are required, further work should examine which function (or whether a completely different function) most closely approximates public desires surrounding waste. Metric definitions should also be considered carefully (and others tested), because the heat-based metric of VFH1 showed that certain formulations can place a natural weighting on a specific waste type.

If, as in this thesis, qualitative insights about decision problems are more valuable than precise quantitative decision outcomes, exploring further variations of this set of metrics and value functions is probably not germane. One potential change to the value function, which more fundamentally alters its structure, is examined in section 6.2.

6.1 Key Takeaways from Sensitivity Analysis of Various Waste Metrics and Utility Functions

Changing the structure of the waste metric and/or utility function does not change the set of basic qualitative conclusions presented in Chapter 5. The biggest change in results occurs when the utility function changes from a linear mapping to a "diminishing returns" function over the amount of waste. The diminishing returns function emphasizes the desirability of waste mitigation, especially at low waste amounts, so rapidly building TFRs becomes more desirable.

6.2 Sensitivity to an Interaction Term in the Value Function

Multi-attribute methodologies require analysts to assume (or to define) each attribute to be independent of the others. The metrics capturing waste and cost attributes in Chapter 4 were defined as much as possible to preserve independence. The cost metric, for example, does not include the price of waste management and disposal: this is assumed to be part of the "social" cost which is captured in non-dollar denominations by the waste metric. Yet the assumption of complete independence between these two attributes is not perfectly accurate. For example, the cost of the nuclear fuel cycle system may affect public opinion of nuclear power overall, which in turn will affect the social disutility for nuclear waste.

Clemen (1996) suggests several ways to adjust value functions so that they account for the non-independence of attributes. One option, which is relatively straightforward for two-attribute functions, is to define a utility surface over the two-dimensional space of the attributes. The advantage of this method is that it does not require any assumptions whatsoever about attribute independence. A significant disadvantage, however, is that properly eliciting the structure of the utility surface is computationally intensive, requiring a decision maker to establish preference-indifference curves by comparing a large number of attribute lotteries.

If the two attributes are determined to be mutually utility independent, attribute utility functions are separable and the value function can be expressed as:

$$U(x,y) = k_x U_x(x) + k_y U_y(y) + \left(1 - k_x - k_y\right) U_x(x) U_y(y)$$

where U(x,y) is the value of having the attributes at level x and level y, and k_x and k_y are weighting constants. Determining the appropriate utility equations for a separated utility function is simpler than determining a joint utility surface.

Attributes are mutually utility independent if a decision maker determines that his or her preferences for uncertain choices involving different levels of Y are independent of the value of X, and vice versa. A thought experiment helps to determine whether cost and waste are utility independent in the context of the U.S. nuclear fuel cycle. If cost is utility independent of waste, the decision maker should have a certainty equivalent for e.g. a lottery where there is a 50% chance the fuel cycle will produce 2 repositories worth of waste over the century, and a 50% chance the cycle will produce enough to fill only one repository: this certainty equivalent should not depend on whether the fuel cycle is expensive or cheap. For some decision makers, this will be true. The decision maker might prefer the certainty of 1.5 repositories' worth of waste to the lottery above, regardless of the nuclear system cost. Similarly, a lottery over costs does not necessarily depend on the amount of waste the system generates: lower cost is always preferred. Political decision makers, however, might be more risk averse about waste if nuclear system costs are high: justifying the need for a new repository site might be made more difficult if nuclear power is expensive compared to other sources of electricity. Further work should explore the true utility independence of these attributes for different decision makers, and could explore whether decisions change substantially given different utility surfaces. Here, mutual utility independence is assumed (and indeed may hold for some decision makers) in order to see the effect of adding one layer of complexity to the value function.

The waste metric is calculated as before, using an additive multi-attribute function for the waste types. The waste utility function is still applied to each type individually, such that

$$WasteUtility_i = wSNF * U_{SNF}(SNF_i) + wTRU * U_{TRU}(TRU_i) + wFP * U_{FP}(FP_i)$$

The overall value function will now be:

$$ScenarioValue_i = wC * U_c(cost_i) + wW * [WasteUtility_i] + (1 - wC - wW) * U_c(cost_i) * [WasteUtility_i]$$

Note that we no longer require that wC + wW = 1 (and indeed, if wC and wW satisfy this condition, the equation simplifies to the original form with no interaction term). If the coefficient of the interaction term is positive, the two attributes are considered to be complements (such that increasing the disutility of the attributes entails an extra penalty: having high cost and a large amount of waste, in a sense, is thus highly catastrophic). A negative coefficient, however, is more likely. A negative coefficient is reasonable if the attributes are substitutes, such that higher cost-scenarios are more tolerable if they produce less waste. The "true" weights could be elicited formally from a decision maker, in order to determine which characterization of the attributes makes more sense. TreeAge, however, allows exploration of an extensive range of the weights, so that we do not need to assume either case and can instead present decision results as a sensitivity to a range of weights.

The figures below show that we receive very little new insight by adding an interaction term to the value function. The decision result depicted in Figure 6-5 is calculated for the five-option tree, which allows the decision maker to build TFRs or EUFRs at 10% of their allowed amount in the first period (compare to the results of section 5.3). All probabilities for growth and cost are set to 0.5, and wC and wW are each varied from 0 to 1. This means that in the upper-right region of the graph, the attributes act as complements, while in the lower-left, they act as substitutes.

The results are fairly intuitive: if cost weights are low and waste weights are high, fast reactors are desirable. The graph does show that for this range of cost and waste weights, the decision is more sensitive to the cost weight than to the waste weight. One possible explanation relates to an insight from chapter five: waiting to build TFRs significantly reduces cost, because of discounting, whereas strong waste benefits can still be gained even if TFRs are built later.

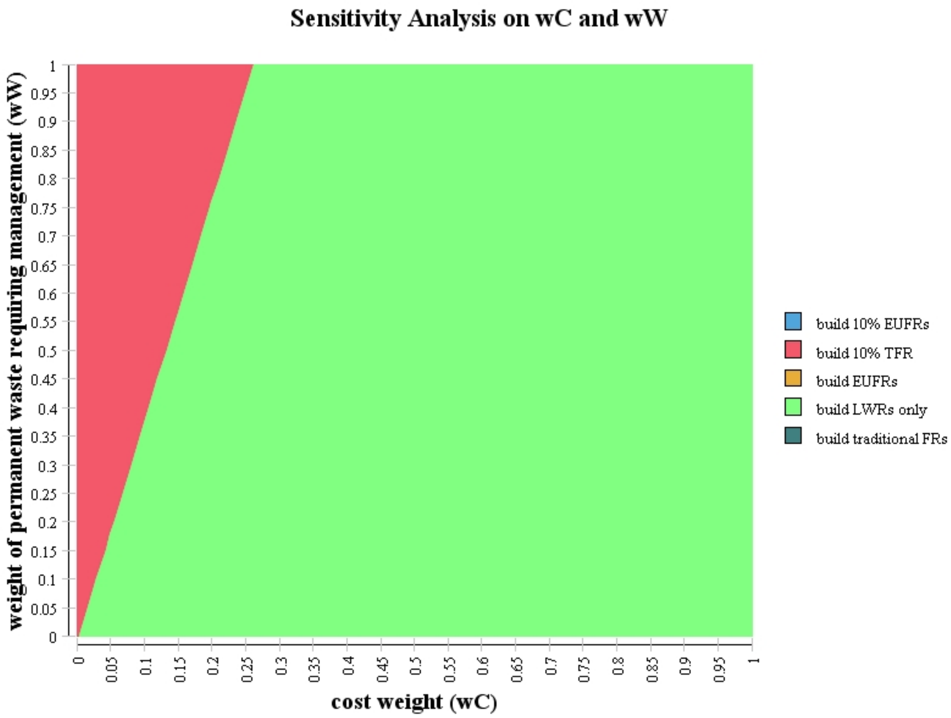

Figure 6-5: Cost and waste weight tradeoff for a value function with a cost-weight interaction term

Figure 6-5 does not change substantially when other parameters are shifted. For example, increasing the probability of high growth in period one simply makes TFRs more attractive, shifting the pink space to the right (as we would expect, given that for the 10% TFR decision, higher growth makes TFRs more attractive). Similarly, graphs comparing the sensitivity of wC to P1 with wW set at various values look almost exactly the same as the graph of the basic 10% decision presented in section 5.3, with the wW determining the thickness of the 10% TFR desirability band.

Overall, the simple addition of an interaction term does not significantly impact the results presented in Chapter 5. It may be, however, that consideration of a wildly different utility surface over the attributes would change the results more substantially (but this surface would need to accurately represent the desires of decision makers, and would not likely depart far from the linear or exponential curves explored in this thesis).

Rather than concentrate too heavily on the independence of the attributes and structure of the value function, analysts conducting further work would probably do better to examine the effect of a completely different value function with new metrics. For example, if decisions were made that rendered volume to be the only relevant waste metric (e.g. because deep "hot" borehole disposal were deemed the best pathway for geologic management), this might substantially alter the results. Alternatively, adding metrics or replacing waste or cost with other attributes deemed to be more important also could have a large effect. Safety, for example, is paramount in the public's post-Fukushima thoughts on nuclear power; further work could examine the effect of incorporating a safety metric into the value function.

6.2 Key Takeaways from the Sensitivity Analysis on the Addition of an Interaction Term
Adding an interaction term to the value function does not change the qualitative outcomes observed so far. We do see that when we separate the waste and cost preference weights so that they are independent of one another, the results are somewhat more sensitive to the cost weight. The decision results generally respond to uncertain parameters as before.

6.3 Sensitivity to the Financial Discount Rate

A major challenge for any analysis with a financial investment component is selection of a discount rate to represent the time-value of money. For all analyses so far, a discount rate of 7% has been used. This follows DeRoo and Parsons, who employ 7% annually compounded in their analysis.(De Roo & Parsons, 2009)[2] The MIT Future of Nuclear Power study estimated the rate for nuclear power plants to be 10%.(J. Deutch, Moniz, & et al., 2003) The Boston Consulting Group, however, in its analysis of the cost of reprocessing nuclear fuel, used a discount rate of 3%.(Boston Consulting Group, 2006) This low figure represents a "risk-free" government rate, which may or may not be applicable to the fuel cycle system. Because the possibility exists that the government would build and operate fuel cycle facilities, and because 3% represents one of the largest departures from an estimated 7% discount rate, 3% is chosen as a rate to explore the sensitivity of decision results.

The new discount rate of 3% is applied to the total cost of the nuclear system, calculated by year. The lower discount rate means that costs do not decrease as quickly with time, so that the present value of system costs in later years is higher than it otherwise would be. The new discount rate is tested using the decision tree of section 5.2, which includes three options for the

[2] DeRoo and Parsons explain that the 7% discount rate is the sole discount they apply to cash flows, and thus the figure subsumes rates related to the financial structure of the project in question. It therefore most closely resembles a Weighted Average Cost of Capital in their analysis; here, it is applied broadly to all expenditures required in future years to build up and sustain the nuclear fuel cycle system.

decision maker (TFR, EUFR, or LWR – all at 100% of the allowed building rate) and two periods.

Figure 6-6 shows how the decision result changes with the new discount rate. As expected, the lower discount rate increases the space for which TFRs are the desirable first-period option. This happens because the advantage of waiting until the second period to build TFRs is no longer so great, given that costs do not fall off as rapidly with time.

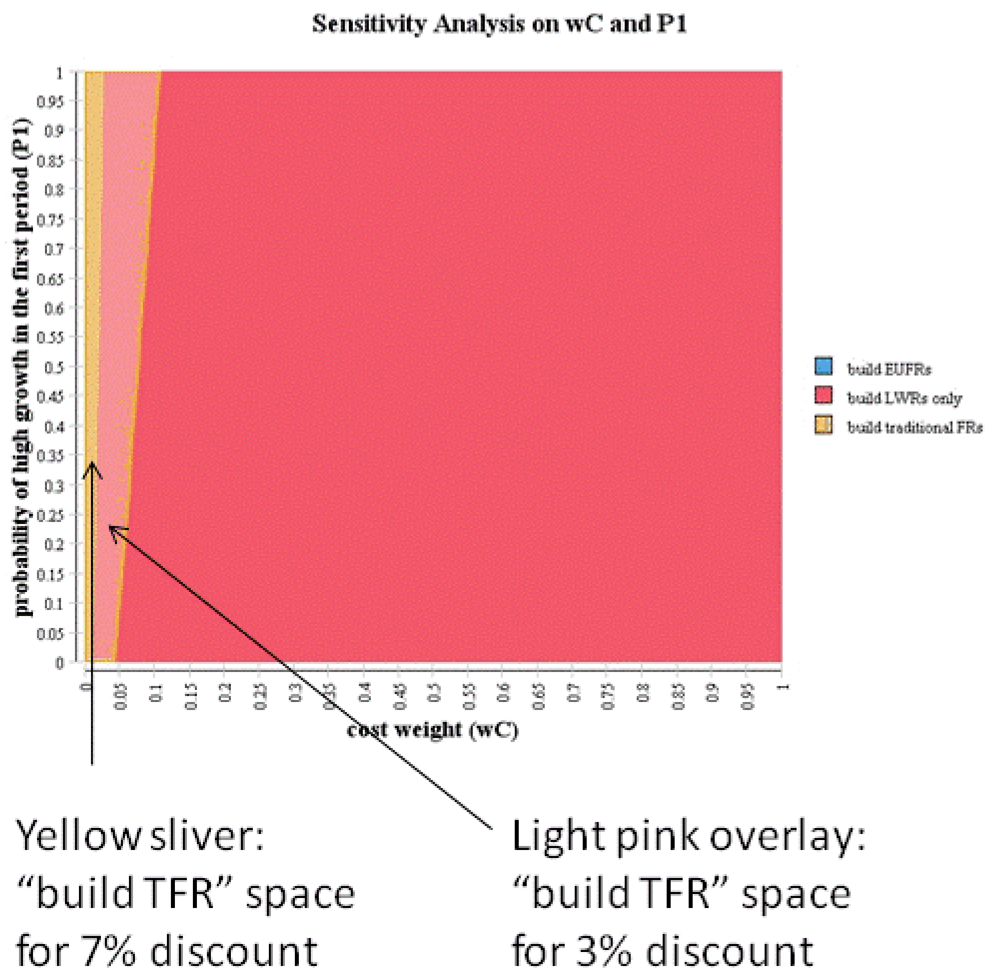

Figure 6-6: Sensitivity of the decision space to a low discount rate

Similar effects occur for the second-period decision. Figure 6-7 shows desirable choices for the second period if "build TFRs" is chosen in the first period. Stopping TFRs becomes more desirable, again because costs are calculated to be higher with the lower discount rate in the later period.

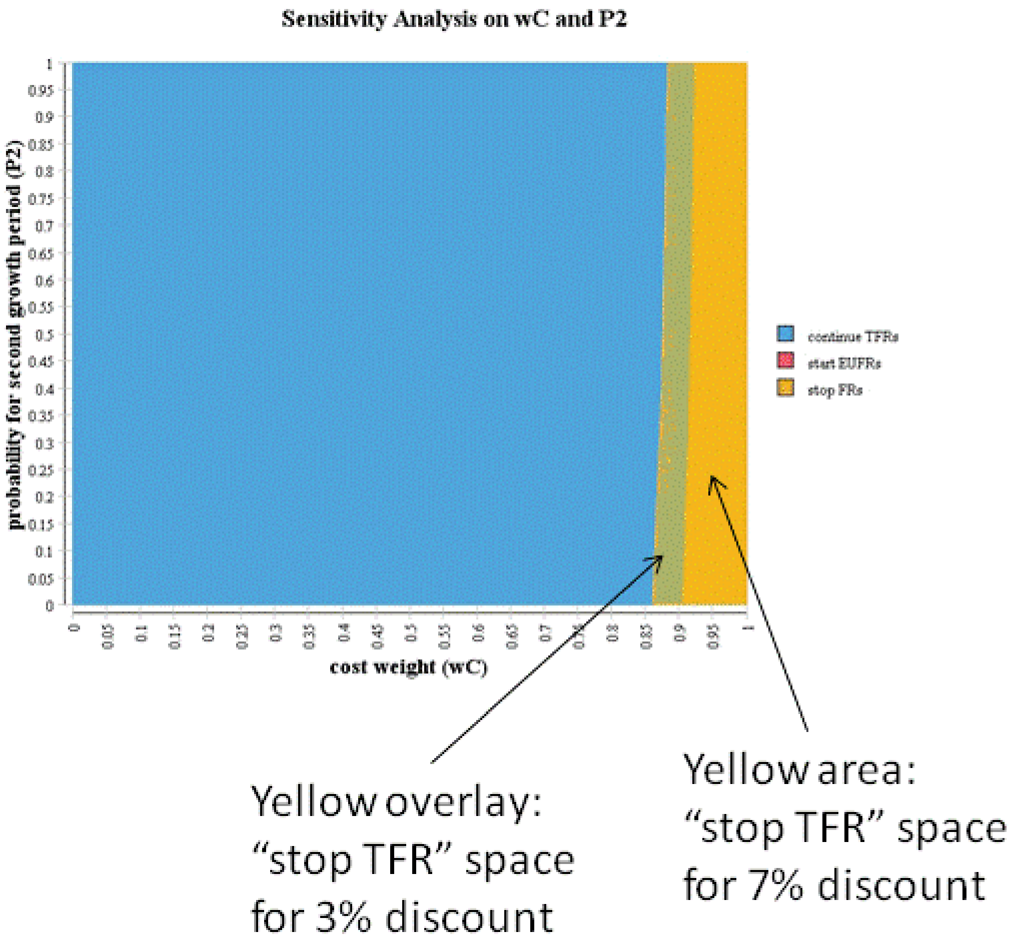

Figure 6-7: Sensitivity of the second-period decision to a lower discount rate

Both the first and second-period decisions exhibit very little change when a new discount rate is applied. This, combined with the fact that the small changes occur in the direction we would expect, increases confidence in results described in Chapter 5. Within the range of plausible discount rates (2-12%), the decision outcomes in this framework are not likely to change qualitatively; analysts should therefore concentrate efforts on narrowing the uncertainty around the other decision parameters.

6.3 Key Takeaways from the Sensitivity Analysis on the Financial Discount Rate

Changing the discount rate applied to yearly system costs from 7% to 3% has the expected effect: building traditional fast reactors in the first period becomes more desirable. This happens because the cost benefit of waiting until later in the century is minimized by the lower discount rate. The effect, however, is not large, and qualitative conclusions about the relationships between the various parameters and their effects on the decision space do not change.

6.4 Waste Liability

Chapter 4 introduced the concept of a "cost liability" intended to capture the financial commitment associated with plants operating at the end of the simulation. Because those plants will presumably continue producing electricity for their full 60-year lifetimes, a fair scenario comparison accounts for the future expenditures that the scenario will require. Similarly, we might consider a "waste liability:" the plants present at the end of the simulation will continue to produce waste, and this is unavoidable if we assume a normal next century where the plants continue to operate.

Calculating the waste the plants will produce is straightforward. For the LWRs operating at the end of the simulation, the total number of remaining reactor-years is calculated and multiplied by the amount of SNF discharged each year. The FRs are assumed to recycle their own fuel year after year, but at the end of their lives, they will discharge their entire core (which will then be managed either in a repository or perhaps in a new FR). The future waste for each FR is thus calculated to be its total core mass. Note that both future losses from reprocessing and the final full-core discharge for LWRs are ignored; both are relatively small amounts, and the waste liability is intended to be a very rough approximation.

Far less straightforward is determining an appropriate discount rate to apply to this waste which will be discharged over a century from now. So far, waste has not been discounted at all within the 100 years of the simulation: waste that appears 100 years from now is as undesirable as waste generated next year. By so far ignoring the waste liability, the discount rate used in the above analyses could be described as a step function, which discounts simulation waste at 0% and post-simulation waste at 100%.

Caplin and Leahy (2000) argue that where societal welfare is concerned, a smaller discount rate should be used than the private (financial) discount rate in order to properly account for society's preferences.(Caplin & Leahy, 2000) Government, says Caplin and Leahy, should pursue future-oriented policies. Given that financial discount rates tend to be on the order of 10% and lower, the discount rate applied to future nuclear waste could be justified at a very low number. This section explores the effect of discounting the waste liability at 0%.

The waste liability is first applied to the results for the simple, three-decision, two-period tree with cost uncertainty (see Figure 5-5). The decision is between building TFRs, EUFRs, and LWRs at 100% their possible rates in each of two periods. Growth can be high or low for either period.

Figure 6-8 shows the resulting decision space for sensitivity over the cost weight and probability of high growth in the first period. This image should be compared with Figure 5-6, for which the yellow band is much thinner: when the waste liability is added in, TFRs become more attractive. This makes intuitive sense, because LWRs will generate waste every year throughout their lifetimes whereas FRs discharge only a single core's worth of fuel when they are decommissioned. The more FRs present at the end of the simulation, the lower the waste liability, so FR scenarios get a boost from considering the waste liability.

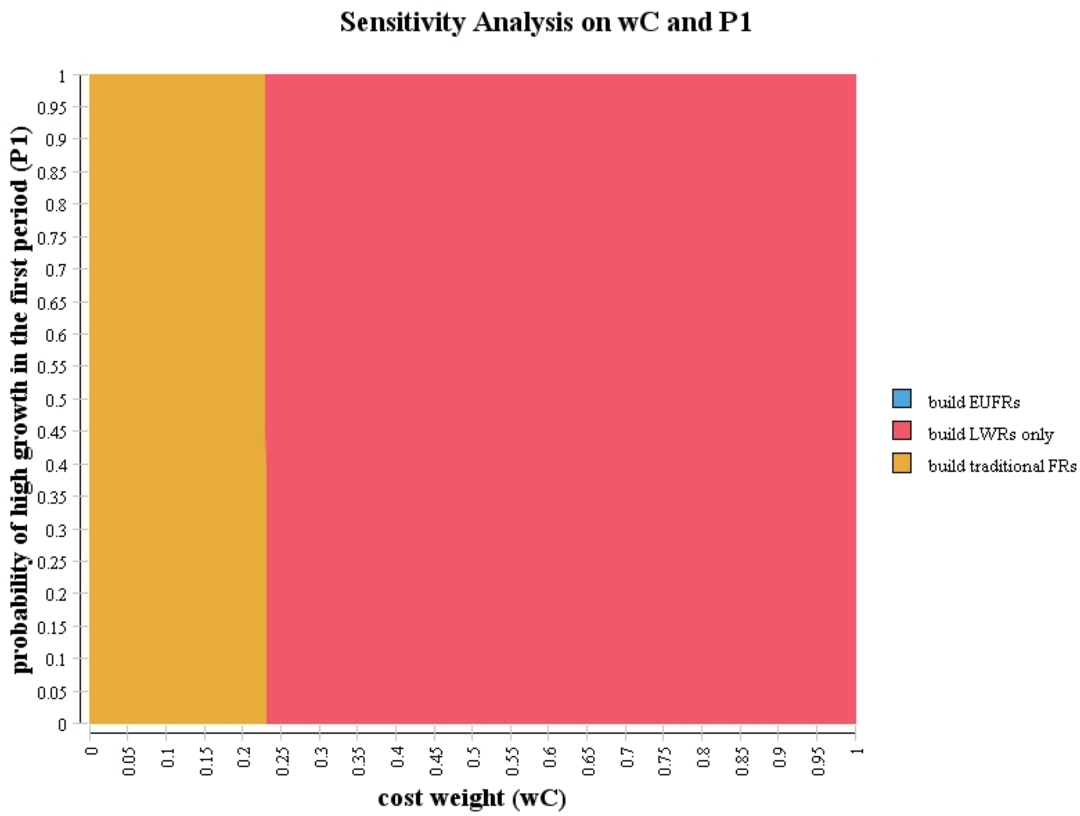

Figure 6-8: Desirable decisions for the two-period tree with waste liability incorporated

EUFR scenarios, of course, entail building the most FRs of all because they do not rely on continued building of LWRs for reactor startups.[3] Building EUFRs does appear as a desirable option if fast reactors are particularly inexpensive, as shown in Figure 6-9. EUFRs are the most expensive option if the cost premium on FRs is high. If, however, the cost premium is low, EUFRs actually are cheaper because the difference in reactor construction is no longer as important, and EUFRs obviate the need for LWR SNF recycle. Note, however, that at very low cost weights (where the waste weight is high), TFRs are still the desirable option because starting with them, using some LWR SNF, and then moving to EUFRs is better from a waste perspective than building EUFRs all the way through.

[3] Note that in reality, it is possible to build fewer EUFRs in addition to some LWRs (the 10% and 25% startup scenarios begin to address these options). In general, those scenarios do not reduce the waste burden enough or come in at low enough cost to be desirable (the EUFRs are "middled-out": if cost matters, LWRs are the best option, whereas TFRs are the best option if waste is important). There may be an optimal mix of EUFRs and LWRs such that cost is not too high and enough waste is obviated by EUFRs, but finding such optima was not the aim of this analysis. The following discussion does show that increasing the importance of obviated waste favors EUFRs.

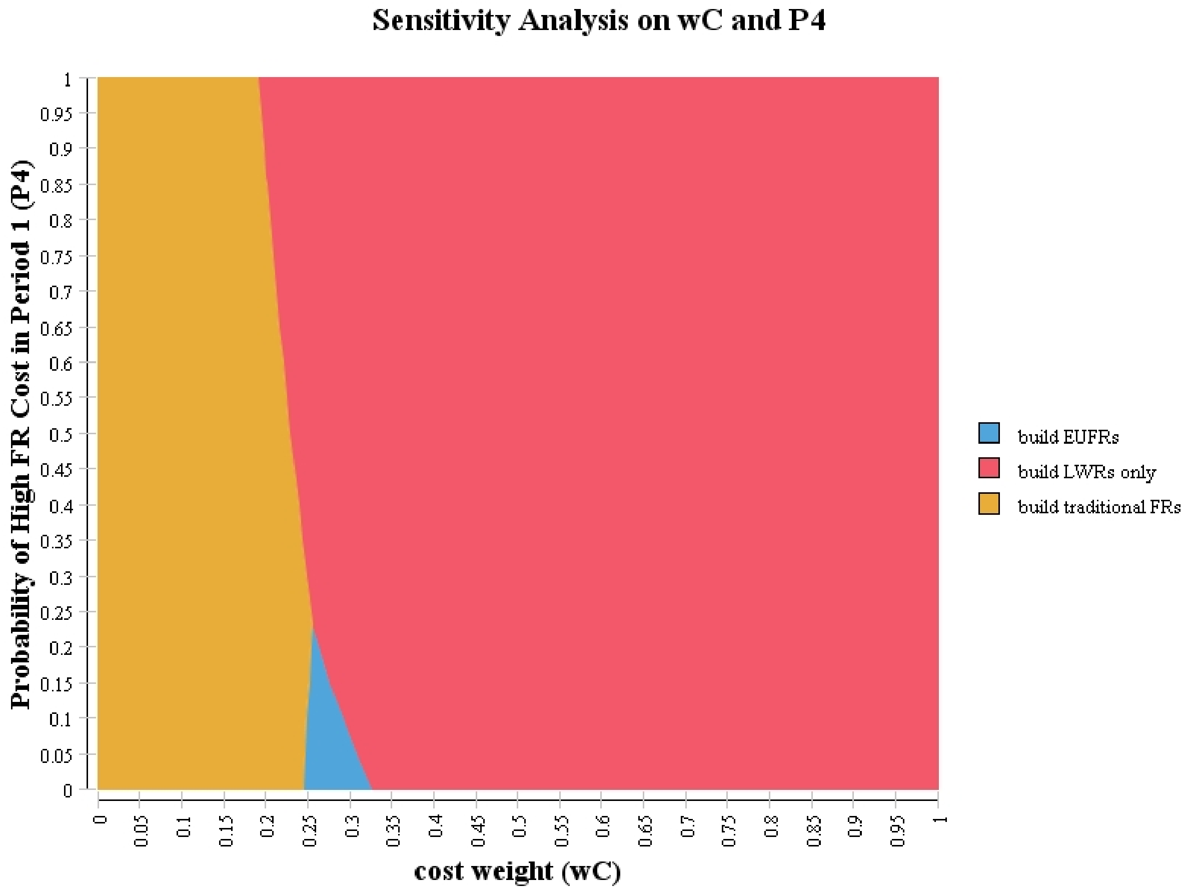

Figure 6-9: Decision sensitivity to fast reactor cost for the two-period tree when waste liability is added

The optimal second-period decision for this tree is almost always to build EUFRs. Figure 6-10 shows the sensitivity of the decision over cost weight and probability of high growth in the second period if "build LWRs" is chosen first; the "build EUFR" space is even larger if TFRs are the chosen first-period option. Predictably, LWRs are most desirable when the cost weight is high. The stark sensitivity to growth in period 2 between the TFR and EUFR options occurs because at low growth, the same number of FRs are built regardless of which scenario is chosen (LWRs no longer limit the building of FRs in the TFR scenario: only FRs are built in the second period). The TFRs, however, are using old spent LWR fuel to start up while EUFRs are not, meaning that there is more leftover LWR SNF from early LWRs contributing to the waste liability for the EUFR scenarios. This gives TFRs the edge at low growth. At high growth, by contrast, TFRs are limited by available SNF, so LWRs must be built alongside the TFRs. This contributes to the waste liability throughout the second period because more LWRs are present at end, and gives EUFRs the advantage from the waste perspective.

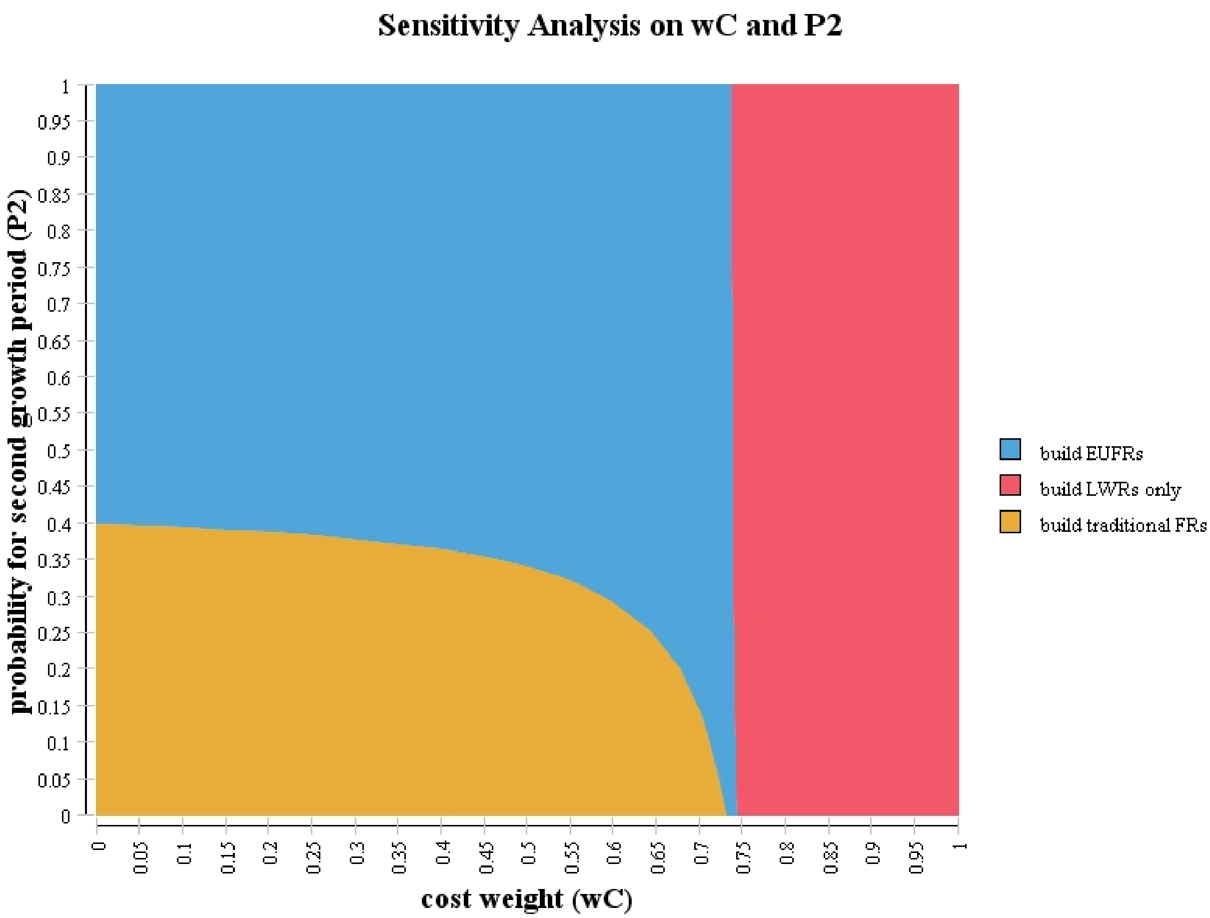

Figure 6-10: Second-period decision for simple two-period tree with waste liability

Considering the waste liability in the context of the five-option tree (where the decision maker can start with 10% TFRs or 10% EUFRs in addition to the options above) creates a more variegated decision space with similar trends. Figure 6-11 shows the decision sensitivity to the probability of high growth in period one and to the cost weight. The “build EUFRs at 10%” decision now overtakes some of the space for which LWRs were the optimal decision in the three-option tree. This happens because the EUFRs can replace some LWRs that would be built early on, reducing the ultimate waste liability for relatively little cost. The other sensitivities for the five-option tree look similar to the three-option tree in shape, with the 10% options gaining some traction for the middle range of cost weights.

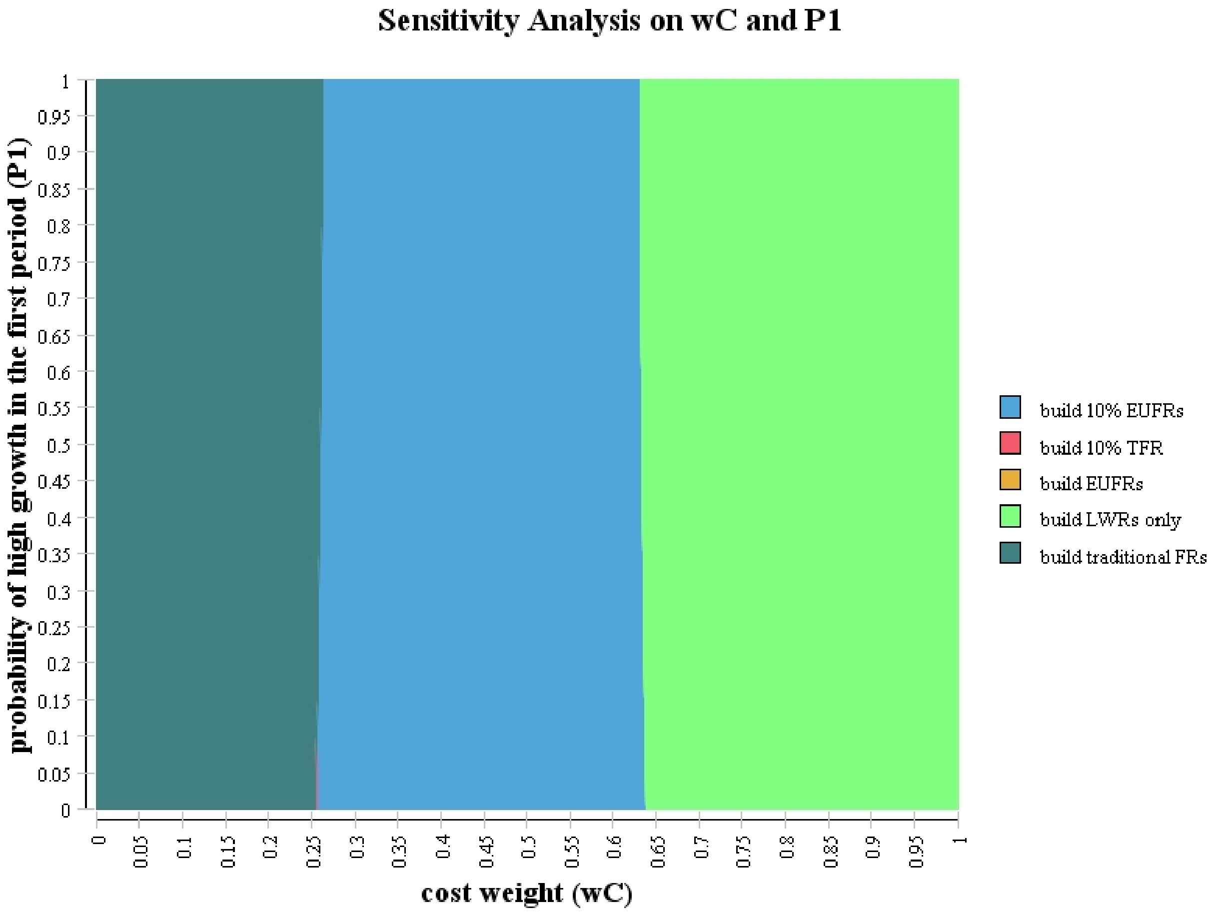

Figure 6-11: Decision sensitivity to growth probability and cost weight for 5-option tree with waste liability

Overall, inclusion of the waste liability has a strong effect on the decision results. This indicates that it is very important for decision makers to consider what the waste management strategy will be. It also means that the time horizon or discounting applied to the waste buildup will dramatically impact choices between fuel cycles. This is especially true for high-growth scenarios, where the number of reactors built, and thus the waste liability, is extremely large by the end of the century.

6.4 Key Takeaways from Sensitivity to the Nuclear Waste Liability

Considering the "waste liability," or the final disposal of all system wastes at some point in the future when (or if) the entire nuclear system is decommissioned, has a dramatic effect on the decision results. Avoiding LWRs becomes very important, because any LWRs operating at the end of the scenario will continue to produce a large amount of waste throughout their lifetimes.

Enriched-uranium fed fast reactors become a desirable option, because they are the only manner of ensuring that few (or no) LWRs exist by the end of the simulation.

6.5 The Impact of Placing an Industrial Capacity Limit on the Pace of Fast Reactor Builds

The comparison of FANTSY with four more sophisticated fuel cycle codes showed that FANTSY produces a more “bumpy” boom-and-bust pattern of fast reactor builds than the other codes (see Chapter 3 for a summary and Appendix B sections 4 and 5 for more full discussion of the build results). As mentioned before, the pattern results in FANTSY because there is no limit on the number of FRs that can be built, so the code builds FRs rapidly and then must stop for a period of time while stocks of LWR SNF are rebuilt. Though each code calculates similar numbers of fast reactors built by the end of the century, a sensitivity analysis is desirable to determine if the difference in *when* fast reactors are built has an impact on decision outcomes. Note that LWRs could also face industrial capacity restrictions over the timeframe of the simulation, but because the LWR build pattern is equal to the demand minus the FR build pattern, it is only necessary to explore restrictions on FR builds.

The pattern of FR builds is determined in CAFCA, VISION, and FANTSY by the industrial limits placed on the number of FRs that can be built in a given period, if any, by the available reprocessing capacity (and hence feedstock available for FR startup), and by the utility “ordering rules” for each new FR. FANTSY includes no limit on FR numbers, uses a set of reprocessing capacity restrictions similar to CAFCA’s, and employs the FR ordering mechanism used by CAFCA. Two of the three parameters are easy to adjust: restrictions on build rates for FRs, and limits on reprocessing capacity additions. Of those two, a cap on the number of allowable FR builds each year produces more dramatic changes in the FR build pattern. This sensitivity analysis thus analyzes the impact of introducing an FR build restriction into FANTSY on decision results.

Two possible sets of build restrictions are shown in Table 6-2. These values are based on the peak build rate for LWRs in the U.S.: about 10-12 reactors per year. Each profile ramps up gradually, reflecting the assumption that it will take industry some time to develop to the “LWR bandwagon” market capacity or beyond. These cases were also chosen because they produce significant differences in the pattern of nuclear reactor builds (and, in the harsh restriction case, the total number of FRs built by the end of the century).

Table 6-2: Restrictions on national industrial capacity for building FRs

PERIOD	MAX FR BUILDS	
	Medium Restriction	Harsh Restriction
2035-2040	5/yr	3/yr
2040-2050	8/yr	4/yr
2050-2065	10/yr	5/yr
2065-2110	15/yr	8/yr

Figure 6-12 shows the impact of build restrictions on the profile of operating fast reactors at high growth. The green lines show the impact of no restriction ("Zero") vs. medium and harsh (or "Hi") restriction for building 10% FRs at the beginning of the simulation and 100% in the second period. The blue lines show the same, for building 100% TFRs the whole way through. One can see that for the 100% TFR all-through case, the total number of reactors operating at the end of 100 years is similar regardless of the restrictions placed on the build pattern. This is not so for the 10% case: the ultimate number of operating FRs is very different between build restriction scenarios. The industrial capacity limit thus has an effect beyond just smoothing the pattern of FR builds for some cases. To truly isolate the impact of a smoother profile while preserving the same number of reactors ultimately built, the FR ordering algorithm would need to be adjusted toward those of the advanced fuel cycle simulations. For now, we continue with the analysis of sensitivity to an FR build restriction, because it could represent a physical industrial reality that both smoothes the building curve and restricts the number of reactors built within the century timeframe.

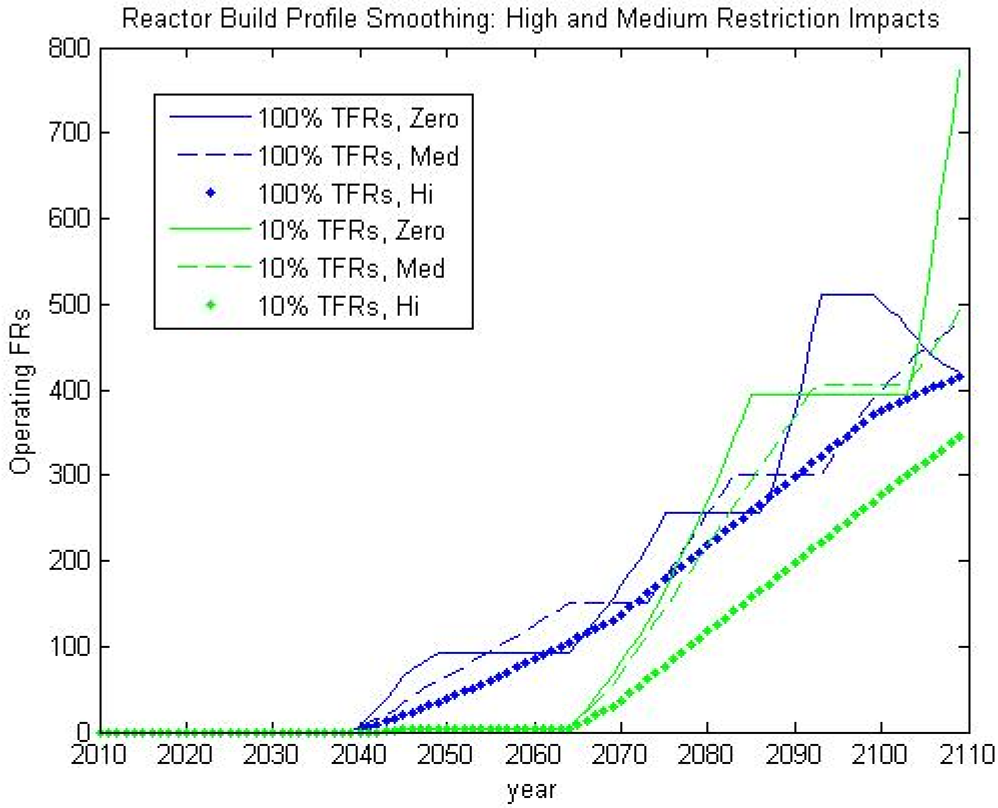

Figure 6-12: Impacts of FR build restrictions on the operating FR profile

The fast reactor fleet grows larger under high growth scenarios, so build restrictions have a much bigger impact on the construction pattern when growth is high. If growth is low throughout the entire simulation (1.2% per year until 2040 and 0.5% per year thereafter), the maximum number of fast reactors ever built in a single year (toward the end of the simulation) is 16. Nearly every single year under low growth, the number of fast reactors required is less than allowed by the medium build restriction. The harsh restriction only has a very small effect on the FR build profile at low levels of demand.

The analysis here therefore concentrates on build restrictions that apply at high growth only. In practice, industrial capacity restrictions on the number of fast reactors could exist at any growth level, but the assumption simplifies the decision tree exhibited in Figure 6-13. The tree is identical to that of section 5.3, except that an extra uncertainty node has been added for conditions of high growth. If growth is high in the same period that fast reactors have been chosen, an uncertainty node reflects the possibility that build restrictions for FRs could be harsh, medium, or non-existent. A new probability parameter, P5a, P5b, and P5c, tracks the probability with which each restriction comes to pass.

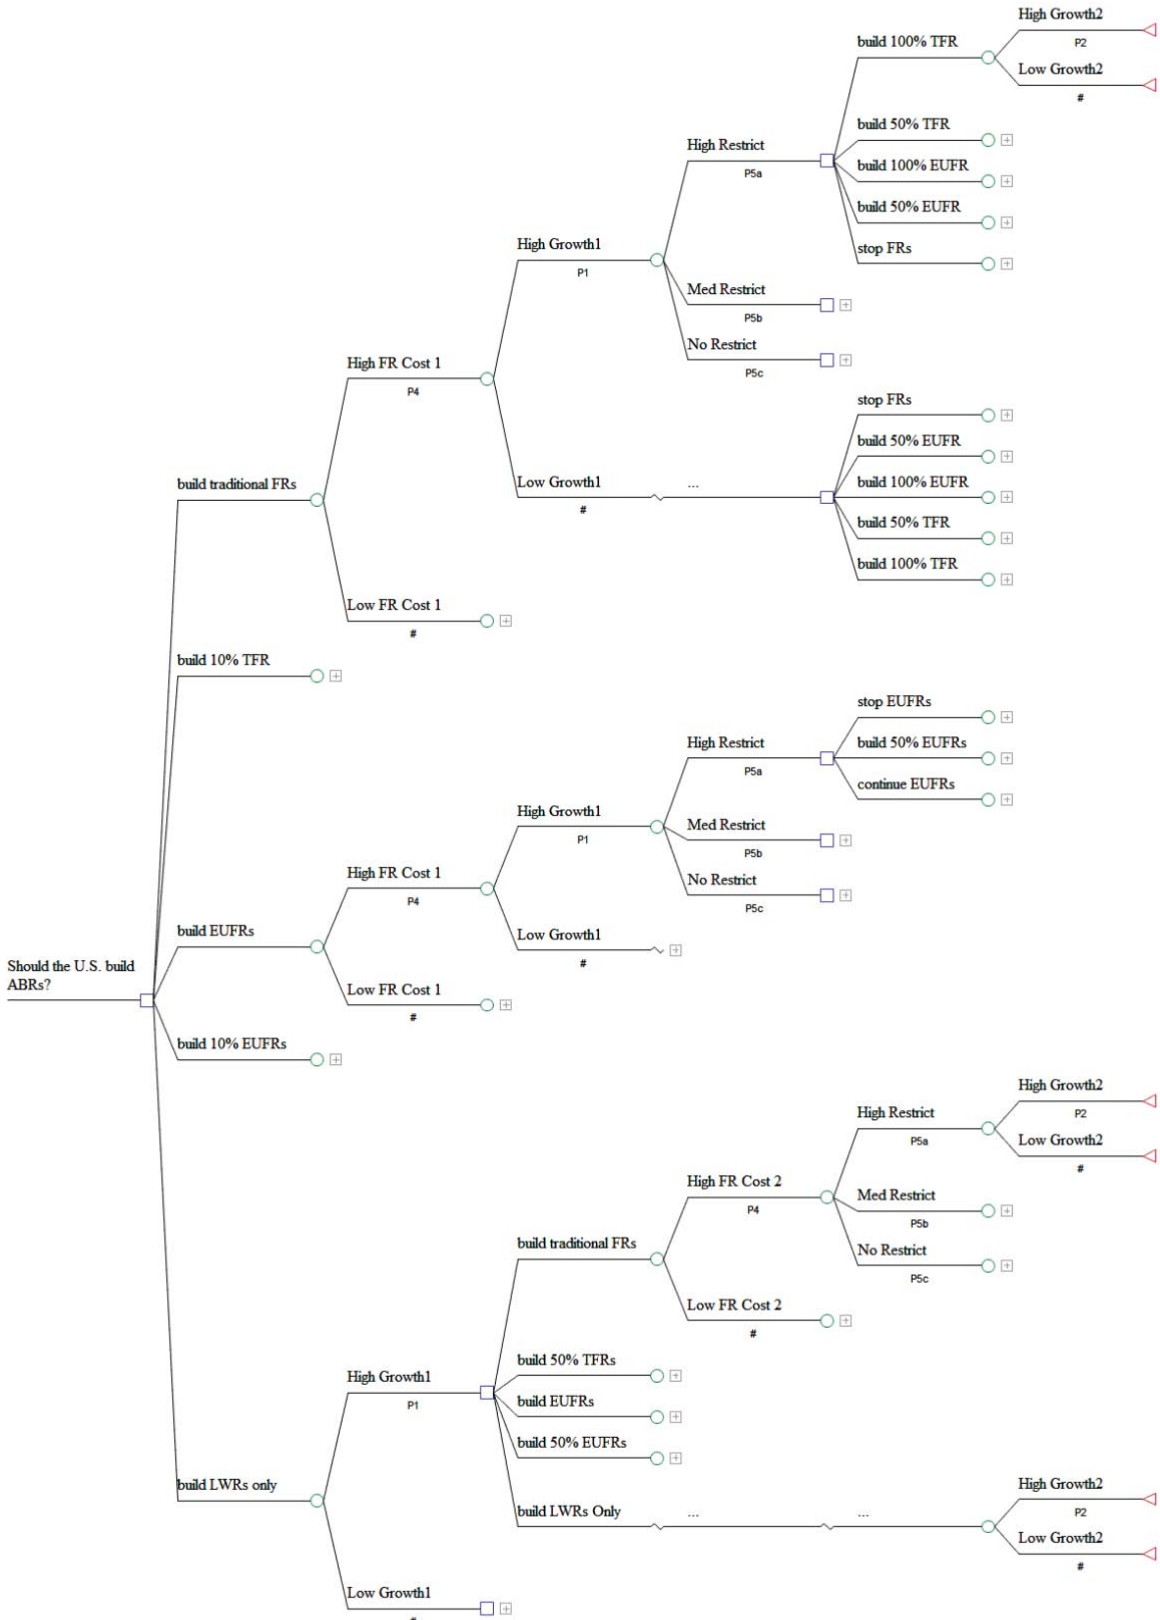

Figure 6-13: Decision tree for industrial capacity restrictions on FRs

When P5c, the probability of zero restriction, equals 1, the tree simplifies to one very similar to that in section 5.3. The only difference is that more options have been calculated for the terminal nodes, causing slight shifts in the utility values. The decision graph for the initial node comparing the probability of high growth to the cost weight is effectively the same as the one from section 5.3 (Figure 6-14). As before, below a cost weight of 0.2, building 10% TFRs in the first period is desirable, and becomes very slightly more so with increasing probability of high growth.

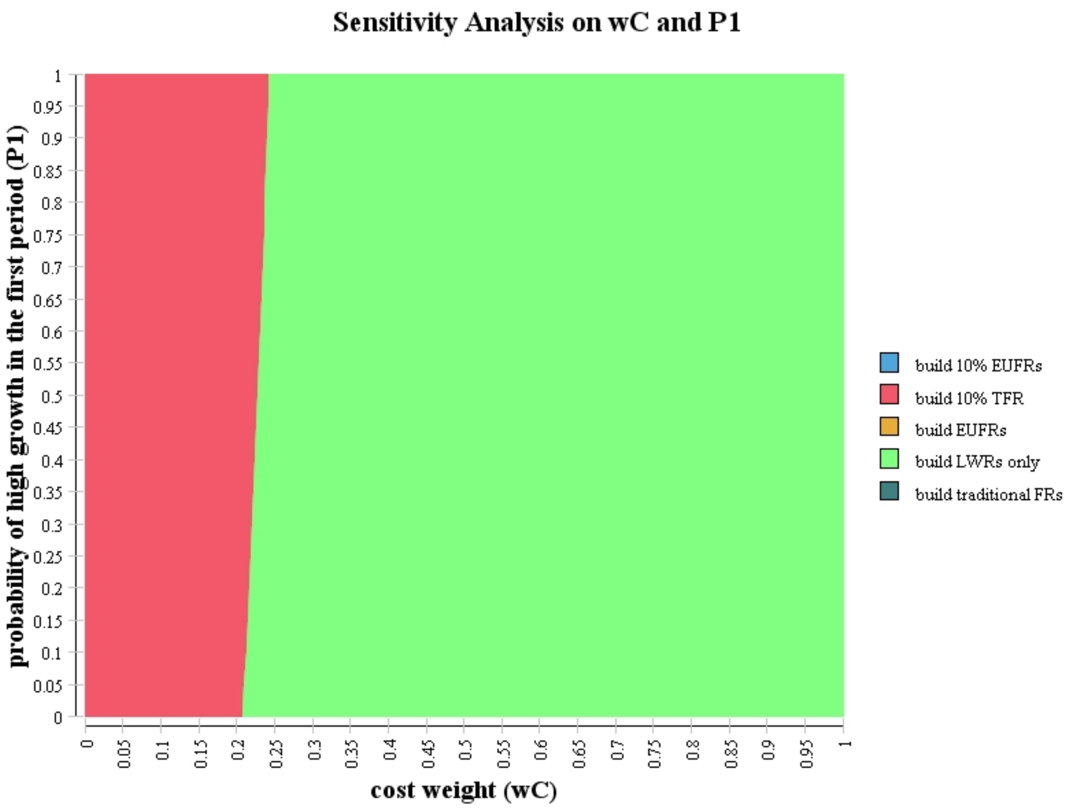

Figure 6-14: Desirable decisions for standard tree with uncertainty in FR building capacity at high growth

Changing the probability of the industrial capacity restrictions has almost no effect on the decision space. Figures Figure 6-15 and Figure 6-16 show the miniscule effects of the restrictions, demonstrated by setting the harsh and then medium restriction probabilities to one, and comparing each to the above scenario where the probability of zero restriction is one. Recall that

P5a is the probability of harsh restrictions plaguing industry, while P5b is the probability of a medium restriction and P5c is the probability of no restriction at all.

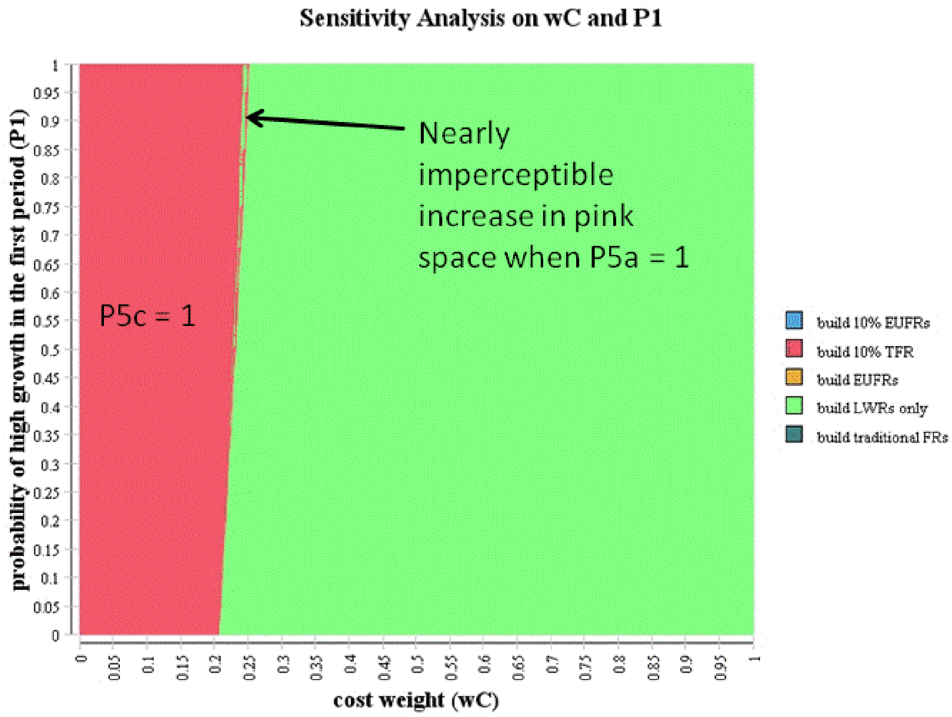

Figure 6-15: Comparison of zero and harsh industrial capacity restriction effects on decision space

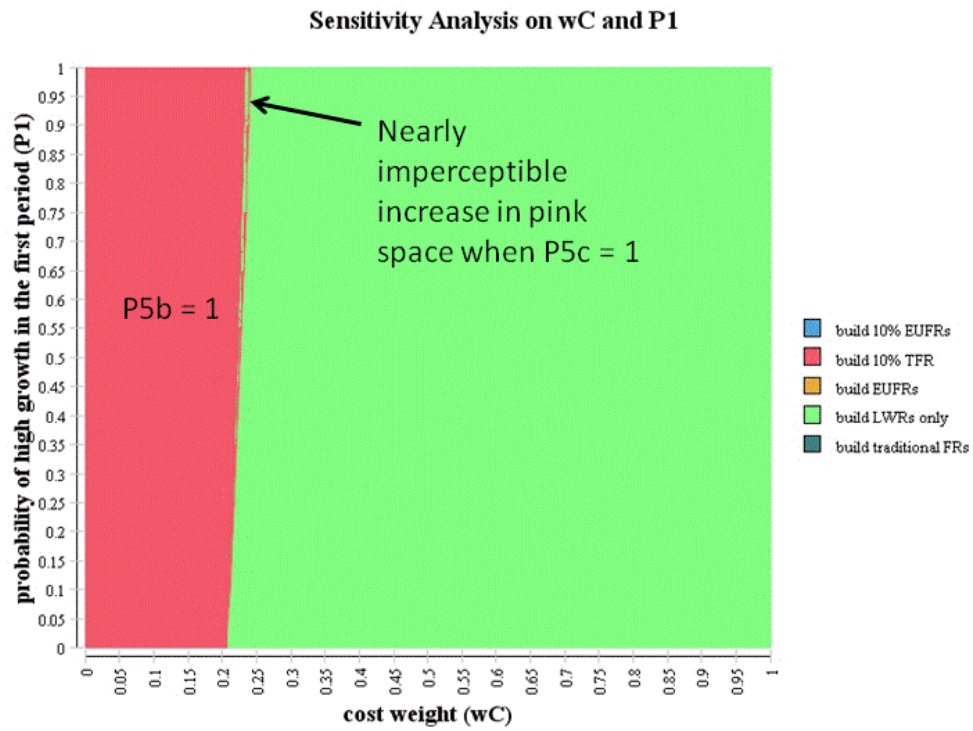

Figure 6-16: Comparison of medium and zero industrial capacity restrictions on decision space

Of note is the fact that the harsh restriction very slightly *increases* the desirability of the 10% fast reactor decision, while the medium restriction decreases its desirability. The reason for this is that both restriction levels achieve similar levels of waste benefit in terms of decreasing the SNF stock, but the medium restriction involves earlier (i.e. less discounted) FR builds.

The primary conclusion to be drawn from this analysis is that restricting the amount of FRs that can be built in a given year has very little impact on the decision space (the decision space for the second period decision, not shown, confirms this conclusion). This is true even though there are some significant differences in the FANTSY results for the build pattern of FRs. Given that FANTSY produces build numbers and patterns fairly similar to those from other codes with different FR ordering algorithms, adjusting the FANSTY ordering algorithm is likely to continue producing similar results with no ultimate effect on the decision space. Further analysis of sensitivity to the drivers for the FR build pattern is therefore unnecessary.

6.5 Key Takeaways from the Sensitivity Analysis to Industrial Capacity Restrictions for Building Fast Reactors

Adding an industrial capacity restriction, allowing only a few fast reactors to be built per year, has a nearly imperceptible impact on the decision space. Neither the smoother build profile nor the decrease in ultimate numbers of reactors built changes the desirability of any of the options evaluated, at least in the absence of a learning curve assumption. The differences between codes' fast reactor build profiles may thus be of relatively little importance to high-level fuel cycle decisions.

Chapter 7: Incentivizing Industry – A Different Decision Maker Perspective

Chapters 5 and 6 identify and analyze desirable pathways for nuclear fuel cycle system evolution, under the assumption that an "ideal" or powerful decision maker will be able to enact decisions at each node. The U.S. government most closely approximates such a decision maker, because federal legislative and executive bodies will be able to make policy to radically alter the structure of the national fuel cycle system. Government, however, is unlikely to act without strong industry collaboration. Rather than mandate the building of certain types of reactors and fuel cycle facilities, government will likely create a system of incentives and penalties to encourage industry to invest in new technologies. This section explores the implications of a different perspective, where government offers incentives to industry in order to steer nuclear fuel cycle evolution.

7.1 A Government-Industry Interaction Model

Private industry currently controls and operates all front-end and electricity production elements of the open fuel cycle, including mining, milling, conversion, enrichment, fabrication, and nuclear reactor construction and operation. Government is charged with waste management, and is supposed to take title to the spent LWR fuel and dispose of it in a permanent repository. A model of interaction in the context of a closed fuel cycle is proposed that preserves this basic division of labor: industry continues to hold purview over the LWR portion of the cycle and operates FRs, while a government entity manages reprocessing and waste disposal. This is a notional institutional arrangement which is neither predicted nor advocated; rather, it serves as one example for a way to set up the problem. Analysts could insert any number of potential arrangements between nuclear fuel cycle actors, and could potentially use the framework to compare their impacts.

Figure 7-1 graphically shows how responsibility would be allocated for this notional structure. Utilities operating LWRs would continue to pay a waste fee to the government in exchange for the government managing spent LWR fuel. The government would use this revenue to operate a repository and recycling facilities, and would offer fast reactor fuel free of charge (and would not charge a waste fee) to any utility operating a fast reactor.

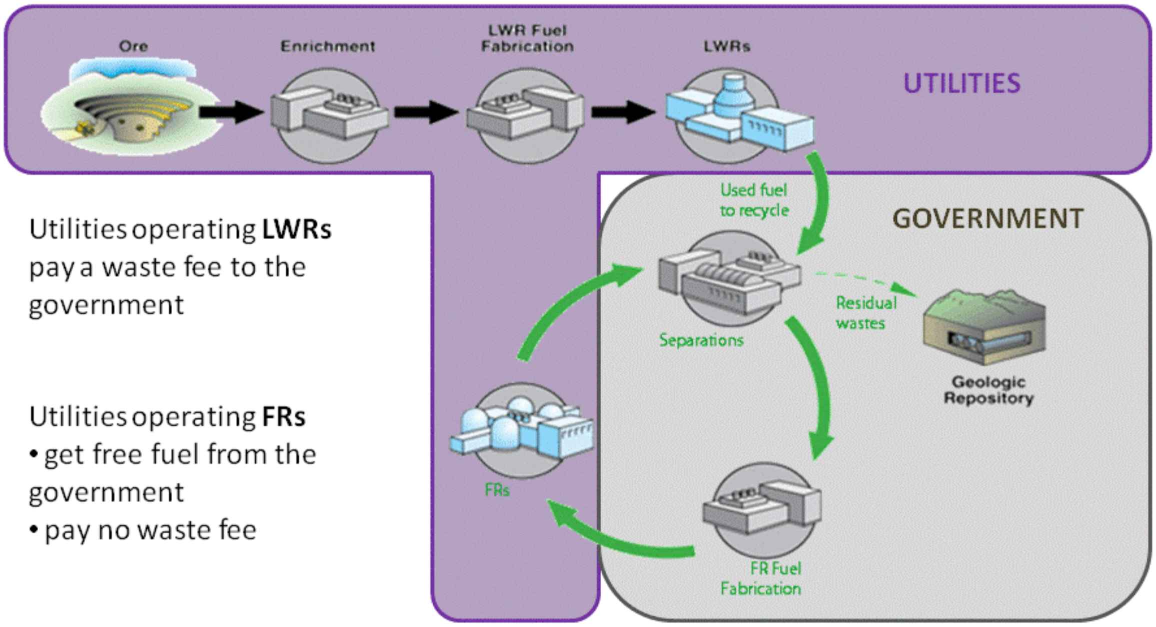

Figure 7-1: Division of labor between industry and government for a closed fuel cycle

Many different versions of the institutional arrangement governing the closed nuclear fuel cycle are possible. For example, a separate governmental corporation could build and operate fast reactors, or government could find ways to incentivize industry to build and operate reprocessing facilities. Changing the institutional structure entails changing where and how much money is exchanged between entities, but the overall system costs will remain the same. The division of labor illustrated in Figure 7-1 was chosen to minimize changes from the current once-through system, and is used to demonstrate one example of how government-industry interaction could work and to evaluate the implications for fees and fuel cycle outcomes. Note that this structure assumes that fast reactors will always be more expensive than LWRs; if FRs become cheaper, the "free fuel" and "no waste fee" policies for FRs would require reevaluation.

When making decisions to invest in new sources of electricity generation, utilities choose based on the technology that will provide the best returns. This often amounts to choosing the cheapest project (with availability and reliability considerations factored into a risk-adjusted cost). Assuming that fast reactors are always more expensive to construct and operate than LWRs, industry will require incentives in order to ever choose to build an FR over a LWR. One way for government to spur a shift in investment is to charge a higher waste fee on LWRs, such that FRs with no waste fee and with free fuel become cheaper.

Figure 7-2 shows the net present cost of building and operating FRs and LWRs for 60 years. The net present cost of an LWR project is $6,515/kWe in 2007 dollars (calculation from (Du & Parsons, 2009), adjusted from a 40-year to a 60-year plant lifetime). Because we do not know what the FR overnight cost premium will be, the NPV of FR cost is presented for a range of cost premiums from 0% to 55%; note that the same premium is applied to decommissioning costs, and FR operating and maintenance costs always have a 20% premium over LWR O&M. The graph clearly demonstrates that at a low cost premium, an FR project is actually cheaper than an LWR project, because LWR operators pay for both fuel and for a 1 mill/kWh waste fee whereas FR operators do not. At 10%, FRs and LWRs have the same cost, and above that, FRs are more expensive.

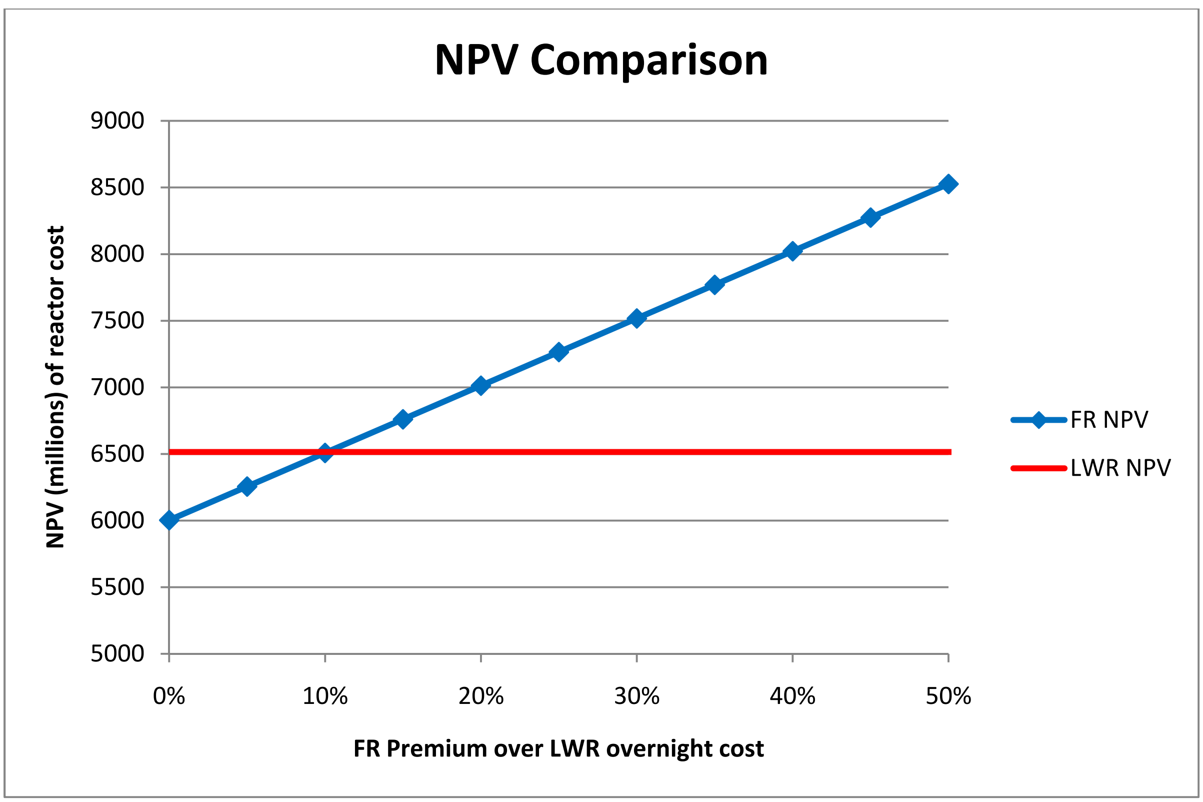

Figure 7-2: NPV of FR and LWR lifetime costs under proposed government/industry interaction structure

Figure 7-2 demonstrates that for any FR cost premium above 10%, the incentives of free fuel and a zero waste fee will not be sufficient to incentivize industry to build FRs. LWRs are cheaper and will be the consistent choice in that range. Government decision makers have many policy options to combine with the suggested regime and further spur FR investment: one includes a steeper waste fee for LWR SNF.

Figure 7-3 shows how much the waste fee would have to increase from its present value at 1 mill/kWh in order to drive industry to build FRs. Below an FR cost premium of 10%, no increase would be required: FRs would be the cheapest option. Above a 10% cost premium, the waste fee climbs precipitously such that if FRs are 55% more expensive to build than LWRs, the fee would need to increase by 5 ¢/kWh. This is an extremely high burden to place on nuclear electricity, given that current levelized nuclear electricity costs are in the range of 8.4 ¢/kWh, and that nuclear electricity has to compete with coal costing closer to 6 ¢/kWh,(Du & Parsons, 2009).

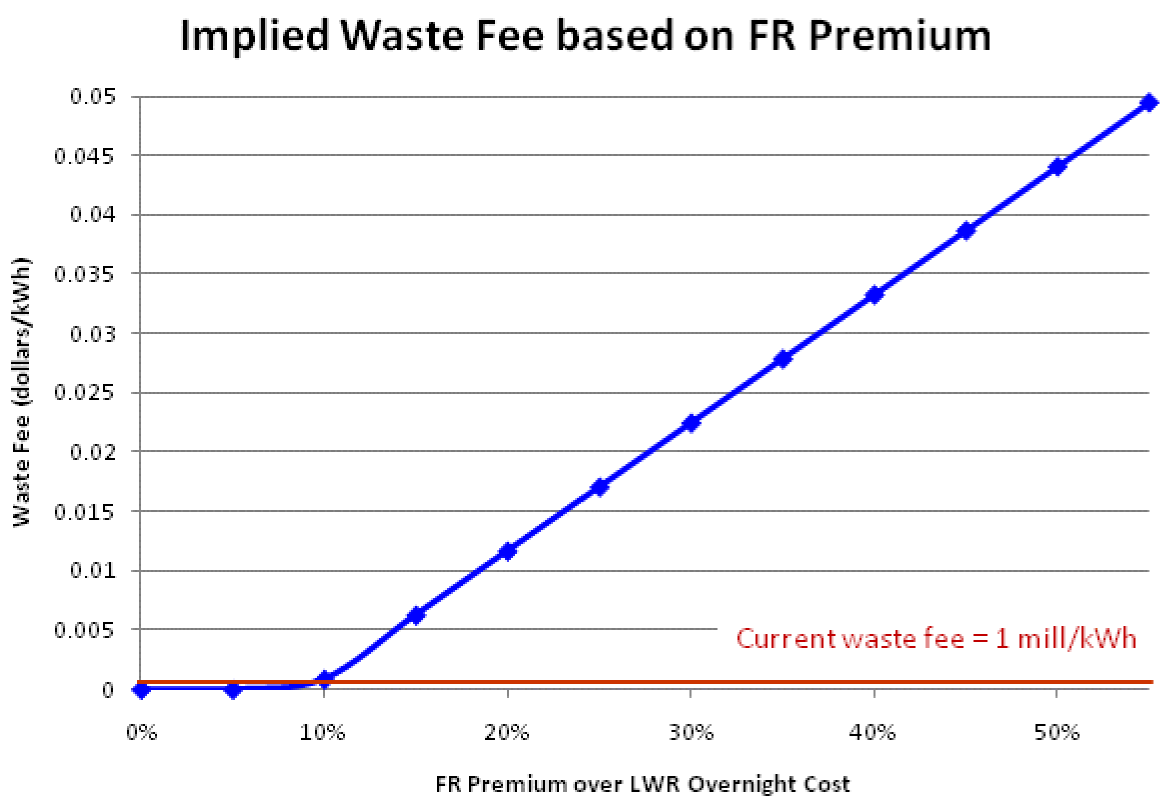

Figure 7-3: LWR waste fee required to incentivize industry builds of FRs

7.2 System Impacts of Raising the Waste Fee

If government were to raise the LWR SNF charge, industry might or might not ultimately respond by building fast reactors at a desirable time or pace. This section assesses whether the pathways identified in Chapter 5 can be achieved through a fee-hike policy instrument.

Figure 7-4 shows, in decision tree format, the likely result of the government increasing the LWR waste fee to a very high level. If the waste fee were far enough above the minimum required to spur FR investment, every utility would always to choose to build FRs over LWRs and the resulting pathway would be "build TFRs at 100% their allowed rate" throughout the century.

The analyses in Chapters 5 and 6, however, showed that building 100% TFRs throughout the entire simulation is almost never a desirable pathway. The five-option, two period tree showed that the generally favored options are to build either LWRs or a small number of TFRs in the first period, depending on the relative weighting between cost and waste. Neither pathway is achievable with a high and static waste fee: industry will only and always build TFRs.

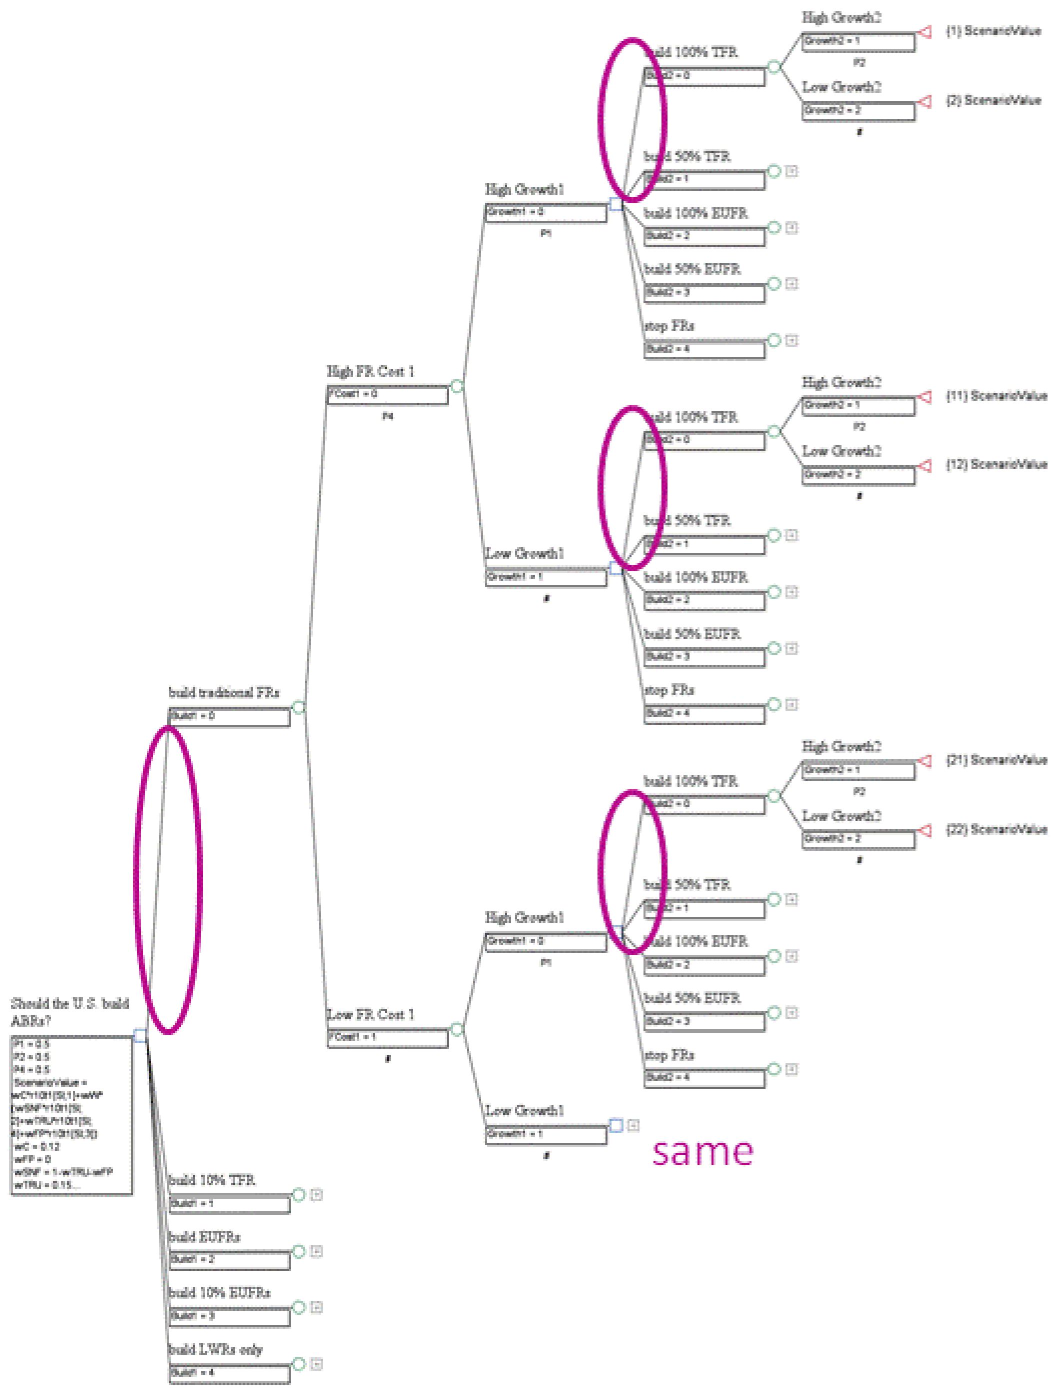

Figure 7-4: Industry response to a high waste fee

Enacting any increase in the waste fee will require a monumental political struggle, especially in the current environment where utility groups are suing the government for *removal* of the 1 mill/kWh waste fee.(World Nuclear Association, 2011a) It is highly unlikely that government will be able to raise the waste fee so substantially that all utilities in all markets will always choose FRs over LWRs.

Far more probable is that government will raise the waste fee by some amount, before clear information is available on the cost of fast reactors, and then utilities will respond in a variety of ways with some choosing LWRs while others choose FRs or neither. Clearly, if the waste fee is too low, zero utilities will build LWRs, and if the waste fee is to high, all utilities will build FRs. If the government agency responsible for the fee hike manages to tag the LWR SNF fee so that it makes FRs just barely more attractive than LWRs, utilities will make choices according to state-level regulations, risk, current generation portfolios, and likely electricity prices. No two utilities will perform the same calculation.

A comprehensive analysis of the U.S. utility industry and the range of conditions companies face is beyond the scope of this thesis. In principle, a detailed model of market conditions and industry decision-making could be incorporated into the decision tree model, in order to capture the uncertain response industry will have to a raised waste fee. This is left to further work.

Instead, a simple decision tree is examined, where the industry decision to build FRs vs. LWRs is a deterministic function of the waste fee and the FR cost premium. If government raises the waste fee, and FR costs turn out to be "high," utilities will build zero fast reactors. By contrast, if government raises the waste fee and costs are "low," utilities will build TFRs at 100% the allowed pace. At each decision node, the government now has the option to increase, keep, or lower the waste fee. Lowering the waste fee will cause industry to stop building FRs. The tree is illustrated in Figure 7-5.

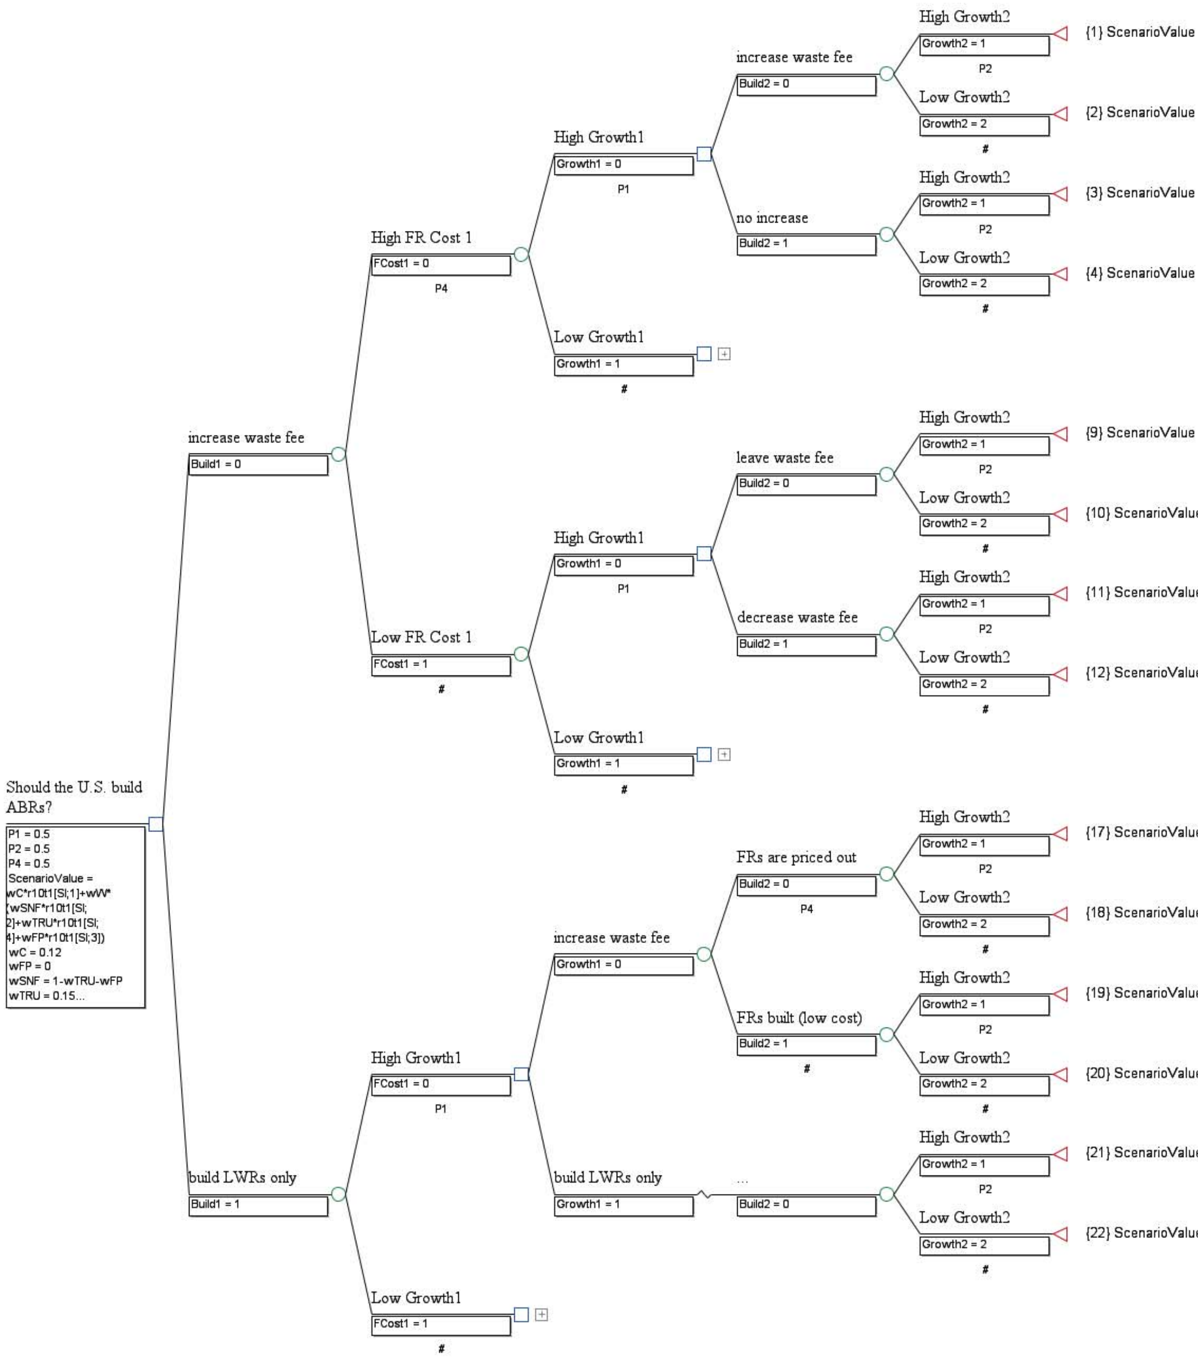

Figure 7-5: Decision tree for government waste fee decision

The result of a sensitivity analysis between the probability for high FR cost and the cost weight, shown in Figure 7-6, is somewhat counterintuitive. The higher the probability that FRs will be expensive, the more attractive it is for government to raise the waste fee in order to incentivize fast reactor investment. This is exactly opposite to the result from section 5.2, where for a 2-period tree with similar options, a higher probability of high FR cost makes FRs *less* desirable.

The phenomenon is similar to that seen in sections 5.4 and 5.5, where the possibility that no fast reactors can be built in later years makes building them sooner the better option. We see it here because of the particular structure of the tree: if government chooses not to increase the waste fee, locking in LWRs only for the first period, a waste fee hike in the second period will still produce no fast reactors if FR costs are high. Therefore, if we think FR costs will be high, we should increase the waste fee in the first period in order to have the opportunity to raise it again in the second period and build FRs despite their high costs.

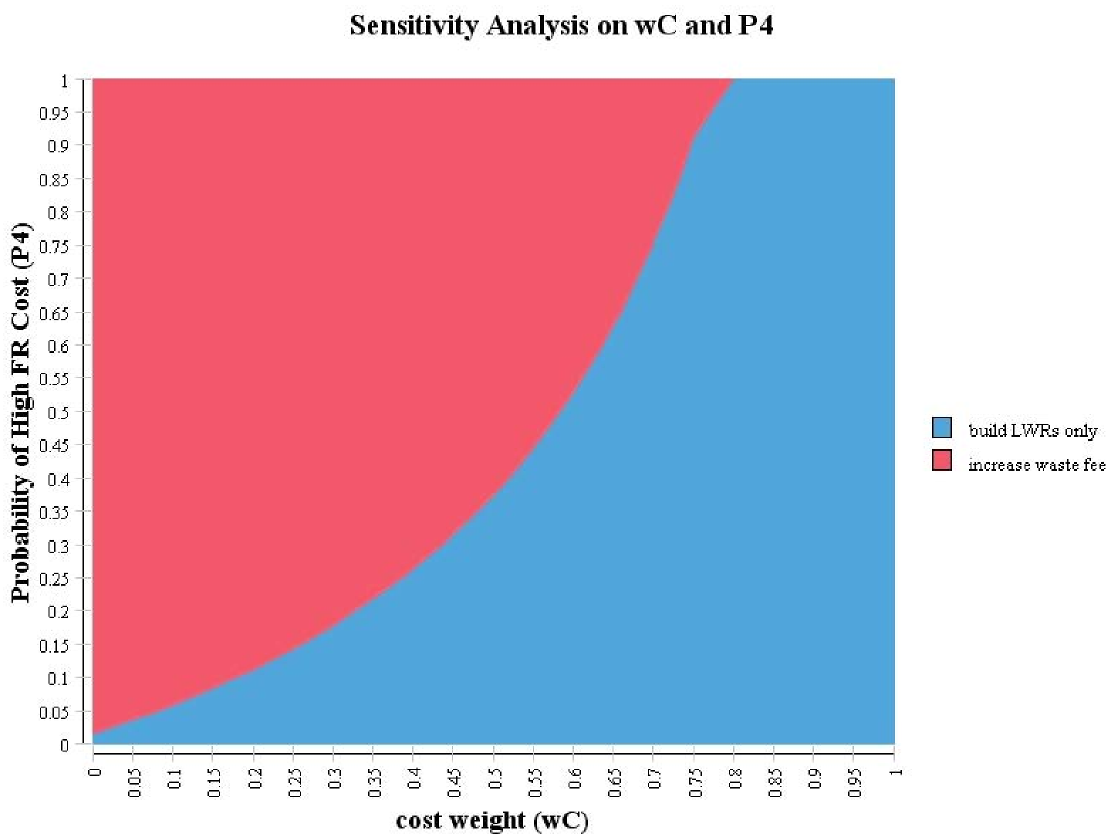

Figure 7-6: Desirable government waste fee decisions for a range of cost weights and cost probabilities

Adding an extra period, and therefore an extra opportunity to raise the waste fee even if LWRs are chosen in the first period, decreases the region of desirability for a waste fee hike but maintains the same trends as the two-period analysis. Figure 7-7 shows the three-period waste fee tree. The periods are the same as for the three-period analysis in section 5.5, with the first decision at 2025, the second at 2050, and the third at 2075.

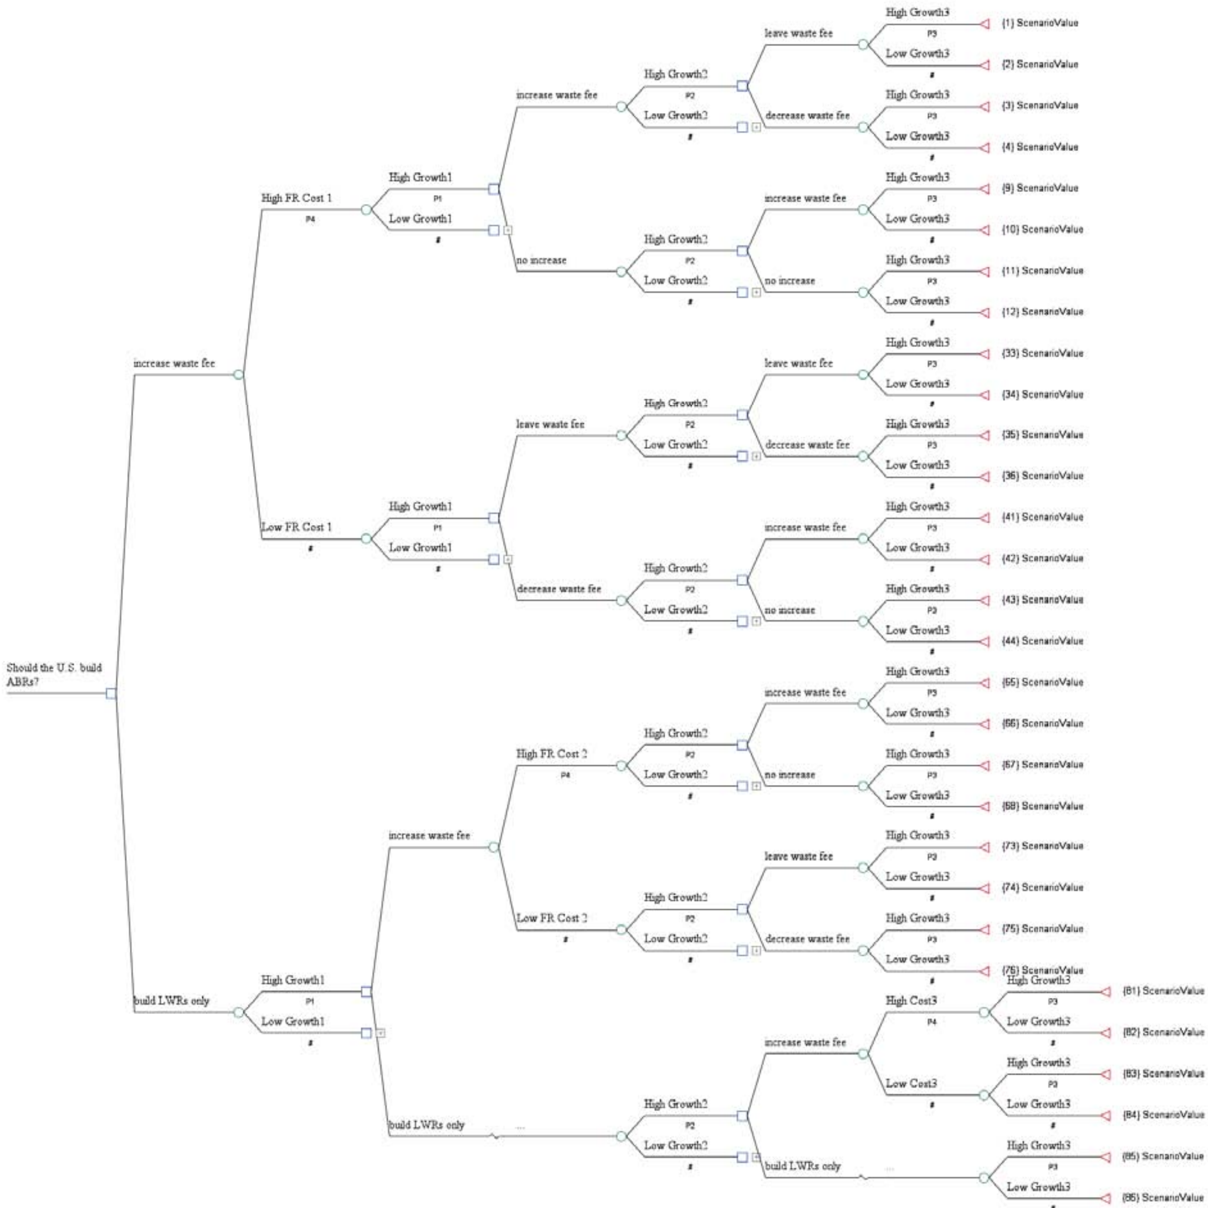

Figure 7-7: Three-period tree for waste fee analysis

The decision tradeoff between the cost weight and probability of high FR cost premium is shown in Figure 7-8. The region of desirability for increasing the waste fee has decreased dramatically, which makes sense now that there will be more opportunities to change course later in the century. Interestingly, however, a high probability of FR cost still makes raising the waste fee more, rather than less, attractive. At low- to mid-range estimates for P4, the results echo those of section 5.2 (see Figure 7-9 and compare to Figure 5-6). At those values of FR cost probability, waiting until later to increase the waste fee makes more sense. But if we think FRs may be expensive, the analysis tells us to increase the waste fee sooner so that we will have more opportunity later on to adjust the fee.

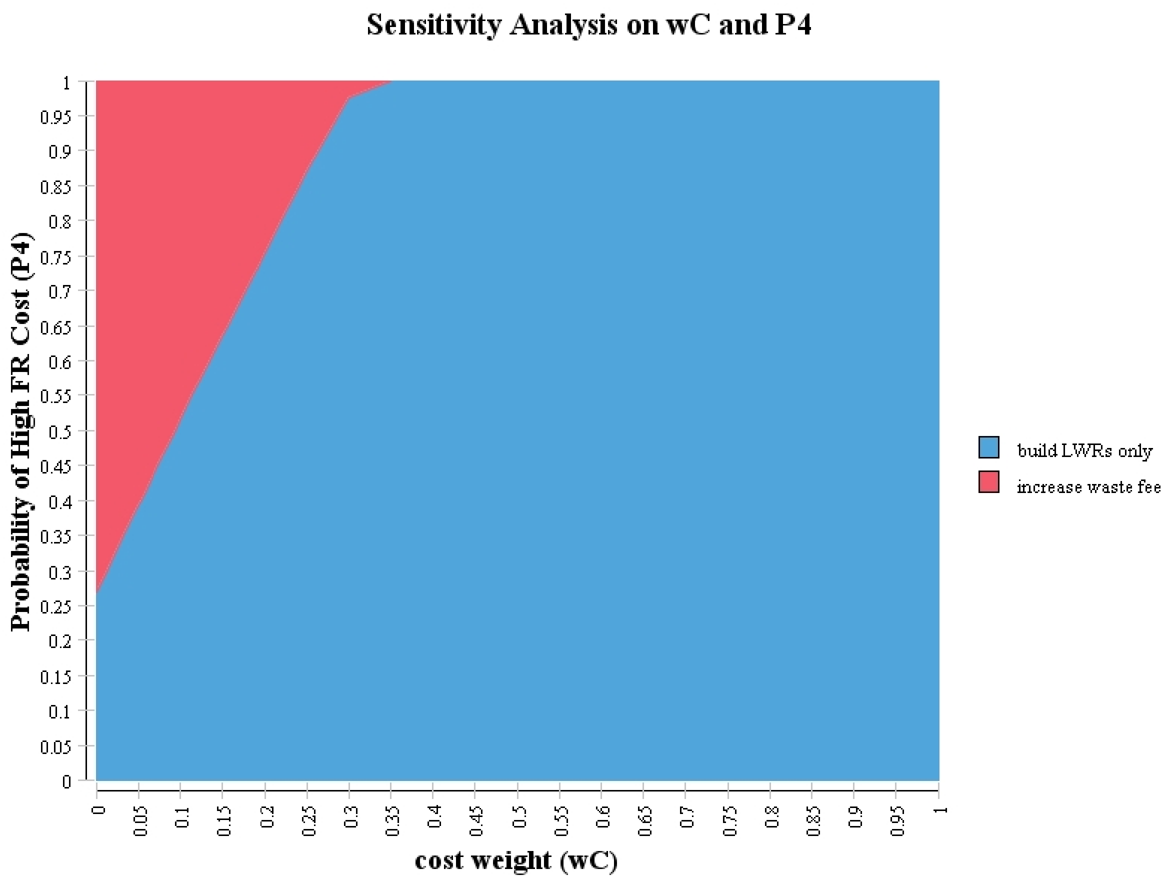

Figure 7-8: Sensitivity to cost weight and FR cost of increasing the waste fee in the first period

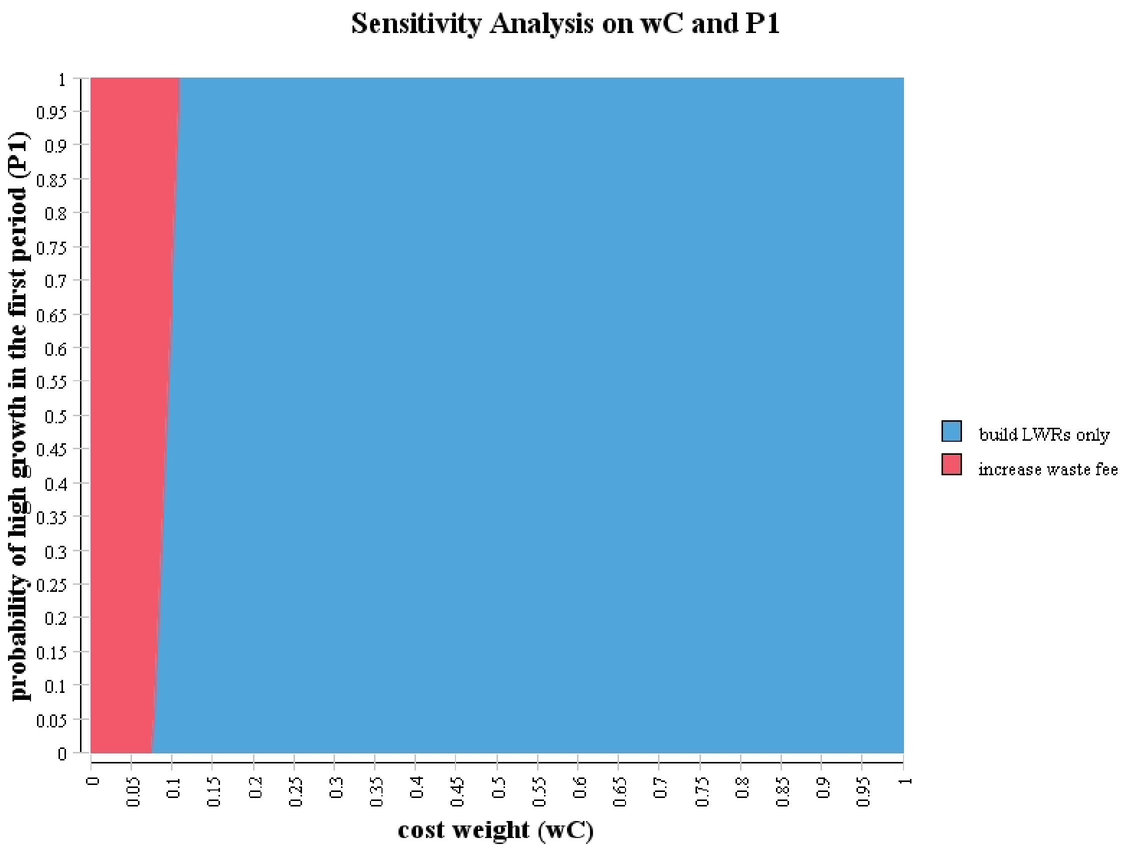

Figure 7-9: Decision sensitivity to first-period growth and cost weight for 3-period waste fee analysis

The takeaway from this analysis is not so much about true political decisions (e.g. that expensive FRs means government should actually move to increase the waste fee) as it is about learning. We understand from the analyses in Chapter 5 that under most circumstances, waiting until later in the century to build fast reactors makes sense. We also discovered that building a few fast reactors early on is attractive, even before we consider the fact that we will learn about how these reactors operate and at what cost. Here, we see that learning about FR costs is very important: if FRs are expensive, the only way they will get built is if the government increases the waste fee twice. In essence, the first increase is required to learn about the reactors and what they cost, and then the second increase actually spurs the investment. Acknowledging that the first policy attempts to incentivize fast reactor builds may not work means that trying earlier to learn about them becomes more important.

In sum: waiting until later in the century to build FRs is desirable, but we should work hard now to understand how they work and how much they will cost so that they are truly deployable. We need to know enough so that the policy instruments used to bring them up to fleet scale will be effective when and if the time comes.

Chapter 8: Conclusions

The work presented in the previous seven chapters contributes to knowledge about the use of canonical decision analysis models applied to complex systems, and increases understanding about the nuclear fuel cycle. The following three sections detail conclusions drawn from this thesis, limitations to the work and opportunities for further research, and contributions to the fields of complex system and nuclear fuel cycle analysis.

8.1 Thesis Conclusions

Conclusions from this work fall into two categories. The first body of conclusions concerns the methodology explored and tested in the thesis, which includes a hybrid decision analysis and system dynamics model applied to a complex system. The second body of conclusions offers new knowledge on evolution pathways for the U.S. nuclear fuel cycle.

8.1.1 Methodological Conclusions

Perhaps the most valuable conclusion to be drawn from this thesis is that a decision-analytic approach to synthesizing system dynamics results can help illuminate the parameters that are important decision drivers, and teach the analyst about the system. For example, it is very apparent throughout Chapters 5, 6, and 7 that decision maker preferences are a major driver of the decision outcome. The decision space is almost always more sensitive to the weighting for cost (which represents [1- the weight for waste]) than it is to the other parameters.

The parameter ranges were purposefully chosen to span nearly all plausible values, so that scenario outcomes would represent the full range of possibilities. We can thus conclude, for example, that over the range of possible cost premiums for fast reactors, decisions about closing the fuel cycle are not as sensitive to the rate of nuclear power growth as they are to the relative preference between cost and waste minimization. Similar conclusions can be drawn with confidence for several other parameters examined in Chapters 5 and 6. The particular combination of a system dynamics model that calculates outcomes combined with an orderly exploration of the parameter space via decision analysis allows us to quickly identify the parameters of significance to the decision.

A second major conclusion is that we gain important insights by considering both an "ideal" and a more realistic decision maker when modeling the system. Neither perspective should necessarily be considered in the absence of the other. From the "ideal" decision maker analysis, we learned about the effects on general welfare of various uncertainties and new structures for alternatives. These represent the best all-around decisions if a decision maker has full power over the system. Changing the perspective slightly, so that the decision maker only has power over a piece of the system and outcomes become more uncertain, shows that different decisions may need to be taken in order to get closest to the "best" decision outcomes. For example, the nuclear fuel cycle example of Chapter 7 tells us that the uncertainty in industry response indicates an earlier start to fast reactor builds than otherwise might be prescribed if the government decision maker had full control. The methodology could be useful for comparing a range of ways in which the decision maker could interact with the system (i.e. different policy structures) in order to learn about the most robust ways to move toward the "best" decisions, given limited decision-maker authority.

Two further conclusions concern the incorporation of multi-attribute utility analysis into the hybrid decision analysis-system dynamics methodology. Multi-attribute utility theory is controversial, because existing methods for eliciting proper utility curves and attribute weights may not actually account properly for decision maker preferences.(McCord & de Neufville, 1983)[4] Indeed, so far, hybrid DA-SD analyses have avoided consideration of more than one attribute (see section 2.3). But this analysis has shown that many of the general trends and conclusions are not highly dependent on the utility function shapes or scenario value function structure. The most important difference occurred when the utility function for waste changed from linear to a diminishing returns function, but conclusions about which parameters are important decision drivers and which decisions are robust to a range of preferences remained generally the same. Analysts should certainly tread carefully when choosing a utility function shape, especially when advising final decisions, but can still draw solid conclusions from this style of analysis.

---

[4] McCord and de Neufville in fact question utility theory even more deeply, asserting that the commonly-held axiom of decision maker utility curves existing independently of situational probabilities is false. If, however, one believes that a utility curve at least *exists* to describe preferences in the case of a decision tree, the fact that very disparate curves produce similar results, as demonstrated in Chapter 6, may mean that the results are believable despite this limitation.

The challenges associated with weight elicitation were avoided by employing sensitivity graphs to present the analysis results. In this way, readers see the best decision outcomes over the whole range of possible decision maker preferences. This approach has two advantages: decision makers can see which choices are robust, and arguments about the proper way to elicit weight values (or whether it can or should be done) are avoided. One major challenge is that interpreting results can be difficult, because explaining the slope of a line in the decision space requires simultaneous consideration of many different parameters. Overall, the benefits of showing decision robustness outweigh the extra burden on the analyst in explaining decision results.

Overall, the methodology shows promise for providing insights on the evolution of complex systems. Similar decision-making problems, involving high levels of interconnectedness and inertia and multiple, ill-defined decision-making groups, may benefit from this framework. Potential examples of systems to study include the U.S. electricity system as a whole, U.S. transportation infrastructure, and climate change mitigation systems and policy.

8.1.2 Conclusions for U.S. Nuclear Fuel Cycle Evolution

The most important fuel cycle conclusion drawn from this work is that the option to change course later in the century has a dramatic impact on the desirability of closing the fuel cycle now. Indeed, decision makers can wait until later in the century to significantly deploy traditional, self-sustaining fast reactors, and doing so will have cost benefits and few drawbacks in terms of waste buildup by the end of the century. Waiting to deploy large amounts of TFRs is the best decision, as long as the only alternatives are LWRs or EUFRs, and as long as the amount of SNF in the temporary stockpile is not an important factor.

Building a few fast reactors as early as 2040, however, is also highly desirable, and this conclusion is confirmed by both the “ideal” and more “realistic” decision maker perspectives. From an ideal perspective, building a few fast reactors early will entail a modest waste benefit by the end of the century and comes at relatively little added cost. It also provides for a smooth building curve (if there are no natural or artificial constraints on fast reactor builds), which could allow companies to take better advantage of learning. From a more realistic perspective, government will not be able to mandate precise numbers of fast reactors be built by industry, but will be able to offer incentives to bring TFRs into the system. This means that the government will want to learn quickly about FR costs, so that incentives can be set properly and early enough to allow for adjustments enabling a desirable build pattern for fast reactors. Overall, it makes sense to build a small fleet early on, and to reserve judgment about larger-scale deployment until later in the century.

A second conclusion is that the desirability of building FRs in the first period (in this case, at about 2040) decreases if nuclear power growth is likely to be low. Though the results of sections 5.4 and 5.5 demonstrated that building FRs in the second period will be attractive for low growth cases in order to get a waste benefit, a new decision will need to be made when those later periods are reached. The conclusion holds that a low likelihood for significant LWR waste generation now decreases the urgency of building TFRs.

A final conclusion regarding the nuclear fuel cycle is that for high-level decisions on whether to employ advanced recycling, the price of uranium, the achievable separations efficiency for spent fuel, and the shape and size of the cost learning curve are relatively unimportant parameters. All three fuel cycle system parameter sets had very little effect on the decision results, even when the parameters were varied across their full plausible ranges. This model fails, however, to account for some considerations related to the variables, including the desirability of uranium/fuel security of supply, the safety implications of different separations efficiencies, and others. Under the current paradigm, where the uranium market appears stable and safety issues surrounding recycling appear soluble, explicit consideration of the set of related issues is unlikely to produce a significant change in the results. This could change, however, and further work should explore the impact of these and other considerations of potential importance.

8.2 Limitations and Future Work

The methodology explored in this thesis has several limitations; decision results should thus be used carefully in the context of a deliberative-analytic process, as a guide and source of information rather than as direct prescription for action. Several of the most important limitations to the methodology are discussed here, with suggestions for future work to understand and alleviate them.

On major limitation of the work presented in Chapters 5, 6, and 7 is that it only considers two system objectives. The nuclear fuel cycle, like many other complex systems, involves more than two system aspects that must be considered and traded off in making decisions about system evolution. This thesis does not consider proliferation resistance or nuclear safety, despite the fact that these are two very important issues for stewards of nuclear power. In principle, further work could continue this analysis and incorporate the full complement of system objectives. Presentation and interpretation of the results, however, would be extremely difficult. Two- and three-dimensional graphs would only show part of the picture, and unpacking the drivers for decision space shapes would be made significantly more complicated with even a few more parameters. Alternatively, the analysis could be done individually with one attribute traded off against another for all possible pairs, but this would be time-consuming and results could be difficult to compare. The methodology might thus be best suited to two-objective systems.

Another limitation of this work is that only bounding point estimates are considered for the uncertainties and decision options. This is likely of greatest importance for the decision alternatives, which are modeled as stylized and discretized versions of the actual options. Future work should relax the discretization assumption, evaluating continuous portfolios of reactor build options. Follow-on studies might also consider full distributions for uncertain parameter values in place of point estimates.

A related limitation of the methodology is that any decision alternatives that are not modeled are not assessed. This is potentially problematic because there may be other decision options that would dominate the choices considered, and this methodology does not help identify what they might be or whether they exist. One example includes the possibility that the best strategy would involve all three options, where traditional fast reactors mitigate some of the existing LWR waste and then the system transitions to exclusively include enriched-uranium fed reactors. The decision options follow the same rigid structure as those of traditional fuel cycle models, where the type of reactor and date of introduction has to be specified before running a scenario. Future work could help alleviate this limitation in two ways. The simplest solution might be to employ creativity and brainstorming tools with a decision-making group in order to identify a wide range of decision options, and then to evaluate them all. Another option includes building a much more complex fuel cycle (or other system) model, upon which e.g. the introduction date for a reactor type could be optimized or decisions made endogenously. With the latter approach, the basic structure of a decision alternative (reactor type, date of introduction, date(s) of change in strategy), would be fixed. Both directions for further research could yield insights.

Another potential limitation of this approach is that it relies heavily on the accuracy of the system dynamics model. This is a challenge often described colloquially as "garbage in, garbage out," indicating that unless the underlying system model adequately represents system impacts, the decision model is useless. In this thesis, the underlying fuel cycle model FANTSY was checked extensively with widely used fuel cycle codes (see Appendix B), but the decision results should nevertheless be confirmed with one of the more sophisticated fuel cycle models. Future analysts should be aware of the importance of system modeling accuracy.

Decision analysis heavily emphasizes the relative values of scenario outcomes (e.g. one scenario scores higher than another) without considering differences in magnitude between scenario scores. Relative scenario magnitude information is used (but not shown) by TreeAge when it calculates and draws the solid-color decision spaces. In general, close to the dividing lines between decisions, scenario outcomes are closer in magnitude than elsewhere on the graph. Further unpacking of the underlying data reveals which parameters drive the magnitude differences, but an understanding of the relative "betterness" of decisions in a certain area is not immediately available with the TreeAge visualization options. Future work on advanced decision software, focusing especially on graphing features, could help make magnitude variations more clear. Use of different colors or contour lines could show places where the scenario value difference grows and shrinks, and creative calculations and uses of color might help identify which parameters are driving the differences between options. The needed information is available, but more research could help illuminate the best ways to display and digest it.

A final limitation of this work is particular to the fuel cycle system: nuclear power growth is modeled as an exogenous variable, so there is no accounting for the potential substitution of other electricity sources. Were nuclear electricity to become expensive, other types of generation would become more desirable and might reduce demand for nuclear power. These relationships are not currently modeled, though the framework does allow examination of various scenarios where there are sudden changes in nuclear demand. A more important implication of this limitation stems from the fact that government-level decisions made about the nuclear fuel cycle will have an impact on the broader electricity market. For example, if the government decides to build a fast reactor demonstration, the funds for the project cannot be applied to research in renewable energy. Similarly, if the government offers support to fast reactors, demand for the nuclear may increase even if it is not the cheapest available option. In essence, the “real” government will be deciding not only on how to evolve nuclear power, but also on which sources of electricity more generally are worth supporting. This broader decision is not considered here. Future work could address the larger question, for example by linking a model of the U.S. electricity market to the nuclear fuel cycle model and then re-framing decision analysis questions.

Despite the limitations listed in this section, the work in this thesis contributes to two areas of knowledge.

8.3 Thesis Contributions

This thesis contributes both to methods for complex system analysis and to an understanding of the U.S. nuclear fuel cycle. The major methodological contribution extends the field of hybrid decision analysis-system dynamics studies outlined in section 2.3, primarily by performing an investigation of a multi-attribute problem. The results in Chapters 5-7 show that a hybrid multi-attribute method can elucidate important drivers of system decisions, as well as highlight certain options that are robust to a range of uncertainties and preferences. The sensitivity analyses in Chapter 6 provide an initial indication that the particular functional form of the value equation may not have an enormous effect on most decision results. More work is needed to confirm this, especially on a theoretical basis. If the sensitivity results hold, this type of multi-attribute approach may prove useful for a wide range of complex system decision-making problems.

A second contribution concerns the challenge of considering multiple decision makers in the context of canonical system analysis frameworks. As discussed in section 2.1, most decision analysis methods were developed to aid a single decision maker. The analytic-deliberative approach to decision-making provides a way to incorporate multiple decision makers, but requires that the important players physically come together as one group with decision-making authority. Game-theoretic models provide another option for analysis with more than one decision maker, but the methods work best when decision makers and system states are well-defined and actions can be clearly assigned to various players. The analysis in Chapter 7, rather than pushing toward a comprehensive, multi-decision maker model, offers an incremental improvement on the single, all-powerful decision maker characterization. Following (Webster, 2008), the government is modeled as only having limited decision-making authority, and the actions of industry are modeled as an uncertainty. The contribution of this work is in showing that *both* the "ideal" and "limited" decision maker characterizations produce interesting and complementary information. Much work remains in exploring the utility and flexibility of a "limited decision maker" approach, but the results here suggest that complex system decision-making could benefit from two separate passes through a canonical system analysis model: one to determine "best" solutions from a general welfare perspective, and one to examine more practical sets of options for a more realistic decision maker.

The final set of contributions lies in the realm of nuclear fuel cycle analysis. Section 2.2 described traditional approaches to the study of the fuel cycle, all of which have provided invaluable knowledge to decision makers and without which this approach would not have been possible. The work in this thesis extends dynamic fuel cycle analyses, asking not just "what will happen" but "which fuel cycle alternatives are more desirable." The decision analysis framework developed here provides a means to synthesize fuel cycle scenario results, examine tradeoffs, and uncover the parameters that matter *in the context of a decision about the system*. None of these insights would be readily obtainable with a system dynamics model alone; indeed, DSARR, an important dynamic fuel cycle analysis, does not address these issues. Though decision rules can be coded into a system dynamics model, and some sort of scoring mechanism applied, inserting a large number of cumbersome if-then statements into a complex nuclear fuel cycle model (e.g. VISION with ~1 million input variables and over 30 system dynamics page views of code) is not likely to be tenable. The decision-analytic hybrid approach allows for a relatively rapid analysis, even with a complex underlying fuel cycle model, that explicitly highlights uncertainties and value tradeoffs.

The MIT Future of the Nuclear Fuel Cycle study (Kazimi et al., 2011) is a good example of a more comprehensive, qualitative + quantitative analysis of fuel cycle issues, intended for government-level decision makers. The analytic framework and results presented in this thesis complement that type of analysis, helping to highlight some of the most important issues worthy of further study as well as identifying which alternatives may hold the most promise for a range of future scenarios. The results in this thesis are not intended to stand alone as prescriptions, but to help decision makers sort through some of the information provided by the MIT study and others to come. The author's sincere hope is that this thesis provides a strong analytic basis for real decision-maker discussions regarding not only the nuclear fuel cycle, but also other long-lived, complex technological systems that would benefit from a hybrid modeling approach.

References

A Joint Report by the OECD Nuclear Energy Agency and the International Atomic Energy Agency OECD. (2010). *Uranium 2009: Resources, production and demand*. Paris, France: OECD.

Boston Consulting Group. (2006). *Economic assessment of used nuclear fuel management in the united states*

Bunn, M., Holdren, J. P., Fetter, S., & van der Zwaan, B. (2005). The economics of reprocessing versus direct disposal of spent nuclear fuel. *Nuclear Technology, 150*, 209-230.

Caplin, A., & Leahy, J. (2000). *The social discount rate* (working paper. Cambridge, MA: National Bureau of Economic Research.

Clemen, R. T. (1996). *Making hard decisions: An introduction to decision analysis*. Boston: Duxbury Press.

Committee on Risk Characterization, National Research Council. (1996). *Understanding risk: Informing decisions in a democratic society*. Washington, D.C.: National Academies Press.

D. Bell, H. Raiffa, A. Tversky (Ed.). (1988). *Decision making: Descriptive, normative, and prescriptive interactions*. New York, NY: Cambridge University Press.

de Neufville, R. (2010). *Intro to drivers of flexibility: The time value of money*. Unpublished manuscript.

De Roo, G., & Parsons, J. (2009). *Nuclear fuel recycling, the value of separated transuranics and the levelized cost of electricity*. Unpublished manuscript.

Deutch, J., Moniz, E., & et al. (2003). *The future of nuclear power: An interdisciplinary MIT study* No. 0-615-12420-8). Cambridge, MA: Massachusetts Institute of Technology.

Deutch, J., & et al. (2009). *Update of the MIT 2003 future of nuclear power.* Massachusetts Institute of Technology.

Dixon, B., Halsey, B., Kim, S., Matthern, G., Piet, S., & Shropshire, D. (2008). *Dynamic systems analysis report for nuclear fuel recycle*. Idaho Falls, Idaho: Systems Analysis Campaign, Idaho National Laboratory.

Du, Y., & Parsons, J. (2009). *Update on the cost of nuclear power*. Cambridge, MA: Center for Energy and Environmental Policy Research, MIT.

Grout, P. (1982). Welfare economics of decision making with changing preferences. *The Review of Economic Studies, 49*(1), 83-90.

Grubler, A. (2009). *An assessment of the costs of the french nuclear PWR program 1970-2000* No. IR-09-036). Laxenburg, Austria: International Institute for Applied Systems Analysis.

Guerin, L., & et al. (2008). *A benchmark study of computer codes for system analysis of the nuclear fuel cycle* No. MIT-NFC-TR-105). Cambridge, MA: Center for Advanced Nuclear Energy Systems, MIT.

Hammond, P. (1976). Changing tastes and coherent dynamic choice. *The Review of Economic Studies, 43*(1), 159-173.

Hansson, S. O. (2005). *Decision theory: A brief introduction*. Stockholm: Royal Institute of Technology (KTH).

Hess, H., et al. (1957). *The disposal of radioactive waste on land: Report of the committee on waste disposal of the division of earth sciences* No. 519). Washington, D.C.: National Academy of Sciences - National Research Council.

Hewitt, G. (2011, 14 June). No to nuclear and no to Berlusconi. *BBC World News,*

Hoffman, E. A., Yang, W. S., & Hill, R. N. (2008). *Preliminary core design studies for the advanced burner reactor over a wide range of conversion ratios*. Argonne, IL: Argonne National Laboratory.

Hovmand, P., & Ford, D. (2009). Real options approach for innovation implementation strategies. *Conference Proceedings of the 27th International Conference of the System Dynamics Society,* Albuquerque, NM.

Howard, R. (1988). Decision analysis: Practice and promise. *Management Science, 34*(6), 679.

Kazimi, M., Moniz, E., Forsberg, C., & et al. (2011). *The future of the nuclear fuel cycle*. Cambridge, MA: MIT.

Keeney, R. (1982). Decision analysis: An overview. *Operations Research, 30*(5), 803.

Martin, W., Ahearne, J., & and the Nuclear Energy Advisory Committee. (2008). *Nuclear energy: Policies and technology for the 21st century*. Washington, D.C.:

McCord, M., & de Neufville, R. (1983). Fundamental deficiency of expected utility decision analysis. In S. French, R. Hartley, L. C. Thomas & D. J. White (Eds.), *Multi-objective decision making*. New York, NY: Academic Press.

Nuclear energy for the future: Required research and development capabilities - an industry perspective (2008). . Washington, D.C.: Batelle, for the U.S. Department of Energy.

Osgood, N. (2005). Combining system dynamics and decision analysis for rapid strategy selection. *Conference Proceedings of the the 23rd International Conference of the System Dynamics Society,* Boston, MA.

Piet, S., Gibson, P., Joe, J., Kerr, T., & Nitschke, R. (2003). Making sustainable decisions using the KONVERGENCE framework. *Waste Management Conference 2003,* Tucson, AZ.

Piet, S., Matthern, G., Jacobson, J., Laws, C., & Cadwallader, L. (2006). *Fuel cycle scenario definition, evaluation, and tradeoffs*. Idaho Falls, Idaho: Idaho National Laboratory, Fission and Fusion Systems Department.

Pope, M., Driscoll, M., & Hejzlar, P. (2004). *Reactor physics design of supercritical CO_2-cooled fast reactors*. Cambridge, MA: Center for Advanced Nuclear Studies, MIT.

Saleh, J., Hastings, D., & Newman, D. (2001). *Extracting the essence of flexibility in system design.* Unpublished manuscript.

Shropshire, D., Piet, S., Soelberg, N., Cherry, R., Henry, R., Meikrantz, D., et al. (2009). *System losses and assessment trade study*. Idaho Falls, Idaho: Advanced Fuel Cycle Initiative, U.S. Department of Energy.

Sussman, J., Dodder, R., McConnell, J., Mostashari, A., & Sgouridis, S. (2007). *The CLIOS process: A user's guide.* Unpublished manuscript.

Tan, B., Anderson, E., Dyer, J., & Parker, G. (2010). Evaluating system dynamics models of risky projects using decision trees: Alternative energy projects as an illustrative example. *System Dynamics Review, 26*(1), 1.

Tang, V. (2006). Corporate decision analysis: An engineering approach. (Ph.D., MIT).

The University of Chicago. (2004). *The economic future of nuclear power*

U.S. Department of Energy. (2011). *Fuel cycle research and development: Fact sheet*. Washington, D.C.: Office of Nuclear Energy, U.S. Department of Energy.

Webster, M. (2008). Incorporating path dependency into decision-analytic methods: An application to global climate-change policy. *Decision Analysis, 5*(2), 60.

Weil, R. and Apostolakis, G. (2001). A methodology for the prioritization of operating experience in nuclear power plants. *Reliability Engineering & System Safety*, *74*(1), 23.

Wigeland, R., & Bauer, T. (2004). *Repository benefits of AFCI options*. Argonne, IL: Nuclear Engineering Division, Argonne National Laboratory.

World Nuclear Association. (2011a, March 10, 2011). Double attack on U.S. nuclear waste fees. *World Nuclear News,*

World Nuclear Association. (2011b). *Nuclear power in france.* Retrieved April 11, 2011, from http://www.world-nuclear.org/info/inf40.html

Yucca mountain cost estimate rises to $96 billion. (2008, August 6). *World Nuclear News.*

Appendix A: The Flexible Advanced Nuclear Technology Simulation by Year (FANTSY) Source Code

The FANTSY source code is written in MATLAB, saved as .m files. The wrapper code (rapper.m) is called directly from the MATLAB command line, which in turn calls the function sbF one time for each scenario simulated. The results from each scenario are then written to an Excel file called ToTree.xlsx.

The particular version of the code presented below simulates a two-period tree with three reactor options (LWR, TFR, and EUFR) and two growth rates per period. The wrapper code is adjusted substantially depending on the particular needs of the tree. In this case, several “blank” scenarios are coded into the Excel file, in order to more accurately reflect the scenario indexing present in the tree. Note also that the wrapper script calls a function xlswrite1; this is distinct from xlswrite only in that Excel is opened at the beginning of the write functions and closed at the end, rather than opened and closed for each instantiation of xlswrite. This saves hours of computing time.

rapper.m

```
% this is the wrapper script which will run through iterations of the fuel
% cycle model, each time with different inputs corresponding to different
% branches of a tree

% constants
initialCR = 1.0;

%DECISION AND UNCERTAINTY VALUE LISTS: CHANGE MANUALLY
U1 = [0.025, 0.005]; % growth rates in decimal
CS = [2200]; % centralized storage open
D1 = {'TFR', 'EUFR', 'LWR'}; % reactor type in period 1
LS = [0.01]; % loss fractions for TRU/FP
U2 = [0.04, 0.005]; % growth rates for 2nd period
D2 = {'TFR', 'EUFR', 'LWR'}; % reactor type in period 2
P1 = [1, 0.1]; %percent of possible FRs built during period 1
P2 = [1, 0.5]; % percent of possible FRs built in period 2

ActiveIndices = 1;
%VECTOR LENGTHS FOR DECISIONS AND UNCERTAINTIES
U1num = length(U1);
if U1num > 1; ActiveIndices = ActiveIndices + 1; end
CSnum = length(CS);
```

```
if CSnum > 1; ActiveIndices = ActiveIndices + 1; end
D1num = length(D1);
if D1num > 1; ActiveIndices = ActiveIndices + 1; end
LSnum = length(LS);
if LSnum > 1; ActiveIndices = ActiveIndices + 1; end
U2num = length(U2);
if U2num > 1; ActiveIndices = ActiveIndices + 1; end
D2num = length(D2);
if D2num > 1; ActiveIndices = ActiveIndices + 1; end
P1num = length(P1);
if P1num > 1; ActiveIndices = ActiveIndices + 1; end
P2num = length(P2);
if P2num > 1; ActiveIndices = ActiveIndices + 1; end
%maxScenarios = U1num*D1num*U2num*LSnum*D2num*CSnum*P1num...
  %*P2num;

maxScenarios = 120;
%SCENARIO TRACKING
IndexList = cell(maxScenarios,ActiveIndices); %rows are scenarios
IndexTitles = {'First Reactors', '% Build1', ...
  'Growth1', 'Second Reactors', '% Build2', 'Growth2'};

CostHeader = {'NatU', 'SWU', 'LWRfuelfab', 'FRfuelfab', ...
  'LWRsConstructed', 'FRsConstructed', 'LWRsOperating', ...
  'FRsOperating', 'ElectricityLWR', 'ElectricityFR', ...
  'SNFreprocessed', 'FRreprocessed', 'LWRsDecommd', 'FRsDecommd', ...
  'LWRdryStor', 'FRdryStor', 'CSbuilt', 'CSoperating', 'CStransport'};
WasteHeader = {'SNF', 'FP', 'HLW','SNF-scaled'};
CostOutput = cell(maxScenarios,19); % rows are scenarios, cols are
  % variables (17 vars needed for cost model), 100 years for each var
WasteOutput = zeros(maxScenarios,4); % waste is a simple matrix
  % with scenarios as rows and 4 variables assessed at 2109 (SNF, HLW,
  % FP, SNF stored in CS)
Output = cell(maxScenarios,20); % final cell is waste output

for index1 = 1:D1num
  for index2 = 1:P1num
    for index3 = 1:U1num
      for index4 = 1:D2num
        for index5 = 1:P2num
          for index6 = 1:U2num
          scenario = 48*(index1-1) + 24*(index2-1) + ...
            12*(index3-1) + 4*(index4-1) + 2*(index5-1)...
            + index6;
          if (index1 == 3 && index2 == 2), break, end
          if ~(index1 == 2 && index4 == 1)
          Output(scenario,:) = ...
            sbF(initialCR,U1(index3),CS,D1(index1),...
              LS,U2(index6),D2(index4),P1(index2),...
              P2(index5));
          WasteOutput(scenario,:) = Output{scenario,20}();
          IndexList(scenario,:) = {scenario ...
            D1(index1) P1(index2) U1(index3) ...
            D2(index4) P2(index5) U2(index6)};
          display(scenario);
          else
```

```
            blank = zeros(100,1);
            for yo = 1:20
             Output{scenario,yo} = blank;
            end
            sim = sprintf('blank %d',scenario);
            display(sim);
            Output{scenario,20} = [0 0 0 0];
            IndexList(scenario,:) = {scenario 0 0 0 0 0 0};
          end
         end
       end
     end
   end
  end
end


CostOutput = Output(1:maxScenarios,1:19);
years = (2010:2109)';
scs = (1:120)';
WasteOutput = horzcat(scs,WasteOutput);

Excel = actxserver ('Excel.Application');
File='C:\Users\Lara\Documents\MATLAB\ToTree.xlsm';
if ~exist(File,'file')
  ExcelWorkbook = Excel.workbooks.Add;
  ExcelWorkbook.SaveAs(File,1);
  ExcelWorkbook.Close(false);
end
invoke(Excel.Workbooks,'Open',File);

endrow = 0;
for is = 1:maxScenarios
  for iv = 1:19
    startrow = endrow + 2;
    headerrow = endrow +1;
    headspot = sprintf('B%d',headerrow);
    yearspot = sprintf('A%d',startrow);
    xlswrite1(File, CostHeader, 'Data', headspot);
    xlswrite1(File, years, 'Data', yearspot);
    col = 65 + iv; % ASCII dec code for letters, A is 65
    column = char(col);
    spot = sprintf('%s%d',column,startrow);
    xlswrite1(File, CostOutput{is,iv}(), 'Data', spot);
  end
  endrow = startrow + 100;
end

xlswrite1(File, WasteHeader, 'WasteData','B1');
xlswrite1(File, WasteOutput, 'WasteData', 'A2');
xlswrite1(File, IndexTitles, 'Scenarios', 'B1');
xlswrite1(File, IndexList, 'Scenarios', 'A2');
for il = 1:120
  r1 = IndexList{il,2};
  r2 = IndexList{il,5};
  ill = il+1;
```

```
  place1 = sprintf('B%d',ill);
  place2 = sprintf('E%d',ill);
  xlswrite1(File, r1, 'Scenarios', place1);
  xlswrite1(File, r2, 'Scenarios', place2);
end
xlswrite1(File, initialCR, 'Data', 'A1');
invoke(Excel.ActiveWorkbook,'Save');
Excel.Quit
Excel.delete
clear Excel
```

The code below (sbF.m) is also often adjusted to reflect the particular needs of a decision tree. The basic structure, however, is the same for all studies presented in this thesis.

sbF.m

```
function ToPass = sbF(initialCR,U1,CS,D1,LS,U2,D2,P1,P2)
% function to calculate 100-year mass flows for the U.S. nuclear fuel cycle
% sbF (base = sbF14)

per1start = 2040;
per2start = 2065;

year1 = per1start;
year2 = per2start;
yrChange1 = year1 - 2010;
yrChange2 = year2 - 2010;
growth0 = 0.012; %growth rates in fractional form
growth1 = U1;
growth2 = U2;

% set starts and stops to end of simulation (aka never): they
% will only change if the scenario calls
EUFRstartyear = 2200;
FRstartyear = 2200;
justfastreprocstart = 2200;
reprocstart = 2200;

if ~strcmp(D1,'EUFR') && ~strcmp(D2,'EUFR')
  EUscenario = 0;  % no EUFRs are built
elseif ~strcmp(D1,'TFR') && ~strcmp(D2,'TFR')
  EUscenario = 1;  % all FRs that start up are EUFRs
elseif strcmp(D1, 'TFR') && strcmp(D2, 'EUFR')
  EUscenario = 2;  % some FRs are EUFRs, some are traditional;
else
  display('Bad EUFR/TFR scenario');
end

% OT or FR flag
if strcmp(D1,'LWR') && strcmp(D2,'LWR')
  OTflag = 1;
else
  OTflag = 0; % if zero, this is an FR run; if 1, this is an OT run
end
```

```
% FR start flags
if EUscenario == 0 && OTflag == 0 % condition for a TFR-only run
  if strcmp(D1,'TFR')
    FRstartyear = per1start;
  elseif strcmp(D2,'TFR')
    FRstartyear = per2start;
  else
    display('Flag error: FR start flags');
  end
  if strcmp(D2,'LWR')
    FRstopyear = per2start;
  else
    FRstopyear = 2200;
  end
  FRstart = FRstartyear - 2010 - 4;
  reprocstart = FRstart - 5;
  FRend = FRstopyear - 2010;
end

% EUFR start flags - set the general FR start flags with EUFR start times
if EUscenario == 1 && OTflag == 0 % condition for all-EUFR run
  if strcmp(D1,'EUFR')
    EUFRstartyear = per1start;
  elseif strcmp(D2,'EUFR')
    EUFRstartyear = per2start;
  else
    display('Flag error: EUFR start flags');
  end
  if strcmp(D2,'LWR')
    EUFRstopyear = per2start;
  else
    EUFRstopyear = 2200;
  end
  FRstart = EUFRstartyear - 2010 - 4;
  EUFRstart = FRstart;
  justfastreprocstart = FRstart - 5;
  reprocstart = 1000;
  FRend = EUFRstopyear - 2010;
elseif EUscenario == 2 % condition for TFR switch to EUFR (can't reverse)
  FRstartyear = 2040;
  FRstart = FRstartyear - 2010 - 4;
  EUFRstartyear = 2065;
  EUFRstart = EUFRstartyear - 2010 - 4;
  reprocstart = FRstart - 5;
  justfastreprocstart = EUFRstartyear - 2010;
  FRend = 1000; % for EUscenario 2, FRs go to end of century always
end

fractionbuild = zeros(1,106); %fraction of possible FRs actually built
for i = 1:29
  fractionbuild(i) = 0; % 2010-2039
end
for i = 30:54
  fractionbuild(i) = P1; % 2040-2064
end
```

```
for i = 55:106
  fractionbuild(i) = P2; % 2065-end
end

% CS flag - reads in the start year for centralized storage of LWR SNF
% also initialize CS-related parameters
CSstart = CS-2010;
CSbuilt = zeros(1,100);
CSoperating = zeros(1,100);
opCS = 0;
spaceCS = 0;
CSsnftransported = zeros(1,100);
CSamountstored = zeros(1,100);
removalflagoff = 1;

% initialize loss-related parameters
switch LS
  case 0.1
    lossfraction = 0.1;
    squish = 5;
  case 0.01
    lossfraction = 0.01;
    squish = 20;
  case 0.001
    lossfraction = 0.001;
    squish = 42;
  otherwise
    display('bad loss fraction');
    return
end

% constants: these do not change for any run
legacySF = 53000; % MT
reactors = [0  0  0  0  0  0  0  0  0  0  0  0  0  0  0 ...
  0  0  3  3  2  6  11 13 2  7  3  3  0  3  4  5  4 ...
  8  6  9  5  1  2  1  0  0  1  0  1  0  0  0  0  0 ...
  0  0  0  0  0  0  0  0  0  0  1]; % with legacies
LWRcapfac = 0.9;
FRcapfac = 0.85;
capacity2010 = 104*LWRcapfac; % GWe - 100 for benchmark, 104 otherwise
SNFfpfraction = 0.053; % Hoffman 100s col D
EUFRenrichment = 0.195;
EUFRcoresize = 45.5; % commensurate with CR = 1.0 FR
CR = initialCR;

if CR == 1.0
  coreFR = 45.5;
  yearlyFRneed = 11.19; % total fuel needed yearly (TRU + other)
  truFractionFR = 0.141; % fraction of FR SNF that is TRU - Hoffman
  UtruFractionFR = 0.922; % fraction of FR SNF that is HM - Hoffman
  Ffpfraction = 0.078; % Hoffman 600s col AP
  fuelperFR = 45.5; % MT per year - from CAFCA 'FR' Fuel Core Mass
  fuelperyearFR = 11.19;
  truFracFresh = 0.139;
  rpnCap = 500; % nominal reproc plant capacity in MTHM/year, acc to CAFCA
elseif CR == 0.5
```

```
  coreFR = 25.66;
  yearlyFRneed = 6.19;
  truFractionFR = 0.270;
  UtruFractionFR = 0.859; % fraction of FR SNF that is HM - Hoffman
  Ffpfraction = 0.142; % Hoffman 600s col AD
  fuelperFR = 25.66;
  fuelperyearFR = 6.19;
  truFracFresh = 0.333;
  rpnCap = 200;
else
  display('Bad Conversion Ratio');
end

% initializations of temporary and output vectors
demandzero = capacity2010;
demand = zeros(1,106);
underconstruction = zeros(1,5);
ReactorsConstructed = zeros(1,100);
ReactorsOperating = zeros(1,100);
ReactorsDecommissioned = zeros(1,100);

FRsConstructed = zeros(1,100);
FRsOperating = zeros(1,100);
FRsDecommissioned = zeros(1,100);
underconstructionFR = zeros(1,5);
makeupfuel = zeros(1,100);
operatingFRs = 0;
fastreactors = zeros(1,60);

underconstructionEUFR = zeros(1,5);
EUFRsConstructed = zeros(1,100);
NatUforEUFRs = zeros(1,100);
SWUEUFR = zeros(1,100);

TotalFRsConstructed = zeros(1,100);

% demand vectors
Forecast = zeros(1,100);
LWRdemand = zeros(1,100);

%%%% REPROCESSING CAPACITY %%%%
maxSNFyearlyadd = 500; % amount of MT reproc capacity that can be built
yearly
maxFRyearlyadd = 50;
toProcThermal = 0;
toProcFast = 0;
reprocPlantLife = 50; % years reprocessing plants will be in operation
tRC = 0;
fRC = 0;
thermalRP = zeros(1,100); % reprocessing capacity by year in MTHM/year
fastRP = zeros(1,100);
FP = zeros(1,100);
ThFP = zeros(1,100);
FFP = zeros(1,100);
```

```
FRfuelfab = zeros(1,100);

%%%% FUEL MODULE %%%%

lwrenrichment = 0.042;
xf = 0.0071; % feed enrichment (natural uranium at 0.71%)
xt = 0.0025; % tail assay (0.25%)
fuelperLWR = 87.76; % MT/reactor
fuelperyearLWR = 19.73; % MT/reactor/year


lwrTRUfraction = 0.0128; % fraction MTIHM - from CAFCA 'LWR':M11
coolingpool = ones(1,5)*2000; % size of this vector is the wet cooling time
coolingpoolFR = zeros(1,2);

% yearly front-end output data
FuelConsumed = zeros(1,100); % MTHM
NaturalUraniumConsumedlwr = zeros(1,100);
SWUlwr = zeros(1,100);
SWU = zeros(1,100); % kg-SWU
NaturalUraniumConsumed = zeros(1,100); % MT
LWRfuelfabricated = zeros(1,100);
FRfuelfabricated = zeros(1,100);

electricityLWR = zeros(1,100);
electricityFR = zeros(1,100);

electricityOT = zeros(1,100);
electricityTFR = zeros(1,100);
electricityEUFR = zeros(1,100);

% yearly back-end output data
SNFcooled = zeros(1,100);
FRcooled = zeros(1,100);
TRUlwr = zeros(1,100);
TRUfr = zeros(1,100);
TRUlosses = zeros(1,100);
TRUstock = zeros(1,100);
SNFstock = zeros(1,100);
FRstock = zeros(1,100);
SNFinventory = legacySF;
FRinventory = 0;
TRUinventory = 0;
SNFreprocessed = zeros(1,100);
FRreprocessed = zeros(1,100);


% nuclear electricity growth and forecasting
% run through 100 years to create demand vector
for year = 1:106
  if year == 1
    demand(year) = demandzero;
  elseif year <= yrChange1
    demand(year) = demand(year - 1)*(1+growth0);
  elseif year <= yrChange2
```

```
    demand(year) = demand(year - 1)*(1+growth1);
  else
    demand(year) = demand(year - 1)*(1+growth2);
  end
end

for year = 1:100 % 100 year simulation, 2010 - 2109
    if year > 1
      EUFRstarting = EUFRsConstructed(year - 1);
    else
      EUFRstarting = 0;
    end

    forecast = demand(year + 5);
    ops = sum(reactors(6:60));
    ReactorsOperating(year) = sum(reactors);
    generation = ops*LWRcapfac + sum(fastreactors(6:60))*FRcapfac...
      + EUFRstarting*FRcapfac;
    futuregen = sum(underconstruction(1:4))*LWRcapfac + ...
      sum(underconstructionFR(1:4))*FRcapfac + ...
      sum(underconstructionEUFR(1:4))*FRcapfac;
    forecastedneed = forecast-(generation+futuregen);
    Forecast(year) = forecastedneed;

% build LWRs
  if OTflag
    reactorsneeded = forecastedneed/(LWRcapfac);
    if reactorsneeded >= 1
      underconstruction(5) = floor(reactorsneeded);
    else
      underconstruction(5) = 0;
    end
  end

% FRONT-END FUEL CALCS
   FuelConsumed(year) = fuelperLWR*reactors(60) + fuelperyearLWR*...
     sum(reactors(1:59));
   NaturalUraniumConsumedlwr(year) = FuelConsumed(year)*(lwrenrichment-...
     xt)/(xf-xt);
   vp = (2*lwrenrichment-1)*log(lwrenrichment/(1-lwrenrichment));
   vf = (2*xf-1)*log(xf/(1-xf));
   vt = (2*xt-1)*log(xt/(1-xt));
   SWUlwr(year) = (FuelConsumed(year)*vp +
(NaturalUraniumConsumedlwr(year)...
     - FuelConsumed(year))*vt - NaturalUraniumConsumedlwr(year)*vf)*1000;
   LWRfuelfabricated(year) = FuelConsumed(year);

   % fuel EUFRs that are starting up
   if EUscenario == 1 || EUscenario == 2
     if year > EUFRstart
      fuelEUFR = EUFRcoresize*EUFRstarting;
      NatUforEUFRs(year) = fuelEUFR*(EUFRenrichment - xt)/(xf-xt);
      vpEU = (2*EUFRenrichment-1)*log(EUFRenrichment/(1-EUFRenrichment));
      SWUEUFR(year) = (fuelEUFR*vpEU + (NatUforEUFRs(year)...
        - fuelEUFR)*vt - NatUforEUFRs(year)*vf)*1000;
      LWRfuelfabricated(year) = LWRfuelfabricated(year) + fuelEUFR;
```

```
      end
    end

  % totals
  NaturalUraniumConsumed(year) = NaturalUraniumConsumedlwr(year) + ...
    NatUforEUFRs(year);
  SWU(year) = SWUlwr(year) + SWUEUFR(year);

% CALCULATE ELECTRICITY FOR THIS YEAR
electricityLWR(year) = (sum(reactors)*LWRcapfac)*365*24*10^6; % kWh produced
electricityFR(year) = ((sum(fastreactors) + EUFRstarting)*FRcapfac)...
  *365*24*10^6;

   %%% REPROCESS AND FUEL UP %%%

% REPROCESSING - note all FRSNF/SNF has same losses
if EUscenario ~= 1 && year > reprocstart
  % build reprocessing capacity
  if year < 40
    buildtRC = 500; %MT/year that can be added each year (see Benchmark)
    buildfRC = 50;
  elseif year < 55
    buildtRC = 1000;
    buildfRC = 50;
  else
    buildtRC = 1000;
    buildfRC = 150;
  end
end

if year > justfastreprocstart || EUscenario == 1
  % build JUST FAST reprocessing capacity
  if year < 40
    buildtRC = 0; %MT/year that can be added each year (see Benchmark)
    buildfRC = 50;
  elseif year < 55
    buildtRC = 0;
    buildfRC = 100;
  else
    buildtRC = 0;
    buildfRC = 500;
  end
end

if year > reprocstart || year > justfastreprocstart
  influxCapT = (SNFcooled(year-1)/1000)*(reprocPlantLife/50);
  influxCapF = (FRcooled(year-1)/rpnCap)*(reprocPlantLife/50);
  inventoryCapT = SNFinventory/(50000); % 50 = instantaneous depletion T
  inventoryCapF = FRinventory/(rpnCap*50);

  thNeeded = (influxCapT + inventoryCapT)*1000 - tRC;
  fNeeded = (influxCapF + inventoryCapF)*rpnCap - fRC;

  thBuilt = min(buildtRC,thNeeded);
  fBuilt = min(buildfRC,fNeeded);
```

```
  if thBuilt < 0, thBuilt = 0; end
  if fBuilt < 0, fBuilt = 0; end

  tRC = tRC + thBuilt;
  fRC = fRC + fBuilt;

  thermalRP(year) = tRC;
  fastRP(year) = fRC;

  if (EUscenario == 2 && year > EUFRstart + 4) || year > FRend
    toProcThermal = 0;
    toProcFast = min(fRC, FRinventory);
  else
    toProcThermal = min(tRC, SNFinventory);
    toProcFast = min(fRC, FRinventory);
  end

  SNFreprocessed(year) = toProcThermal;
  FRreprocessed(year) = toProcFast;
  ThFP(year) = toProcThermal*SNFfpfraction;
  FFP(year) = toProcFast*Ffpfraction/UtruFractionFR;
  FP(year) = ThFP(year) + FFP(year);

  TRUlwr(year) = lwrTRUfraction*toProcThermal;
  TRUfr(year) = truFractionFR*toProcFast;
  SNFinventory = SNFinventory - toProcThermal;
  FRinventory = FRinventory - toProcFast;

  if toProcThermal > (SNFinventory - CSamountstored(year-1))
    removalflagoff = 0;
    CSamountstored(year) = max(0,(CSamountstored(year-1) - ...
      (toProcThermal - (SNFinventory - CSamountstored(year-1)))));
  end

  TRUlosses(year) = lossfraction*(TRUfr(year)+TRUlwr(year));
  TRUlwr(year) = (1-lossfraction)*TRUlwr(year);
  TRUfr(year) = (1-lossfraction)*TRUfr(year);
end

if OTflag == 0
% FUEL FAST REACTORS
  TRUinventory = TRUlwr(year) + TRUfr(year) + TRUinventory;
  fuelneededFR = (sum(fastreactors(1:59))*yearlyFRneed + fastreactors(60)*...
    coreFR)*truFracFresh; %fastreactors(60) is zero for EUFR scenarios
  FRfuelfabricated(year) = fuelneededFR/truFracFresh;

  if fuelneededFR > TRUinventory
    makeupfuel(year) = fuelneededFR - TRUinventory;
    TRUinventory = 0;
  else
    TRUinventory = TRUinventory - fuelneededFR;
    FRfuelfab(year) = fuelneededFR/truFracFresh;
  end

%%%%% BUILD FRs & LWRs %%%%%
```

```
  unprocessedTRU = (SNFinventory - toProcThermal)*lwrTRUfraction*0.99...
    + (FRinventory - toProcFast)*truFractionFR*0.99;
  coolingstorage = sum(coolingpool)*lwrTRUfraction*0.99 ...
    + sum(coolingpoolFR)*truFractionFR*0.99;

% put EUFRs into normal fast reactor cycle
fastreactors(60) = fastreactors(60) + EUFRstarting;

% FORECAST FRs
 if year > FRstart && year < FRend
  lifetimeTRUneed = (fuelperFR + fuelperyearFR*59)*truFracFresh;
  fleetforecast = 60*(TRUinventory)...
    /lifetimeTRUneed;
  operatingFRs = sum(fastreactors);
  reactorsneeded = forecastedneed/(FRcapfac);
  if EUscenario == 0
    tobuild = min(fleetforecast - operatingFRs, reactorsneeded);
  elseif EUscenario == 1
    tobuild = -1;
    tobuildEUFR = reactorsneeded;
  elseif EUscenario ==2 && year < EUFRstart
    tobuild = min(fleetforecast - operatingFRs, reactorsneeded);
    tobuildEUFR = -1;
  elseif EUscenario == 2 && year >= EUFRstart
    tobuild = -1;
    tobuildEUFR = reactorsneeded;
  else
    display('Unknown EU Scenario');
    return
  end
 else
   tobuild = -1;
   tobuildEUFR = -1;
   operatingFRs = sum(fastreactors);
 end

 FRsOperating(year) = operatingFRs;

% BUILD FRs
 if tobuild >= 1
   actualbuild = tobuild * fractionbuild(year);
   underconstructionFR(5) = floor(actualbuild);
 else
   underconstructionFR(5) = 0;
 end

 if tobuildEUFR >= 1
   actualbuildEUFR = tobuildEUFR * fractionbuild(year);
   underconstructionEUFR(5) = floor(actualbuildEUFR);
 else
   underconstructionEUFR(5) = 0;
 end

 totalconstructionFR = underconstructionFR(5) + underconstructionEUFR(5);
```

```
% build remaining LWRs
  remainingneed = forecastedneed - totalconstructionFR*FRcapfac;
  LWRdemand(year) = remainingneed;
  reactorsneeded = remainingneed/(LWRcapfac);
    if reactorsneeded >= 1
      underconstruction(5) = floor(reactorsneeded);
    else
      underconstruction(5) = 0;
    end

end

    %%% DISCHARGE AND DECOMMISSION %%%

% LWR SNF
  snfdischarged = reactors(1)*fuelperLWR +...
    sum(reactors(2:60))*fuelperyearLWR;
  SNFcooled(year) = coolingpool(1);
  coolingpool = circshift(coolingpool,[0, -1]);
  coolingpool(5) = snfdischarged;
  SNFstock(year) = SNFinventory + SNFcooled(year);
  SNFinventory = SNFstock(year);

% put SNF into CS
%{
if year >= CSstart && removalflagoff == 1
  if SNFstock(year) > 20000 && spaceCS < 3000 % build CS facilities
    CSbuilt(year) = 1;
    opCS = opCS + 1;
    spaceCS = spaceCS + 15000;
  end
  CSoperating(year) = opCS;
  if spaceCS > 15 && SNFstock(year) > 15  % transport SNF
    CSsnftransported(year) = min(2500,SNFstock(year));
    spaceCS = spaceCS - CSsnftransported(year);
    CSamountstored(year) = CSamountstored(year-1) + CSsnftransported(year);
  end
end
%}



% FR SNF
  frdischarged = fastreactors(1)*fuelperFR + ...
    sum(fastreactors(2:60))*fuelperyearFR;
  FRcooled(year) = coolingpoolFR(1);
  coolingpoolFR = circshift(coolingpoolFR,[0, -1]);
  coolingpoolFR(2) = frdischarged;
  FRstock(year) = FRinventory + FRcooled(year);
  FRinventory = FRstock(year);

  TRUstock(year) = TRUinventory;

% step reactor arrays
   FRsConstructed(year) = underconstructionFR(1);
   FRsDecommissioned(year) = fastreactors(1);
```

```
    fastreactors = circshift(fastreactors, [0,-1]);
    fastreactors(60) = underconstructionFR(1);
    underconstructionFR = circshift(underconstructionFR, [0, -1]);

    EUFRsConstructed(year) = underconstructionEUFR(1);
    underconstructionEUFR = circshift(underconstructionEUFR, [0, -1]);

    TotalFRsConstructed(year) = FRsConstructed(year) + EUFRsConstructed(year);

    ReactorsConstructed(year) = underconstruction(1);
    ReactorsDecommissioned(year) = reactors(1);
    reactors = circshift(reactors, [0,-1]);
    reactors(60) = underconstruction(1);
    underconstruction = circshift(underconstruction, [0, -1]);

end

scaledSNF = sum(SNFreprocessed)/squish;

Waste = [SNFstock(100), sum(FP), sum(TRUlosses), scaledSNF];

ToPass = {NaturalUraniumConsumed', SWU', LWRfuelfabricated', ...
  FRfuelfabricated', ReactorsConstructed', TotalFRsConstructed', ...
  ReactorsOperating', FRsOperating', electricityLWR', ...
  electricityTFR', SNFreprocessed', FRreprocessed', ...
  ReactorsDecommissioned', FRsDecommissioned', SNFcooled', ...
  FRcooled', CSbuilt', CSoperating', CSsnftransported', ...
  Waste};
```

Appendix B: Comparison of FANTSY with the MIT Center for Advanced Nuclear Energy Systems Benchmark of CAFCA, VISION, DANESS, and COSI

The Flexible Advanced Nuclear Technology Simulation by Year (FANTSY) code is a dynamic 100-year simulation of the U.S. nuclear fuel cycle. It contains a very simple set of rules, which initialize the current reactor fleet and then build light water reactors (LWRs) and fast reactors (FRs) according to a nuclear electricity demand profile. The code does not simulate any facilities other than reactors, but does contain calculations to limit the amount of reprocessing capacity available to the system. One difference between the FANTSY code and the benchmark study conducted by MIT CANES is that FANTSY is run for 100 years, from 2010 to 2109, while the benchmark study required codes to run from 2005 to 2100.

B.1 Initial Comparison: Benchmark Scenario 3A

The CANES benchmark exercise includes several different scenarios. Scenario 3A calls for zero increase in electricity demand over the next 100 years (meaning that reactors are only built in order to replace decommissioned ones). Fast reactors with conversion ratio 1.0 are introduced in 2040. The reprocessing capacity profile is shown below in Figure B-1; this is the profile that was implemented by all codes for runs of scenario 3A in the benchmark study. The profile was generated by CAFCA.

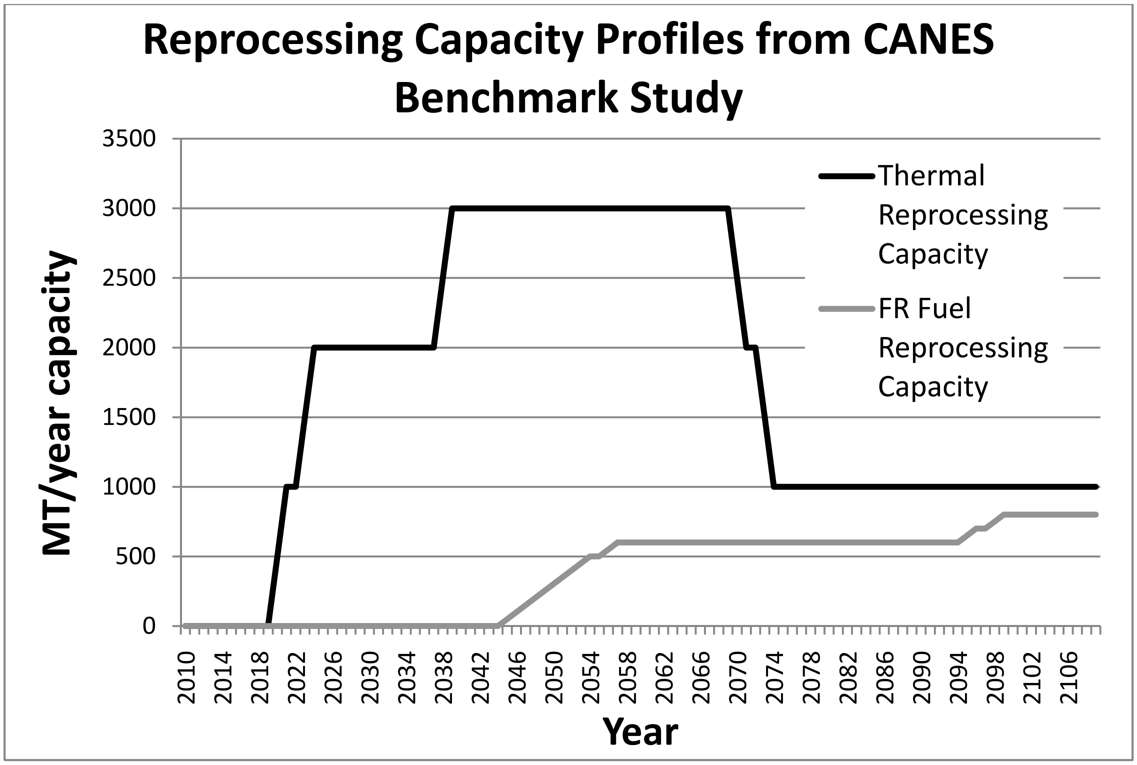

Figure B-1: Reprocessing capacity profile for scenario 3A, provided by CAFCA to the other codes for benchmarking purposes in the CANES benchmark study

FANTSY provides simulation results that track very well with the other codes in this benchmark exercise.

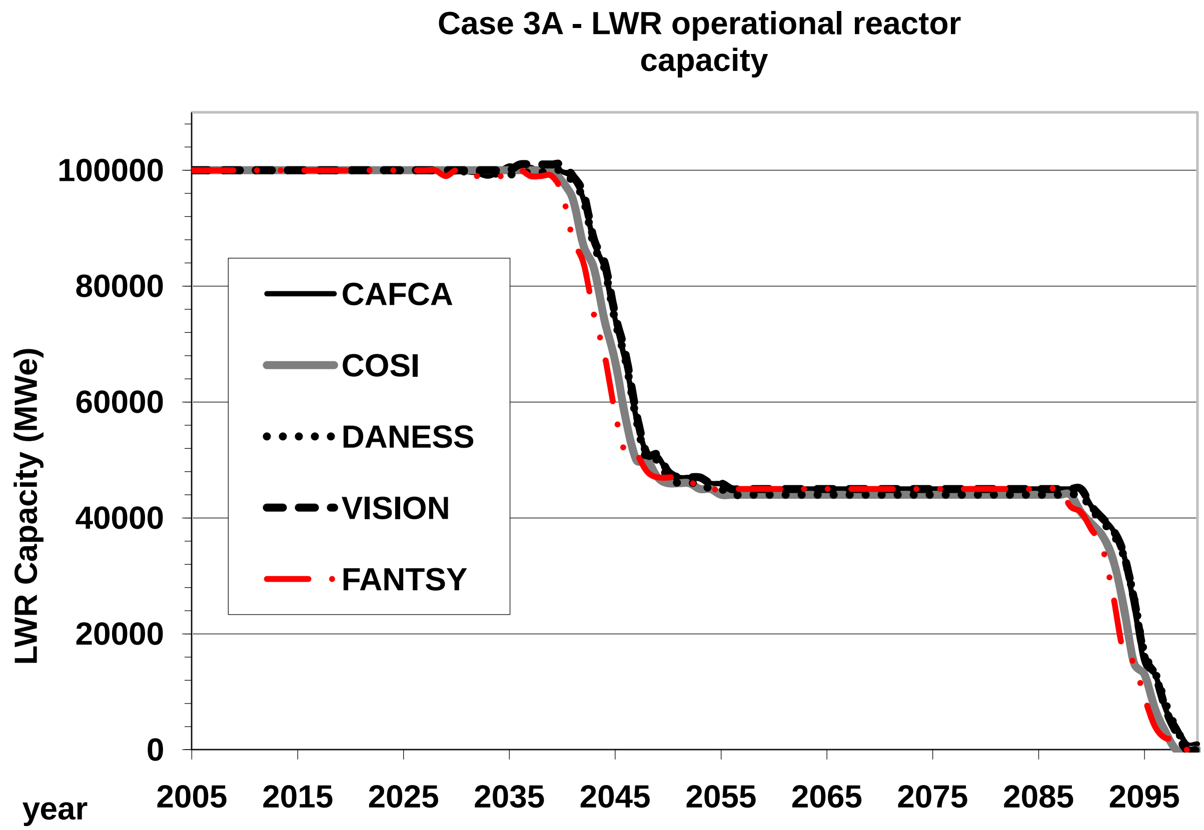

Figure B-2: Operating reactors by MWe capacity (Case 3A)

Figure B-2 shows the profile of operating light-water reactors over the course of the simulation, comparing FANTSY to the results from the four-code benchmark. Figure B-3 and Figure B-4 show the FR and total capacity profiles generated by all five codes. All figures demonstrate a leveling of LWR capacity at about 45 GWe, at the same time FRs levelize at about 60 GWe. The number of reactors increases slightly with the introduction of FRs because FRs have a lower capacity factor than LWRs (85% vs. 90%).

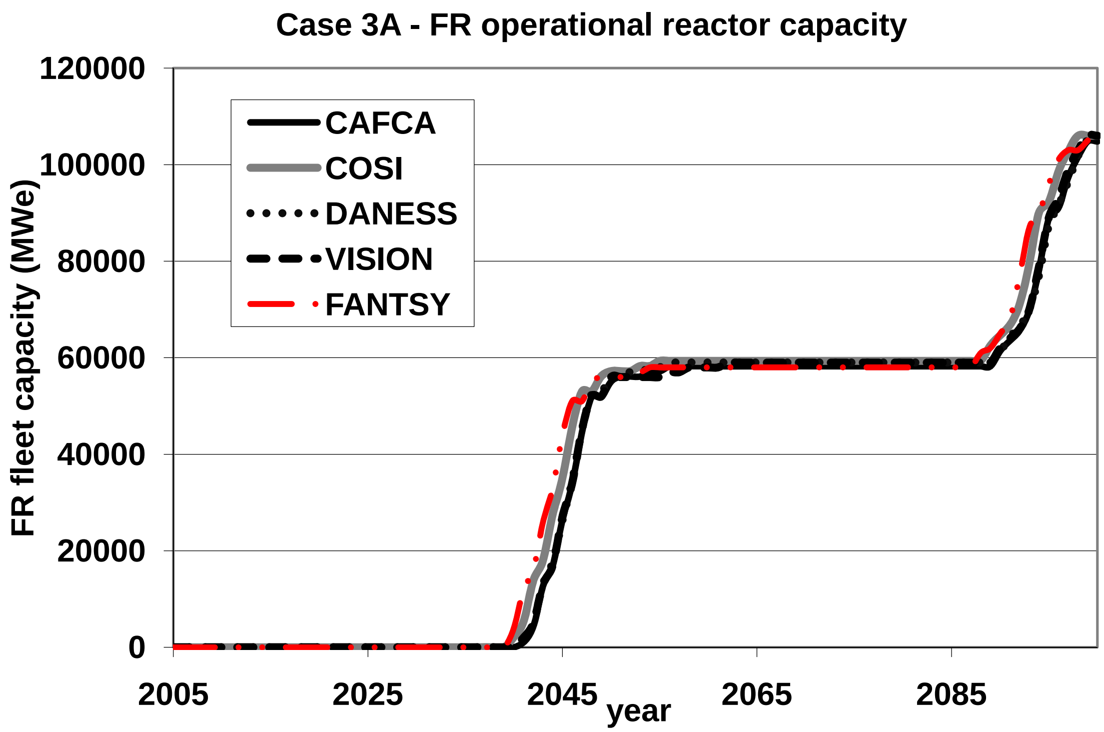

Figure B-3: FR operational reactor capacity (Case 3A)

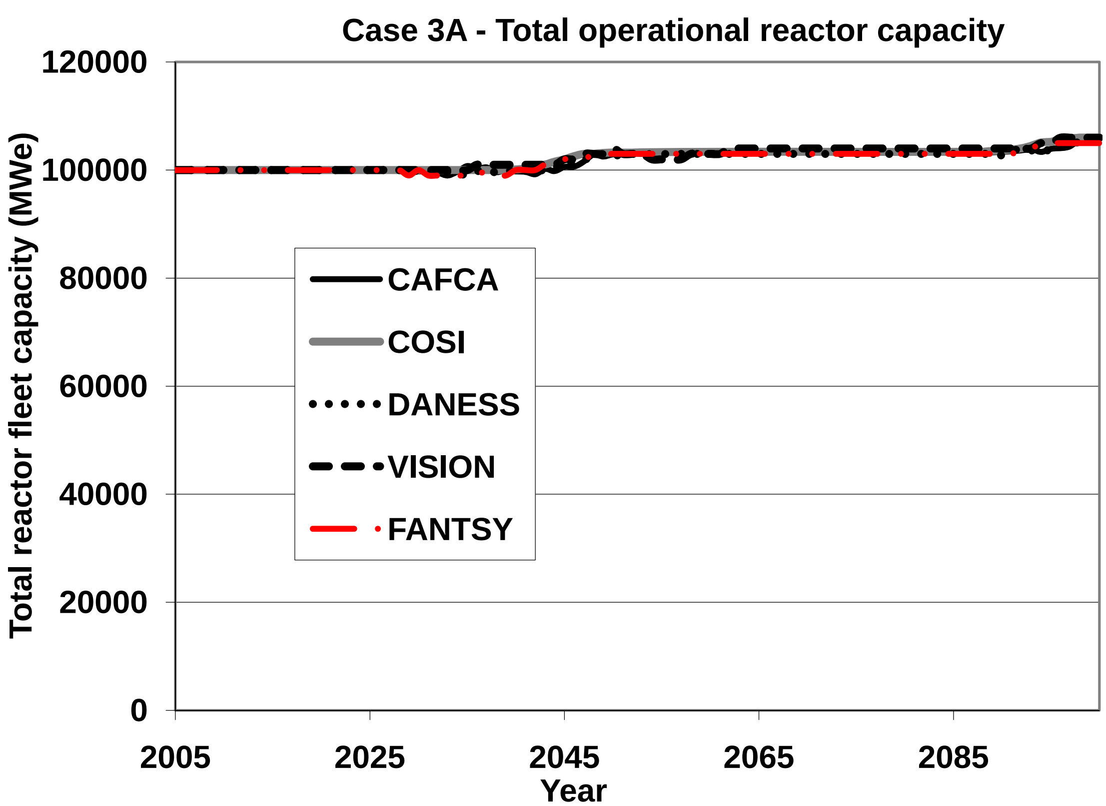

Figure B-4: Total operational reactor capacity (Case 3A) in benchmark exercise

FANTSY results are next compared to the benchmark runs for one front-end and one back-end variable. The yearly natural uranium consumption is calculated by FANTSY, and presented in Figure B-5 alongside the results from the other four codes. The profile closely matches those from the benchmark study. FANTSY tracks best with CAFCA and VISION, though the high peak of uranium consumption reached by FANTSY much more closely matches the CAFCA high value. This makes sense because the fast reactor ordering rule in FANTSY mirrors CAFCA's, and other modeling choices made in FANTSY explicitly reflect the CAFCA calculation.

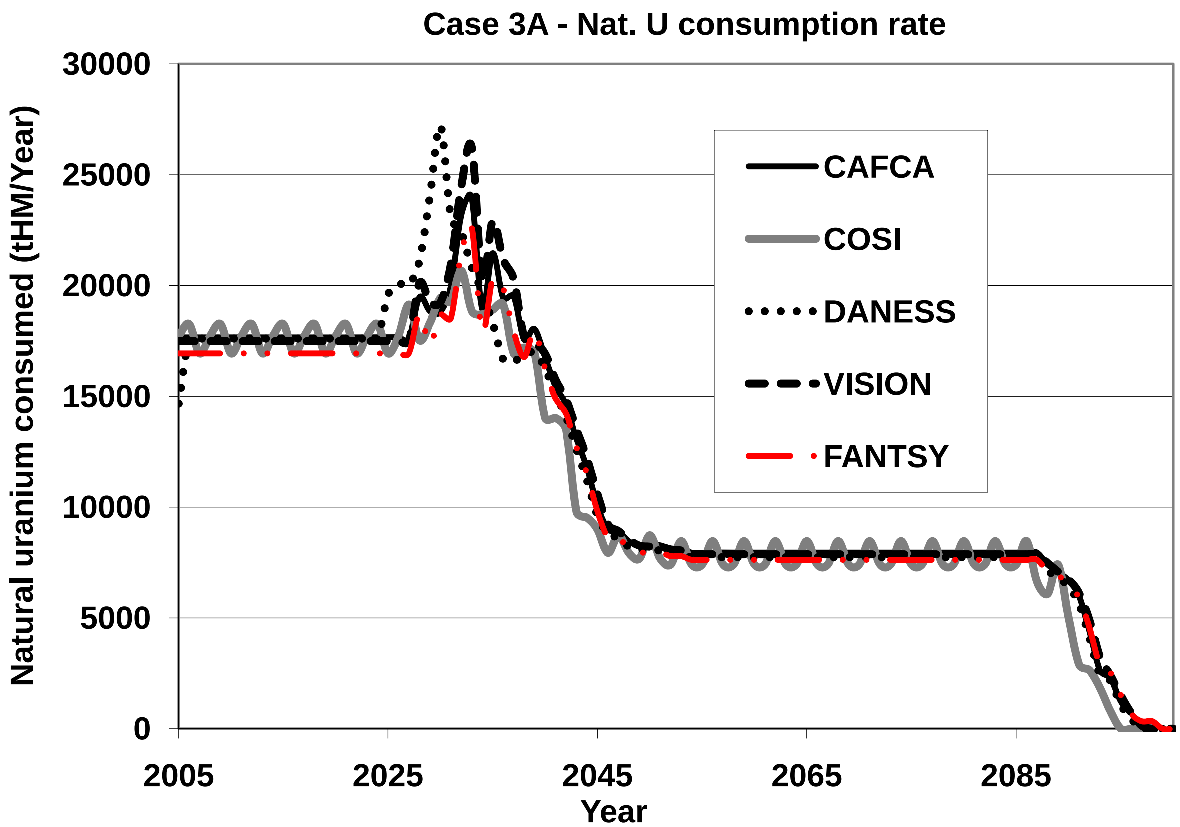

Figure B-5: Yearly natural uranium consumption rate (Case 3A)

The buildup of separated transuranic material (TRU) similarly shows good agreement between FANTSY and the benchmarking exercise. Figure B-6 shows the TRU buildup for FANTSY and the benchmarked codes. Note that for this variable, FANTSY tracks better with VISION, DANESS, and COSI than CAFCA. This is because CAFCA immediately fabricates any separated TRU as long as FRs are in operation. After FRs are introduced in 2040, CAFCA no longer maintains any separated TRU in storage. The design choice in FANTSY was to match the non-CAFCA fabrication rules, such that a buildup of TRU happens even after FRs are deployed.

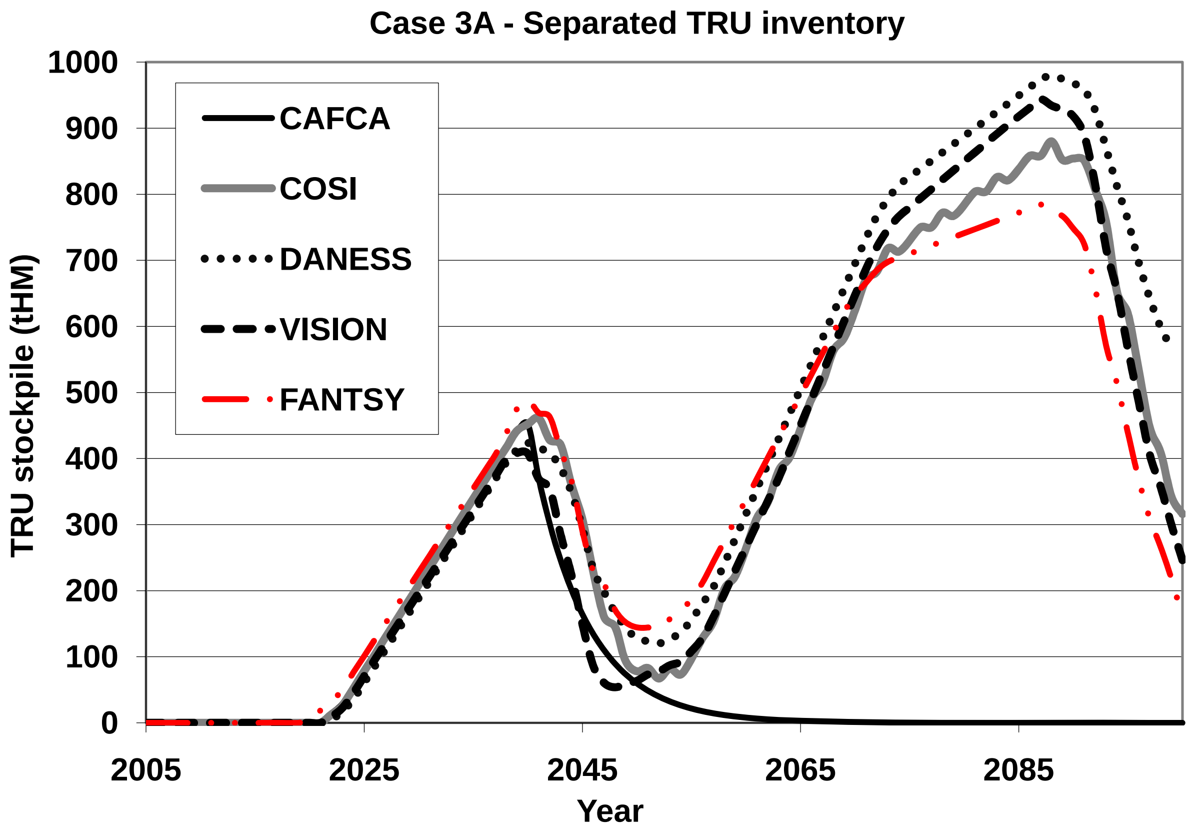

Figure B-6: Amount of separated TRU in storage (Case 3A)

FANTSY does generate more separated TRU during the earlier periods (about 2035-2065), and generates less TRU during the second major buildup. FANTSY nevertheless exhibits the same general behavior as the other codes, showing deltas roughly equivalent to those between the benchmarked scenarios.

B.2 Second Comparison: Case 3B

Case 3B of the benchmark report involves fast reactors with a conversion ratio of 1.0, and a growth rate of 1.5% per year in nuclear electricity demand. For case 3A above, the benchmark specifically stated that all four codes did several runs to carefully tune the results to one another in order to get very high agreement. For case 3B, the codes were not so carefully tuned and therefore display a range of results.

The FANTSY code is no different. Indeed, FANTSY was designed expressly to conform to results from the benchmark case 3A, and similarly shows behavior different from the other codes as soon as a growth rate is introduced. Some of the difference could be because a single timestep for FANTSY is one year, whereas for the other codes it may be smaller (for example, VISION and CAFCA perform calculations over month-long timesteps). In general, the results provided by FANTSY seem to be within the range of deltas between the four originally benchmarked codes.

For the benchmark, the reprocessing capacity profile was not directly specified as it was in case 3A (codes were allowed to optimize reprocessing capacity additions). The results for case 3B are thus shown for two profiles incorporated into FANTSY: one mirrors the profile CAFCA used in the benchmark study, and one is adjusted to show the effect of tuning the profile toward the eventual design decisions made for FANTSY's reprocessing rules. Figure B-7 shows the adjustment made for the thermal reprocessing capacity: the profile is shifted so that reprocessing begins two years later, in 2012. This is a very crude way of implementing a 2-year waiting period between initializing reprocessing and using reprocessed fuel (some of the benchmark codes assume 1 year for reprocessing and 1 year for fuel fabrication before use of the fuel, such that 2012 marks the first year reprocessed fuel is available).

Figure B-8 shows the fast reprocessing profile. In the benchmark study, fast reactor fuel reprocessing begins in 2045. A design decision was made for FANTSY to begin building a small amount of FR fuel reprocessing capacity in 2035, five years before the introduction of CR 1.0 reactors which will recycle their own fuel. This way, technology and design issues would presumably have a greater chance of being resolved by the time the first spent FR fuel is available in 2041. The orange line represents this adjusted reprocessing profile.

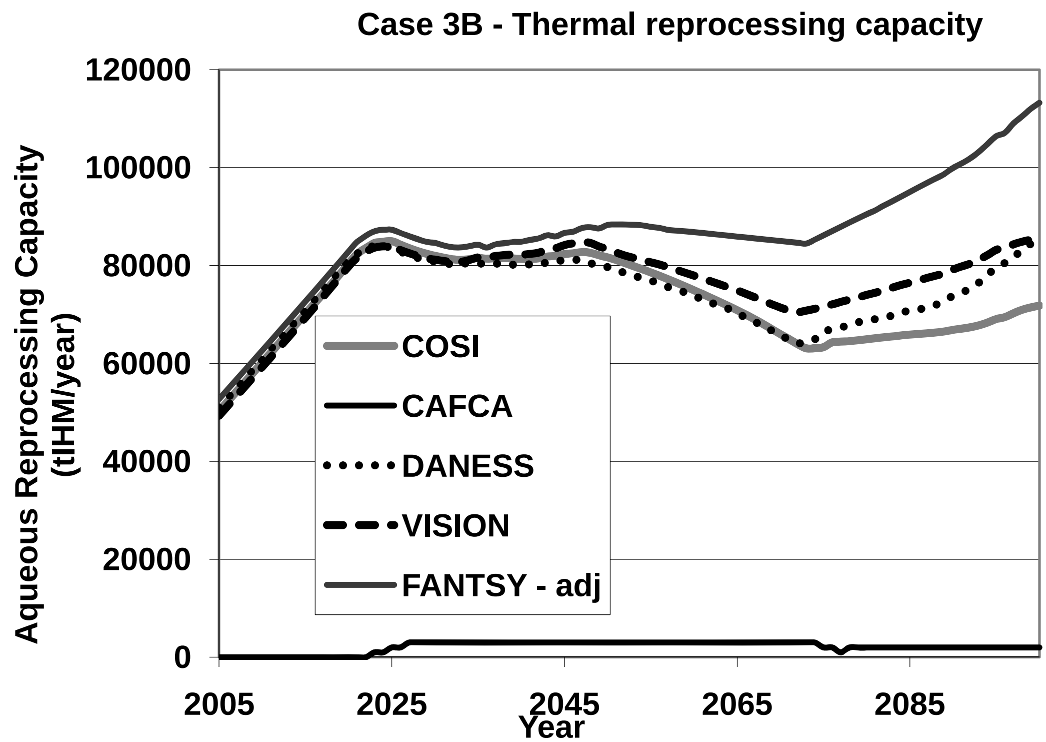

Figure B-7: Thermal reprocessing capacity: FANTSY adjustment to allow for two-year delay between start of reprocessing and use of reprocessed fuel

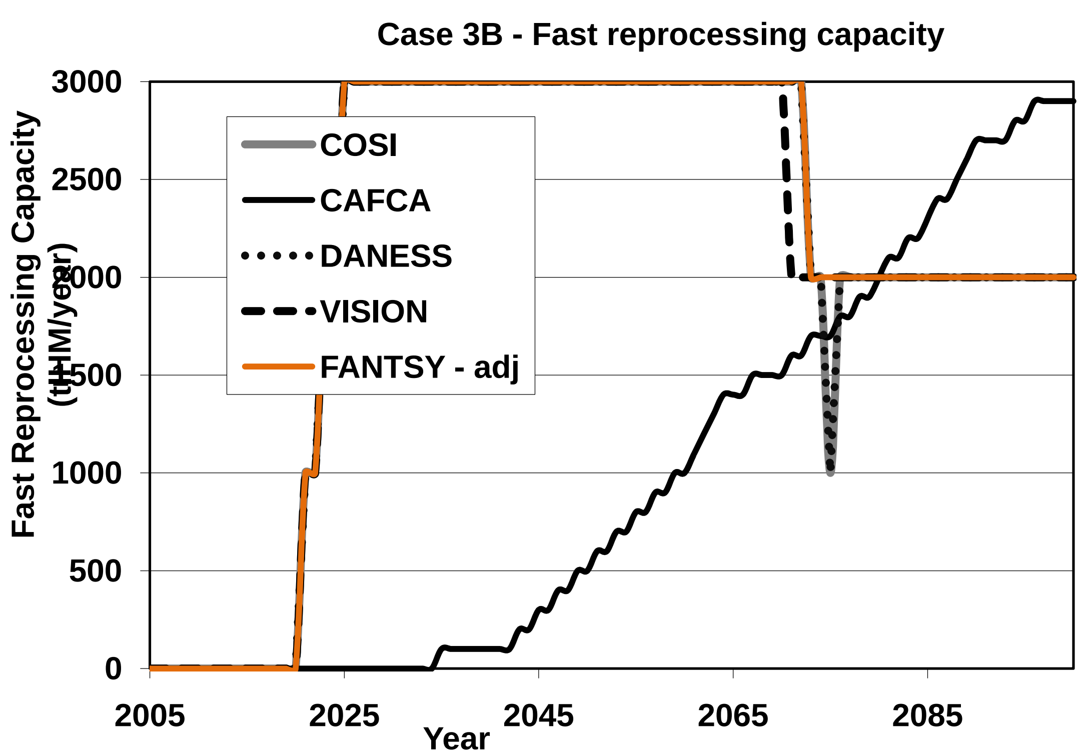

Figure B-8: Fast reprocessing capacity adjustment: FR reprocessing starts in 2041

Figure B-9 and Figure B-10 show profiles for the operating fleet of reactors. These are the most important parameters for the decision analysis study, because both the system cost and amount of waste discharged are depend very heavily on the number of reactors. There are some significant differences between the FANTSY results and the benchmarked code behaviors.

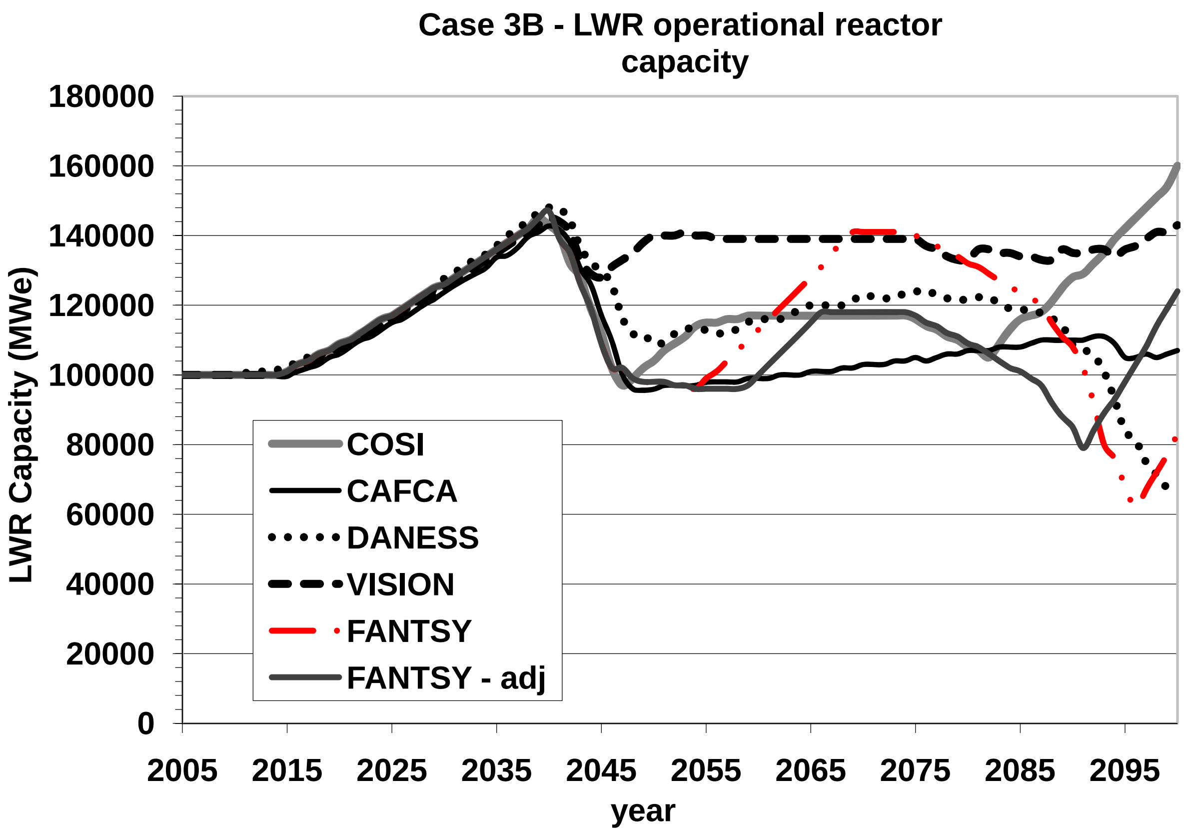

Figure B-9: Installed LWR Capacity by year (Case 3B)

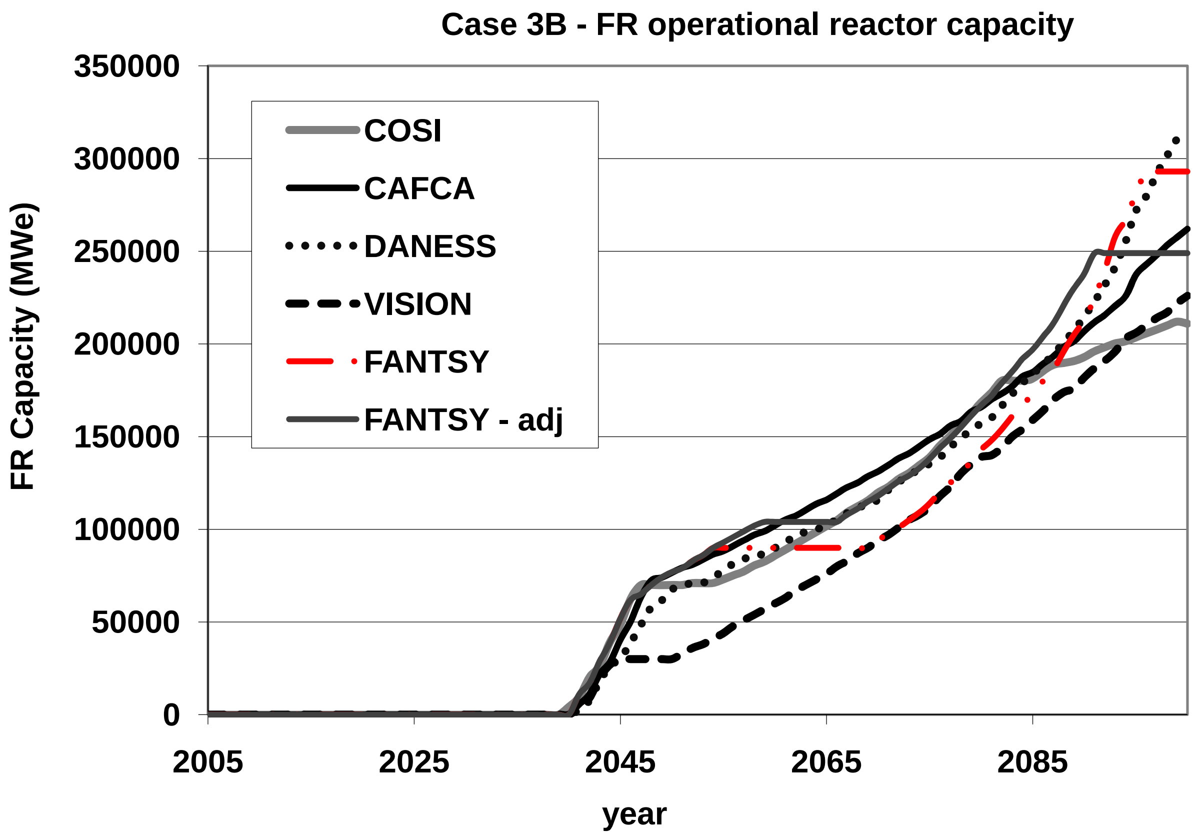

Figure B-10: Installed FR capacity by year (Case 3B)

The first difference is that the fast reactor buildup in the benchmarked cases is nearly always continually increasing (although COSI and VISION demonstrate short periods where FR capacity flattens). The number of reactors operating in any given year, however, is commensurate with the general pattern exhibited in the benchmark scenarios. Differences are most likely due to differences in the ordering algorithm for fast reactors. DANESS, for example, shoots for a certain percentage of fast reactors by the end of the century, and only requires a minimum of five years of TRU fuel be available. FANTSY, by contrast, will only build fast reactors if there is a lifetime supply of TRU available for a fast reactor AFTER all FRs in a given year have been fueled. This constraint is similar to the one used in CAFCA, and it leads to overbuilding when stocks of potential fuel are high and plateaus later.

The differences seen in the number of LWRs are a direct result of the different patterns of FR deployment. Because the slope of FR builds is so steep from ~2072-2090, the LWR capacity dips extremely low for the FANTSY case.

As expected, adjusting the profile for building reprocessing capacity has an impact on the pattern of FR vs. LWR builds. The adjustment for FANTSY (orange curves) involves building thermal reprocessing one year later and fast reprocessing earlier. This enables a greater buildup of feedstock for FRs, so that the first plateau is higher when thermal reprocessing comes later. Readily available FR spent fuel reprocessing then enables a steeper FR buildup later in the century.

Figure B-11 shows the inventory of separated TRU for FANTSY and for the benchmarked codes. One can see that the TRU stock for FANTSY more or less fits within the realm exhibited by the benchmarked codes, except for a dramatic rise in TRU inventory when CAFCA's reprocessing profile is used (red dotted curve). Several factors contribute to FANTSY's higher buildup of TRU later in the century. One is that FANTSY (like COSI, but not the other codes) includes separated TRU from FR SNF reprocessing. This is a relatively small but observable effect. The bigger reason is that the plateaus for FR builds are longer in FANTSY, so TRU stocks build up significantly after an overbuild and a lag (five years of construction time) once they recover from depletion. This effect is similar in CAFCA, but CAFCA continues to fabricate any separated TRU regardless of need, so the TRU stock never builds up after the middle of the century. Finally, FANTSY always separates as much TRU as possible (using 100% of the available reprocessing capacity as long as feedstock exists). This allows for conservatism in the code: given that separated TRU stocks are generally undesirable, FANTSY calculates the upper limit.

Most importantly, Figure B-11 shows that the FANTSY TRU stock profile is consistent with the FANTSY reactor builds, given a delay in construction once TRU stocks rise. It also shows differences between FANTSY and the other codes of a similar magnitude to the between-code benchmark differences. Note that by providing a less restrictive reprocessing profile, the FR plateaus would shorten and the second peak would decrease in height (see Figure B-12). Furthermore, though this parameter exhibits greater differences between FANTSY and the benchmark than other parameters, the amount of TRU in storage does not currently factor into the decision framework. Future analyses which take TRU stockpiles into account would require a more sophisticated version of the code.

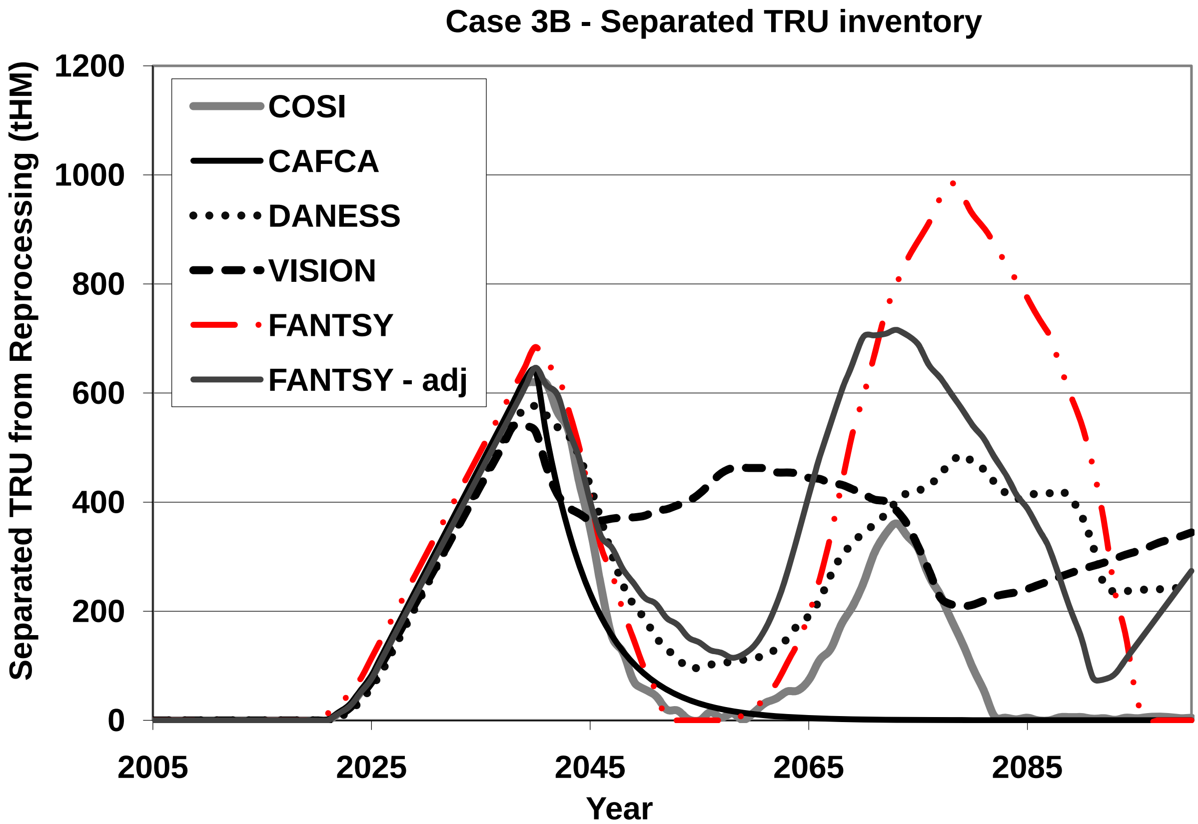

Figure B-11: TRU inventory by year (Case 3B)

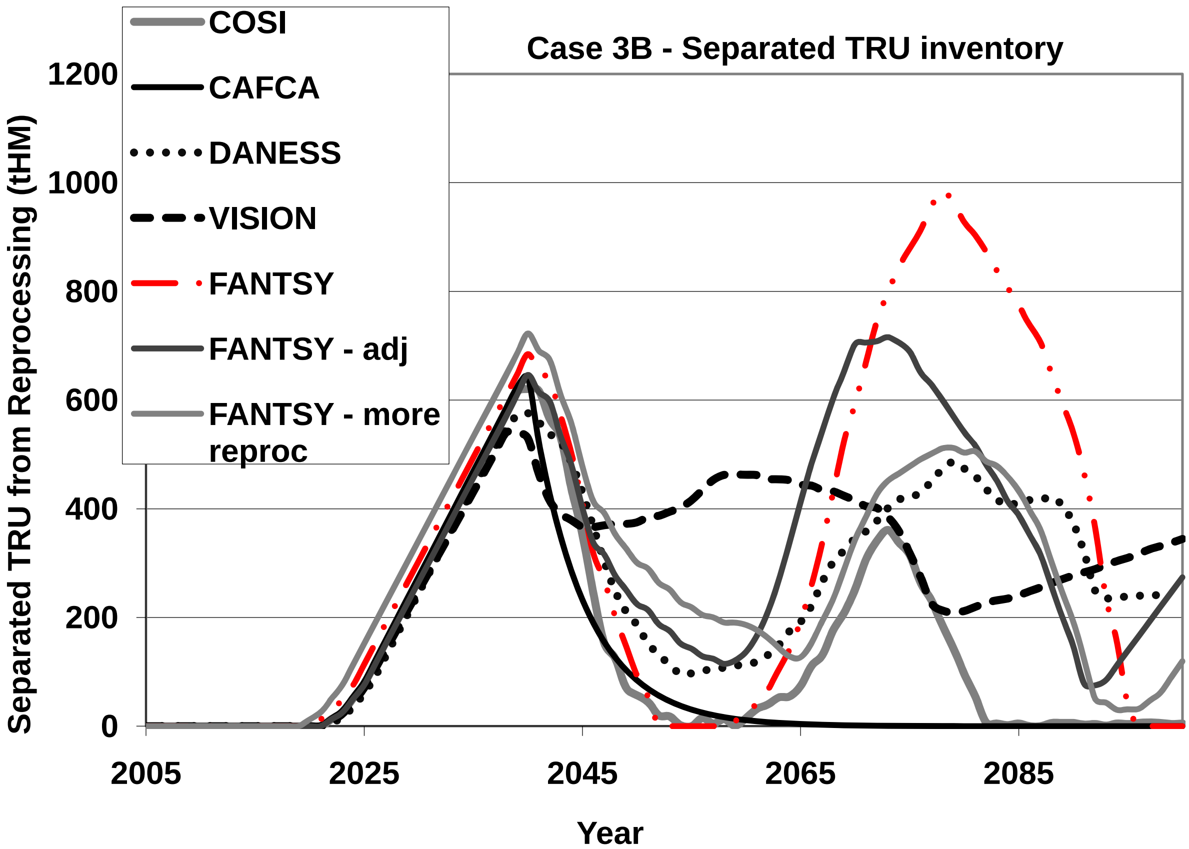

Figure B-12: TRU stockpile including an agressive reprocessing build profile (lessens second buildup)

The natural uranium consumption rate for FANTSY and the benchmarked codes shows good agreement, as demonstrated in Figure B-13. The SWU and UO2 fabrication rates closely mirror natural uranium consumption, so they are not shown here. Some of the dips shown for FANTSY are more dramatic than those for the benchmarked codes. This is due to the rapid buildup of FRs in the middle of the FANTSY simulation.

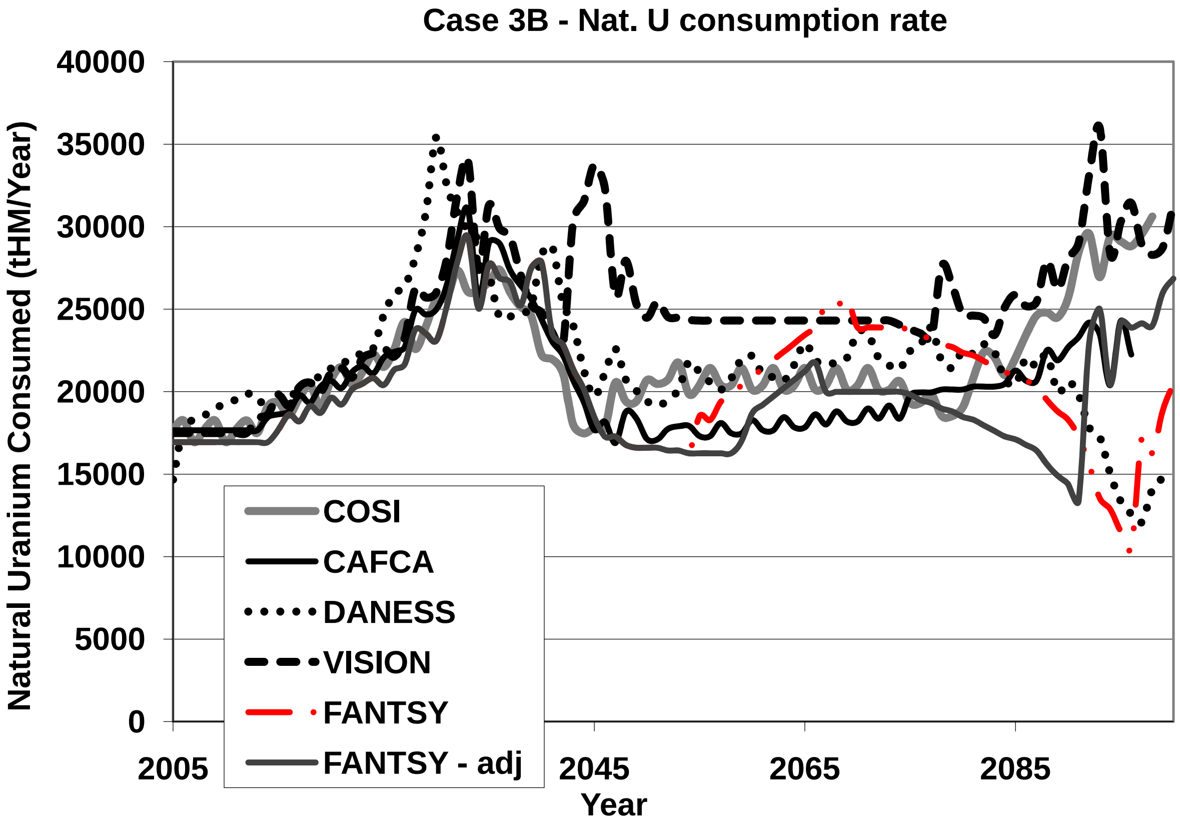

Figure B-13: Natural uranium consumed by LWRs (Case 3B)

Figure B-14 shows the FR fuel fabrication rates. Yet again, one can see plateaus in the rate exhibited by the FANTSY code, because fuel is fabricated as needed by the operating FRs (which plateau for reasons explained above). In general, the FANTSY results align well with the other codes.

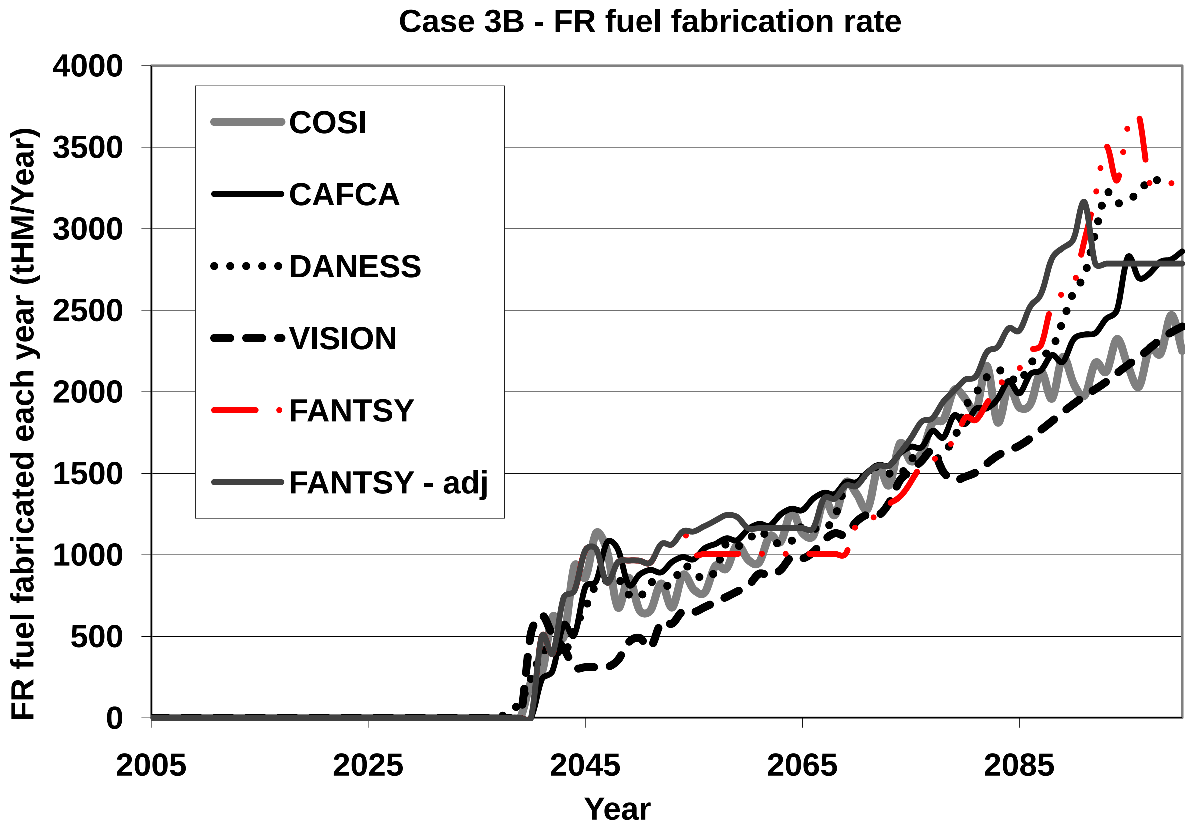

Figure B-14: FR fuel fabrication by year (Case 3B)

The final two figures show the cooled UO2 and FR spent fuel stocks over the course of the simulation. The UO2 stocks match fairly closely. FANTSY initializes in 2010, and assumes 63,000 MT spent UO2 in that year (close to the actual amount reported by industry: 62,500 in May 2010 by the Nuclear Energy Institute). (Nuclear Energy Institute, 2010) This is slightly higher than the amount originally extrapolated by the other codes in 2005. Note also that the adjusted reprocessing profile involves a slightly higher buildup of spent UO2, because reprocessing is delayed.

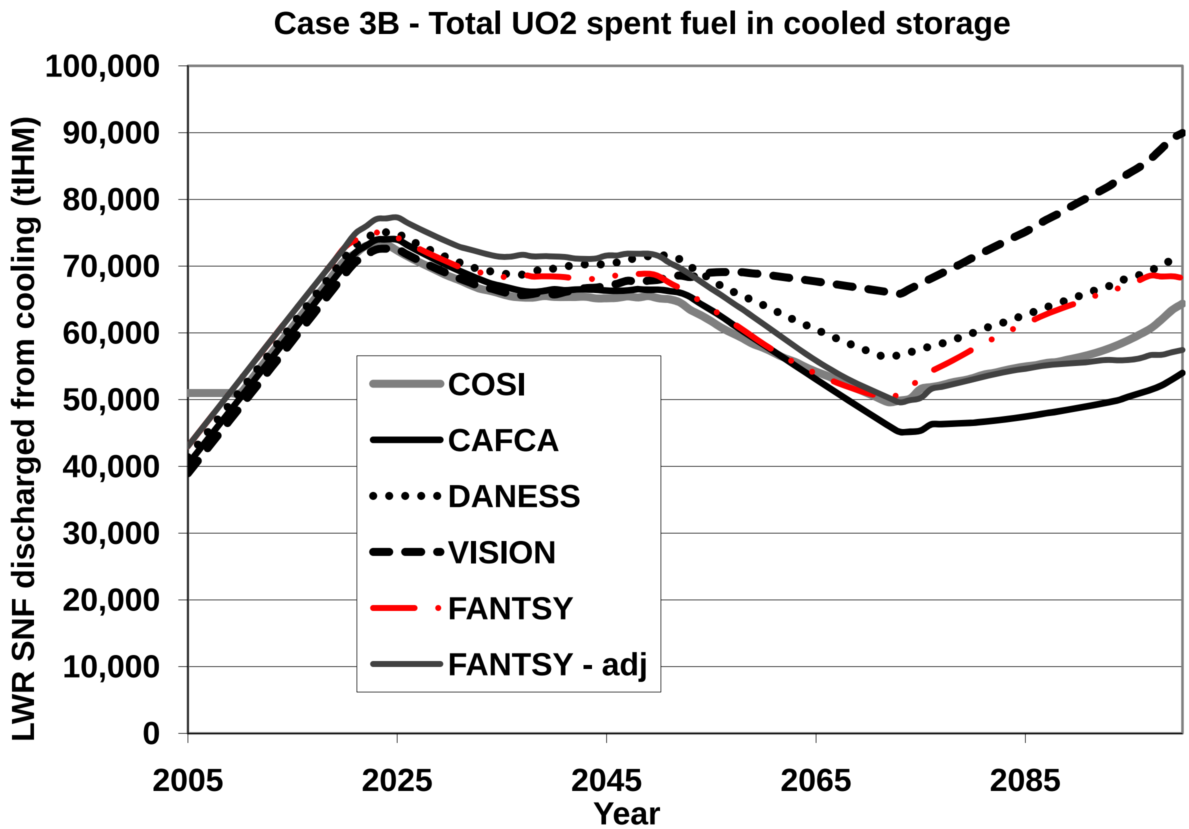

Figure B-15: Spent LWR fuel in cooled storage (Case 3B)

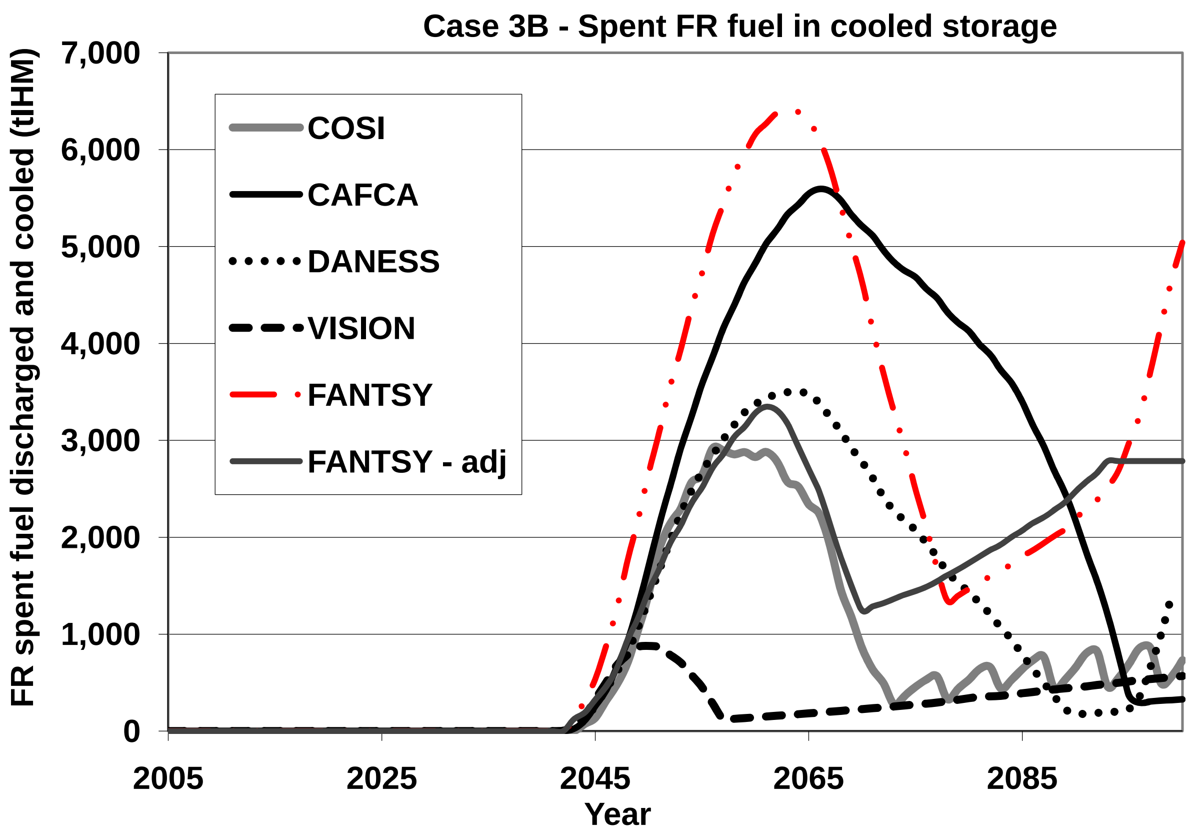

Figure B-16: Spent FR fuel in cooled storage (Case 3B)

The FR spent fuel stocks show dramatic variation between all codes. The adjusted FANTSY run shows very good agreement with COSI and DANESS, producing a smaller buildup of spent FR fuel because of a more aggressive FR reprocessing profile. The CAFCA-based FANTSY scenario (red) includes a bigger FR buildup because of the FR plateau combined with a more restrictive profile for spent FR fuel reprocessing.

B.3 Third Comparison: Scenario 1B

Scenario 1B consists of the same nuclear electricity demand increase (1.5% per year), but now with the introduction of conversion ratio 0.5 (burner) reactors rather than self-sustaining. As for case 3B, two reprocessing profiles are employed: the red lines result from CAFCA's reprocessing profile, and the adjusted profile (orange data) involves slightly later thermal reprocessing, and an earlier and more aggressive build pattern for FR spent fuel reprocessing capacity.

The results of the comparison are similar to those for scenario 3B. This time, however, the immediate buildup of FRs is even more rapid for the FANTSY code than for the benchmarked codes. This leads to an extensive time of zero FR builds, followed by a catching-up toward the end of the simulation. About the same number of FRs is eventually built for the standard CAFCA reprocessing profile, but the adjusted profile allows for building many more FRs due to the greater amount of FR reprocessing capacity. In both cases, the FANTSY build pattern is significantly different (see Figure B-17 and Figure B-18).

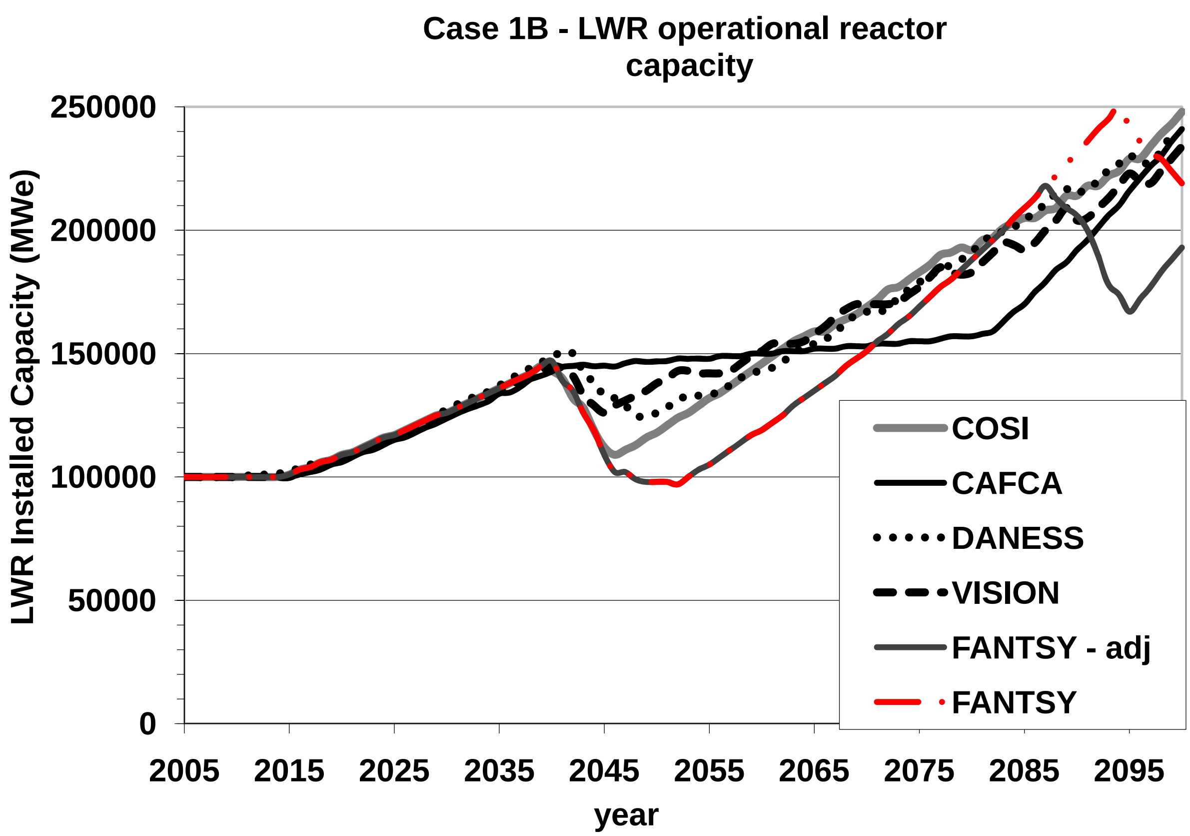

Figure B-17: Operating LWRs (Case 1B)

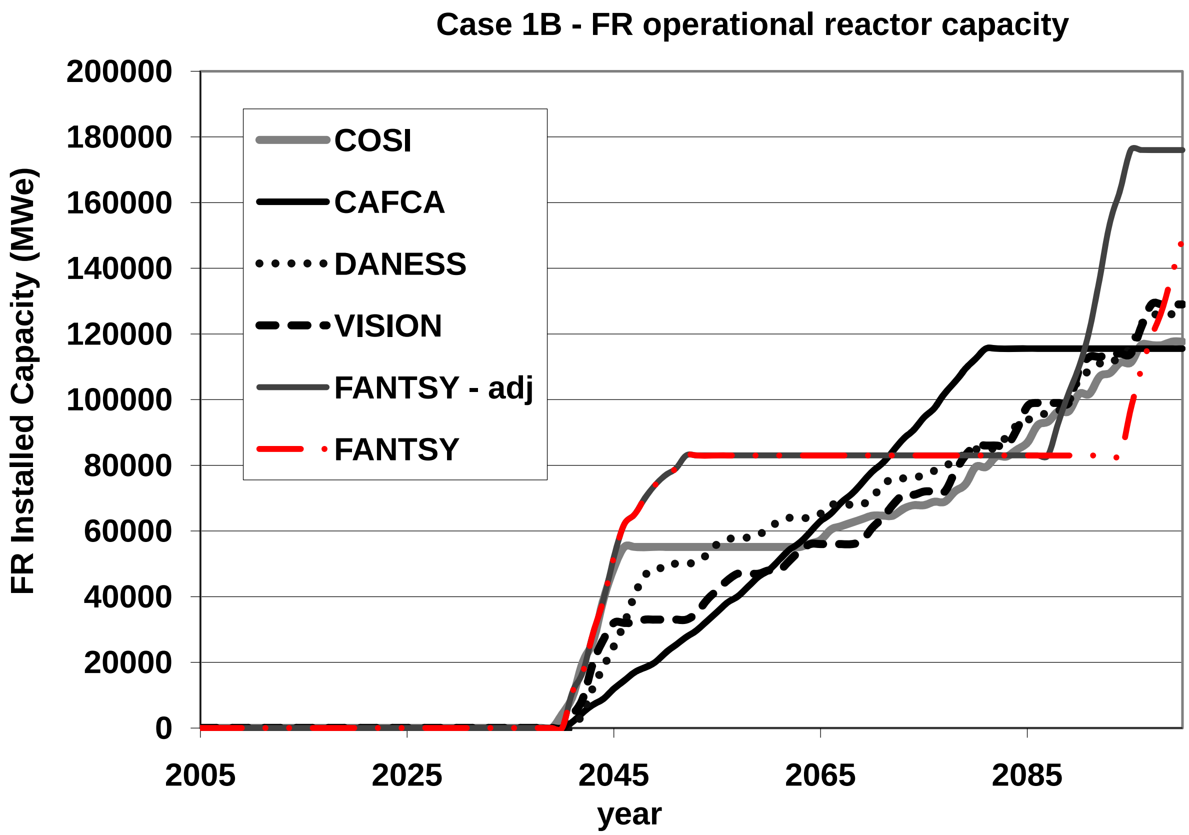

Figure B-18: Operating FRs (Case 1B)

The two reprocessing profiles produce similar results for FANTSY until the end. This is because the greatest difference between them is a much more aggressive FR spent fuel reprocessing capacity build (more aggressive than for all other codes), indicating that appropriately restricting the FR capacity is important. As for case 3B, the rises and plateaus happen because of overbuilding and time lags in FANTSY. The TRU stockpile follows expectation for FANTSY, given the build profile for fast reactors (Figure B-19).

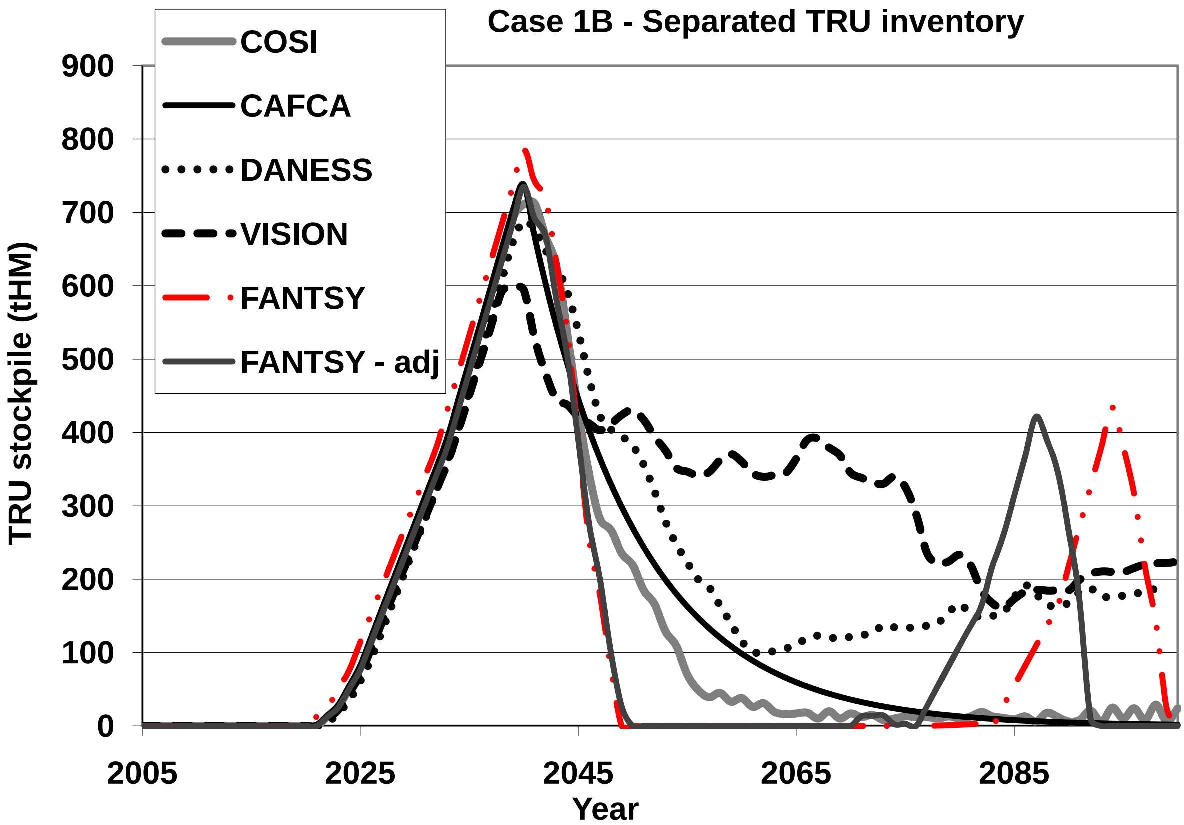

Figure B-19: Separated TRU (Case 1B)

The overbuilding and plateau effects in FANTSY occur because there are no restrictions on the number of fast reactors that can be built in any given year. The other codes also tend to implement restrictions about the size of reprocessing facilities, and require that they maintain a high capacity factor while they are operating. While these assumptions may be well-founded and representative of how industry will perform its capacity planning, it is also possible that delays and overbuilds may characterize fleet construction. FANTSY has the ability to simulate both possibilities.

In order to make the results from FANTSY smooth, like those in the benchmarking exercise, a limit was imposed on fast reactor builds of 2 reactors per year. When applied over the entire length of the simulation, the FANTSY code exhibited reactor build behavior similar to CAFCA but at a shallower slope and without a resulting plateau (see Figure B-20).

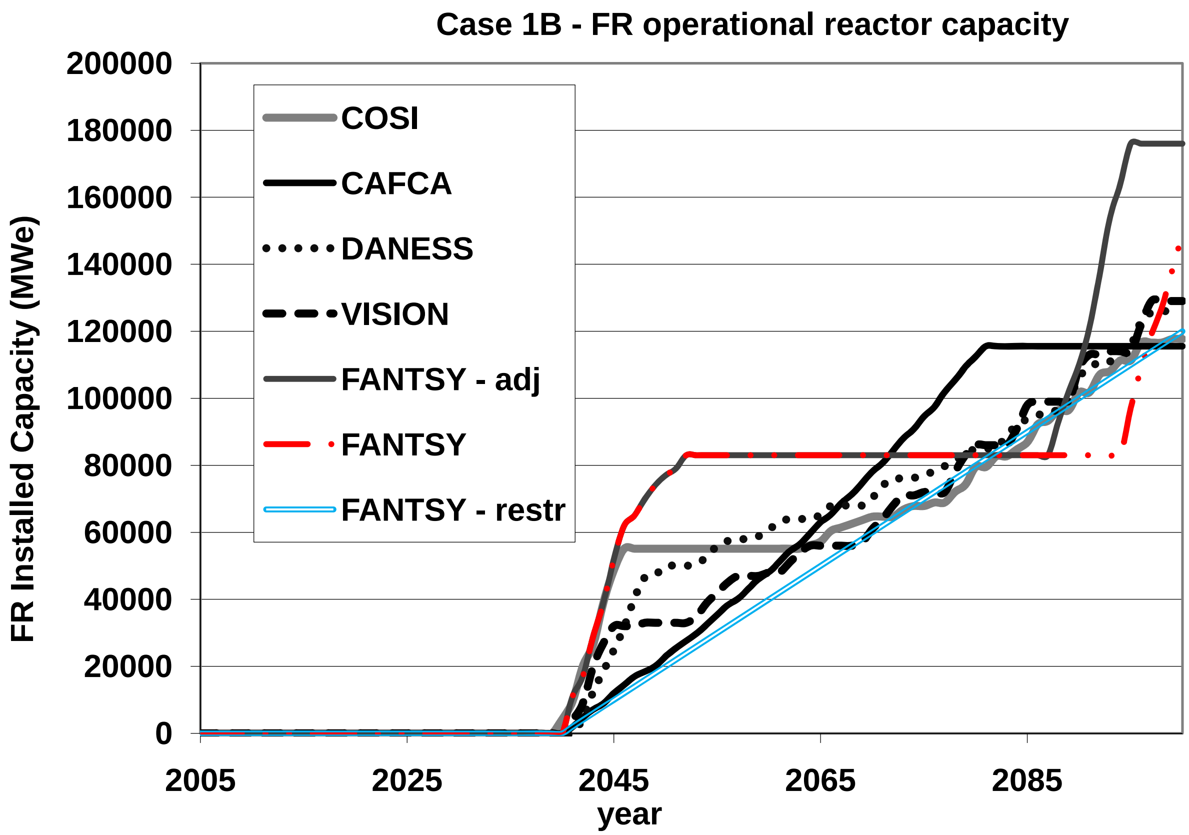

Figure B-20: FR capacity, including a restricted FANTSY build (maximum 2 FRs per year)

Interestingly, this artificial smoothing of the FR builds enables much more predictable management of the TRU stocks, which never reach zero. In turn, this means that FR operators would never need to turn to enriched uranium to fuel their reactors. This is not the case when there is no limit on the pace of FR builds. For the cases above, when the FR operating profile is flat, for many of those years there is not enough TRU from LWRs and FRs to completely fuel the whole FR fleet. This shows that the system properties are extremely sensitive to the build rate of FRs, and whether any industrial constraints exist on building them or reprocessing plants.

B.4 Master Code: Design Choices and Results

In the above examples, the reprocessing plant build profile was specified year by year (a profile was calculated by CAFCA at some point early in the benchmarking process, and then this profile was given as input to all four codes). Rather than use an artificial, static profile for reprocessing capacity expansion, the master code operates its own capacity calculation based on that employed by CAFCA.

The code assumes that until 2050, a maximum of 500MT/year thermal reprocessing capacity can be added each year. After 2050, this increases to 1000MT/year. Fast reprocessing capacity is restricted to 50MT/year built each year until 2065, after which 150MT/year capacity can be added yearly.

The needed reprocessing capacity is calculated based on the influx of spent nuclear fuel to cooling storage and the current inventory. Figure B-21 shows the existing reprocessing capacity each year. Figure B-22 shows the resulting profile of reactor builds (note that for the master code, the initial fleet is changed to reflect 104 reactors and the associated initial nuclear capacity).

A factor known in CAFCA as the "instantaneous depletion time" actually determines the level of conservatism in building reprocessing plants. A high instantaneous depletion time means that more TRU is required in the stocks and flows of spent fuel in order to allow the building of a given number of reprocessing plants. Changing this depletion time from 30/33 (thermal/fast) to 50 (the lifetime of the reprocessing plants) changed the number of thermal reprocessing plants built (see Figure B-23) but not the reactor build profile for this particular case. In the end, 50 was chosen in order to be conservative and slow the building of fast reactors to make FANTSY results closer to those of other codes.

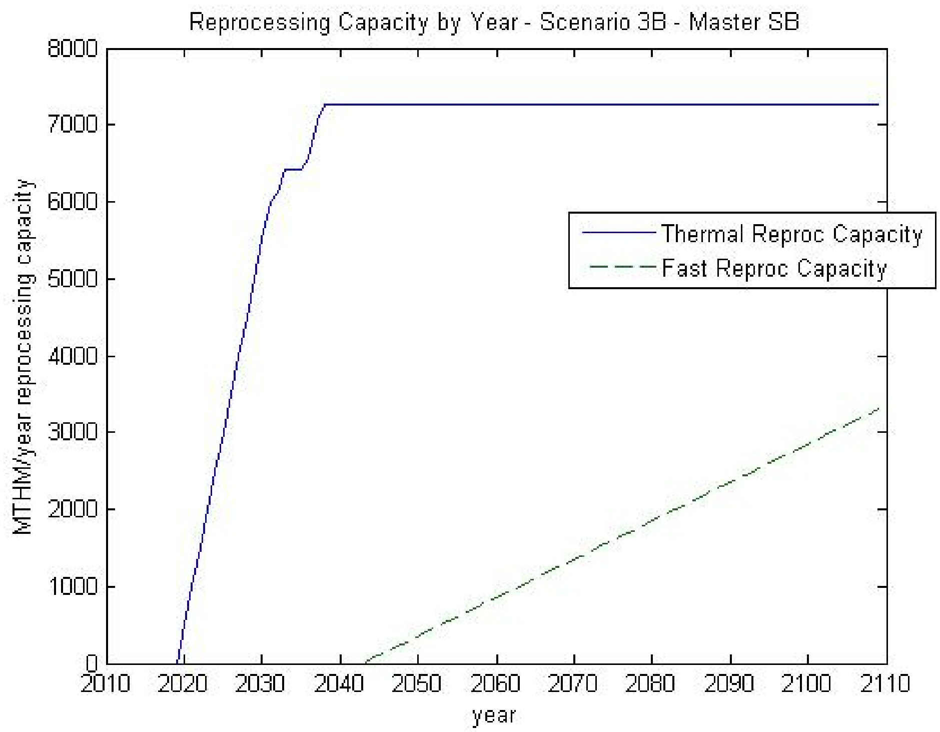

Figure B-21: Reprocessing Capacity as modeled by the FANTSY master code for scenario 3B

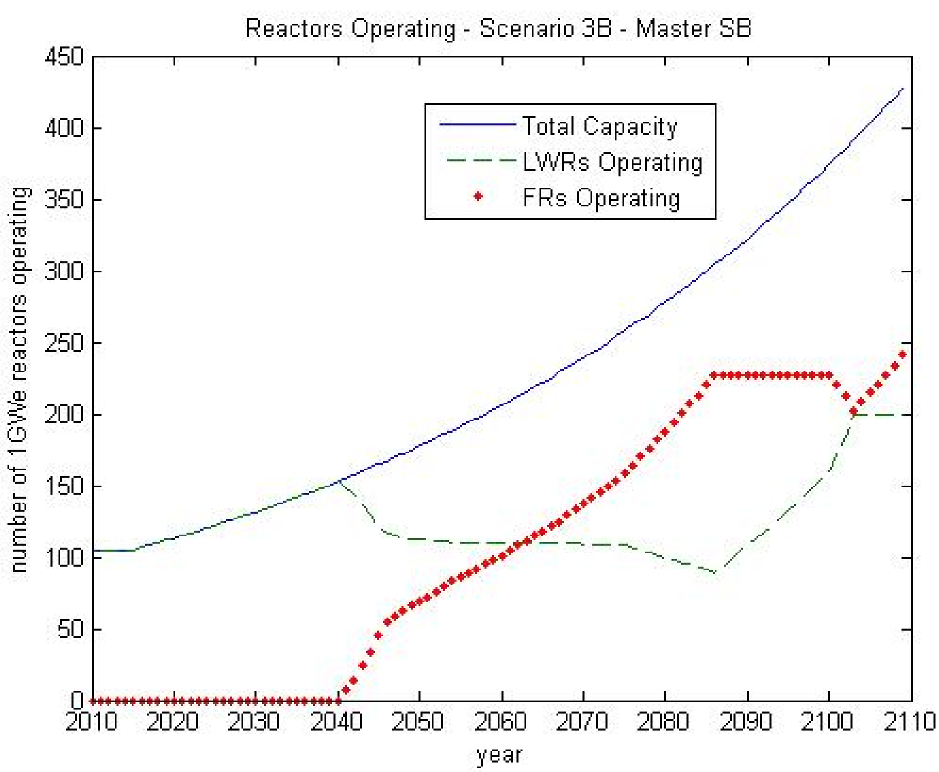

Figure B-22: Operating reactors generated by FANTSY master code for scenario conditions 3B

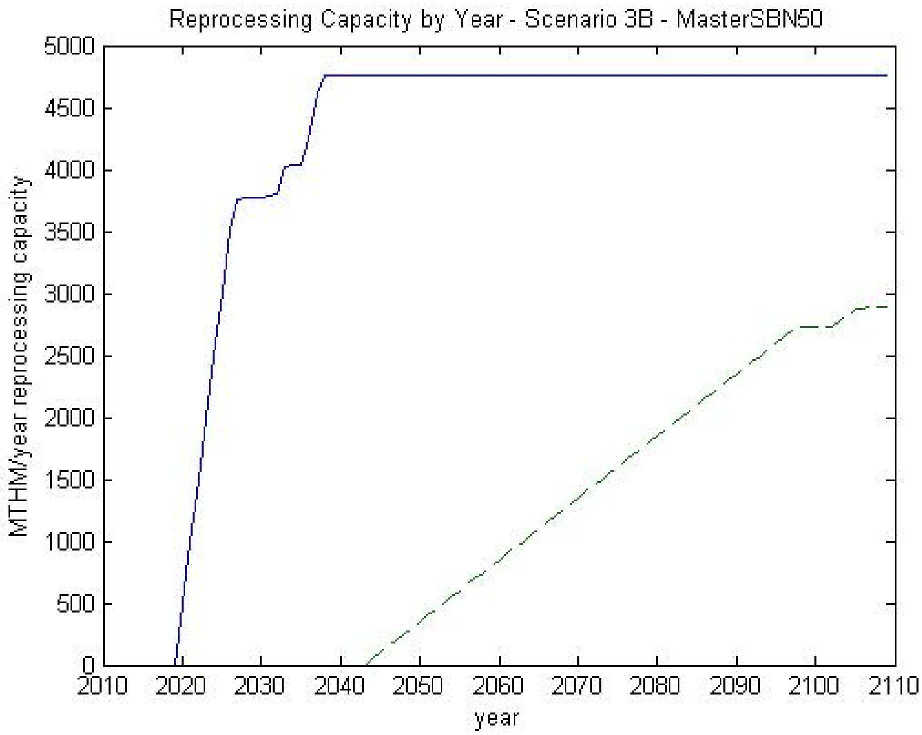

Figure B-23: Reprocessing Capacity as modeled by the FANTSYn50 master code for scenario 3B, with a 50-year "instantaneous depletion time"

B.5 Sensitivities of the Fuel Cycle Code

The FANTSY fuel cycle code is sensitive to several parameters. The above analyses already demonstrated some sensitivity of the results to the profile for building reprocessing plants. The results were well-aligned with the benchmarked codes for the CAFCA-esque reprocessing curve, so rules somewhat similar to CAFCA were employed in the final code. Given that fairly significant adjustments in the reprocessing profile still produce similar results (other than the pattern of FR builds), these rules are used consistently throughout the thesis and sensitivity analysis is performed on the FR build pattern.

The results depend much more strongly on the specific year in which FRs are introduced. Figure B-24 shows the operating reactor profiles when FRs are introduced in 2060 rather than 2040 (using the assumptions of case 3B). The later introduction date actually allows about 75 more reactors to be built by the end of the simulation, because the LWR fleet builds up to provide more fuel initially for a rapid FR takeover.

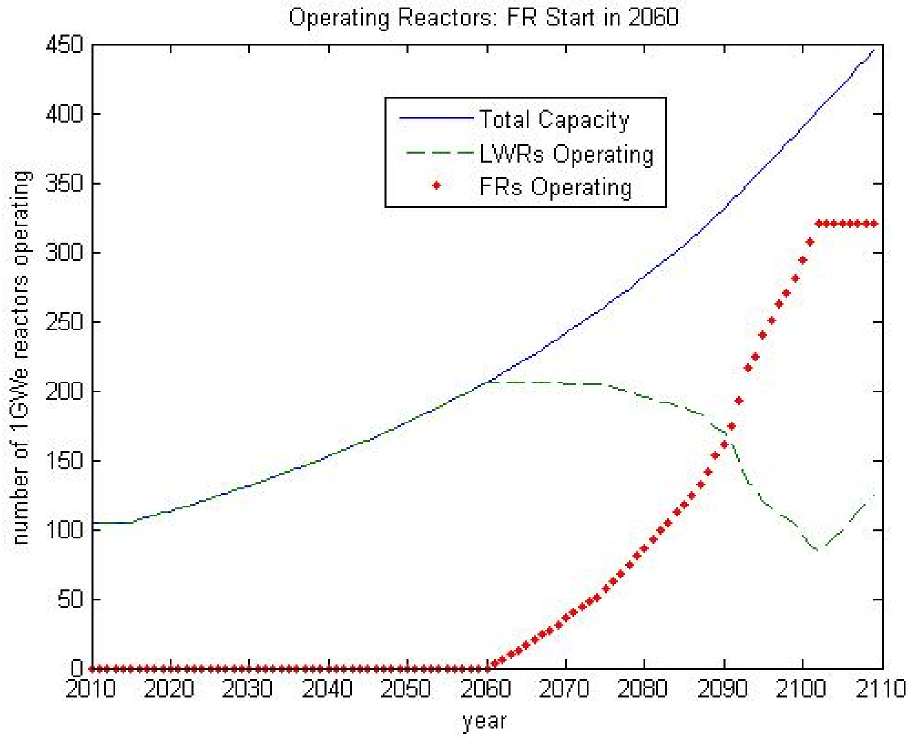

Figure B-24: Operating reactor profiles for LWRs and FRs given a 2060 start date for FRs rather than 2040

The profile of TRU in storage differs considerably between the two simulations with different FR introduction dates (see Figure B-25 and Figure B-26). Interestingly, the amplitude of the TRU stock is roughly the same for both cases; in neither scenario does the TRU stock go above 1500 MTHM. There is, however, a stark difference in the timing of maximum TRU in storage.

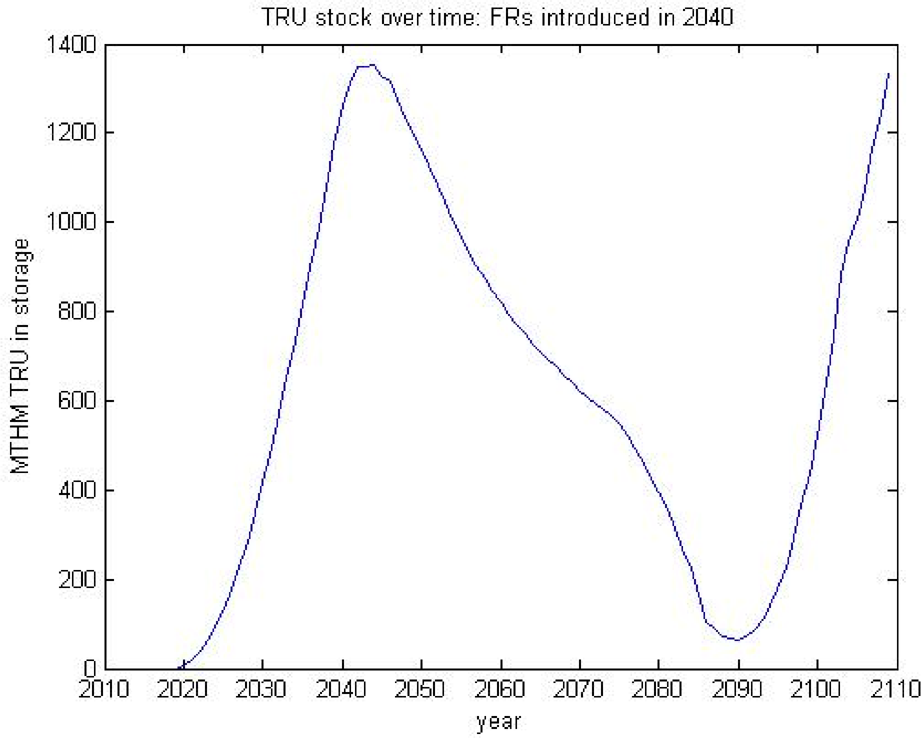

Figure B-25: TRU stocks over time when FRs are introduced in 2040

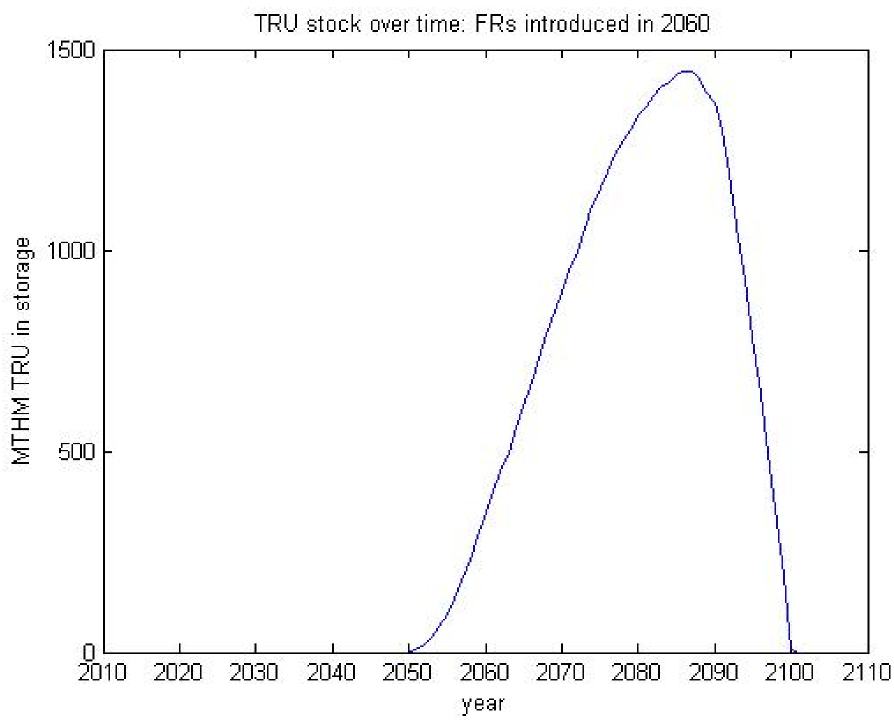

Figure B-26: TRU stocks over time when FRs are introduced in 2060

The results are also extremely sensitive to any restrictions on industrial capacity for building reactors. If one limits the number of reactors that can be built per year, and if this number is lower than the total number of reactors demanded, the results will show dramatic differences in the buildup of TRU and the resulting reactor mix. The difference is not so stark for a limit of 10 LWRs built per year (see Figure B-27, compare to Figure B-22).

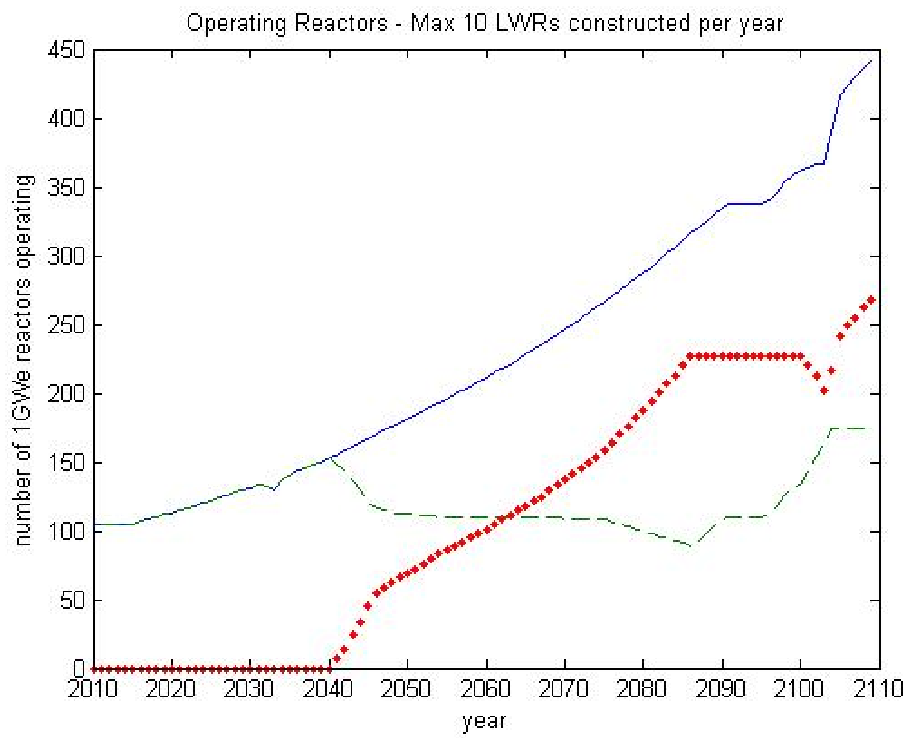

Figure B-27: Operating reactors over time with a 10-LWR per year construction limit

The difference is quite large, however, between the base master case and a case which employs a more complex, gradually expanding industrial capacity. The statement added to the code controls the maximum amount of each type of reactor that can be built per year:

```
if year < 2040
  LWRmax = 5;
  FRmax = 2;
elseif year < 2045
  LWRmax = 8;
  FRmax = 5;
elseif year < 2060
  LWRmax = 10;
  FRmax = 8;
```

```
else
  LWRmax = 13;
  FRmax = 10;
end
```

Results for this case are shown in Figure B-28. Many fewer fast reactors are built with this constraint, indicating that the operating profile is especially sensitive to FR build restrictions.

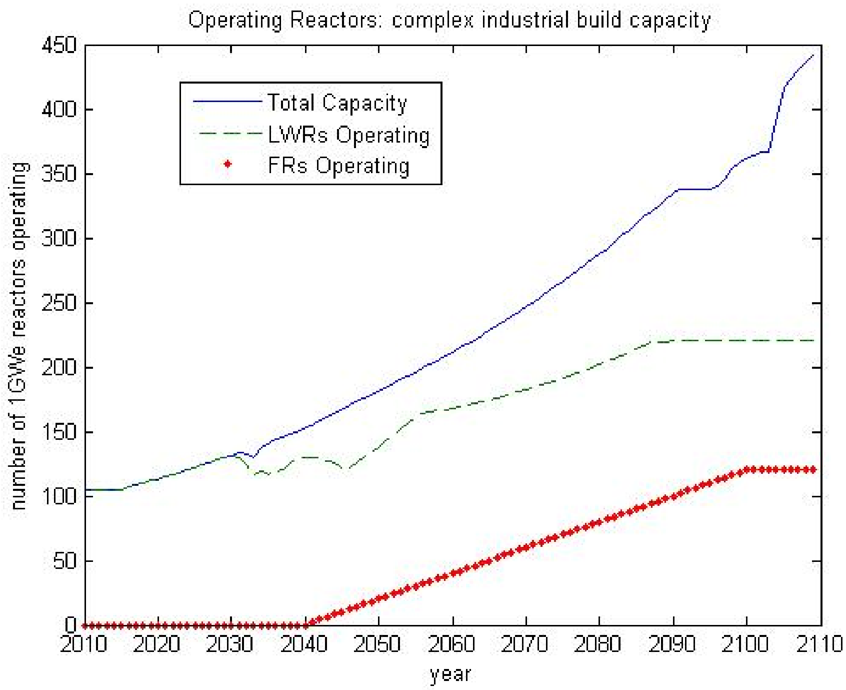

Figure B-28: Operating reactors over time with complex yearly expanding construction limits on LWRs and FRs

The number of reactors built is NOT very sensitive to the amount of legacy fuel in the system.

B.6 Conclusions

Overall, the most important conclusion from this benchmarking exercise is that FANTSY produces results within similar ranges to those of four more sophisticated fuel cycle simulations.

Each of the five codes described here has advantages and disadvantages; VISION and COSI, for example, are very detailed and offer an enormous amount of flexibility in simulating different scenarios (e.g. with up to 10 different reactor types built simultaneously into the fleet for VISION). FANTSY's primary advantage is its simplicity, allowing for rapid adjustments to the code structure and fast runtime (microseconds versus 5 minutes in VISION for a single scenario). Rapid runs are essential for the high numbers of scenario iterations demanded by decision trees in this thesis.

There are some differences between all codes in the amounts and timing of reactor builds and material stockpiles. In general, the differences between FANTSY and the other codes occur because of the *pattern* of FR builds, but the ultimate vales attained (of FRs built and related parameters) are generally equivalent. The magnitude of difference between FANTSY and other codes tends to be the same as the differences among the other four codes, further increasing confidence that FANTSY provides valid results.

Section B6 discussed the primary sensitivities of FANTSY, exploring the reasons for the difference in FR build patterns: restricting the industrial building capacity for fast reactors smoothes the pace of FR builds, as does increasing the date of FR introduction. These parameters are studied explicitly in the body of the thesis to see the effect of the FR build pattern on fuel cycle decision making. In this way, the thesis explores a system with delays (the natural "start-stop" pattern of building exhibited by FANTSY) and a more smoothed introduction of fast reactors (generally typified by the other codes). Both patterns of building FRs into the reactor system are plausible.

Choosing rules for FR ordering and material stockpiling is partly an art. It is not clear, for example, how industry will assess the market for TRU to start up and fuel a fast reactor when making construction decisions. It is also unclear whether market restrictions may limit the pace of building fast reactors. Each of the five codes handles these issues somewhat differently. Ultimately, however, the codes can be tuned to produce nearly identical results, and otherwise give qualitatively similar solutions. FANTSY can thus be trusted to provide a reasonably consistent and valid picture of the nuclear fuel cycle.

B.7 The Initial FANTSY Code

Note that several versions of the MATLAB code were created during the benchmarking process. The code below represents the most "finely tuned" to the CANES benchmarking specifications. Like the CANES benchmark, this version of the simulation has a static, built-in reprocessing capacity profile. This code mimics the benchmarking scenario 3A.

```
% single-run core file
% FANTSY7 (base = FANTSY6): restricts the building of reprocessing plants, to
try
% to get TRU profile in line with benchmark
% this file has a "hard constraint" that includes exactly the reproc build
% profile specified in the benchmark for case 3A

% NOTE: this also fixes the fact that the decision to build FRs is based on
% TRU stocks *AND* SNF and FR fuel stocks that have not undergone reproc

% OT or FR flag
OTflag = 0; % if zero, this is an FR run; if 1, this is an OT run

% when to start and stop FR builds
FRstartyear = 2040; % type in the year to start building FRs; 0 is never
start
FRstopyear = 0; % type in the year to stop FRs; zero is never stop

% constants: these do not change for any run
legacySF = 63000; % MT
% reactors = [0 0  0  0  0  0  0  0  0  0  0  0  0  0  0 ...
 % 0 0  3  3  2  6  11 13 2  7  3  3  0  3  4  5  4 ...
  % 8 6  9  5  1  2  1  0  0  1  0  1  0  0  0  0  0 ...
  %0 0  0  0  0  0  0  0  0  0  1]; % with legacies
reactors = [0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 3 2 2 4 8 10 2 6 4 ...
  3 0 2 3 8 4 9 7 10 7 0 3 1 0 0 1 0 1 0 0 0 0 0 0 0 ...
  0 0 0 0 0 0 0 0 0]; % benchmark version
LWRcapfac = 0.9;
FRcapfac = 0.85;
capacity2010 = 100*LWRcapfac; % GWe - 100 for benchmark, 104 otherwise

TRUinventory = 0; % assume we will NOT recycle legacy SF; otherwise this >0
CR = 1.0;

if CR == 1.0
  lifetimeTRUneed = 705.71*0.139; % MTHM, given
  coreFR = 45.5;
  yearlyFRneed = 11.19;
  truFractionFR = 0.141; % fraction of FR SNF that is TRU - Hoffman
  fuelperFR = 45.5; % MT per year - from CAFCA 'FR' Fuel Core Mass
  fuelperyearFR = 11.19;
  truFracFresh = 0.139;
elseif CR == 0.5
  lifetimeTRUneed = 390.87*0.333; % MTHM, given
  coreFR = 25.66;
  yearlyFRneed = 6.19;
  truFractionFR = 0.270;
  fuelperFR = 25.66;
  fuelperyearFR = 6.19;
  truFracFresh = 0.333;
else
  display('Bad Conversion Ratio');
end
```

```
% parameters that change by run
year1 = 2015;
year2 = 2040;
year3 = 2065;
yrChange1 = year1 - 2010;
yrChange2 = year2 - 2010;
yrChange3 = year3 - 2010;
growth1 = 0.0; %growth rates in fractional form
growth2 = 0.0;
growth3 = 0.0;
growth4 = 0;
% FR start flags
switch FRstartyear
  case 0
    FRstart = 1000;
    reprocstart = 1000;
  otherwise
    FRstart = FRstartyear - 2010 - 5;
    reprocstart = 2020 - 2010;
end
switch FRstopyear
  case 0
    FRend = 1000;
    reprocend = 1000;
  otherwise
    FRend = FRstopyear;
    reprocend = FRstopyear;
end

% initializations of temporary and output vectors
demandzero = capacity2010;
demand = zeros(1,100);
underconstruction = zeros(1,5);
ReactorsConstructed = zeros(1,100);
ReactorsOperating = zeros(1,100);
ReactorsDecommissioned = zeros(1,100);

FRsConstructed = zeros(1,100);
FRsOperating = zeros(1,100);
FRsDecommissioned = zeros(1,100);
underconstructionFR = zeros(1,5);
makeupfuel = zeros(1,100);

fastreactors = zeros(1,60);

% demand vectors
Forecast = zeros(1,100);
LWRdemand = zeros(1,100);
LWRdemandfilled = zeros(1,100);
FRdemandfilled = zeros(1,100);

%%%% REPROCESSING CAPACITY %%%%
maxSNFyearlyadd = 500; % amount of MT reproc capacity that can be built
yearly
maxFRyearlyadd = 50;
```

```
thermalReprocCapacity = [0 0 0 0 0 0 0 0 0 0 500 1000 1000 1500 2000 2000 ...
2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2500 ...
3000 3000 3000 3000 3000 3000 3000 3000 3000 3000 3000 ...
3000 3000 3000 3000 3000 3000 3000 3000 3000 3000 3000 ...
3000 3000 3000 3000 3000 3000 3000 3000 3000 2500 2000 ...
2000 1500 1000 1000 1000 1000 1000 1000 1000 1000 1000 ...
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000 1000 ...
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000 1000 ...
1000 1000 1000 1000 1000];
frReprocCapacity = [0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 ...
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 50 100 150 200 250 ...
300 350 400 450 500 500 550 600 600 600 600 600 600 ...
600 600 600 600 600 600 600 600 600 600 600 600 600 ...
600 600 600 600 600 600 600 600 600 600 600 600 600 ...
600 600 600 600 600 600 650 700 700 750 800 800 800 ...
800 800 800 800 800 800 800 800];


%%%% FUEL MODULE %%%%

lwrenrichment = 0.042;
xf = 0.0071; % feed enrichment (natural uranium at 0.71%)
xt = 0.0025; % tail assay (0.25%)
fuelperLWR = 87.76; % MT/reactor
fuelperyearLWR = 19.73; % MT/reactor/year


lwrTRUfraction = 0.0128; % fraction MTIHM - from CAFCA 'LWR':M11
coolingpool = ones(1,5)*2000; % size of this vector is the wet cooling time
coolingpoolFR = zeros(1,2);

% yearly front-end output data
FuelConsumed = zeros(1,100); % MTHM
SWU = zeros(1,100); % kg-SWU
NaturalUraniumConsumed = zeros(1,100); % MT

% yearly back-end output data
SNFcooled = zeros(1,100);
FRcooled = zeros(1,100);
TRUlwr = zeros(1,100);
TRUfr = zeros(1,100);
TRUlosses = zeros(1,100);
TRUstock = zeros(1,100);
SNFstock = zeros(1,100);
FRstock = zeros(1,100);
SNFinventory = legacySF;
FRinventory = 0;
TRUinventory = 0;


for year = 1:100 % 100 year simulation, 2010 - 2110
% nuclear electricity growth and forecasting
  if year == 1
    demand(year) = demandzero;
    forecast = demand(year)*(1+growth1)^5;
```

```matlab
  elseif year < yrChange1 - 5
    demand(year) = demand(year - 1)*(1+growth1);
    forecast = demand(year)*(1+growth1)^5;
  elseif year < yrChange1 && year > yrChange1-5
    demand(year) = demand(year - 1)*(1+growth1);
    forecast = demand(year)*(1+growth2)^5;
  elseif year < yrChange2 - 5
    demand(year) = demand(year - 1)*(1+growth2);
    forecast = demand(year)*(1+growth2)^5;
  elseif year < yrChange2 - 5 && year < yrChange2
    demand(year) = demand(year - 1)*(1+growth2);
    forecast = demand(year)*(1+growth3)^5;
  elseif year < yrChange3 - 5
    demand(year) = demand(year - 1)*(1+growth3);
    forecast = demand(year)*(1+growth3)^5;
  elseif year < yrChange3 - 5 && year < yrChange4
    demand(year) = demand(year - 1)*(1+growth3);
    forecast = demand(year)*(1+growth4)^5;
  else
    demand(year) = demand(year - 1)*(1+growth4);
    forecast = demand(year)*(1+growth4)^5;
  end

    ops = sum(reactors(6:60));
    ReactorsOperating(year) = sum(reactors);
    generation = ops*LWRcapfac + sum(fastreactors(6:60))*FRcapfac;
    futuregen = sum(underconstruction(1:4))*LWRcapfac + ...
      sum(underconstructionFR(1:4))*FRcapfac;
    forecastedneed = forecast-(generation+futuregen);
    Forecast(year) = forecastedneed;

% build LWRs
  if OTflag
    reactorsneeded = forecastedneed/(LWRcapfac);
    if reactorsneeded >= 1
      underconstruction(5) = floor(reactorsneeded);
      LWRdemandfilled(year) = underconstruction(5)*0.9;
    else
      underconstruction(5) = 0;
    end
  end

% FRONT-END FUEL CALCS
   FuelConsumed(year) = fuelperLWR*reactors(60) + fuelperyearLWR*...
     sum(reactors(1:59));
   NaturalUraniumConsumed(year) = FuelConsumed(year)*(lwrenrichment-...
     xt)/(xf-xt);
   vp = (2*lwrenrichment-1)*log(lwrenrichment/(1-lwrenrichment));
   vf = (2*xf-1)*log(xf/(1-xf));
   vt = (2*xt-1)*log(xt/(1-xt));
   SWU(year) = (FuelConsumed(year)*vp + (NaturalUraniumConsumed(year)...
     - FuelConsumed(year))*vt - NaturalUraniumConsumed(year)*vf)*1000;

   %%% REPROCESS AND FUEL UP %%%

% REPROCESSING - note all FRSNF/SNF has same losses
```

```
if year > reprocstart
  % build reprocessing capacity

  tRC = thermalReprocCapacity(year);
  fRC = frReprocCapacity(year);
  toProcThermal = min(tRC, SNFinventory);
  toProcFast = min(fRC, FRinventory);

  TRUlwr(year) = lwrTRUfraction*toProcThermal;
  TRUfr(year) = truFractionFR*toProcFast;
  SNFinventory = SNFinventory - toProcThermal;
  FRinventory = FRinventory - toProcFast;

  TRUlosses(year) = 0.01*(TRUfr(year)+TRUlwr(year));
  TRUlwr(year) = 0.99*TRUlwr(year);
  TRUfr(year) = 0.99*TRUfr(year);
end

if OTflag == 0

% FUEL FAST REACTORS
  TRUinventory = TRUlwr(year) + TRUfr(year) + TRUinventory;
  fuelneededFR = (sum(fastreactors(1:59))*yearlyFRneed + fastreactors(60)*...
    coreFR)*truFracFresh;

%%%%% BUILD FRs & LWRs %%%%%
  unprocessedTRU = (SNFinventory - toProcThermal)*lwrTRUfraction*0.99...
    + (FRinventory - toProcFast)*truFractionFR*0.99;

  if year > FRstart && year < FRend
% FORECAST FRs
  fleetforecast = 60*(TRUinventory)/lifetimeTRUneed;
  operatingFRs = sum(fastreactors);
  FRsOperating(year) = operatingFRs;

  reactorsneeded = forecastedneed/(FRcapfac);

  tobuild = min(fleetforecast - operatingFRs, reactorsneeded);

  if fuelneededFR > TRUinventory
    makeupfuel(year) = fuelneededFR - TRUinventory;
    TRUinventory = 0;
  else
    TRUinventory = TRUinventory - fuelneededFR;
  end

% BUILD FRs
    if tobuild >= 1
      underconstructionFR(5) = floor(tobuild);
      FRdemandfilled(year) = underconstructionFR(5)*0.9;
    else
      underconstructionFR(5) = 0;
    end
  end
```

```
% build remaining LWRs
  remainingneed = forecastedneed - underconstructionFR(5)*FRcapfac;
  LWRdemand(year) = remainingneed;
  reactorsneeded = remainingneed/(LWRcapfac);
    if reactorsneeded >= 1
      underconstruction(5) = floor(reactorsneeded);
      LWRdemandfilled(year) = underconstruction(5)*0.9;
    else
      underconstruction(5) = 0;
    end

end

    %%% DISCHARGE AND DECOMMISSION %%%

% LWR SNF
  snfdischarged = reactors(1)*fuelperLWR +...
    sum(reactors(2:60))*fuelperyearLWR;
  SNFcooled(year) = coolingpool(1);
  coolingpool = circshift(coolingpool,[0, -1]);
  coolingpool(5) = snfdischarged;
  SNFstock(year) = SNFinventory + SNFcooled(year);
  SNFinventory = SNFstock(year);

% FR SNF
  frdischarged = fastreactors(1)*fuelperFR + ...
    sum(fastreactors(2:60))*fuelperyearFR;
  FRcooled(year) = coolingpoolFR(1);
  coolingpoolFR = circshift(coolingpoolFR,[0, -1]);
  coolingpoolFR(2) = frdischarged;
  FRstock(year) = FRinventory + FRcooled(year);
  FRinventory = FRstock(year);

  TRUstock(year) = TRUinventory;

% step reactor arrays
   FRsConstructed(year) = underconstructionFR(1);
   FRsDecommissioned(year) = fastreactors(1);
   fastreactors = circshift(fastreactors, [0,-1]);
   fastreactors(60) = underconstructionFR(1);
   underconstructionFR = circshift(underconstructionFR, [0, -1]);


   ReactorsConstructed(year) = underconstruction(1);
   ReactorsDecommissioned(year) = reactors(1);
   reactors = circshift(reactors, [0,-1]);
   reactors(60) = underconstruction(1);
   underconstruction = circshift(underconstruction, [0, -1]);

end
```

Appendix C: A Medium-Cost, Waste-Neutral Technology Option

Throughout the thesis, nuclear electricity demand is treated as an exogenous variable: nuclear power growth may be high or low, depending on the availability and cost of other energy generation technologies. This enables an exclusive focus on nuclear technology decisions, but in reality, such decisions will be made in the context of the broader energy mix. One natural question that arises, and is not easily answered within the framework presented here, is whether nuclear technology decisions will be different if an option exists with a cost between LWRs and FRs and which neither mitigates nor produces waste. We can imagine, for example, that coal plants with carbon capture and sequestration achieve costs below those for FRs, but are still more expensive than LWRs. Such plants would not produce any nuclear waste, but also would not help with our existing stocks of spent fuel.

In some sense, EUFRs can approximate the role of a medium-cost, waste-neutral technology if we assume costs between those of TFRs and LWRs. If anything, EUFRs are in fact likely to cost slightly more than TFRs, because the reactor cores will need to accommodate both uranium and transuranic fuels and may include reflectors that would be unnecessary in TFRs. The EUFR assumption is further imperfect because EUFRs do produce a small amount of nuclear waste during TRU processing and at the decommissioning stage. But this amount of waste is very small compared to the waste produced by LWRs, and the cost assumption, while likely inaccurate, enables the EUFR to act as a place-holder for a truly waste-free and medium cost electricity option.

This exploration uses the same tree employed in section 5.3, reprinted as Figure C-1. Now, however, the cost premium applied to EUFRs that are built will be different from the cost premium applied to TFRs.

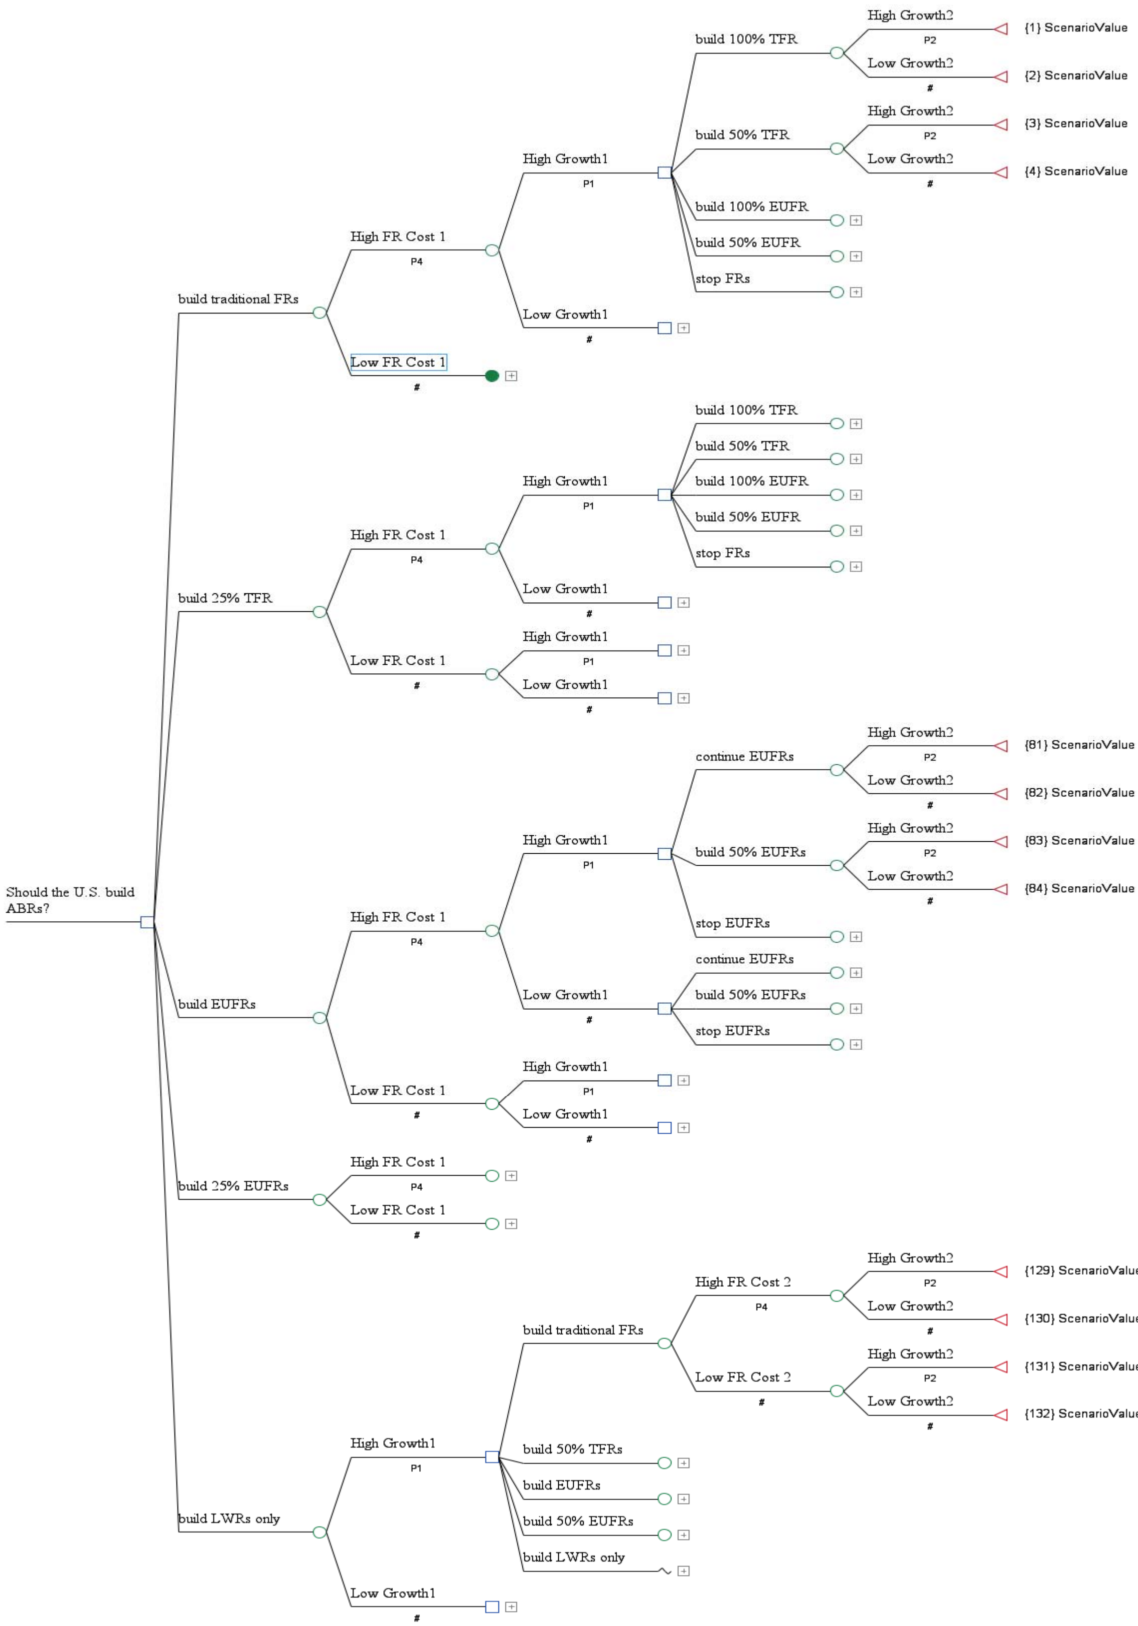

Figure C-1: Two-period tree with partial build options

Two sets of cost premiums are explored for EUFRs (Table C-1). TFRs maintain the traditional cost premiums, at 55% above LWR costs when FR costs are high and 5% when FR costs are low. The first set of percentage values were chosen to essentially split the TFR premium in half, and the second were chosen to see if discounting EUFRs even more would cause them to appear in the decision space.

Table C-1: EUFR cost premiums over LWR cost

% premium over LWR cost	EUFRs		TFRs
	Set 1	Set 2	
High Cost	25%	10%	55%
Low Cost	2%	1%	5%

The basic result for set 1 is shown in Figure C-2. Changing the cost premium for EUFRs has virtually no effect on the decision space. There is a very slight change: building 10% TFRs becomes less desirable by a miniscule amount compared to the result from section 5.3. The reason for this shift is that the EUFR scenarios become less expensive, so the range of cost values is now smaller, leading to minor changes in the relative utilities of the various options. Other decision spaces for this scenario (e.g. for the second-period decisions) exhibit the same behavior. Changing the EUFR cost premiums to set 2, making them even more inexpensive, merely shifts the decision result very slightly in the same manner (Figure C-3).

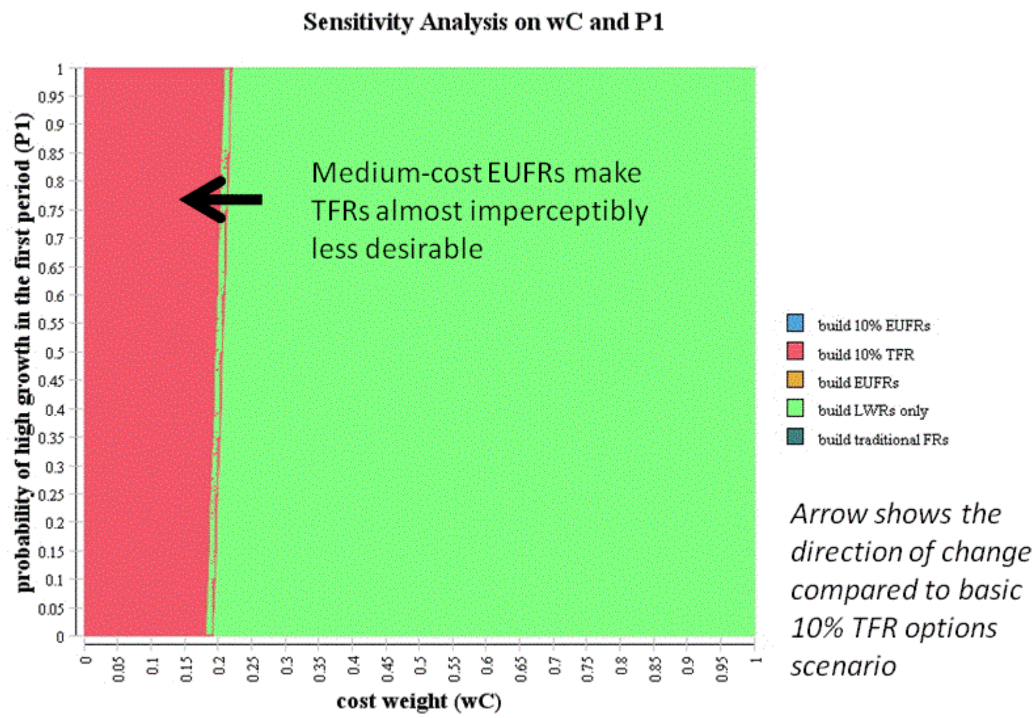

Figure C-2: Desirable decisions when EUFRs are a medium-cost option

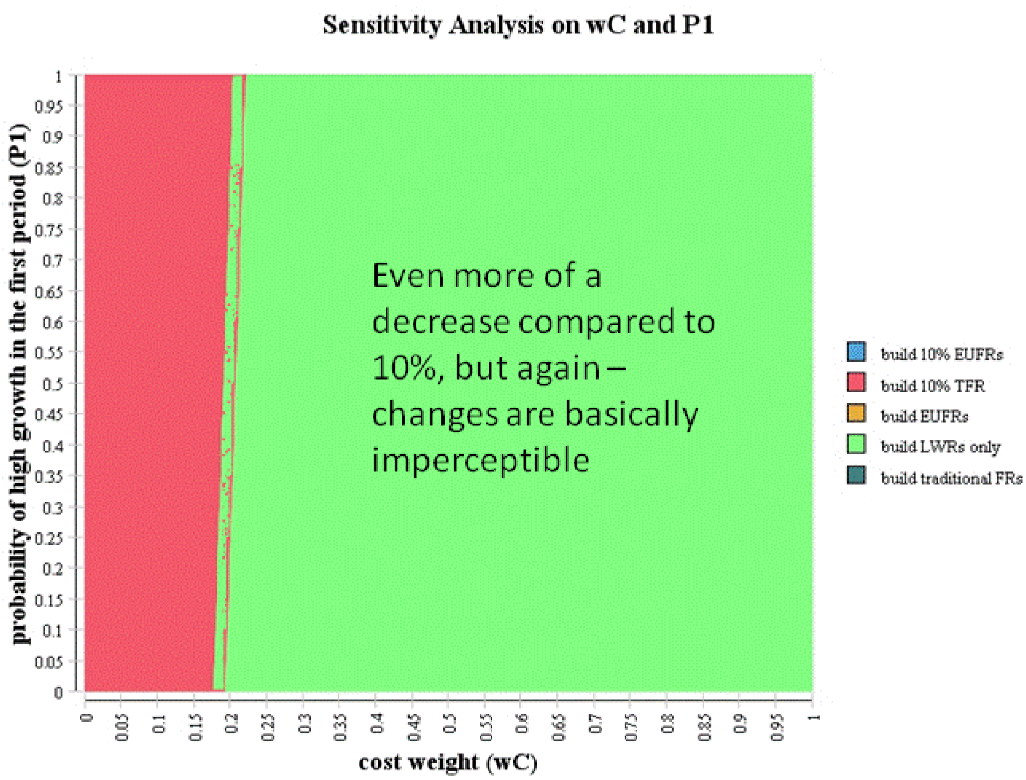

Figure C-3: Desirable decisions when EUFRs are even lower-cost

The reason we don’t see a dramatic takeover by EUFRs is that even with a lower cost premium, EUFR scenarios are expensive because so many fast reactors are built. Figure C-4 shows how several of the decision options impact the stockpile of spent LWR fuel. If EUFRs are built to take over only 10% of the fleet (expanding to 50% later in the century), many LWRs are still built and an increasing SNF stockpile results (dashed green line). Building a 100% EUFR fleet, however, provides a significant waste benefit (solid green line). But the 10% → 50% TFR scenario provides an even greater waste benefit (blue line), and entails building far fewer fast reactors than the 100% EUFR scenario. Even with stark differences in the cost of TFRs and EUFRs, the TFR scenario remains cheaper because of that vast difference in total number of FRs. EUFRs are not able to compete.

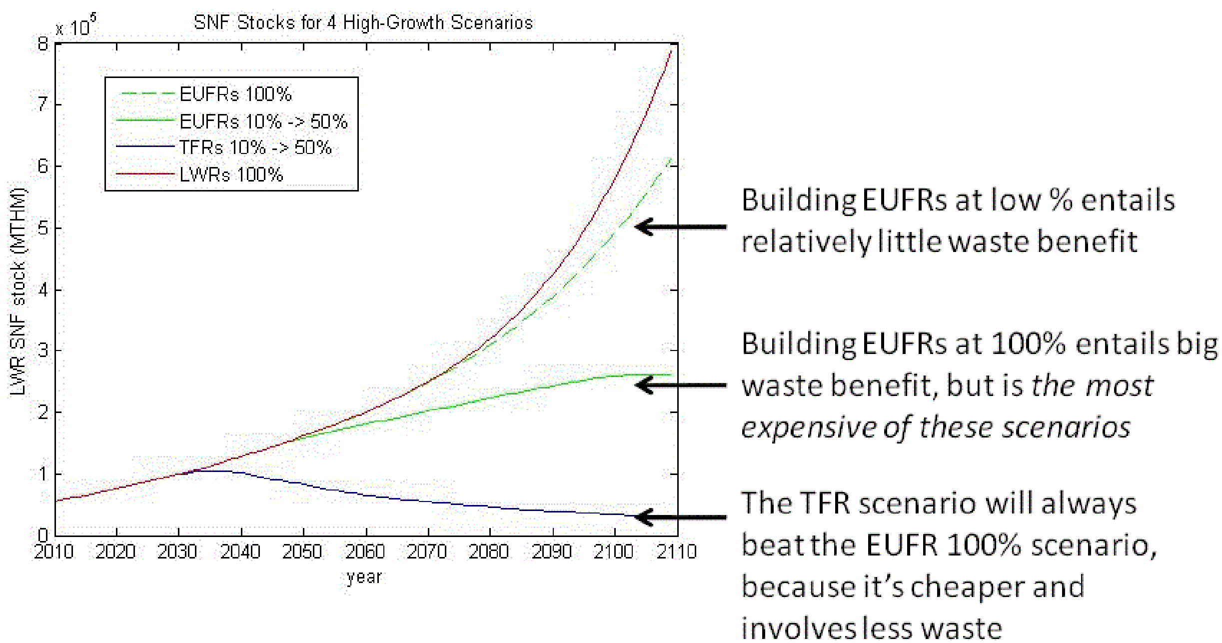

Figure C-4: Impact of different fast reactor decisions on stocks of SNF

Of course, as before, if we consider the waste liability (i.e. what happens if reprocessing eventually ceases and all reactor waste must then be managed), EUFRs become a desirable option because they can phase out LWRs (see section 6.4). If EUFRs are *cheaper* than TFRs, they become even more desirable than before, when same-cost TFRs and EUFRs were assumed and the waste liability was taken into account. Figure C-5 shows that EUFRs take over significantly more of the desirable decision space once they are cheaper than TFRs (large graph, compare to small graph), and that the decision "build 100% EUFRs" (yellow) now appears. With a medium cost, EUFRs are overall much more desirable than they were before.

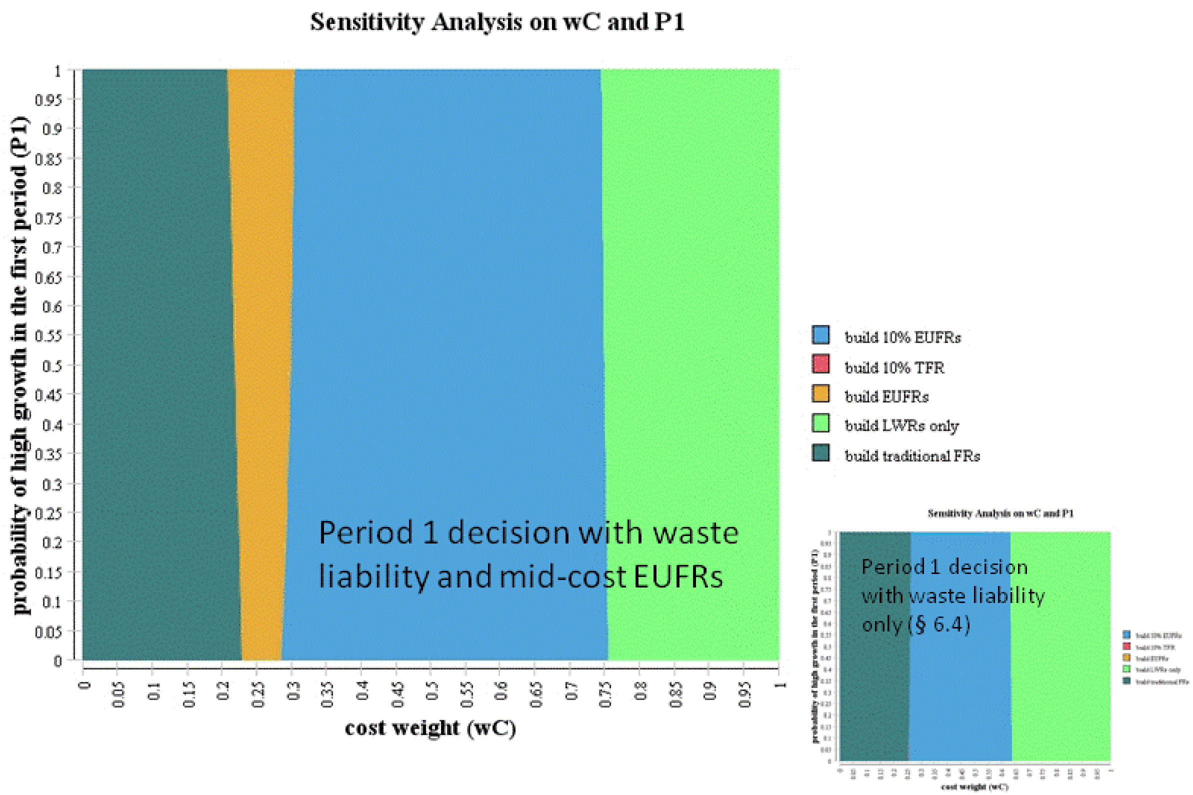

Figure C-5: Desirable decisions when waste liability is considered

Several lessons can be drawn from this analysis. One is that if reducing stocks of LWR spent fuel is of significant concern, TFRs remain the best option because large amounts of LWR SNF can be mitigated by building relatively few TFRs. The picture changes somewhat, however, when the waste liability is considered. "Mitigating" LWR SNF actually means placing waste that would otherwise have gone to a repository into fast reactors: as long as reprocessing plants and fast reactors continue operating, spent fuel nuclides stay out of the waste system. If society continues reprocessing indefinitely (or at least far enough into the future that wildly different waste management strategies become available, such that waste disposal is no longer problematic or undesirable), TFRs handle waste well for the near term. If, however, final waste disposal of all nuclear fuel a century or so in the future is a concern (i.e. if we imagine the eventual end of the nuclear system), EUFRs become desirable, and even more so if they cost less than traditional fast reactors.

Ultimately, a medium-cost, nuclear waste-free technology is not a desirable option under the base assumptions of the thesis analysis: TFRs are better in modest quantities, because they reduce existing waste relatively cheaply. If, however, other considerations like the ultimate waste liability are important, a mid-cost technology that produces little waste makes sense. Furthermore, the current framework does not directly account for other factors affecting the competitiveness of nuclear power. For example, if the mid-cost technology were non-nuclear, and public tide turned (again) strongly against nuclear power, such a technology could handily become more desirable than TFRs or LWRs in the electricity mix. More work is needed to determine the desirability of different nuclear technologies under competition with other types of electricity generation.

One final dimension is very important to questions about a mid-cost, zero-waste competitive technology in this framework: the time of technology introduction. Developing and deploying fuel cycle and reactor systems for TFRs (or for EUFRs) is likely to take decades, and will require sustained R&D funding and political support. The licensing process for such wildly new nuclear technologies will be challenging and lengthy. Given the barriers to rapid revolutionary change in the nuclear industry, it is possible that a non-nuclear technology producing no nuclear waste but offering costs just above LWRs will be competitive on the grid before TFRs are even developed. If this happens, by 2040 TFRs may no longer be desirable because we may decide to manage our existing waste (given that we are not generating much more) rather than complete the fast reactor development process. A quantitative analysis of this possibility is beyond the scope of the thesis, but is important and left to future work.

References

Nuclear Energy Institute. (2010). *Safely managing used nuclear fuel: Fact sheet.* Retrieved July 17, 2011, from http://www.nei.org/resourcesandstats/documentlibrary/nuclearwastedisposal/factsheet/safely managingusednuclearfuel/