A Game-Theoretic Approach to Nuclear Fuel Cycle Transition Analysis Under Uncertainty

U. B. PHATHANAPIROM[1,2*], D. A. HAAS[1], B. D. LEIBOWICZ[1]

[1]The University of Texas at Austin, Department of Mechanical Engineering, 204 East Dean Keeton Street, Austin, TX 78712

[2]Oak Ridge National Laboratory, PO Box 2008 MS-6165, Oak Ridge, TN 37831-6165

E-mail: [1]bphathanapirom@ornl.gov

Abstract

We present a novel methodology for optimizing nuclear fuel cycle transitions that incorporates a game-theoretic approach and captures interactions among multiple decision makers. The methodology is demonstrated using a two-person sequential game with uncertainty, where the two players represent a policy maker and an electric utility company, though the method generalizes to any number and type of individual decision making entities. Coupled with a sophisticated nuclear fuel cycle simulator, rich transition scenarios may be analyzed to identify robust transition strategies. These strategies explicitly treat uncertainties using a stochastic programming approach, devising optimal near-term hedging strategies that simultaneously consider all possible states of the world, maintaining flexibility to allow for intelligent recourse decisions once uncertainties are resolved. In the demonstration game, reactor technology and fuel cycle scheme adopted by the electric utility are shown to depend on both the policy maker's decisions and the distributions over uncertain technological and economic outcomes.

KEYWORDS
decision making under uncertainty; game theory; nuclear fuel cycle transition analysis; sequential games against nature

A Game-Theoretic Approach to Nuclear Fuel Cycle Transition Analysis Under Uncertainty

U. B. PHATHANAPIROM[1,2*], D. A. HAAS[1], B. D. LEIBOWICZ[1]

[1]The University of Texas at Austin, Department of Mechanical Engineering, 204 East Dean Keeton Street, Austin, TX 78712

[2]Oak Ridge National Laboratory, PO Box 2008 MS-6165, Oak Ridge, TN 37831-6165

E-mail: [1]bphathanapirom@ornl.gov

1. Introduction

The U.S. Department of Energy, Office of Nuclear Energy (DOE-NE) has been researching technology to shape the nuclear fuel cycle to better balance the need for energy, economic and proliferation security; environmental sustainability; and the risks associated with the development and deployment of new fuel cycle technologies. Previous fuel cycle analysis tools have been constrained by human and computer resources. Many studies have opted for a less detailed, but more expansive view of the fuel cycle because of the trade-off between depth and breadth of modeling. As a result, some key features in fuel cycle transition analysis have suffered, including: (1) treatment of transients, instead examining the fuel cycle operating at equilibrium where facilities are continually built, operated and decommissioned as needed, (2) simultaneous optimization across multiple criteria and objectives, instead focusing on specific areas of interest to the researchers conducting the study, and (3) explicit modeling of decision making under uncertainty, instead examining uncertainties through sensitivity or scenario analysis that simply varies parameter values within a deterministic model.

This paper presents a novel methodology that incorporates a game-theoretic approach for optimizing nuclear fuel cycle transitions by modeling fuel cycle decision making as a sequential game against nature featuring a policymaker and electric utility. Here, the policymaker acts first to influence the state of the world that will prevail, and as a response, the electric utility chooses to continue operating the U.S. fuel cycle as is, or close it. This approach captures the strategic interest of each player, as they are not necessarily aligned, and the subsequent effects of these players' interactions. Further, the effects of uncertainty on each player's decision making are captured, leveraging the capabilities of nuclear fuel cycle simulators that model time-dependent processes in the nuclear fuel cycle.

The methodology is presented in generic form; in that form, it can be applied to any fuel cycle simulator and with any character archetype of an individual decision making entity, each of which chooses their decisions based on different decision criteria and objectives. Previous transition optimizations have assumed the perspective of a single "benevolent dictator" shaping the future of the nuclear fuel cycle. Instead, in the new methodology presented, each participant in a nuclear project shapes a fuel cycle transition based on their unique decision criteria, and each responds uniquely to uncertainty resolutions as well as to the other participant's previous decisions as time moves forward. Here the policymaker simultaneously considers economic, waste management and proliferation resistance metrics, whereas the electric utility's sole concern is the economics of nuclear power. Incorporating these players and a *Nature* player that moves randomly brings autonomous decision making into the fuel cycle simulator.

A game-theoretic approach to energy systems modeling is not in itself new, nor is the explicit treatment of uncertainties in nuclear fuel cycle transition analysis through a stochastic programming approach. The novelty of the work presented here is their merger, coupled with a sophisticated fuel cycle simulation tool, casting the fuel cycle transition as a sequential game against nature. Compared to past literature in fuel cycle transition analysis, this work moves away from the idea of a single decision maker, instead modeling the strategic interactions between two decision makers whose objectives are not necessarily aligned. Coupling to a fuel cycle simulator allows for multiple, intelligent decisions to be made over the time horizon of the transition, and these decisions are chosen optimally through an iterative hedging algorithm, removing the modeler's responsibility to choose these decisions. This hedging algorithm explicitly incorporates uncertainties into its solution by examining the available decisions over the time horizon of the transition. Further, the work leverages past uses of fuel cycle simulators to calculate time-evolving fuel cycle metrics, and uses these metrics to depict distinct character archetypes by assigning importance weightings of these metrics to the individual decision makers. This approach brings novel richness into the analysis of fuel cycle transitions, addressing previously identified shortcomings of past fuel cycle analyses.

2. Background

2.1 Uncertainty in Nuclear Fuel Cycle Transition Analyses

More comprehensive nuclear fuel cycle transition studies have recently been made possible through the use of complex fuel cycle simulators coupled with enhanced technological capabilities. All nuclear fuel cycle transition studies determine the natural resource and technology requirements for changing over from one nuclear fuel cycle to another. For instance, many transition studies have examined the changeover from the current U.S. open fuel cycle consisting of a light water reactor (LWR) fleet to a closed fuel cycle comprising fast reactors (FRs) burning recycled used fuel from LWRs (Yacout et al., 2004; Dixon et at., 2009; Djokic et al., 2015; Feng et al., 2016; Bae et al., 2016). Uncertain parameters abound in these transitions—technology costs and availability dates, demand growth for nuclear electricity, and the potential for government loan guarantees and tax credits, to name a few. Previously, nuclear fuel cycle transition studies have handled these uncertainties using sensitivity and scenario analysis. Newer work has taken a stochastic programming approach to explicitly optimize decisions made under uncertainty, formulating robust transition strategies (Carlsen, 2016; Phathanapirom and Schneider, 2016; Pierpoint, 2017). Kann and Weyant (2000) offer a thorough description of these approaches to uncertainty analysis, which are summarized here.

Sensitivity analysis may help attribute uncertainty in model output to different sources of uncertainty in its input. Sensitivity analysis is performed by recalculating outcomes of the model while varying uncertain input parameters over their possible ranges. When variation of an input parameter produces relatively small changes in model output, that output is considered *robust*; whereas if a large variation is observed, the output is considered *sensitive*. Given its straightforward nature requiring zero modifications of the model, this type of analysis is commonly used. Although simple, sensitivity analysis is useful in that it allows increased understanding of relationships between model inputs and outputs; it may aid in future investigations by reducing computational burden via identifying inputs that cause larger uncertainty in model outputs, which should be the focus of future investigations (Bistline, 2013; Wian, 2013).

Scenario analysis is roughly similar to sensitivity analysis in that no model modifications are required, and input parameters are varied across their possible ranges. The crucial distinguishing feature of scenario analysis is its construction of different states of the world through some combination of uncertain parameter values, which represent a plausible description of how the system as a whole and its driving forces may develop in the future (Walker et al., 2003; Gabbert et al., 2010). Solutions of a deterministic optimization model are unique to each individual scenario and offer a set of coherent, internally consistent futures.

Sensitivity and scenario analysis treat uncertainties as exogenous to a deterministic model. A key implicit assumption in these approaches is that decision makers have perfect information about the state of the world that will prevail. In reality, decision makers must act before uncertainties are resolved and, in many situations, act *to* resolve those uncertainties.

Stochastic programming explicitly handles uncertainties by simultaneously considering all possible states of the world, offering a systematic approach to decision making under uncertainty. Stochastic programming requires a decision maker to make some decision now that minimizes the (usually) expected cost or consequence of that decision. Considering stochastic programs in this way gives rise to a *recourse* model, in which information available to the decision maker is updated in each sequential stage (Golub et al., 2014; Leibowicz, 2018). The simplest form of a stochastic program is the two-stage linear program with recourse; however, because of the complex, dynamic nature of the nuclear fuel cycle, linear programming is insufficient to represent the most important associated regulatory, economic, and technological issues. Instead, the nuclear fuel cycle simulator is treated as a tool that a solution algorithm can invoke to obtain an objective function value.

Pierpoint (2017) constructs an optimization wrapper that invokes the FANTSY fuel cycle simulator to examine how uncertainties in the nuclear power demand growth rate and reactor capital costs affect the decision to close the fuel cycle. The wrapper enumerates each branch of the decision tree and obtains objective function values for each branch. Once each branch is scored, the method of backward induction is used to find optimal hedging strategies. Similarly, Carlsen (2016) examines time-wise uncertainties such as disruption in fuel supply to devise optimal hedging strategies for reactor deployment. Because of the fidelity of the Cyclus simulator (Huff et al., 2016), coupled with the fine decision space examined, Carlsen relied on a custom particle swarm optimizer to obtain an approximate solution. Phathanapirom and Schneider (2016) developed a progressive hedging algorithm to find optimal hedging strategies for use in considering uncertain radioactive waste disposal costs.

2.2 Extension to Game Theory

Binsbergen and Marx (2007) found that some sequential games can be analyzed by considering multiple decision trees, with each tree corresponding to an individual player. In sequential games, one player chooses their action and it is observed by the other player before they choose their action. If the second player has no information about the first's action, the effect is equivalent to the players choosing their actions simultaneously. Table 2.1 identifies the components of a no-data problem and its corresponding counterparts in a game against nature, both of which find their roots in decision making under uncertainty. These trees account for the dependence of the payoffs on the actions of the other players. While the two may be mathematically equivalent in some cases, insights into the strategic interaction between decision makers can be gained by examining a nuclear fuel cycle transition through a game-theoretic lens.

These sequential decision making problems may be solved using a recursive method known as backward induction or dynamic programming (Rust, 2006).

Table 2.1: Key elements of decision analysis and corresponding counterparts in game theory.

Decision Analysis	Game Theory
Set of alternatives	Strategy set
Chance and unknown events	Moves of *Nature*
Results	Payoff mapping
Solution concept	Equilibrium concept

Pierpoint (2011) proposed and briefly examined a government–industry interaction model. Pierpoint examined industry response to an increase in the 1 mil per kWh nuclear waste fee under the Nuclear Waste Policy Act of 1982. In her work, Pierpoint assumed the obligation of the government stops at waste management, and as a consequence the government absolutely benefits from the transition to FRs. The variable waste fee is exogenously applied in the simulation. Under a different guise, the variable waste fee could be viewed as a first-stage strategy.

Resource allocation in safeguards and security applications to nuclear facilities have been examined using a game-theoretic approach. Avenhaus (2013), Butler et al. (2013), and Ward and Schneider (2016) each examine a facility operating with static material flows, independent of the dynamics of material flow resulting from the nuclear fuel cycle system in which facilities are continually being built and retired. However, in addition to individually examining an enrichment and reprocessing facility, Ward explores a systems approach to optimization across the two facilities. Ward considers a simultaneous game, aimed at optimal safeguards such as random inspections against a proliferation scenario. An “efficient frontier” is identified that depicts payoff as a function of budget. Butler et al. incorporate uncertainties in a sequential decision making model, though they do not examine recourse decisions following uncertainty resolutions, which is equivalent to solving for the expected value solution.

Key to a game-theoretic approach in nuclear fuel cycle transition applications is the consideration of external costs—those costs that are paid by society as a whole rather than exclusively by consumers of nuclear power. Each decision maker in a nuclear fuel cycle transition may select their decisions based on a unique set of decision criteria and weightings, a process that gives rise to the interaction among individual decision makers. The challenge in examining these externalities arises in deriving a mechanism for estimating the costs of the impacts and identifying appropriate importance weightings.

Table 2.2 summarizes concepts presented in this section. The role of uncertainty in fuel cycle decision making has primarily been investigated using sensitivity and scenario analysis, although a few recent works have employed stochastic programming to explicitly consider hedging decisions in a setting where uncertainties unfold dynamically over time. The novelty of this paper is casting the fuel cycle transition as a sequential game against nature, in which multiple players take turns making strategic decisions over time in response to, and in anticipation of, each other's decisions and the uncertain moves of the random nature player.

Table 2.2: Approaches to fuel cycle transition analysis and their corresponding features.

Approach	Summary	Meaningful States of the World?	Hedging Decisions?	Multiple Strategic Players Interacting?
Sensitivity Analysis	Varies parameters over their possible ranges	No	No	No
Scenario Analysis	Varies parameters that are chosen to represent plausible, coherent states of the world	Yes	No	No
Stochastic Programming	Considers plausible, coherent states of the world and evolving information sets as subsequent decisions are made	Yes	Yes	No
Sequential Games Against Nature (Our Approach)	Considers plausible, coherent states of the world and evolving information sets as subsequent decisions are made by multiple agents	Yes	Yes	Yes

3. Methodology

Fig. 3.1 depicts a visualization of the two-stage decision tree arising from a sequential game against nature with two players. In Fig. 3.1, decision nodes are those in which a decision is made, and are represented by squares and colored according to the acting decision maker; chance nodes are the roots of branches in the tree where stochastic parameter outcomes are realized, and are represented by circles; and end points (nodes) show the final outcome of a decision path and are represented by triangles. The cost (or objective function or metric value) $F(d_1, \omega, d_2)$ is history-dependent on the first-stage decision d_1, the outcome of the stochastic parameter ω, and the second-stage decision d_2. We find the optimal recourse decision d_2^* in response to d_1 using

$$d_2^* = \underset{d_2}{\operatorname{argmin}}\, F(d_1,\omega,d_2) \tag{3.1}$$

which chooses from available d_2 after realizing the value of ω such that $F(d_1, \omega, d_2)$ is minimized. The first-stage decision (d_1) represents a near-term hedging strategy that is chosen optimally using

$$d_1^* = \underset{d_1}{\operatorname{argmin}} \sum_{\omega} \mathcal{P}_\omega F(d_1,\omega,d_2^*) \tag{3.2}$$

which minimizes the expected cost over the available d_1 which is a probability-weighted sum of $F(d_1, \omega, d_2^*)$ given $\mathcal{P}_\omega$, the probability of realizing some value of ω. Typically, the assumption is that the decision maker's objective is to choose a strategy that minimizes the expected costs, although that is not necessarily the case. For instance, a decision maker could aim to minimize the maximum possible cost, referred to as a "minimax problem," which is an example of robust optimization. The costs (or gains) in decision theory are equivalent to the notion of payoffs in game theory.

In Fig. 3.1, two decision makers act at d_1 and d_2, represented by the differently colored decision nodes. In this two-stage game, or any sequential decision making problem, a bilevel program arises wherein the upper optimization problem is constrained by the lower optimization problem, with the decision maker in the upper problem able to anticipate the subsequent optimal reaction by the decision maker in the lower problem. The sequential game is solved through the method of backward induction, a recursive method that finds an optimal decision rule. In the context of the nuclear fuel cycle, fuel cycle simulators must be employed to calculate a payoff. Although the interactions between fuel cycle facilities are formulated as a systems model (consisting of levels and flow between objects), the highly nonlinear problem effectively restricts its solution to an algorithm that calls the fuel cycle simulator as a "black-box" model to be invoked.

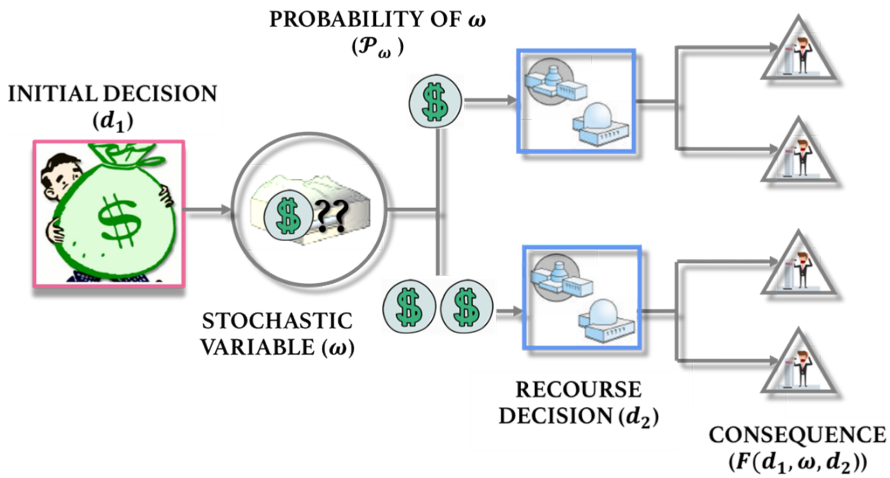

Figure 3.1: Two-stage decision tree for a sequential game against nature with two players. Depiction includes an initial decision of one player (pink square), a recourse decision of a second player (blue square) with an associated cost (grey triangle) and a single stochastic parameter (ω).

For the general *N*-stage program, determination of the optimal hedging strategies for each stage $n = 1, 2 \ldots N$ is carried out recursively with the aid of a function named *Hedge* depicted in Fig. 3.2. *Hedge* accepts two arguments: k, the stage for which possible hedging strategies are being enumerated, and n, the stage for which strategies are being selected. *Hedge* is always initialized at the first stage, $k = 1$, and iterates through all J_k strategies in stage k:

$$d^k_{j=1,2\ldots J_k} \tag{3.3}$$

For all J_k strategies in stage k, *Hedge* increments k by 1 and calls itself:

$$Hedge(k_{,} = k + 1, n) \tag{3.4}$$

This process is repeated until k reaches the value of n. In this way, all histories:

$$\left(d_{j_k}^{k}, \omega\right) \forall\, k, j_k,\omega \tag{3.5}$$

are enumerated. These histories are all the permutations of the strategies for previous stages and all outcomes of the stochastic parameters. For all histories, an optimal hedging strategy at stage n is found using:

$$h_n^* = \underset{h_n}{\operatorname{argmin}}\, F\left(\left\{\left(d_j^k,\omega\right)\forall k < n, j_k, \omega\right\}, h_n, h_{n+1\ldots N}^* \right) \tag{3.6}$$

where F is the cost function described previously. Here, *Hedge* considers each available hedging strategy in stage n, and calculates the associated cost of choosing that strategy given the history leading up to that decision, and all downstream optimal hedging strategies for stages $n + 1$ to N determined previously.

The sequential two-person game in Fig. 3.1 is analyzed by considering multiple decision trees. Each decision maker has a unique objective function $F_p(d_1, \omega, d_2)$ that determines the cost of a path from root to leaf through the decision tree. Then, the costs can be represented by a bimatrix, consisting of values $F_p(d_1, \omega, d_2^*)$, where p indicates the cost to a player p. The *Hedge* algorithm determines all $F_p(d_1, \omega, d_2)$ to determine the optimal hedging strategy for each stage based on the objective of the player who is acting at that stage.

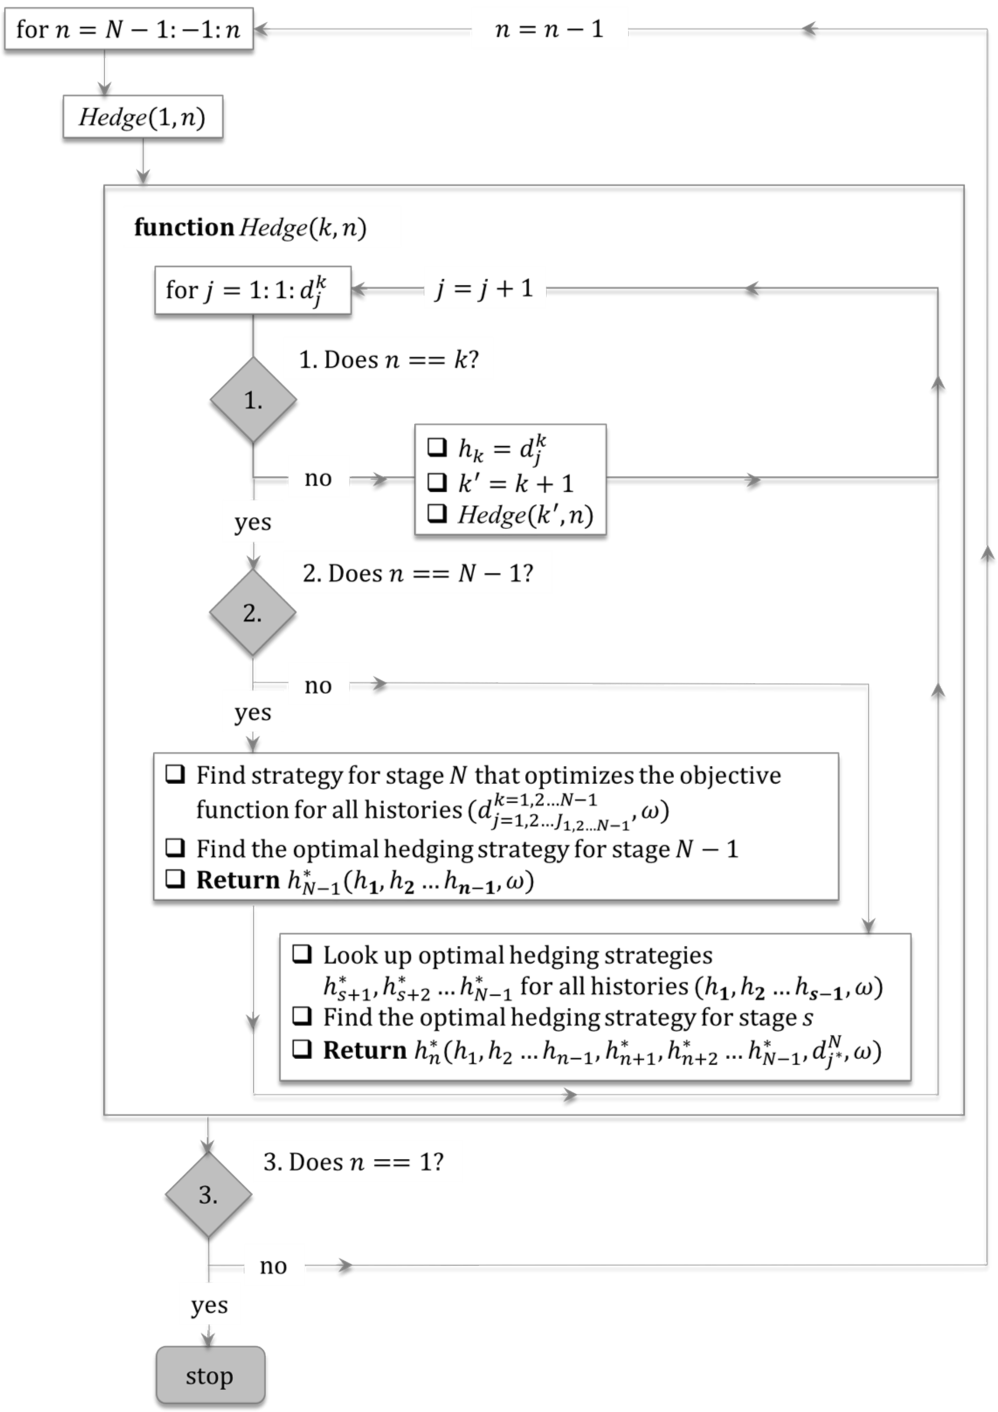

Figure 3.2: Recursive hedging algorithm via a backward induction method for the general *N*-stage decision problem.

4. Demonstration

A two-stage demonstration game is illustrated in Fig. 3.1. The game (scenario) includes two players whose payoffs have a net result greater than or equal to zero (general-sum) where gains by one player do not necessarily correspond to losses for the other. These players represent two of the prime participants in a nuclear project: a government entity (Player *Government*) and a utility company that generates electricity (Player *Utility*). The strategic interactions of these players pose a major challenge in the development of nuclear infrastructure. Players act sequentially with complete information, that is, each knows the other's available strategies, probabilities of *Nature*'s moves, and the corresponding payoffs. Thus, *Government* chooses his decision anticipating the influence of that decision on *Utility*'s subsequent decision by examining *Utility*'s available decisions and corresponding payoffs. *Utility* responds with the knowledge of *Government*'s initial decision and the outcome of any upstream stochastic parameters. These stochastic parameters are represented as choices of a *Nature* player who moves randomly with no concept of a payoff. Moves by *Nature* define the state of the world $\omega \in \Omega$, about which *Government* and *Utility* are assumed to have full distributional knowledge of the probabilities of all states occurring. Calculations of the payoffs for the demonstration game are summarized here, although the specifics are omitted.

4.1 Two-stage Game Description

The two-stage game is informed by results from the VEGAS nuclear fuel cycle simulator that calculates a material- and technology-constrained mass balance (Schneider and Phathanapirom, 2016). VEGAS's unique mass balance calculation ensures that each reactor's fuel demands are met throughout its entire lifetime by utilizing a roll-back feature that returns the simulation to the year in which a reactor that violates material- or technology-constraints was added and removes it from the simulation. A summary timeline of the reference transition scenario and key VEGAS simulation input parameters is shown in Figure 4.1. Each VEGAS simulation begins in 2018 with an initial 100 GW_e reactor fleet of LWRs and ends in 2160, allowing the simulation to run beyond the decision making period by an additional lifetime of the longest-operating facility to ensure that liability costs are accounted for. During the simulation, a nuclear electricity demand growth rate of 2.3% per year is assumed (WNA, 2017). Legacy used fuel is assumed to be directly disposed of in a geologic repository; consequently, it has no effect on the results presented here.

In the two-stage game, *Government* plays first and chooses whether to invest in waste disposal or reprocessing research and development (R&D). *Waste Disposal R&D* affects the probability distribution of possible waste disposal cost outcomes, whereas *Reprocessing R&D* deterministically chooses the reprocessing technology cost. Based on the waste disposal cost outcome, influenced by *Government*'s chosen R&D strategy, and the reprocessing cost, *Utility* chooses a fuel cycle scheme, consisting of either (1) the once-through fuel cycle with direct disposal of used fuel discharged from thermal reactors or (2) the closed fuel cycle with continuous recycle of used fuel discharged from thermal reactors in FRs. Technology costs are available in Appendix A.

Prior to 2035, electricity demand targets are met by building new LWRs, with the original fleet of LWRs retiring by 2040. In 2035, *Utility* chooses a fuel cycle strategy based on the outcome of *Government*'s stage-one decision and has three available options to choose from: continue building LWRs, build high-temperature gas-cooled reactors (HTGRs) that benefit from increased fuel utilization, or build sodium-cooled FRs (SFRs) that recycle used fuel discharged from thermal reactors, closing the fuel cycle. If *Utility* decides to build SFRs, VEGAS attempts to build SFRs to satisfy demand; but if available used fuel for recycle is insufficient to continue steady supply of fresh fuel for an SFR during its lifetime, VEGAS's roll-back feature is implemented. The roll-back feature moves the simulation clock back to the year of construction for the SFR for which insufficient fuel is available, replacing it with a lower-tier reactor (LWR or HTGR) that requires enriched natural uranium, of which there is an unlimited supply.

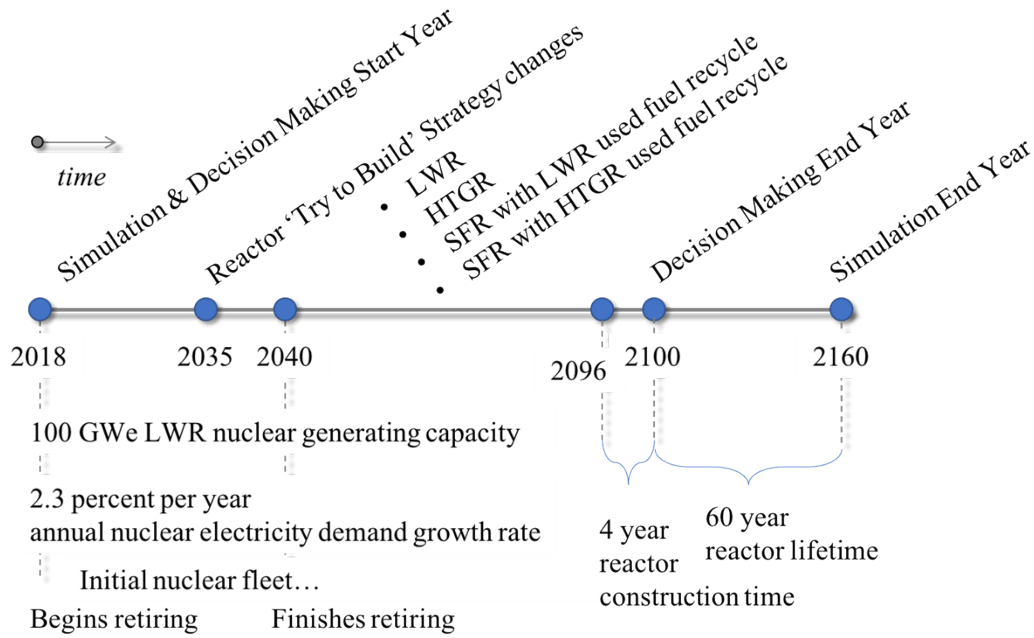

Figure 4.1: Timeline of reference fuel cycle transition scenario with key events and VEGAS simulation input parameters.

4.2 Payoff calculation

Each unique VEGAS simulation is determined by *Government*'s decisions, *Utility*'s decisions, and *Nature*'s moves; these collectively correspond to a path from root to leaf of the decision tree, defined as a fuel cycle transition scenario, such as that depicted in Fig. 3.1. The desirability of each scenario is evaluated by *Government* and *Utility* based on their decision criteria, each of which is measured through some fuel cycle metric quantified through their objective function values. Each scenario results in unique "consequences". Here, the three metrics of interest to the two players are the total cost of electricity, the cumulative heat load to the repository, and the nuclear security measure during the fuel cycle simulation period. These fuel cycle metrics compete directly, as fuel cycles that rely on separating actinides for recycle in FRs generally are low in proliferation resistance, although they result in a lesser waste repository heat burden. The opposite trend is true for fuel cycles that directly dispose used fuel. Costs for both advanced technologies are uncertain. The values of these metrics in a given scenario are calculated using the output mass balance from the VEGAS simulator. These calculations are described briefly.

4.2.1 Cost of Electricity

The levelized cost of electricity (LCOE) is the constant dollar price of electricity that would be necessary over the economic life of a plant to provide an acceptable return on equity for investors. The LCOE consists of two components: (1) C_{FC}, front- and back-end fuel cycle charges that are calculated by applying unit costs (typically in dollars per kg U, initial heavy metal (IHM), or separative work units (SWU)) to the mass balance that quantifies the flow of materials between fuel cycle processes as kg U, IHM, or SWU and (2) $C_{reactor}$, reactor charges that are calculated from the total overnight capital cost of the plant, amortized under assumptions regarding the discount rate and construction time. The annual LCOE is calculated using

$$\mathrm{LCOE} = \frac{C_{FC} + C_{reactor}}{E} \quad (4.1)$$

where E is the total energy in kWh_e generated in that year. The total cost of electricity (COE) is then found using

$$\mathrm{COE} = \sum_{t=t_0}^{T} E_t \mathrm{LCOE}_t \quad (4.2)$$

where E_t is the total electricity generated each year in the VEGAS simulation.

4.2.2 Decay Heat

For a given design and location of a geologic waste repository, the size of the repository is controlled by the decay heat (Hardin et al., 2011; Wigeland et al., 2006). The total cumulative heat load to the repository can be calculated using:

$$D = \sum_{t=t_0}^{T} \sum_r M_{r,t} \left\{ (1 - r_{r,t}) d_r^{SNF} + y_{FP_r} r_{r,t} d_r^{FP} \right\}, \quad (4.3)$$

where

$M_{r,t}$	=	mass throughput of reactor type r in year t
$r_{r,t}$	=	fraction of fuel discharged from reactor type r in year t ultimately reprocessed
y_{FP_r}	=	output FP mass fraction from reactor type r
d_r^{SNF}	=	decay heat constant of reactor type r from SNF
d_r^{FP}	=	decay heat constant of reactor type r from FPs

The decay heat intensities per tIHM[1] of spent nuclear fuel (SNF), d_r^{SNF}, and high-level waste (HLW) resulting from reprocessing of discharged used fuel, d_r^{HLW}, are applied to the mass balance, giving the total heat load to the repository over the lifetime of the VEGAS simulation. The mass balance gives $M_{r,t}$, the total mass of fuel (tIHM) from reactor type r in year t, and $r_{r,t}$, the fraction of that fuel ultimately reprocessed. In Eq. (3), $M_{r,t}(1 - r_{r,t})$ is the quantity of SNF disposed of, which, when multiplied by d_r^{SNF} , yields the total heat load to the repository resulting from SNF. In Eq. (3), $r_{r,t}$ is the fraction of discharged used fuel reprocessed, and y_{FP_r} is the mass fraction of fission products (FPs) in the resulting HLW. Then $M_{r,t}r_{r,t}y_{FP_r}$ is the quantity of FPs in HLW producing heat, at a rate d_r^{FP}, with $M_{r,t}r_{r,t}y_{FP_r}d_r^{FP}$ being the total heat load to the repository resulting from disposal of FPs in HLW. The decay heat intensities per tIHM or SNF and HLW for LWR and HTGR technologies is available in Appendix A.

4.2.3 Proliferation Resistance

The proliferation resistance of a fuel cycle is a measure of its ability to resist the illicit diversion of material for the production of weapons-usable material through both intrinsic (to the material or process) and extrinsic (or engineering) barriers. Charlton et al. (2007) provide a thorough methodology for calculating the dynamic proliferation resistance of a fuel cycle.

---

[1] The calculation of d_r^{SNF} and d_r^{HLW} involves reactor fuel depletion and decay calculations using the Oak Ridge Isotope Generation (ORIGEN) code included in the SCALE 6.2 package (Rearden and Jessee, 2016).

The total nuclear security measure, defined by Charlton as the total time-and mass-weighted average of the static proliferation resistance values PR_i for a fuel cycle consisting of $i = 1, 2, ..., I$ processes, is found using:

$$NS = \frac{\sum_{i=1}^{I} m_i \Delta t_i PR_i}{\sum_{i=1}^{I} m_i \Delta t_i} \tag{4.4}$$

where

m_i	=	amount of material in process i in significant quantities
Δt_i	=	time material is in process i at the static proliferation resistance value for process i
PR_i	=	static proliferation resistance value for process i

The static proliferation resistance for process i is found using:

$$PR_i = \sum_{j=1}^{J} w_j u_j(x_{ij}) \tag{4.5}$$

where

w_j	=	weight for attribute j
u_j	=	utility function for attribute j
x_{ij}	=	input value for the utility function for attribute j in process i

which is a weighted sum of a score (translated from a utility function u_{ij} based on a quantitative or qualitative measure) from j different attributes of the ability of the process ability to impede proliferation.

The average nuclear security measure during a given VEGAS simulation is used here to measure the overall proliferation resistance of the fuel cycle, which can be calculated using:

$$\mathrm{NS} = \frac{1}{T} \sum_{t=t_0}^{T} NS_t \tag{4.6}$$

NS_t	=	total nuclear security measure in year t
NS	=	average nuclear security measure for a given fuel cycle transition path during the fuel cycle simulation period

Each fuel cycle metric is translated into a score using the min-max normalization, in which the least favorable metric value across all scenarios receives a score of 0 and the most favorable metric value receives a score of 1, with all intermediate metric values normalized using linear interpolation. For the total COE and the cumulative heat load to the repository, the score is found using:

$$S_c^{\vartheta} = \frac{f_c^{\vartheta} - \min_{\Theta} f_c^{\vartheta}}{\max_{\Theta} f_c^{\vartheta} - \min_{\Theta} f_c^{\vartheta}} \tag{4.7}$$

where f_c^{ϑ} is the metric value for path ϑ through the decision tree, with the set $\vartheta \in \Theta$ producing the entire decision tree depicted in Fig. 3.1. The score of the proliferation resistance metric is:

$$\mathrm{S}_c^{\vartheta} = 1 - \frac{f_c^{\vartheta} - \min_{\Theta} f_c^{\vartheta}}{\max_{\Theta} f_c^{\vartheta} - \min_{\Theta} f_c^{\vartheta}} \tag{4.8}$$

The evaluation criteria are indexed over c, with the overall payoff for path ϑ calculated using:

$$P^{\vartheta} = \sum_c w_c S_c^{\vartheta} \,. \tag{4.9}$$

Once payoffs are determined for each $\vartheta \in \Theta$, the *Hedge* algorithm depicted in Fig. 3.2 is used to select the decision that maximizes the expected payoff.

Government's and *Utility*'s assumed decision criteria weightings are given in Table 4.2 for the two-stage demonstration game, which are subject to the constraint:

$$\sum_c w_c = 1 \tag{4.10}$$

where c indexes over the cost of electricity, decay heat and proliferation resistance criteria. *Government*'s three criteria are assumed to be of equal importance. Ensuring that nuclear power remains a viable marketplace option along with other electricity generation technologies is advantageous if climate change policy is enacted. The international consensus is that geologic repositories represent the best-known method for permanently disposing of SNF and HLW generated from nuclear power production. Siting, construction, and licensing of a geologic repository is assumed to be a federal responsibility (DOE, 2013). Since the decay heat from SNF and HLW ultimately determines the size of the repository needed, minimizing the heat load is favorable. Finally, from a national security standpoint, the need to develop a proliferation-resistant fuel cycle is apparent. On the other hand, *Utility*'s two criteria include the COE and decay heat, with greater importance placed on the former. Utilities are typical businesses in many respects; and increased electricity sales result in increased revenues and therefore profits, especially in an unregulated market. Historically, under the Nuclear Waste Policy Act of 1982, utilities were charged 1 mill per kWh of nuclear electricity, paid to a Nuclear Waste Fund, which was to fund the development of repositories for disposing of SNF and HLW (DOE, 2004). Yucca Mountain was designated as the first site for a geologic repository for nuclear waste in 1987 and was originally approved in 2002. However, as a result of the DOE shutdown of the Yucca Mountain project in 2010, the federal government has failed to meet its obligation to dispose of nuclear waste, leaving 39 states to store radioactive waste on-site (NEI, 2018). Because of the issues surrounding licensing, constructing, and operating a nuclear waste repository, utilities may be responsible for long-term storage of their nuclear waste. Generally, fuel cycles that directly dispose of discharged used fuel result in a larger repository heat burden, although the proliferation resistance is high because these cycles avoid producing separated actinides. The opposite relationship holds in considering fuel cycles that recycle discharged used fuel.

Table 4.2: Player *Government*'s and *Utility*'s assumed decision criteria weighing for demonstration problem.

	Criterion Weighting ($\boldsymbol{w_c}$)	
Evaluation Criterion ($\boldsymbol{c}$)	Player *Government*	Player *Utility*
Cost of Electricity	$0.\overline{3}$	0.9
Decay Heat	$0.\overline{3}$	0.1
Proliferation Resistance	$0.\overline{3}$	0.0

4.3 Results: Perfect Information and Hedging Strategies

The perfect information strategies represent those strategies that would be optimal *if* the decision maker knew all of *Nature*'s moves (outcomes of stochastic parameters) in advance. Additionally, these strategies require that each player correctly anticipate the moves of the other player. For the two-stage demonstration game, *Government*'s and *Utility*'s decisions when perfect information is available are depicted in Fig. 4.2. In the case where the waste disposal cost outcome is low (Fig. 4.2a), *Utility* chooses thermal reactors, which produce a greater volume of waste, since disposing of used fuel is cheap. *Government* can choose a "do nothing" strategy because the *Reprocessing R&D* strategy affects the reprocessing cost; no reprocessing technology is employed for the once-through fuel cycle; and the *Waste Disposal R&D* strategy affects the probability distribution of the waste disposal cost outcomes, which are already known in the perfect information scenario. If *Government* were to choose the "do nothing" strategy, it is likely that leftover funds from opting out of R&D spending would be used in operation of the waste repository. However, these leftover funds are implemented in the VEGAS simulations as a capital subsidy for LWRs to ensure that all decision alternatives entail the same level of total expenditures, which facilitates the analysis. Alternatively, in the case where the waste disposal cost outcome is high (Fig. 4.2b), *Government* chooses the *Reprocessing R&D* strategy to lower the cost of reprocessing, and *Utility* chooses to transition to a closed fuel cycle by building FRs. As a result, less waste is disposed of in the repository, which has the effect of minimizing the fuel cycle cost and heat load to the repository.

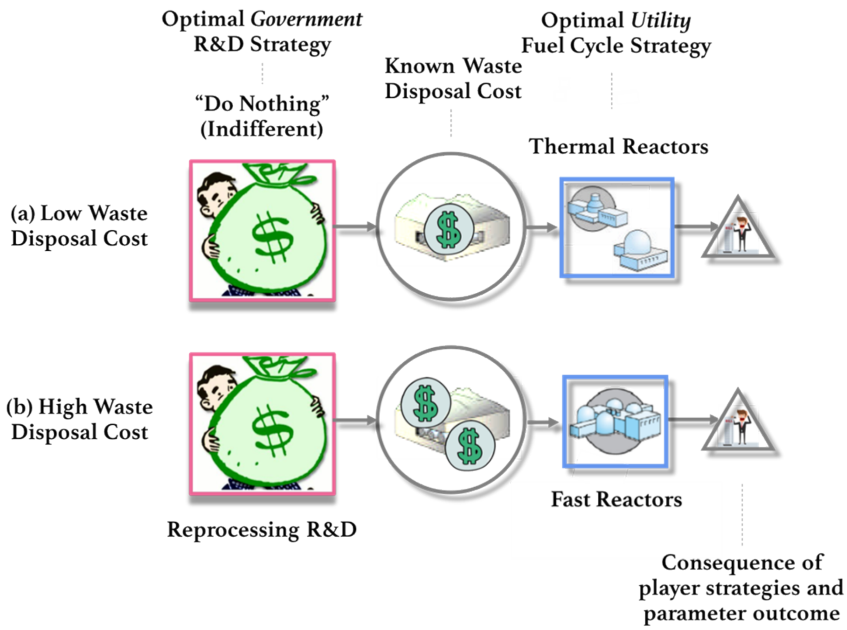

Figure 4.2: Player *Government*'s optimal R&D strategy and Player *Utility*'s optimal fuel cycle strategy (thermal reactors: open fuel cycle and fast reactors: closed fuel cycle) when both players operate with perfect information, knowing that in (a) the waste disposal cost is low and in (b) the waste disposal cost is high.

When *Government* and *Utility* hedge optimally under uncertainty, a transition to a closed fuel cycle is never observed. *Government* hedges by choosing the *Waste Disposal R&D*, since the expected state of the world (overall disposal cost outcomes) is unfavorable to recycling of used fuel. Then, in all cases, *Utility* builds thermal reactors and directly disposes of discharged reactor fuel. On the contrary, if *Government* instead chooses the *Reprocessing R&D* strategy and the waste disposal cost outcome is high, then *Utility* is observed to switch the strategy toward a closed fuel cycle, building FRs. These behaviors are depicted in Fig. 4.3a, in which *Government* and *Utility* hedge optimally, and Fig. 4.3b, in which *Government*'s decision is fixed as *Reprocessing R&D* and the waste disposal cost outcome is high.

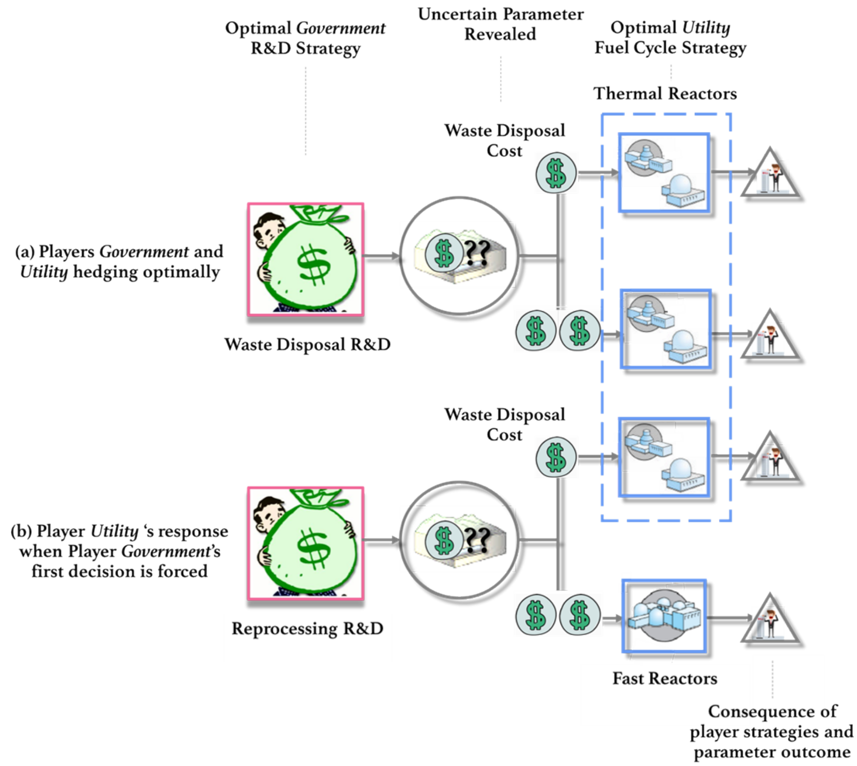

Figure 4.3: Player *Utility*'s optimal recourse decision (thermal reactors: open fuel cycle and fast reactors: closed fuel cycle) (a) when Player *Government* chooses his optimal hedging R&D strategy and (b) when *Government*'s R&D strategy is fixed. In (a), *Government* chooses his strategy before knowing the state of the world that will prevail, and in both (a) and (b) *Utility* chooses his recourse decision after the state of the world is revealed.

The strategies chosen by each player—both perfect information and hedging strategies—are contingent on their decision criteria weightings. Fig. 4.4 illustrates how *Utility*'s decision to pursue an open fuel cycle (blue) or closed fuel cycle (red) depends on the weightings of the three decision criteria (w_{COE}, w_{PR}, w_{DH}) for (cost of electricity, decay heat and proliferation resistance). Fig. 4.4a represents cases where *Government* chose *Waste Disposal R&D*, whereas Fig. 4.4b represents cases where *Government* chose *Reprocessing R&D*. The quantiles of Fig. 4.4a and 4.4b are then shaded based on *Utility*'s recourse decision (opening or closing the fuel cycle) after having observed *Government*'s decision.

Examining the bottom-left quantile of both Fig. 4.4a and 4.4b, *Utility*'s optimal recourse decision is shaded blue to indicate the decision to keep the fuel cycle open, building thermal reactors and directly disposing of the resulting used fuel. This decision is made when *Utility* is concerned only about lowering the COE, regardless of whether *Government* chooses *Waste Disposal* or *Reprocessing R&D*. Compare that scenario with the bottom-right region shaded orange, in which *Utility*'s optimal recourse decision is now to close the fuel cycle, based only on lowering the heat load to the repository.

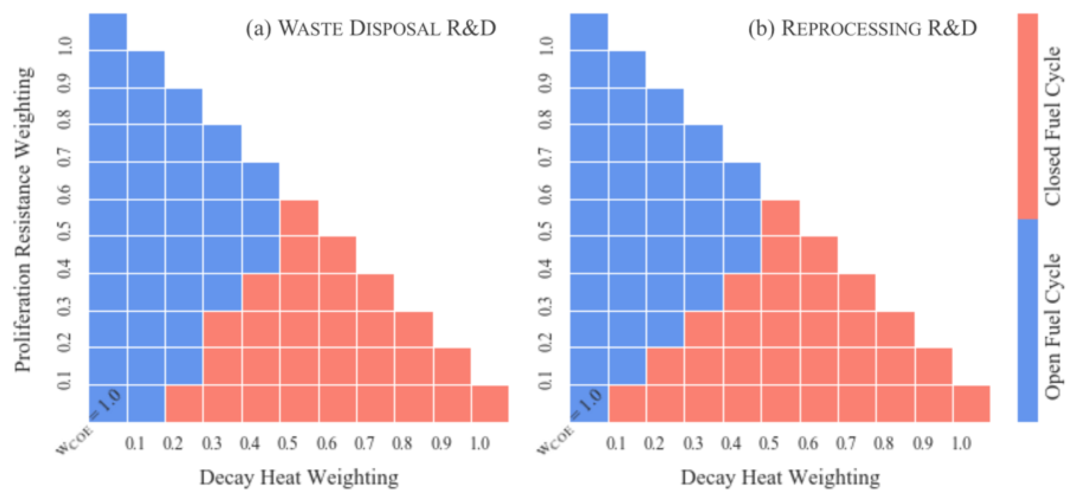

Figure 4.4: Player *Utility*'s weighted decision matrices (with $w_{COE} = 1\text{-}w_{PR}\text{-}w_{DH}$) for selecting an optimal recourse decision based on Player *Government*'s first-stage strategy being (a) *Waste Disposal R&D* or (b) *Reprocessing R&D* and a high waste disposal cost outcome.

Utility's recourse decision is most heavily influenced by the tradeoff between proliferation resistance and decay heat. However, the effect of the cost of electricity can be observed when considering the transition of *Utility*'s weightings of (w_{COE}, w_{DH}, w_{PR}) from (0.0, 0.5, 0.5) to (0.2, 0.4, 0.4). Here, the ratio between w_{DH} and w_{PR} remains the same, but w_{COE} increases from 0.0 to 0.2. Using Fig. 4.3a, we observe that *Utility*'s optimal recourse decision changes from a closed fuel cycle to an open fuel cycle.

The threshold observed in Fig. 4.3, in which *Utility* changes his decision based on *Government*'s stage-one strategy, can be seen by comparing the quantile of row one and column two of *Utility*'s decision matrices in Fig. 4.4. In this quantile, *Utility* has a baseline criteria weighting in Table 4.2 (0.9, 0.1, and 0.0 for the COE, decay heat and proliferation resistance, respectively). Here, *Government*'s choice between *Waste Disposal R&D* and *Reprocessing R&D* results in a different response by *U*. These inflection points at which *Government* is able to influence *Utility*'s recourse decision can be seen by comparing *Utility*'s decision matrices in Fig. 4.4a and 4.4b, and where *Utility*'s recourse decision is altered by his own criteria weighting by comparing the individual quantiles in Fig. 4.4a or 4.4b.

5. Conclusions

This paper presents a novel methodology for optimizing nuclear fuel cycle transitions that captures interactions between self-interested agents. The methodology is demonstrated in an example two-person, two-stage problem, in which players represent a policy maker and an electric utility company. While the example scenario features just two players and two stages, the hedging algorithm presented generalizes to any number of stages and agents. Results from the example scenario identify a near-term hedging strategy that maintains flexibility to adapt a transition strategy or policy choice based on new information. These adapted decisions are termed "recourse decisions." This strategy shift in response to information gained during the transition illustrates the importance of coupling a fuel cycle simulator with a model of autonomous decision making.

The work presented in this paper addresses some shortcomings of past fuel cycle analysis studies, including treatment of transients, a multi-agent decision making process (versus the "benevolent dictator" norm), and explicit handling of uncertainties. Many components of the analysis may be developed further in future work. Briefly, some expansions may be the inclusion of simultaneous decisions when multiple utility companies compete to fulfill a fixed demand for nuclear electricity, different scalarization methods for multi-objective optimization, and inclusion of other objectives in addition to the three considered here.

Beyond the realm of fuel cycle transition analysis, the work presented here has further applications in nuclear safeguards and security, where a fuel cycle simulator may be used to identify vulnerabilities in the fuel cycle. A novel coupling of a fuel cycle simulator to an adversarial game offers the ability to more realistically calculate decision criteria such as the time requirement for significant diversion of special nuclear material, idle enrichment or reprocessing capacity, and quantities of stockpiled separated actinides. A temporal cross section in a fuel cycle simulation containing quantities and qualities of material circulating and in stockpiles, as well as available capacities of nuclear technologies may give initial conditions for a breakout scenario in which an aggressive strategy to produce or divert large quantities of high-value special nuclear material is pursued or, at worst, a regime change whereby a civilian nuclear power program is abandoned in favor of a nuclear weapons production program.

References

[Avenhaus, 2013] Avenhaus, R. (2013). Safeguards Systems Analysis: With Applications to Nuclear Material Safeguards and Other Inspection Problems. Springer Science & Business Media.

[Bae et al., 2019] Bae, J. W., Peterson-Droogh, J. L., & Huff, K. D. (2019). Standardized verification of the CYCLUS fuel cycle simulator. *Annals of Nuclear Energy*, *128*, 288-291.

[Binsbergen and Marx, 2007] Van Binsbergen, J. H., & Marx, L. M. (2007). Exploring relations between decision analysis and game theory. *Decision Analysis*, *4*(1), 32-40.

[Bistline, 2013] Bistline, J. E. (2013). Essays on Uncertainty Analysis in Energy Modeling: Capacity Planning, R & D Portfolio Management, and Fat-tailed Uncertainty (Doctoral dissertation, Stanford University).

[Butler et al., 2013] Butler, J. C., Cronin, P. M., Dyer, J. S., Edmunds, T. A., & Ward, R. M. (2013). *Decision Analysis Methods For the Analysis of Nuclear Terrorism Threats with Imperfect Information* (No. LLNL-TR-635767). Lawrence Livermore National Laboratory.

[Carlsen, 2016] Carlsen, R. W. (2016). *Advanced Nuclear Fuel Cycle Transitions: Optimization, Modeling Choices, and Disruptions* (Doctoral dissertation, The University of Wisconsin-Madison).

[Charlton et al., 2007] Charlton, W. S., LeBouf, R. F., Gariazzo, C., Ford, D. G., Beard, C., Landsberger, S., & Whitaker, M. (2007) Proliferation resistance assessment methodology for nuclear fuel cycles. *Nuclear Technology, 157,* 143-156.

[Dixon et al., 2009] Dixon, B., Kim, S., Shropshire, D., Piet, S., Matthern, G., & Halsey, B. (2008). *Dynamic Systems Analysis Report for Nuclear Fuel Recycle* (No. INL/EXT-08-15201). Idaho National Laboratory (INL).

[Djokic et al., 2015] Djokic, D., Scopatz, A., Greenberg, H. R., Huff, K. D., Nibbleink, R. P., & Fratoni, M. (2015). *The Application of CYCLUS to Fuel Cycle Transition Analysis* (No. LLNL-CONF-669315). Lawrence Livermore National Laboratory.

[DOE, 2004] U.S. Department of Energy. (2004) *Nuclear Waste Policy Act (Amended with appropriations acts appended)*. Office of Civilian Radioactive Waste Management.

[DOE, 2013] U.S. Department of Energy. (2013). Strategy for the management and disposal of used nuclear fuel and high-level radioactive waste.

[Feng et al., 2016] Feng, B., Dixon, B., Sunny, E., Cuadra, A., Jacobson, J., Brown, N. R., ... & Gregg, R. (2016). Standardized verification of fuel cycle modeling. *Annals of Nuclear Energy*, *94*, 300-312.

[Gabbert et al., 2010] Gabbert, S., Van Ittersum, M., Kroeze, C., Stalpers, S., Ewert, F., & Olsson, J. A. (2010). Uncertainty analysis in integrated assessment: the users' perspective. *Regional Environmental Change*, *10*(2), 131-143.

[Golub et al., 2014] Golub, A., Narita, D., & Schmidt, M. G. (2014). Uncertainty in integrated assessment models of climate change: Alternative analytical approaches. *Environmental Modeling & Assessment*, *19*(2), 99-109.

[Hardin et al., 2011] Hardin, E., Blink, J., Greenberg, H., Sutton, M., Fratoni, M., Carter, J., ... & Howard, R. (2011). *Generic Repository Design Concepts and Thermal Analysis*. (SAND2011-6202). Sandia National Laboratory.

[Huff et al., 2016] Huff, K. D., Gidden, M. J., Carlsen, R. W., Flanagan, R. R., McGarry, M. B., Opotowsky, A. C., ... & Wilson, P. P. (2016). Fundamental concepts in the Cyclus nuclear fuel cycle simulation framework. *Advances in Engineering Software*, *94*, 46-59.

[Kann and Weyant, 2000] Kann, A., & Weyant, J. P. (2000). Approaches for performing uncertainty analysis in large-scale energy/economic policy models. *Environmental Modeling & Assessment*, *5*(1), 29-46.

[Leibowicz, 2018] Leibowicz, B. D. (2018). The cost of policy uncertainty in electric sector capacity planning: Implications for instrument choice. *The Electricity Journal*, *31*(1), 33-41

[NEI, 2018] Nuclear Energy Institute. (2018) https://www.nei.org/advocacy/make-regulations-smarter/used-nuclear-fuel. Accessed July 2018.

[Phathanapirom and Schneider, 2016] Phathanapirom, U. B., & Schneider, E. A. (2016). Nuclear fuel cycle transition analysis under uncertainty. *Nuclear Science and Engineering*, *182*(4), 502-522.

[Pierpoint, 2011] Pierpoint, L. M. (2011). *A Decision Analysis Framework for the U.S. Nuclear Fuel Cycle* (Doctoral dissertation, Massachusetts Institute of Technology).

[Pierpoint, 2017] Pierpoint, L. M. (2017). Illuminating fuel cycle decision drivers using a decision analysis framework. *Nuclear Science and Engineering*, *186*(1), 66-82.

[Rearden and Jessee, 2016] Rearden, B. T. & Jessee, M. A. (2016). *SCALE Code System,* (ORNL/TM-2005/39, Version 6.2). Oak Ridge National Laboratory.

[Rust, 2006] Rust, J. (2006). Dynamic programming. *New Palgrave Dictionary of Economics*.

[Schneider and Phathanapirom, 2016] Schneider, E. A., & Phathanapirom, U. B. (2016). VEGAS: A Fuel Cycle Simulation and Preconditioner Tool with Restricted Material Balances. *Nuclear Technology*, *193*(3), 416-429.

[Shropshire et al., 2009] Shropshire, D. E., Williams, K. A., Boore, W. B., Smith, J. D., Dixon, B. W., Dunzik-Gougar, M., ... & Schneider, E. (2009). *Advanced fuel cycle cost basis*. (INL/EXT-07-12107) Idaho National Laboratory.

[Walker et al., 2003] Walker, W. E., Harremoës, P., Rotmans, J., van der Sluijs, J. P., van Asselt, M. B., Janssen, P., & Krayer von Krauss, M. P. (2003). Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support. *Integrated Assessment*, *4*(1), 5-17.

[Ward and Schneider, 2016] Ward, Rebecca M., and Erich A. Schneider. A game theoretic approach to nuclear safeguards selection and optimization. *Science & Global Security* 24.1 (2016): 3-21.

[Wian, 2013] Tian, W. (2013). A review of sensitivity analysis methods in building energy analysis. *Renewable and Sustainable Energy Reviews*, *20*, 411-419.

[Wigeland et al., 2006] Wigeland, R. A., Bauer, T. H., Fanning, T. H., & Morris, E. E. (2006). Separations and transmutation criteria to improve utilization of a geologic repository. *Nuclear Technology*, *154*(1), 95-106

[Wigeland et al., 2014] Wigeland, R., Taiwo, T., Ludewig, H., Todosow, M., Halsey, W., Gehin, J., Jubin, R., Buelt, J., Stockinger, S., Jenni, K., & Oakley, B. (2014). Nuclear fuel cycle evaluation and screening – final report. Idaho National Laboratory, INL/EXT-14-31465, FCRD-FCO-2014-000106

[WNA, 2017] World Nuclear Association. (2017). *Nuclear Power Economics and Project Structuring*. (No. 2017/001).

[Yacout et al., 2004] Yacout, A. M., Hill, R. N., Van Den Durpel, L., Finck, P. J., Schneider, E. A., Bathke, C. G., & Herring, J. S. (2004). Dynamic analysis of the AFCI scenarios. *PHYSOR 2004*, 25-29.

Appendix

Values of key parameters used throughout the analysis of the reference transition scenario are given in the table below. Some commentary on these values follows:

Waste Disposal Costs: Obtained from the 2009 Advanced Fuel Cycle Cost Basis report (Shropshire et al., 2009).

Reactor Capital Costs: SFR capital costs were collected from fuel cycle expects of the Advanced Fuel Cycle Initiative Economic Working Group, who also informed updated LWR capital costs from their 2009 estimates.

Decay Heat Coefficients: Reactor fuel recipes are approximated from the DOE's Evaluating and Screening Study (Wigeland et al., 2014). Fuel burnup and depletion calculations were performed using the Oak Ridge Isotope Generation (ORIGEN) code included in the SCALE 6.2 package (Rearden and Jessee, 2016). Post processing of ORIGEN results yield actinide and fission product decay heat intensities.

Static Proliferation Resistance Values: Similar calculations to decay heat coefficients, though at a finer isotopic scale. Calculation of nuclear security attributes for material in their forms during individual fuel cycle processes which contribute to the overall proliferation resistance value follow the methodology presented by Charlton et al. (2017). Values in the table are presented in units of utility per unit mass through the fuel cycle process.

Table A.1: Key parameter values for analysis of reference transition scenario.

Waste Disposal Costs		
	Low Cost Outcome	High Cost Outcome
SNF Disposal ($/kg IHM)	602	987
HLW Disposal ($/kg FP in IHM)	4,133	8,795
Reactor Capital Costs		
	LWR	SFR
Total Overnight Capital Cost ($/kWe)	4,177	4,155
Decay Heat Coefficients		
	LWR	SFR
SNF Disposal (watts per tIHM)	1.849E+03	8.929E+03
HLW Disposal (watts per tFP in HLW)	1.322E+03	2.738E+03
Static Proliferation Resistance Values		
	LWR	SFR
Uranium Mining (per kg U as U_3O_8)	0.791	0.791
Conversion (per kg U as U_3O_8)	0.791	0.791
Enrichment (per SWU)	0.788	1.000
Fuel Fabrication (per kg IHM)	0.807	0.575
SNF Storage (per kg IHM)	0.854	0.815
Reprocessing (per kg IHM)	0.583	0.599
SNF Disposal (per kg IHM)	0.825	0.838
HLW Storage (per kg IHM)	0.914	0.916
HLW Disposal (per kg IHM in HLW)	0.919	0.917