Modeling the Transition from Fossil Fuel to Renewable Electricity Generation Using a Finite Horizon Markov Decision Process

Alaisha K. Sharma

Harvard College
Computer Science 91r

15 May 2020

Contents

1 Abstract

Modeling the transition to majority renewable energy systems is becoming increasingly critical as countries around the world race to meet ambitious decarbonization targets set for 2030, 2050, and even 2100. One approach to this problem is to capture as many different components and layers of the energy system as possible. This produces much more comprehensive models but at the expense of computational simplicity and ability to generalize results across different regions. We propose an energy system transition model based on a finite horizon Markov Decision Process (MDP) that considers a simplified problem space with four main components. Using this model we quantify how, over a fixed period of time, combining a) a deterministic carbon tax with b) stochastically decreasing costs of technology would incentivize c) the conversion of fossil fuel electricity plants into renewable energy plants that are d) accompanied by energy storage systems. In accordance with most literature on decarbonizing energy systems, we found that the costs of energy storage are by far the biggest bottleneck to increasing renewable penetration of the electricity system. Motivated by existing carbon taxes around the world, our results also show that the starting price of carbon must increase exponentially to achieve linear increases in renewable penetration. Importantly, this reveals that a carbon tax schedule starting at about 40 USD and increasing by 5% annually (the most widely supported option for a U.S. national carbon tax) would need to almost double its starting price in order to achieve just 50% renewable penetration by 2070. Although our model lacks the complexity to make specific policy recommendations for specific regions, we hope that these results highlight the urgent need to reduce the costs of utility scale energy storage and to implement significantly higher carbon taxes than most countries in the world currently have.

2 Introduction

The concept of an *energy transition* is rapidly gaining traction as global annual carbon emissions are projected to reach 37 billion tons in 2030 and 45 billion tons in 2050 [16]. The devastating impacts of maintaining business-as-usual (BAU) emissions within the current global energy system cannot be understated: record high industrial greenhouse gas (GHG) emissions, peak intensity heatwaves and storms, ocean acidification, and soil deterioration, among countless other factors, are already posing serious threats to rural communities, mostly in warmer developing countries. However, climate research indicates this is just the beginning of an acceleration that could cost the world over 10% in GDP by 2100 in the case of BAU, which projects a global mean temperature increase of between 3 and 5 degrees Celsius [3]. As a result, political leaders are starting to recognize the urgency of a long term, global transition to a sustainable, carbon-free or at least carbon-neutral energy system.

The staggering complexity of the energy system presents several challenges to modeling and optimizing this transition. Even narrowed to just electricity, determining the optimal transition from 90% fossil fuel to anywhere from 50 - 100% renewable energy based production could involve thousands of variables, from the individual costs of different renewable energy technologies, to the penetration of electric vehicles, to the addition of high-speed transmission cables between disconnected sections of national grids [37]. Many hybrid energy system models disregard time completely by performing overnight optimization. Other models optimize the energy system for some set year in the future but fail to capture the transition itself. Comprehensive models including the LUT Energy System Transition Model and TIMES family models use a high degree of both spatial and temporal granularity and incorporate hundreds of customizable inputs. This, however, makes their results both extremely computationally expensive and difficult to generalize beyond the specific region whose data is used as input to the model.

Despite these differences, an overarching takeaway from modeling the energy transition is that reaching higher levels of renewable penetration involves an extremely high price tag. To reach net zero carbon emissions in the U.S. by 2030, as the Green New Deal[1] proposes, could cost as high as $5.7 trillion, close to a quarter of the national debt [9]. This implies a severe discrepancy between renewable penetration targets for this century and existing economic incentives towards decarbonizing the energy system. For instance, as of 2020 the U.S. still has no nationally implemented carbon tax scheme.

Our model is motivated by the question of how high the penalty of continued fossil fuel generation would need to be in order to incentivize majority renewable penetration on a less ambitious timescale of about 50 years. Instead of a very complex model whose results are necessarily tailored to a specific region, we intentionally kept our model very general. We intend this model to be useful not in designing specific energy transition roadmaps but rather in providing more universal, region-agnostic insights on the challenges of this transition. This paper offers corroborating evidence for the fact that reinforcing renewable energy with sufficient energy storage makes increasing renewable penetration exponentially more difficult as more fossil fuel energy is phased out. Furthermore, although existing carbon taxes in countries such as Sweden and Finland seem exorbitantly high, our results bring into question the ability of carbon taxes below this level to achieve a majority renewable energy based electricity system by 2070.

[1] The Green New Deal is a piece of legislation that argues strongly for transitioning the U.S. to a sustainable energy system by 2030. However, it is less a roadmap than a set of ideas without specific cost estimates.

3 Background

3.1 Electric Grid

The energy system modeled in this paper represents an isolated electric grid. The *grid* is a general term that comprises the entire electricity supply chain, from production by power plants to delivery into residential buildings. The flow of electricity within the grid is driven by supply and demand of electric power. On the consumer side, *demand* refers to the power (in units of kW) required by consumers at some point in time. The *load* refers to the energy (in units of kWh) required by consumers, which is equal to the power required over some period of time. Formally, energy is the integral of power over time. *System load* therefore specifies the total electrical demand over time, usually one year, and a *load profile* refers to the shape of the power curve during a specific time period, usually a day, month, or year. The reliability of an electric grid is often measured as its ability to supply or meet load, and it is particularly challenging for a grid to follow a load profile that has many sharp peaks because it requires rapid ramping up and down of electricity production.

In reality, the electric grid writ large is an interconnected system comprising several smaller sub-grids. These sub-grids are usually connected via high-speed transmission cables that can deliver bulk electricity over hundreds of miles. This allows electricity to be bought and sold across regions and ensures that the suppliers of electricity, power plants, need not produce power exactly equal to demand at all times. For instance, the Western Interconnection is an electric grid that powers from southwestern Canada to Baja California; this larger grid connects state level grids, and is itself connected with the other major electric grids that power the rest of North America [32]. With energy inflows of over 80 million MWh per year, the state of California is the biggest net electricity importer in the U.S. [4] A grid that is disconnected from the main distribution system is known as *islanded*. These grids are closed systems, in that all of the electricity produced must be consumed and if demand exceeds power output, it will not be met. Islanded grids are usually a more favorable option for remote, unelectrified areas in which the community does not have a high power demand.

3.2 Power Plant Economics

There are several types of costs associated with an electricity generating power plant. The first is the *capital cost*, or the one-off, initial cost of construction. Capital costs depend on the size or *nameplate capacity* of the power plant (in units of kW or MW). The second type of costs is the umbrella *operation and maintenance* (O&M) cost, or the continual cost of keeping the plant running. O&M costs can be further divided into *fixed* costs, which also depend on the size of the power plant, and *variable* costs, which depend on how much power is produced (actual power output). An example of fixed O&M costs would be labor costs, and an example of variable O&M cost would be fuel costs, which for renewable energy plants are zero since resources like sun and wind are considered free.

The actual power output of a plant is calculated as the nameplate capacity multiplied by the *capacity factor*, which is defined as the ratio between how much power the power plant actually generates versus the theoretical maximum power it could generate over time. This is especially relevant in the context of renewable energy plants, since their maximum power output is always much more than their actual power output. For example, if the sun were shining constantly, the maximum output of a 100% efficient solar PV plant would be exactly its nameplate capacity. Yet because, on average, only half the day sees sunshine, the average maximum capacity factor of solar PV is 50% (in reality it is lower since PV panels cannot perfect capture all solar energy).

Since the upfront (capital) costs of renewable energy plants are so high, one especially important factor in determining the payoff on investment is the *discount rate*. The discount rate describes how much less a sum of money today would be worth one year in the future. When upfront costs are high, it takes a long time for an investment like a renewable energy plant to pay off; the a higher discount rate, the less valuable revenue made in the future becomes and the less optimal the investment seems. Higher discount rates are typically used for industry projects with short timescales (a decade or two), but there is substantial controversy in using the same rates to make decisions about renewable energy projects since the timescale of climate change is on the order of centuries [49].

3.3 Renewable Energy and Storage

Integrating more renewable energy into the electric grid presents challenges beyond the cost of building renewable energy plants themselves. One of the biggest problems with renewable energy is *intermittency*, which refers to the inherently unpredictable nature of renewable resources such as wind and solar radiation. The power curve of an onshore wind plant, for instance, would have regular peaks in the nighttime but could drop to almost zero in the middle of the day. Furthermore, renewable resources vary significantly with location and season, leading to drastic differences in the minimum and maximum power output of renewable energy plants.

As *renewable penetration* of the grid, or the percentage of electricity generated by renewable energy plants, increases, the power supply becomes increasingly unreliable. Whereas fossil fuel generation can provide a stable *base load*, or minimum power output over a period of time, there is always a risk that the wind may not blow or the sun may not shine, and renewable generation will thus fail to supply the minimum load required. On the other hand, the wind may blow extremely strongly and the sun may shine extremely brightly, but if there is not a high enough load on the grid, this leads to *curtailment* of power plants; they are shut down for periods of time because the electricity being produced has no consumer. This paper considers onshore wind and solar photovoltaic (PV) power generation, both of which are highly susceptible to this issue.

Energy storage is often hailed as the holy grail solution to the challenges of intermittency. When the renewable resource is abundant, electricity produced in excess of demand could be stored, avoiding curtailment. Then later when demand is high but the resource is low, the stored electricity could provide a reliable power supply. The most common form of storage in the United States is hydroelectric or *pumped hydro*. This refers to a system in which excess power is used to pump water uphill, which then can be allowed to flow downhill and produce power as needed. In the U.S., because of its relatively low cost, technological maturity and reliability, pumped hydro accounts for 95% of utility scale energy storage [19]. However, since pumped hydro requires access to an appropriate body of water, and even then comes with negative environmental impacts, it is not a feasible storage solution for most regions wishing to build renewable energy systems [36].

Other common methods of energy storage include flywheels, compressed air, and battery systems. The main downside of these methods is their prohibitive cost, especially storage via battery systems. A 2018 study published in *Energy & Environmental Science* found that meeting 80% of electricity demand in the U.S. with wind and solar power would require either a nationwide high-speed transmission system to balance excess supply and demand over hundreds of miles, or 12 hours of electricity storage for the entire system [39]. The *MIT Tech Review* estimates the cost of the latter option at $2.5 trillion [47].

3.4 CLC Proposed Carbon Tax

A *carbon tax* is one of the most widely proposed economic tools for incentivizing a reduction in carbon emissions. Many environmental economists see carbon taxes as crucial to the low-carbon energy transition, especially in terms of electricity production. Putting a price per ton on CO2 emitted both increases the price of fossil fuel produced electricity, as suppliers tend to absorb the annual cost of paying carbon taxes into their LCOE, and incentivizes research and development of renewable energy, as investors reap the long term benefits of avoiding carbon taxes.

The carbon tax schedule that we focus on in this paper uses $41.84 per ton (2020-USD) as its starting price of carbon, and increases that price by 5% annually. This schedule is based on a carbon tax proposed in 2017 by the Climate Leadership Council (CLC), a non-profit organization bringing together policy and corporate leaders in favor of carbon taxation with dividends[2] implemented at the national level [48]. This tax would start at $40 (2017-USD) and increase at 5% per year above inflation.

Since the 2019 publication of a statement supporting carbon dividends, endorsed by over 3500 economists in the U.S. including all four former Chairs of the Federal Reserve, the CLC proposal has seen increased bipartisan support from many large corporations, environmental NGOs, and political leaders. A sensitivity analysis using the Goulder-Hafstead Energy-Environment Economy E3 CGE model[3] of the U.S. showed that this CLC proposed carbon tax would lead to 50% fewer carbon emissions relative to 2005 levels by 2035 [25].

4 Related Work

4.1 Hybrid Energy System Design

A significant amount of research has been published on integrating various technologies to create a hybrid energy system. These models may combine renewable energy with fossil fuels, as well as heat generation, energy storage, and hydrogen fuel cell technology (including gas storage tanks) in a *hybrid* energy system. One such area focuses on finding the optimal number and size of each component in a theoretical hybrid energy system, according to some cost-based objective function. However, these models do not consider time, so the task of actually transitioning from some current system to the hybrid optimal system remains unresolved.

Fetanat and Khorasaninejad in 2015 used an ant colony optimization for continuous domains based integer programming to find the optimal sizing of wind-battery, solar-battery, and wind-solar-battery systems. [20]. They found that a majority wind mix resulted in the lowest cost because it reduced the number of batteries required. Maleki and Askarzadeh in 2014 modeled both cost and load related objectives to find the optimal sizing for a system comprising PV panels, wind turbines, and fuel cells [34]. Using an artificial bee swarm optimization (ABSO), they found that tightening the constraint on the maximum loss of power supply probability[4] had the biggest impact on sizing fuel cells since the inability to meet demand strongly correlated with how much gas was stored as energy reserves in the hydrogen tanks.

[2]Carbon dividends are based on the idea that revenues generated from a carbon tax be redistributed amongst citizens, with a focus on low-income households.

[3]The E3 CGE model is an intertemporal general equilibrium model based on the U.S. tax system and including international trade [24]. The model tracks interactions representing domestic and foreign producers of various goods and services, a typical domestic and foreign household, and a domestic and foreign government. Equilibrium is reached when supply balances demand, imports balance exports, and government spending balances tax revenue.

[4]The loss of supply probability (LPSP) is measured as a value between 0 and 1. If 0, this means that the system load will never be met, and if 1, that load will always be met. Thus, the closer the LPSP of an energy system to 1, the more reliable its power supply.

A literature review compiled by Sharafi and ELMekkawy in 2014 showed that most optimizations of hybrid energy systems focused on a single objective, in most cases system costs [40]. Sharafi and ELMekkawy propose a novel approach to particle swarm optimization (PSO) that allowed the optimization of multiple objectives: unmet load, fossil fuel emissions, and total cost. They found that the set of non-inferior solutions changed dramatically when each of these objectives were preferred.[5]

Finally, although the uncertainty of renewable energy resources is a major factor in the ability of renewable energy to meet system load, many studies on hybrid energy system design (including all three above) use averaged power curves for resources like wind and solar. Kamjoo et al. in 2016 introduce chance constrained programming (CCP) into a non-dominated sorting genetic algorithm (NDSGA-II) in order to model renewable resource uncertainties [28]. Norbu et al. one year later demonstrated the effectiveness of CPP on optimizing generator size and storage capacity in a standalone renewable energy system, which makes an appealing case for remote, unelectrified regions [35].

4.2 Transition to Renewable Energy

There have also been several studies published on the feasibility of 100% renewable energy systems by 2050. These tend to focus on a single country or region in order to optimize specifically for the renewable energy resources, existing energy infrastructure, and power demand profiles in that region.

The LUT Energy System Transition Model, which offers high temporal and spatial resolution linear optimization, is one of the more widely used models for simulating energy transitions within a specific region. LUT incorporates power, heat, and transportation sectors, as well as additional industrial sectors such as carbon removal and desalination. Child et al. in 2019 demonstrated the feasibility of achieving 100% renewable energy in Europe by 2050 using this model [7]. They found that increased storage capacity, especially for electricity prosumers[6], and adding transmission connections between countries were critical in reducing the total cost of the energy transition. Similar analyses using the LUT model were performed for Turkey and Pakistan, both with the goal of 100% renewable energy by 2050 [30] [38]. For Turkey, optimal scenarios relied on up to 73% of the power supply coming from solar, and for Pakistan this increased to 92 - 96%, demonstrating how sensitive energy transition optimization is to local geography.

Another popular family of models are based on the TIMES economic model generator, which given inputs including energy end usage, existing energy stocks, primary energy sources, and availability of future technology, aims to supply energy services at minimum global cost [33]. Krakowski et al. in 2016 applied a model from the TIMES family to analyze transitions to renewable energy between 40% and 100% in France by 2050 [31]. The simulation relied heavily on the ability to import and export power, but interestingly, found that flexibility in power generation[7] was an even more important factor in total cost than energy storage capacity.

In all of these energy transition simulations, the optimal solutions naturally are extremely dependent on region-specific inputs, including renewable resources, existing electric infrastructure, consumer load profiles, and ability to import/export electricity; therefore, they cannot be easily generalized to other regions. For this reason, most studies do not simulate a region smaller than a single country and must be repeated in full for different regions.

4.3 Energy Storage Schemes

A consistent theme, both in optimal hybrid energy system design and energy transitions that successfully meet some renewable energy goal, is the role of energy storage systems. Since the main purpose of storage is to address grid reliability isssues posed by intermittency, research on optimal storage capacity given some level of renewable penetration indicates that storage is highly dependent on the region-specific renewable energy mix, distribution of renewable resources, and load profiles.

Solomon et al. in 2017 explain how effective storage size depends on both the energy and power capacity of storage, the distribution of local renewable energy resources, and the level of renewable penetration, among several other factors [43]. They found that, although renewable energy universally exhibits intermittency characteristics, comparing region-specific studies did not yield convergence to a global relationship between renewable penetration and storage requirements. They hypothesize that this is due to differences in how the above factors are inconsistently emphasized in modeling storage requirements.

One major discrepancy between studies was the willingness to tolerate curtailment, since the flip side of building storage capacity is building excess renewable energy capacity (which would go to waste except

[5]A preferred objective for this particular paper meant that the objective was explicitly minimized, whereas other objectives were set as constraints. During the optimization, the maximum or minimum permissible levels for these constraints were varied to obtain different non-inferior solutions.

[6]A *prosumer* functions as an entity that can both produce and consume electricity, which means they sell to and buy from the grid as needed. Prosumers usually take the form of residential consumers who install solar PV panels, although this can often lead to much peakier demand curves [27].

[7]The flexibility of a power plant depends on several factors, including ramping capability, minimum load, and start up costs [46]. Conventional power plants were mostly designed to supply a steady load and to run almost 24/7 since powering up and shutting down were expensive operations. Building flexible power plants that can meet a highly variable load with minimum lag time is a major engineering challenge and significantly increases costs.

during peak demand). For instance, Denholm and Hand in 2011 determined using the REFLex model[8] that for the ERCOT grid[9], renewable penetration levels of close to 80% with less than 10% curtailment would require building storage capacity equal to the average demand for one day [12]. In contrast, Solomon et al. found that a renewable penetration level of 85% in California would require building storage capacity of only 22% of the average demand for one day, but would result in 17% energy loss or curtailment [42].

Another variable that could strongly influence storage requirements is demand itself, since higher variance in demand makes it more difficult to supply power reliably and thus would require more storage for the same level of renewable penetration. Lastly, as with renewable energy itself, diversification of storage technologies can significantly increase the overall efficiency of storage and help reduce costs [18]. However, this amplifies the challenges with modeling storage requirements since storage technologies differ significantly in cost, lifetime, efficiency, and other factors.

5 Problem Formulation

The guiding question underlying our model is: how does a 100% fossil fuel energy system transition to a lower carbon (hopefully 100%) renewable energy system over time?

We considered an energy system consisting of some number of power plants that we observe over a fixed time span. We imagined each of these power plants beginning as a generic fossil fuel plant, a theoretical average of coal and natural gas, with the possibility of conversion into a generic renewable energy plant, a theoretical average of onshore wind and solar PV, over time. Since renewable energy is intermittent, it would be impossible to meet the system load without some way of addressing the high variance in renewable energy power curves. Accordingly, we added a storage requirement that depended on the level of renewable penetration. To reflect the variety in available storage technologies, we allowed storage to be built as a combination of battery systems and pumped hydro. We also addressed the phenomenon of technological progress over time; in our system, the costs of building renewable plants and storage decrease with time. Finally, we included a carbon tax applied to the emissions produced by the fossil fuel plants.

In order to narrow the scope of this problem, we simplified the energy system in several ways. Firstly, we considered only the electric grid, ignoring other parts of the the broader energy system, such as include heat generation, transportation, etc. Secondly, we fixed the annual demand on this theoretical grid as constant and kept the grid islanded, assuming that demand is exactly equal to the total power output of plants in the system so there is no curtailment. Finally, we kept the problem entirely driven by plant- and storage-related costs, and did not include factors like subsidies for renewable energy projects, renewable portfolio standards, etc.

We reduced the wide array of costs associated with converting fossil fuel plants into renewable energy plants to the following main categories:

- Maintaining existing fossil fuel plants
- Building new renewable plants and required energy storage
- Maintaining existing renewable plants and storage
- Paying carbon taxes based on the remaining fossil fuel electricity generation

Using this setup, we wanted to test how initial conditions (such as storage costs and carbon tax prices) affected the cumulative cost of transitioning from fossil fuels to renewable energy. Formally, our goal was to compute the optimal schedule for converting fossil fuel plants into renewable energy plants given the following parameters:

- A finite time horizon
- A starting number of fossil fuel plants
- Stochastically decreasing costs of technology
- A way to calculate storage requirements by renewable penetration
- A deterministic carbon tax schedule

We then used this solver to quantify how decreasing costs of storage or different carbon tax schedules would affect the feasibility of specific renewable penetration goals. See Section 6.6 for a full list of assumptions made in our model.

[8]According to NREL, the Renewable Energy Flexbility (REFLex) model is a reduced form dispatch model that evaluates the limits of variable renewable generation as a function of system flexibility. It is an updated version of the PVFlex model [13].
[9]ERCOT refers to the Texas Interconnection, which powers the state of Texas.

6 Model Design

To address this broad question of what the energy transition might look like, we formulated the problem as a finite horizon Markov Decision Process (MDP). Over a fixed time period, the MDP models how many fossil fuel plants are converted into renewable energy plants and in which years. Section 6.5 provides the details of our MDP model; see Appendix A for a full list of model parameters with their baseline values.

6.1 Notation

The following abbreviations may be used throughout the paper as well as in the code repository:

Abbreviation	**Refers To**
FH, fh	finite horizon as used in MDP
FF, ff	fossil fuel plant
RES, res	renewable energy plant
BS(S), bss	battery storage system
PH(S), phs	pumped hydro system

6.2 Tech Stages

One of the most important considerations when trying to model the low-carbon energy transition is the progress of energy technology. Through economies of scale, improvements in equipment efficiency, and increasing support from climate policy, we expect existing energy technologies to decrease in cost over time. The widespread development and subsequent drops in capital costs of solar photovoltaic (PV) panels are an excellent example of all three phenomena [29]. In addition, disruptive energy technologies may appear in the future even though we are unable to conceive of them today.

To capture the general trend of technological advancement with time, we consider discrete "tech stages." We assume that it is only possible to advance to the immediate next tech stage, and we assign some probability `p_adv_tech` of moving on to tech stage $v+1$ given that v is the current tech stage:

$$p_v = \frac{1}{n_v} \tag{1}$$

where p_v is the probability of moving into tech stage $v+1$, a quantity dependent on v, and n_v is the expected number of years that tech stage v should last. In our model we use three tech stages; this is motivated by energy technology projections that generally include data on technology in the present day, estimates for the near future (usually one decade ahead), and estimates for the far future (between two to four decades ahead).

The parameters that we made dependent on tech stage include the capital costs of building renewable plants (`c_res_cap`) and the capital costs of building storage (`c_bss_cap`/`c_phs_cap`). Costs decrease with tech stage to model renewable energy technology becoming more cost-effective over time. We chose not to make O&M costs independent of tech stage, since they are zero for RES plants and negligible compared to capital costs for BSS/PHS. We also chose to make all costs for FF plants independent of tech stage. See Section 6.5.3 for a more detailed description of costs.

6.3 Carbon Tax

In our model, the key incentive towards building RES plants comes from a monotonically increasing carbon tax. For simplicity, we chose to model only a per ton carbon tax as opposed to a cap & trade scheme.[10] The functions we used to calculate a linearly (Equation 2) or exponentially growing (Equation 3) carbon tax at time t given coefficient c are as follows:

$$E_{\text{lin}}(t) = c_{\text{init}} + c_{\text{inc}} \cdot t \tag{2}$$

$$E_{\text{exp}}(t) = c_{\text{init}}(1 + 0.01 c_{\text{inc}})^t \tag{3}$$

where c_{init} = `c_co2_init` and c_{inc} = `c_co2_inc`. Importantly, the price of carbon is deterministic given the time t for it does not depend on tech stage. However, since the number of RES plants (and thus the number of FF plants) does depend on tech stage, and the carbon tax is applied to emissions from the FF plants, the total cost incurred by the carbon tax also depends on tech stage and is thus stochastic.

We also assume that FF plants emit CO2 only, and that we can apply the carbon tax to 100% of these emissions. In reality, however, fossil fuel electricity produces several kinds of pollution, and taxes on pollution (including CO2) can range widely depending on the material and its form when released. Even within carbon taxes, there is usually a distinction between emissions from transport fuels versus other sources [23]. However, since we model all types of FF plants as a single generic plant, we make the simplifying assumption that we can apply the carbon tax universally to all emissions.

[10] A carbon tax theoretically allows an unlimited amount of carbon to be emitted as long as it is taxed. On the other hand, a *cap & trade scheme* creates a market for "carbon credits" under which the total amount of carbon emissions must remain fixed. Organizations can then buy and sell the ability to emit more or less carbon than they were allotted in credits.

6.4 Energy Storage

Since we model RES plants in our system as an average of solar and onshore wind power plants, whose power curves show extremely high variance, we needed to account for increasing stress on the grid as more FF plants are converted into RES plants. Henrik Zsiboracs et al. modeled energy system scenarios for Europe in 2040 with varying levels of renewable penetration, and aggregated estimates for amount of energy required for a given percentage of renewable penetration in different European countries [51]. We used data points for the U.K. from this survey to fit an exponentially increasing function that calculates the required energy storage as a percentage of total system load given the percentage of renewable penetration:

$$S(r) = c_0 \cdot \exp\left(c_1 \cdot 100\frac{r}{n}\right) + c_2 \tag{4}$$

where $\frac{r}{n}$ is the fraction of RES plants out of total plants (including plants to be built), and $S(r)$ is the total storage that must be built to support a r RES plants, expressed as a percentage of total system load. The reason we fit an exponential function for the storage requirement is because introducing more intermittent power sources has a multiplicative rather than an additive effect on power supply reliability. This is supported empirically by the majority of studies surveyed by Zsiboracs et al., which showed a trend of exponentially increasing storage requirements for linearly increasing renewable penetration.

6.5 Markov Decision Process (MDP)

We developed several iterations of a finite horizon MDP model; the version used to obtain all results presented in this paper is MDP V3, although all previous versions can be found in the code repository. The models were built using the MDP Toolbox library in Python.

In the finite horizon case, the MDP is defined by a controlled, dynamic system and a cost or reward structure; this becomes the objective function in the optimization [41]. Starting from some initial state, the "agent" in the model has a choice of actions it can perform in that state. Performing a specific action in a specific state results in some reward. Given an action and the current state, there is some probability of transitioning to any other state.

The goal of the optimization is to determine the optimal policy: a mapping of states to actions that maximizes the expected reward of the agent. In our case, the optimal policy defines a schedule for converting FF plants into RES plants that incurs the lowest total cost.

6.5.1 State Space

In our model, we define state as a combination of any valid values for the following variables:

Variable	Name	Description	Range
t	Time	Timestep in MDP	[0, `n_years`)
v	Tech Stage	Stage of technological development	[0, 2]
r	RES Plants	Number of renewable plants existing at beginning of year	[0, `n_plants`]

Given an optimal policy, both t and r can be determined with certainty, t because it is the current time step and r because actions build RES plants deterministically. The only variable that lends uncertainty to the state is v, since there is a nonzero probability of being in any tech stage at time t.

6.5.2 Actions

The acceptable action that a policy may take in some state (t, v, r) is to convert a fossil fuel (FF) plants into renewable energy (RES) plants where $a \in \{i \in \mathbb{Z} : 0 \leq i \leq \texttt{n\_plants}\}$. Converting a FF plant into a RES plant involves building a new RES plant as well as a required amount of energy storage. There is no action that can reconvert a RES plant into a FF plant, or shut down an existing RES plant.

Since the total number of plants in the system is constant at all times, we model the physical impossibility of building more than $\texttt{n\_plants} - r$ in some state (t, v, r) by placing an infinite cost on any renewable plants built past `n_plants`. In other words, $C(t, v, r, a)$ in Equation 5 below would equal negative infinity for any $a > \texttt{n\_plants} - r$.

6.5.3 Reward Function

We use a negative cost function as the reward function in the MDP. Given a state (t, v, r) and an action a taken in that state, the cost can be calculated as follows:

$$C(t, v, r, a) = \sum_f (C_{\text{co2}}(t) + C_{\text{ff}})) + \sum_r C_{\text{old-res}}(v) + \sum_a C_{\text{new-res}}(v) + C_{\text{old-stor}}(r, a) + C_{\text{new-stor}}(v, r, a) \tag{5}$$

where $f = \texttt{n\_plants} - r$ is the total number of FF plants remaining at the end of time step t, C_{FF} is the cost paid for FF plants, $C_{\text{old-RES}}$ is the cost paid for existing RES plants, $C_{\text{new-RES}}$ the cost paid for RES plants built in the time step t, $C_{\text{old-stor}}$ is the cost paid for existing storage capacity, and $C_{\text{new-stor}}$ is the cost paid for building the required storage to accompany a new RES plants. Each of these cost components is calculated as follows:

$$C_{\text{co2}}(t) = E(t) \cdot e_{\text{ff}} \tag{6}$$

where e_{ff} is the per kWh emissions of a FF plant, `ff_emit`, and $E(t)$ is the carbon tax calculated at time t according to either Equation 2 or Equation 3.

$$C_{\text{ff}} = \frac{c_{\text{ff-cap}}}{l_{\text{ff}}} + c_{\text{ff-fix}} \cdot z_{\text{ff}} + c_{\text{ff-var}} \cdot (z_{\text{ff}} \cdot y_{\text{ff}} \cdot 8760) \tag{7}$$

where $c_{\text{ff-cap}}$ = `c_ff_cap` is the capital cost of building a FF plant, l_{ff} = `ff_lifetime` is the plant lifetime, $c_{\text{ff-fix}}$ = `c_ff_fix` is the fixed O&M cost, $c_{\text{ff-var}}$ = `c_ff_var` is the variable O&M cost, z_{ff} = `ff_size` is the plant size or nameplate capacity, y_{ff} = `c_ff_capacity` is the capacity factor, and 8760 is the number of hours in 1 year.[11]

$$C_{\text{old-RES}}(v) = \frac{c_{\text{res-cap},v}}{l_{\text{res}}} \tag{8}$$

$$C_{\text{new-RES}}(v) = c_{\text{res-cap},v} \tag{9}$$

where $c_{\text{res-cap},v}$ = `c_res_cap` is the capital cost of building a renewable energy plant in tech stage v, and l_{res} = `res_lifetime` is the plant lifetime. We assume that there are no O&M costs for RES plants.[12]

$$C_{\text{old-stor}}(r, a) = S(r+a) \cdot \left(\frac{bc_{\text{bss-fix}} + hc_{\text{phs-cap}}}{8760} + bc_{\text{bss-var}} + hc_{\text{phs-var}} \right) \tag{10}$$

$$C_{\text{new-stor}}(v, r, a) = (S(r+a) - S(r)) \cdot (bc_{\text{bss-cap},v} + hc_{\text{phs-cap},v}) \tag{11}$$

where $c_{\text{bss-cap},v}$ = `c_bss_cap` and $c_{\text{phs-cap},v}$ = `c_phs_cap` are the capital costs of building battery system and pumped hydro energy storage in tech stage v, $c_{\text{bss-fix}}$ = `c_bss_fix` and $c_{\text{phs-fix}}$ = `c_phs_fix` are the fixed O&M costs of battery systems and pumped hydro, and $c_{\text{bss-var}}$ = `c_bss_var` and $c_{\text{phs-var}}$ = `c_phs_var` are the fixed O&M costs of battery systems and pumped hydro, respectively. We assume that fixed and variable O&M costs do not depend on tech stage, mainly because they are negligible compared to the capital costs.

6.5.4 Transition Probabilities

Let E_i refer to the carbon tax in state i and p_{v_i} refer to the probability of moving into tech stage $v_i + 1$. Then given an action a, the transition probability of some state (t_i, v_i, r_i) to any state (t_j, v_j, r_j) in the MDP may be described by the following rules:

- $p_{ij} = 0$ if and only if:
  - $t_j < t_i$ (cannot go backwards in time)
  - $v_j < v_i$ (cannot revert to a previous tech stage)
  - $r_j \neq r_i + a$ (must build number of RES plants exactly equal to action)
- For all other state pairs:
  - $p_{ij} = p_{v_i}$ if $v_j = v_i + 1$
  - $p_{ij} = 1 - p_{v_i}$ if $v_j = v_i$
  - $p_{ij} = 0$ otherwise

6.6 MDP V3 Assumptions

Summarizing from previous sections, the following assumptions were held across all versions of the model:

- The energy system is closed and has a fixed demand over the given time period.
- A RES plant is a 50-50 average of wind and solar costs, and a FF plant is a 60-30[13] average of coal and natural gas costs.
- The size of a RES plant can be exactly determined from the size of a FF plant as well as capacity factors of RES and FF plants.
- It is only possible to advance by 1 tech stage at a time.
- The carbon tax is entirely deterministic given time and number of RES plants out of total plants.

[11] Annual fixed O&M costs of a power plant depend on its nameplate capacity, whereas variable O&M costs depend on the actual power output, which involves the capacity factor and how many hours in the year the plant runs (here we assume it runs for all 8760).

[12] Generally, wind and solar plants do have fixed O&M costs but they are negligible compared to the capital costs. The variable O&M costs are considered to be zero because there are no fuel costs; energy is "free" once the plant is constructed.

[13] This weighted average is based on the proportion of electricity generated globally by coal versus natural gas in 2016 [10].

- O&M costs of FF plants, BSS storage, and PHS storage do not depend on time or tech stage.
- RES plants have zero O&M costs.
- RES construction incurs zero CO2 emissions.

In addition, the following assumptions were incorporated specifically into MDP V3:

- The probability of advancing tech stage depends only on the current tech stage.
- Periodic renewal of both RES and FF plants can be modeled as an additional O&M cost: the capital cost scaled down by the lifetime of the plant.
- For a RES plant, the renewal cost uses the capital cost of the current tech stage, not the tech stage in which the plant was built.
- O&M costs of storage account for degradation, so additional replacement costs can be treated as zero.
- All BSS related costs are based on a four hour lithium-ion battery system. BSS capital cost combines energy and power dependent costs.

7 Results

7.1 Stochasticity

Figure 1 visualizes stochasticity in MDP model V3. Both the tech stage (shown in blue) and the optimal policy (actions shown in dark green) were averaged over 200 runs of the model. We chose `p_adv_tech`, the probability of advancing to the immediate next tech stage, such that, as closely as possible, tech stage 1 would be reached on average by 2030, and tech stage 2 by 2050.

We set this particular tech stage curve for several reasons. Firstly, it aligns with the cost projections for renewable energy and storage that we used as data sources, since the projects gave point estimates for reduced costs in 2030 and in 2050. Secondly, this allows the model on average an additional 20 years in tech stage 2, meaning that each simulation can spend more years when it is cheapest to build RES plants and the accompanying storage (while the carbon tax continues to increase), so this might incentivize building more RES plants. Finally, this somewhat captures technological development in the real world, in that there may be a minimum capital cost that RES plants and storage systems converge to simply because of theoretical efficiency limits and the cost of raw materials. Based on our data sources, it also seems difficult if not impossible to provide accurate cost estimates beyond roughly 30 years.

Figure 1: Average optimal policy and average tech stage over time. MDP V3 was run for 200 iterations with baseline parameters and the starting price of carbon as 99.11 USD/ton.[14] Averaged over all model runs, the current tech stage is shown in blue, and new RES plants built in the current time step and total existing RES plants are shown in dark and light green, respectively.

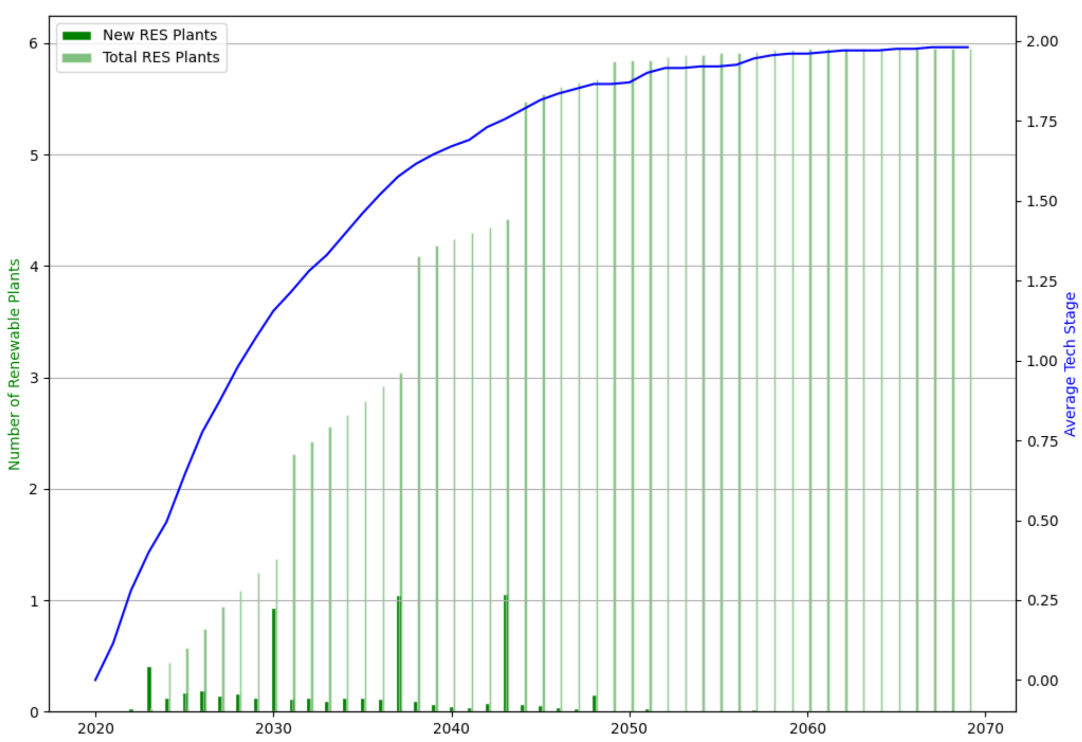

[14]We use 99.11 USD/ton here (the current carbon tax in Switzerland) simply for the purposes of visualizing uncertainty in the optimal policy more clearly. The heights of the dark green bars are much less noticeable with a lower starting carbon price since the optimal policy builds fewer RES plants.

The purpose of Figure 1 is to give a graphic overview of how the optimal policy (in this case for a carbon tax of 99.11 USD/ton increasing by 5% annually) progresses alongside a stochastic tech stage. The blue line represents the current tech stage averaged over all 200 runs of MDP V3, reaching approximately 1.125 in 2030 and 1.875 in 2050. The dark green bars represent the action taken (number of RES plants built in the current time step) and the light green bars represent the cumulative number of existing RES plants (at the beginning of the time step), both averaged over all model runs.

One interesting aspect of the average optimal policy in Figure 1 is the decrease in height of the dark green bars with time, except for single years in which the dark green bar is much higher compared to neighboring years. This can be attributed to how variance in tech stage decreases with time as more the model runs reach tech stage 2. Since a higher tech stage brings down the costs of building RES plants, the earlier a particular run advances in tech stage, the earlier it will build more RES plants. In the early years (steps) of the simulation, there is a higher variance in tech stage among the model runs; in other words, the model runs occupy the state space across all tech stages. However, since tech stage increases monotonically, as time progresses, the tech stage in any given run is more likely to be at stage 2. Thus, as time progresses, more of the model runs reach a states consisting only of tech stage 2. Since a greater percentage of the runs are operating from similar states as time goes on, they make similar policy decisions. In this graph, with the starting carbon price at 99.11 USD/ton, the optimal policy seems to converge to building the last two plants in 2037 and 2043.

7.2 Cost Breakdowns

The purpose of this section is to give an overview of how different components of the system contribute to the total cost under the optimal policy. Figure 2 compares the total annual and cumulative cost incurred by following the optimal policy under two different conditions: a carbon tax schedule with a starting price of 41.84 USD/ton that increases by 5% annually and a carbon tax schedule with a starting price of 123.18 USD/ton that also increases by 5% annually. The former represents the CLC proposed carbon tax for the United States, which we focus on in this paper, and the latter represents the highest carbon tax in the world as of February 2020, currently implemented in Sweden.[15]

Most notably, the cumulative costs under the Sweden-based tax schedule are significantly higher, about 850 billion USD compared to the CLC proposed tax schedule, about 600 billion USD. However, the optimal policy under the Sweden-based tax schedule also achieves more than double the renewable penetration as the optimal policy under the CLC proposed tax schedule (70% compared to 30%), and the only difference is the starting price of carbon. Thus, although the high price leads to an overall more costly optimal policy, it is a 133% increase in renewable penetration for less than a 42% increase in total cumulative cost.

[15]The Swedish carbon tax is for now fixed at the price of 123.18 USD/ton, and does not include a 5% annual increase. This is only added for the purposes of comparing with the CLC proposed carbon tax of 41.84 USD/ton.

Figure 2: Comparison of total annual and cumulative costs, as well as RES penetration, between optimal policies with a 41.84 USD/ton (CLC proposed carbon tax for the United States) and a 123.18 USD/ton (Sweden-based carbon tax) tax schedule that increases at 5% annually. All other parameters follow the baseline parameters and results were averaged over 200 iterations of MDP V3.

(a) Total Annual and Cumulative Costs

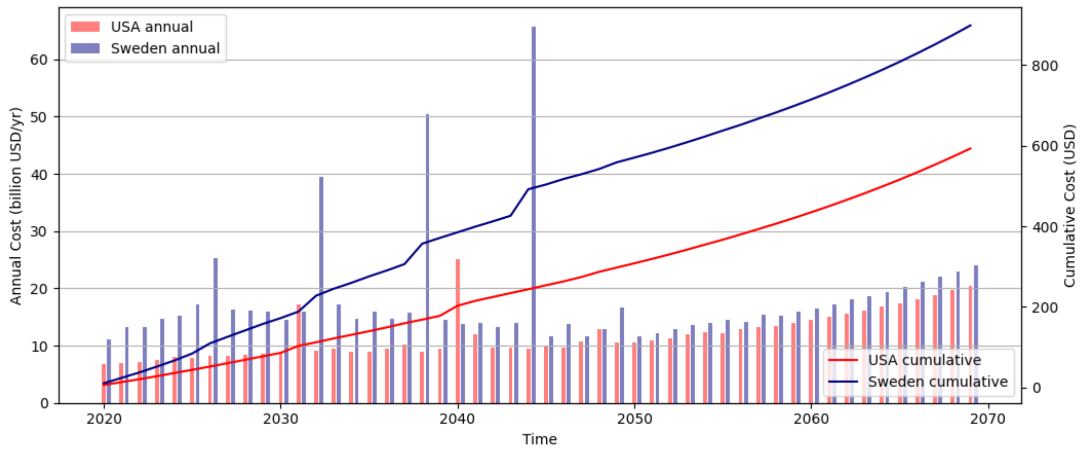

(b) RES Penetration

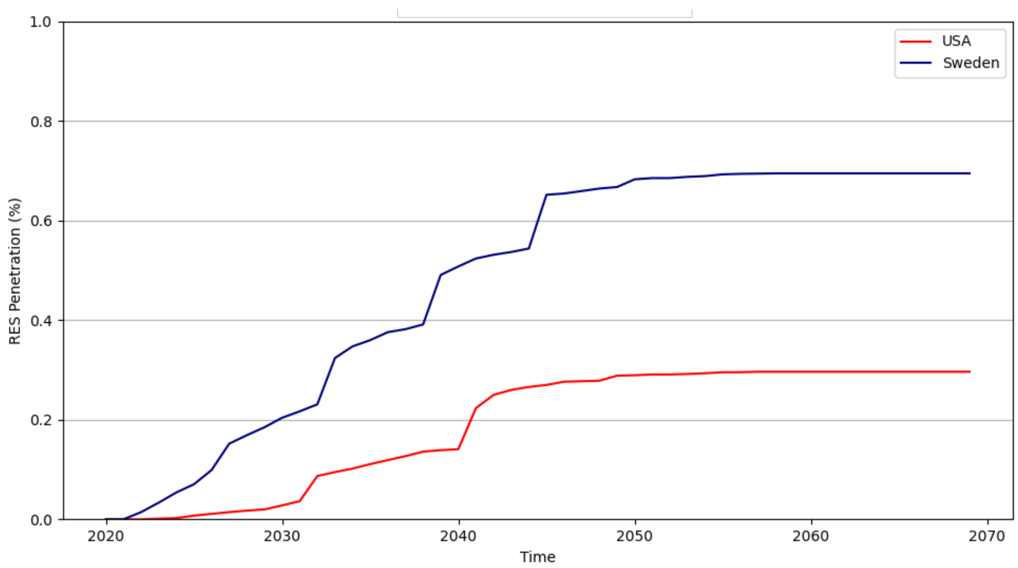

The gains in renewable penetration for comparatively less substantial increases in cost imply that the continued cost of maintaining more FF plants and paying tax on emissions is a substantial burden. In Figure 2, total annual costs for both the CLC proposed and Sweden-based carbon tax schedules reaches a low point at around 2050. Yet for the CLC proposed carbon tax, from this point the exponential increase in carbon tax begins to contribute more to the total cumulative cost than the capital costs of the three RES plants.

Figure 3 displays a more granular cost breakdown, split by tech stage, specifically for the carbon tax schedule of 41.84 USD/ton increasing by 5% annually. The most obvious finding is how small the capital costs of building RES plants are compared to the capital costs of building the accompanying storage. These storage costs are almost entirely responsible for the large spikes in the years when the optimal policy builds a RES plant. Following construction, the annual costs (operation and maintenance) of both RES plants and storage are negligible; the annual costs of fossil fuels are much higher. In any year when no RES plants are built, the carbon tax forms the majority of annual cost, and this majority tends to increase with time.

Figure 3: Total annual cost broken down an absolute terms and as percentages for the optimal policy under a carbon tax schedule starting at 41.84 USD/ton and increasing 5% annually (CLC proposed carbon tax). Plots are separated by tech stage, which is held fixed over time.

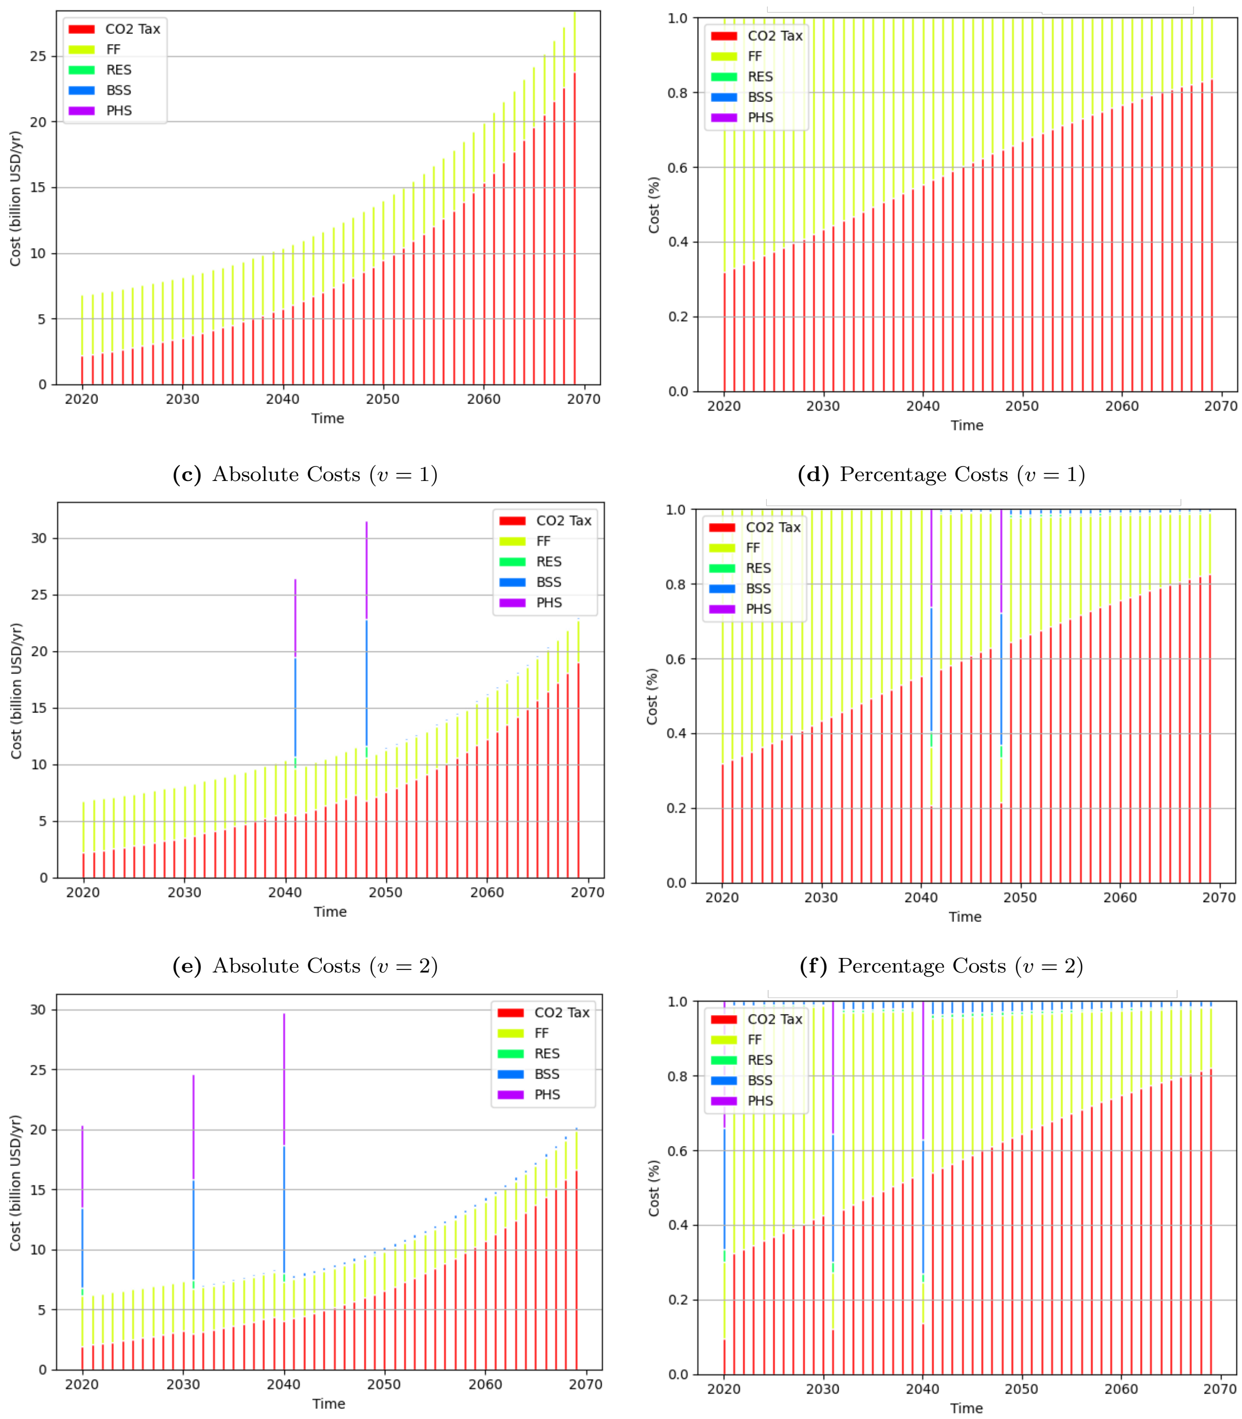

(a) Absolute Costs ($v = 0$)

(b) Percentage Costs ($v = 0$)

(c) Absolute Costs ($v = 1$)

(d) Percentage Costs ($v = 1$)

(e) Absolute Costs ($v = 2$)

(f) Percentage Costs ($v = 2$)

7.3 Storage Cost Reductions

This experiment attempts to answer the following: by how much do storage costs need to be reduced in order to achieve a certain level of renewable penetration? To answer this question, we applied a range of percent reductions to the costs of storage. Since the capital costs form the vast majority of total storage costs, we only applied reductions to the parameters `c_bss_cap` and `c_phs_cap`. We ran our model with these successive reductions in storage costs for two different starting carbon prices: 41.84 USD/ton (CLC proposed carbon tax) and 123.18 USD/ton (Sweden-based carbon tax). For all other parameters, we used the baseline parameters. Note that the baseline parameters specify a 50-50 storage mix of BSS and PHS, so the results below should not be generalized to different storage mixes.

As displayed in Figure 4 (a), linear reductions in storage costs do not have a linear effect on total cumulative cost under the CLC proposed carbon tax. Rather, an interesting tiered behavior emerges amongst reductions of 10% (yellow curve), 20% (green curve), and 30% (blue curve). It appears that a cost threshold for building four RES plants is somewhere between 0 - 10% reduced costs, but the threshold for building five is between 30 - 40% reduced costs. This behavior is consistent with the exponential nature of the storage requirement function (see Equation 4), which makes building renewable plant $r + 1$ exponentially more expensive than building plant r. A linear change in storage costs, for instance from a 10% to a 20% reduction, may simply not be enough to make up for the more-than-linearly-higher cost of building the fifth RES plant; this explains why only the reduction of 40% (purple curve) converges to 50% renewable penetration.

Figure 4: Percent reductions in the capital cost of storage (`c_bss_cap` and `c_phs_cap`) with a carbon tax schedule starting at 41.84 USD/ton with 5% annual increase (CLC proposed carbon tax). All other parameters follow the baseline parameters and results were averaged over 200 iterations of MDP V3.

(a) Effect on total cumulative cost

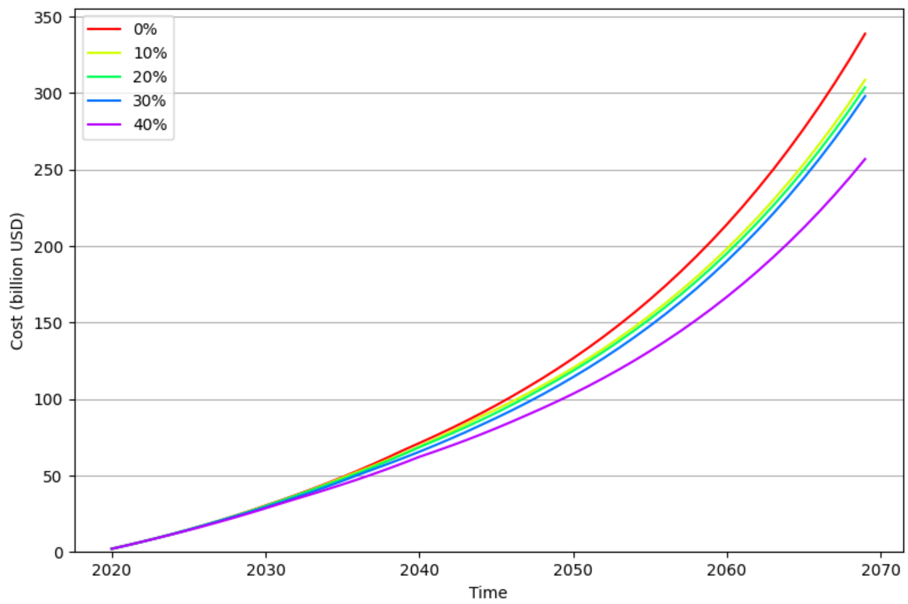

(b) Effect on renewable penetration

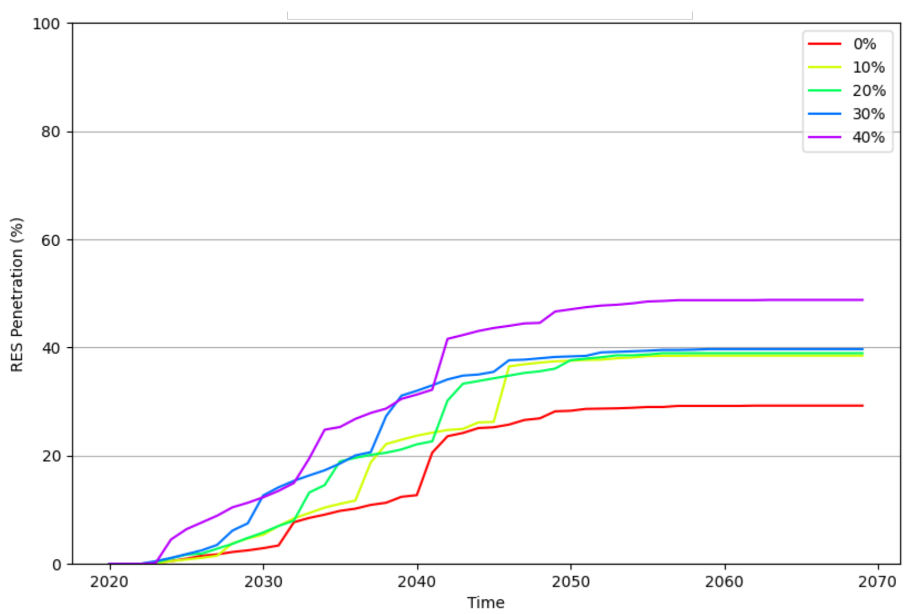

In our model, it is impossible to build a fraction of a RES plant, which would achieve renewable penetration between the tens of percentages. For this reason, we only see a change in end-result renewable penetration when storage costs become low enough to build all of the storage required for a new RES plant, resulting in the the tiered behavior of the red, yellow-green-blue, and purple curves. However, the years before 2050 in Figure 4 (b) reveal some nuance behind this result. Here we see that the optimal policies for a 20% storage cost reduction (green curve) and a 30% reduction (blue curve) converge to the same level of renewable penetration: 40%. This means that both policies build four RES plants out of ten total plants by about 2060; the key difference is that, on average, the optimal policy for the green curve builds the plants later. Graphically we can see this as the green curve tends to lag behind the blue curve, though they eventually converge to the same value. So although reducing storage costs by 20% and 30% appears to achieve the same end-result behavior in terms of number of RES plants built, a 30% reduction induces building the plants earlier. One major benefit of schedules that build RES plants earlier is that they avoid more carbon emissions.

Given the differences in optimal policy schedules, it was somewhat surprising to observe the same grouping of yellow-green-blue amongst the total cumulative cost curves as the renewable penetration curves. We identify a couple factors that may contribute to this behavior. The first is inherent optimization by the MDP: the optimal policy will not build a RES plant (and pay for storage) until the carbon tax grows large enough to justify doing so. The second is the effect of a reduction in storage costs applied statically over the model time horizon, which should induce the optimal policy to build if not more RES plants then at least the same number but earlier. The third factor is more complicated, and has to do with tech stage. As time progresses, the likelihood of reaching later tech stages, and thus of benefiting from tech stage related reductions in storage costs, increases. So an optimal policy with no storage cost reduction that builds RES plants later may pay less because storage costs less in later tech stages, thus bringing its total cost closer to that of an policy with a storage cost reduction that builds RES plants earlier. The similarity between the yellow, green, and blue curves for total cumulative costs supports this concept.

Increasing the starting price of carbon to 123.18 USD/ton, which is the current price of carbon in Sweden and the highest nationally implemented tax in the world as of 2020, we found a similar tiered behavior as for the CLC proposed carbon tax but with different thresholds. Unlike in Figure 4 (a), the total cumulative cost curves in Figure 5 (a) tend to flatten as the storage cost reduction grows. This may be caused by two factors. Firstly, the storage requirements increase exponentially with each new RES plant built, so for policies that build more RES plants, the impact of a storage cost reduction is much greater. Secondly, the carbon tax schedule used in this particular experiment increases exponentially, so the difference in carbon tax cost between optimal policies that build 7 versus 9 RES plants is greater, for instance, than between policies that build 3 versus 5 plants, even though the delta in carbon emissions is the same.

Figure 5: Percent reductions in storage capital cost with a carbon tax schedule starting at 123.18 USD/ton with 5% annual increase.

(a) Effect on total cumulative cost

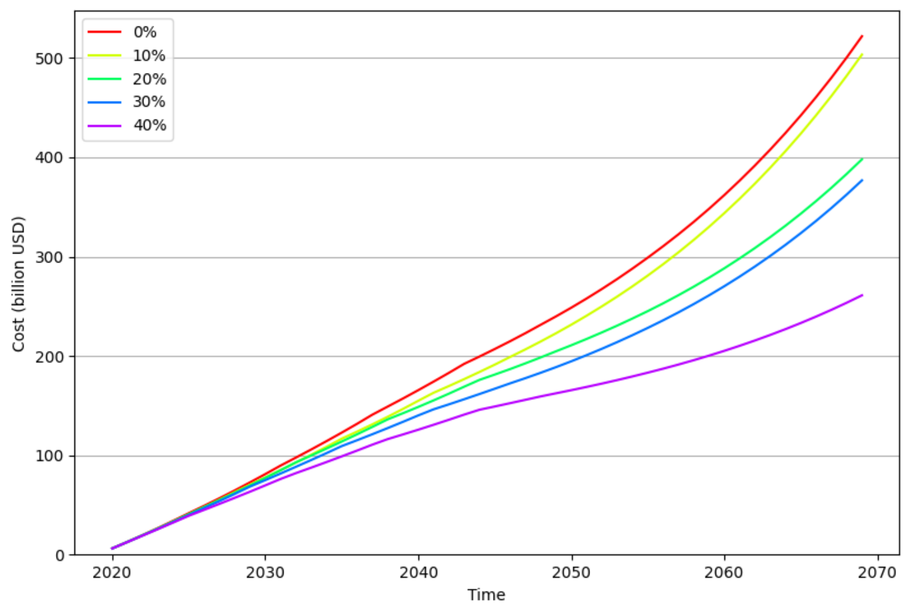

(b) Effect on renewable penetration

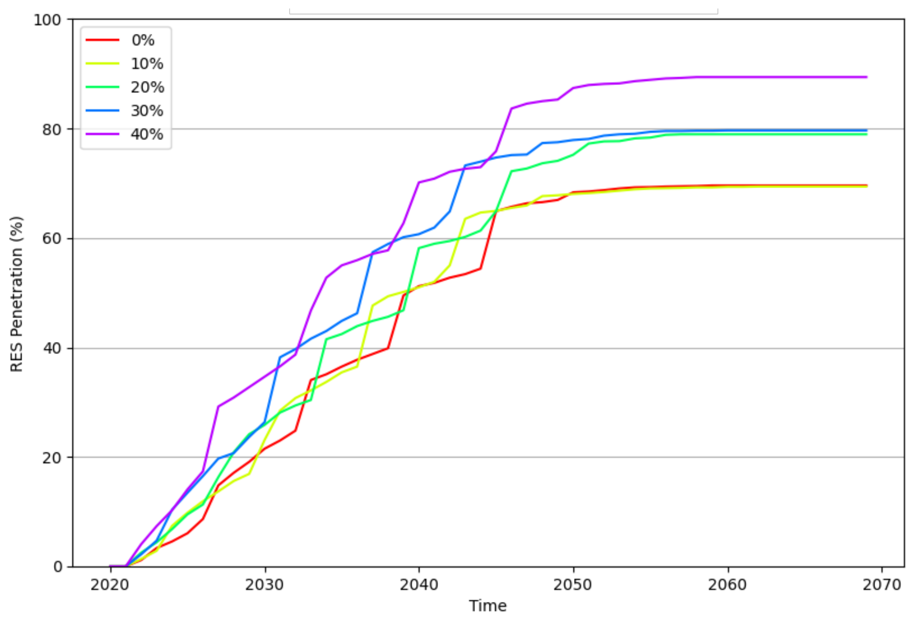

Interestingly, it appears that under the Swedish carbon tax, the 0% (red curve) and 10% (yellow curve) reductions are in one tier and the 20% (green curve) and 30% (blue curve) reductions are in another tier. This implies that since the optimal policy under a higher carbon tax tends to build more RES plants, and storage is increasingly expensive, a 10% reduction does not provide enough incentive to reach the threshold of building an additional renewable power plant. However, like in Figure 4 (b), optimal policies that converge to the same level of renewable penetration still differ in when they build RES plants; the greater the storage cost reduction, the earlier the optimal policy builds RES plants. So even though 0% and 10% reductions result in the same total number of RES plants built, building earlier leads to with fewer carbon emissions and thus less cumulative carbon tax. Although the red and yellow curves appear close to each other graphically, the higher starting carbon price of 123.18 USD/ton actually makes their delta tens of billions in total cumulative cost.

One potential flaw with this experiment is the universal application of the percent reductions in storage costs across tech stage. In other words, a 30% reduction in storage costs in tech stage 0 is perhaps reasonable, since storage costs in tech stage 1 are already set significantly lower. However, a 30% reduction in storage costs in tech stage 2 may not be realistic, since this reduction combined with the built in reduction by tech stage may result in storage costs that surpass theoretical limits on efficiency, even with projections for substantial improvement in storage technology.

7.4 Carbon Tax Schedules

The following two experiments quantify the effects of different carbon tax schedules on both the optimal policy and resulting carbon emissions. The starting prices of carbon are based on national carbon taxes implemented in various countries as of 2020 [6]. The carbon taxes we used for this experiment include:

Table 1: A subset of the national carbon taxes implemented in 2020. We selected countries in order to provide a somewhat even spread of starting carbon prices.

Country	2020-USD/ton
South Africa	8.52
United Kingdom	23.63
United States	41.84
Finland	68.52
Switzerland	99.11
Sweden	123.18

In general, we expected a higher starting price of carbon to lead to the optimal policy building more renewable power plants earlier. However, we saw that different carbon tax schedules can result in significant differences in the level of carbon emissions in the short term (first decade and a half) versus long term (end of model time horizon, five decades).

7.4.1 Effect of Starting Price on Annual Carbon Emissions

For this next experiment, our goal was to quantify how much the starting price of a carbon tax schedule influenced the emissions reductions achieved by the optimal policy. We ran our model using the carbon taxes from Table 1 as starting carbon prices, and applying two specific growth rates. Our hypothesis was that the difference in starting price would have much larger effects for an exponential than a linear growth rate. For all other parameters, we used the baseline parameters. Note that the exponential and linear growth rates used in this experiment may not reflect the actual carbon tax schedules implemented by the countries in Table 1, and are intended simply to standardize across countries for the purposes of comparing the MDP optimal policy.

Figure 6 reveals the tiered behavior of the optimal policy for a linear carbon tax schedule that increases by 10 USD annually. The carbon taxes based on all countries except for South Africa and Sweden seem to converge to a middle tier of annual carbon emissions by about 2055. Within this middle tier, however, the cumulative carbon emissions would still be fewer for Sweden than Finland, for instance, since the optimal policy for Sweden appears to build RES plants sooner and thus avoids more emissions than the optimal policy for Finland. This result echoes the differences we found in the exact schedule of building RES plants, despite the general tiered behavior that emerged when we applied linear reductions in storage costs (see Section 7.3). For carbon emissions in particular, given how soon the policies converge (beyond 2055 no policy seems to build any more RES plants), the tiers may actually be a better indicator of how effective different starting prices of carbon are on avoiding carbon emissions in the long term.

Figure 6: Annual carbon emissions under the optimal policy when the starting price of carbon is varied. All carbon tax schedules increase the price of carbon by 10 USD per year.

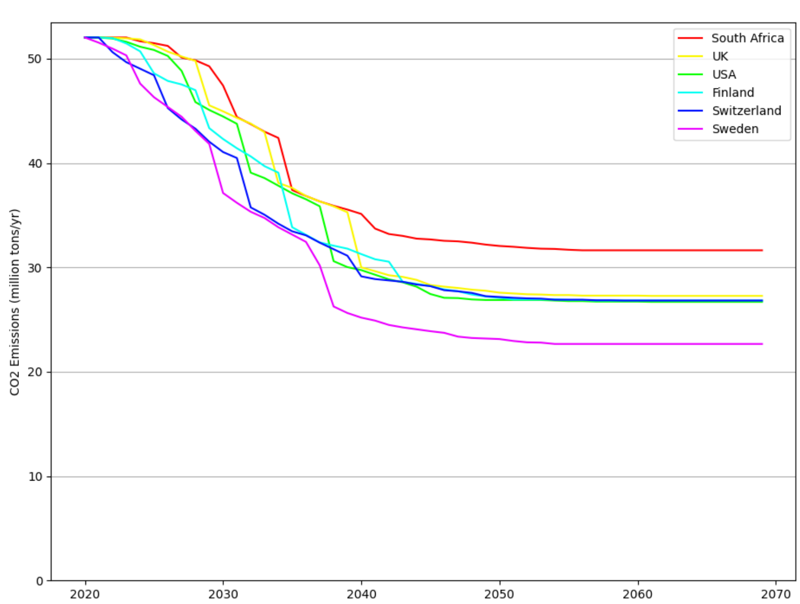

A comparison of Figures 6 and 7 highlights the substantial differences in exponential and linear tax schedules in terms of long term effects on annual carbon emissions. An exponential carbon tax schedule that increases by 5% annually preserves the differences in starting price (not necessarily proportionally) as the simulation converges. In other words, the optimal policies under the exponential tax for different starting prices of carbon do not exhibit tiered behavior. This is likely because the differences in carbon tax are large enough that each increase in the starting price of carbon pushes the corresponding optimal policy past the cost threshold: when building at least one more RES plant costs less than continuing to pay tax on the current amount of carbon emissions.

Figure 7: Annual carbon emissions under the optimal policy when the starting price of carbon is varied. All carbon tax schedules increase the price of carbon by 5% per year.

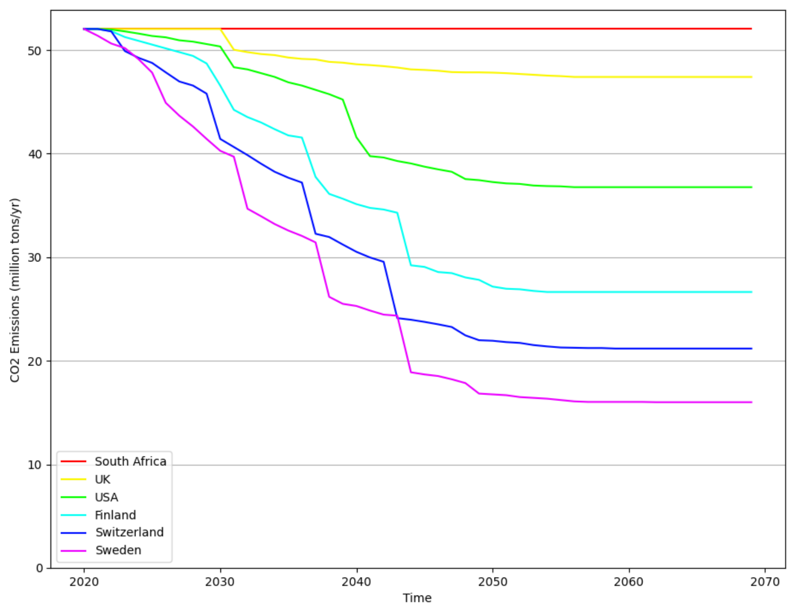

One additional observation is how much more effective the linear tax schedule that increases by 10 USD annually is for the starting carbon prices for South Africa (8.52 USD/ton), the U.K. (23.63 USD/ton) and even the CLC proposed carbon tax for the U.S. (41.84 USD/ton, labeled USA in Figures 6 and 7). Mathematically this seems obvious, since there is a price threshold below which a 10 USD increase annually would result in a higher price of carbon by 2070 than a 5% increase annually. However, it is interesting to compare carbon tax schedules across prices and growth rates, for instance the U.K. starting price with a 10 USD annual increase versus the Finland starting price with a 5% annual increase, both of which converge to about 27 million tons/yr of carbon emissions by 2070.

Figure 8 shows a more detailed comparison of these two tax schedules. Again, the similarity of U.K. schedule (orange) and Finland schedule (dark blue) for cumulative carbon emissions is deceiving; the small delta actually represents a difference of about 50 million tons. In terms of cost, shown in Figure 8 (b), the total cumulative cost for Finland seems to rise more sharply after 2060 than for the U.K. However, the first two spikes in the orange bars (total annual cost for the U.K.), which correspond to building a RES plant and accompanying storage, place a higher cost burden on those single years than the corresponding spikes in the blue bars (total annual cost for Finland), which are spread out over two years each. Except for the last decade, during which exponential growth in the Finland-based carbon tax surpasses linear growth in the U.K. tax (which started at one third the price), the Finland tax schedule may seem favorable because it leads to a slightly lower cumulative cost. However, when considered alongside carbon emissions, the trade-off between optimizing for fewer carbon emissions versus lower total cost becomes clear.

Figure 8: Comparison of total annual and cumulative costs, as well as cumulative CO2 emissions, between optimal policies with tax schedules of 23.63 USD/ton (U.K.-based carbon tax) increasing by 10 USD annually and 68.52 USD/ton carbon (Finland-based carbon tax) increasing by 5% annually. All other parameters follow the baseline parameters and results were averaged over 200 iterations of MDP V3.

(a) Cumulative CO2 Emissions vs. Price of CO2

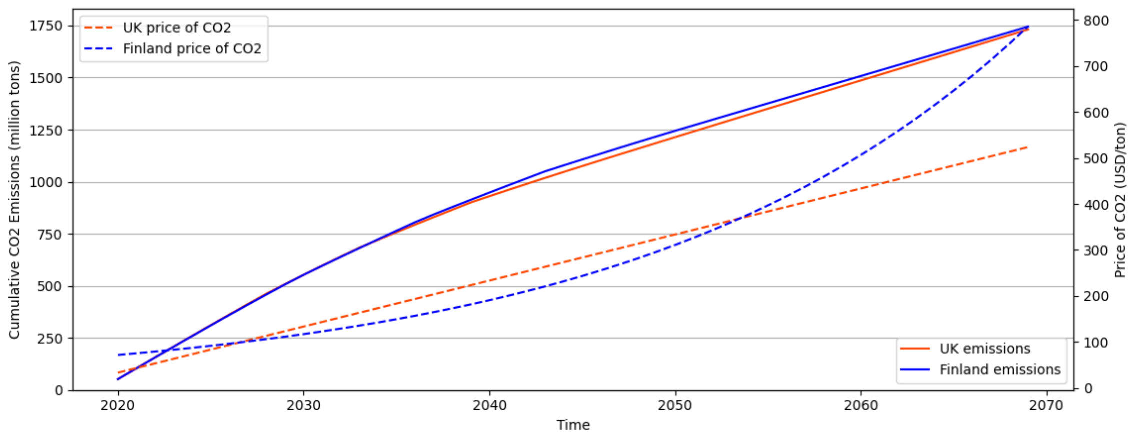

(b) Total Annual and Cumulative Costs

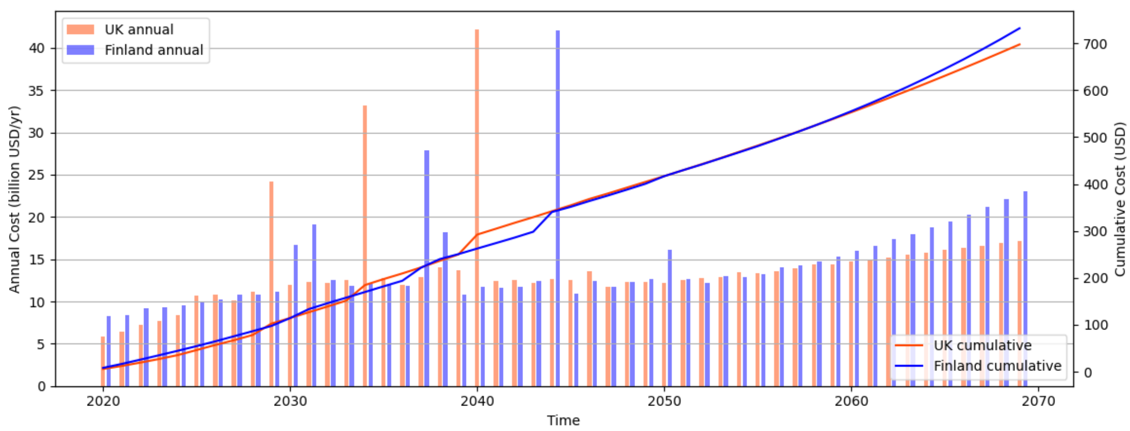

7.4.2 Sensitivity of Renewable Penetration to Starting Price

The final experiment we discuss in this paper attempts to answer the following: what starting price of carbon is required to reach some target renewable penetration by 2070? To answer this question, we performed a sensitivity analysis on the carbon tax schedule, in particular on the starting price of carbon when the tax growth rate is fixed at 5%. This time we ran the model with fixed increments in the starting price of carbon so that we could estimate the threshold beyond which that starting price increasing by 5% annually built the target number of RES plants. For all other parameters, we used the baseline parameters.

Our results indicate that, using the CLC proposed carbon tax of 41.85 USD/ton as a baseline, this value would need to almost double in order to achieve 50% renewable penetration by 2070. Figure 9 compares the optimal policies for different starting prices of carbon. Put in terms of existing carbon taxes, this means that out of the countries in Table 1, only the carbon taxes implemented in Finland,[16] Switzerland and Sweden would be sufficient to reach a target of 50% penetration by 2070, and only in Sweden would be sufficient to reach 70% penetration by 2070. The values of 50% and 70% are motivated in part by the claim that the CLC proposed carbon tax could reduce carbon emissions by 50% in 2035. Despite giving the model an extremely conservative timescale of 50 years to achieve 50% target renewable penetration, still only the three highest carbon taxes in the world were able to meet this target.

[16]The exact value of the Finnish carbon tax is 68.52 USD/ton, but it converges to the same renewable penetration level as 70 USD/ton under a tax schedule that increases by 5%/yr.

Figure 9: Optimal policy schedules for various starting prices of carbon that increase by 5% annually. The optimal policy is shown separately for each tech stage; because the tech stage is held fixed, the optimal policy is deterministic.

(a) Starting price: 40 USD/ton

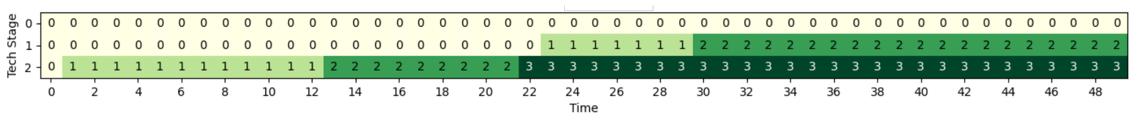

(b) Starting price: 60 USD/ton

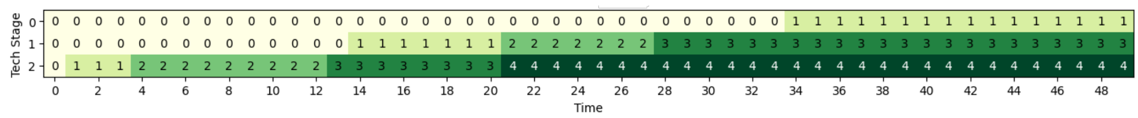

(c) Starting price: 70 USD/ton

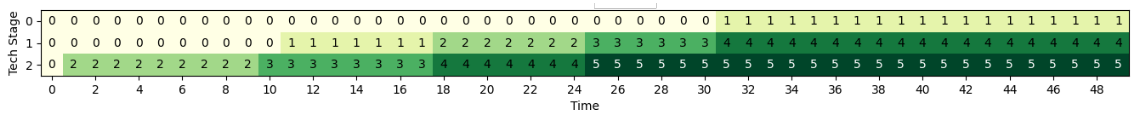

(d) Starting price: 80 USD/ton

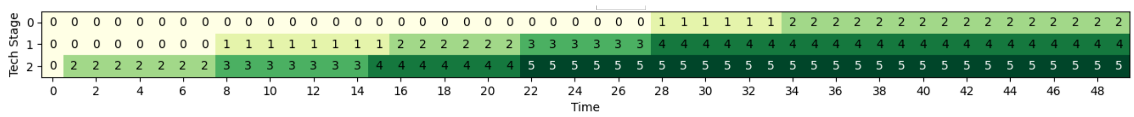

(e) Starting price: 100 USD/ton

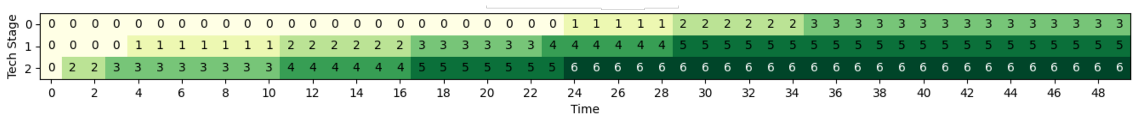

(f) Starting price: 120 USD/ton

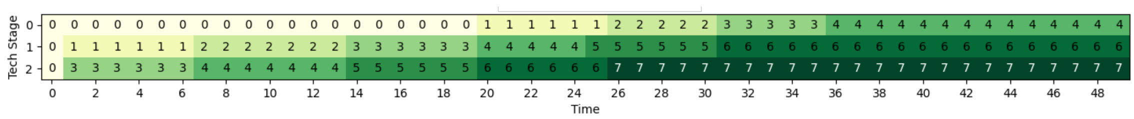

Our estimation of this threshold is based on the assumption that the simulation will reach tech stage 2 by approximately 2050. If this occurs, then we can judge the threshold carbon prices for achieving 50% or 70% renewable penetration by the number of renewable plants that exist by 2070 in tech stage 2. However, Figure 9 demonstrates the drastic differences in the number of RES plants built if the simulation remains in tech stages 0 or 1. This means that even with the most optimal scenario of decreasing costs over time, it still takes a steep starting price of carbon to achieve targets of 50% or higher renewable penetration by 2070.

One final consideration is that since our model implements a finite time horizon, costs after 2070 cease to exist, so this definitely impacts the cost-benefit optimization of building more RES plants. We could imagine that under an infinite time horizon, the optimal policy under any carbon tax would eventually reach 100% renewable penetration to avoid paying carbon tax for an infinite number of time steps. Carbon emissions certainly do not cease to cause harm after 2070, but with a discount rate of just 3% and a duration of 50 years, it seems surprising that the optimal policy under the CLC proposed carbon tax only achieved 30% renewable penetration.

8 Discussion

The main insights from our three experiments using MDP V3 may be summarized as follows:

Firstly, the costs of transitioning from a fossil fuel to a renewable energy system are strongly dominated by the capital costs of energy storage. This aligns with the literature on decarbonizing the electric grid when energy storage is the main method used to compensate for renewable energy intermittency. Achieving target levels of renewable penetration in our model was highly dependent on the simulation reaching tech stage 2, in which storage costs have decreased by over 50%. This implies that without significant technological improvement to reduce the cost of energy storage systems, planners will not be economically incentivized to make the investments necessary for high levels of renewable penetration unless subsidies or other means of long term funding are put in place for renewable energy and storage systems.

Secondly, as supported by the previous point, the optimal policy becomes increasingly sensitive to reductions in storage costs as the number of RES plants built in the base case (zero percent reduction) increases. Linear reductions in storage costs lead to non-linear effects on the total cumulative cost, although they also resulted in tiered behavior among optimal policies: below a certain threshold in the storage cost reduction, all optimal policies eventually built the same number of RES plants. However, within a single tier, the optimal policies with higher reductions in storage cost tended to build RES plants earlier, thus reducing cumulative carbon emissions. This underscores the importance of developing cheaper storage systems,[17] and perhaps indicates that even with optimistic cost projections, BSS/PHS systems may be too expensive to expand to the majority of utility scale storage.

Thirdly, using a linearly increasing versus an exponentially increasing carbon tax has important implications for short term versus long term carbon emissions. Considering the U.K.-based linear tax (23.63 USD/ton increasing by 10 USD/yr) and the Finland-based exponential tax (68.52 USD/ton increasing by 5%/yr), the price of carbon is higher under the linear tax until about 2055. Driven by higher carbon prices, the optimal policy under the linear tax builds RES plants earlier and thus avoids more carbon emissions. However, since waiting increases the chances of seeing RES and storage capital costs decrease, the cumulative cost of this optimal policy is higher. This highlights an important trade-off between placing less economic burden on entities that must convert from fossil fuel to renewable energy and incurring fewer carbon emissions during the energy transition. Furthermore, if the most important goal is reducing carbon emissions, then policymakers could consider a hybrid carbon tax schedule: adjusting the starting price to reach some level of renewable penetration by a target year, the tax begins as linear to incentivize converting to renewable energy sooner, and after passing the target year, switches to exponential in order to penalize emissions remaining above the target more heavily.

Finally, achieving 50% renewable penetration by 2070 would require a starting carbon price of almost 70 USD/ton and an exponential increase of 5% per year. Given the exponentially increasing storage requirement and high capital costs of BS and PH systems, this is somewhat unsurprising. Yet these findings could lead us to question exactly how the CLC proposed carbon tax of 41.84 USD/ton increasing at 5% per year could achieve 50% carbon emission reductions relative to 2005 by 2035. Our model undoubtedly lacks many infrastructural and economic details, but there does seem to be a discrepancy between the low starting price of the CLC proposed tax and its prediction to achieve 50% renewable penetration while maintaining a nationwide reliable power supply.

9 Limitations

The biggest limitation of this model is its inability to capture aspects of the electricity supply chain besides power generation. For instance, we assume an average annual demand for power that will be met exactly by the power output of the plants in the system. However, especially for RES plants, this is unrealistic because the timing (both within a single day and over seasons within a year) of peak power generation for renewables often does not match that of peak demand. The storage requirement is designed to address this, but as literature on renewable energy storage indicates, simply increasing storage is often insufficient to maintain reliability of the electric grid. Flexible power generation, demand-response,[18] and other techniques are as important in increasing renewable penetration as storage capacity.

Another major area of improvement for this model would be enlarging the state space while preserving interpretability. In MDP V3, we kept the state space as small as possible by modeling only the bare minimum of variables as state dependent. Although this allowed the model to run much faster, it also reduced our ability to capture certain aspects of there real world in our model. The largest time sink is actually populating the transition probabilities, which takes $O(n^2)$ time where n is the product of the number of possible values of each state variable. Our model takes on the order of seconds to run, so adding more state variables, especially with limited ranges, should not incur a significant time cost.

Some specific examples of a richer state space include tracking the number of each RES plant built in tech stages 0, 1, and 2 separately, each as a different state variable. Doing so would allow us to make renewal costs depend on the tech stage in which the plant was built rather than the current tech stage; this is a more realistic approach because the equipment in older plants is more difficult and expensive to refurbish. We could also track the age of each power plant as separate state variables. Having access to this information via state would allow us to model more sophisticated plant renewal costs, as well as payment plans spread out over multiple years, since the capital costs of renewable energy projects are financed based on amortized costs rather than paid for entirely up front.

Finally, one specific limitation of this model that may indirectly tie it to a specific geographic region is how we model energy storage. We use E.U. based estimates for calculating the required storage for a given level of renewable penetration, which presumably depend to some degree on the renewable resource distribution and existing electric infrastructure in the E.U. We assumed that since the E.U. encompasses a range of renewable resources (strong solar radiation in the southern nations versus strong wind in the Scandinavian nations) this would lead to an averaging effect on storage requirements, but in order to claim that our model is truly generalizable we would need equivalent estimates based on a global survey of renewable penetration and accompanying storage.

[17] Innovation in energy storage is a rapidly expanding space, especially in terms of developing alternatives to batteries and pumped hydro. One idea is to synthesize ammonia fuel using renewable electricity, since it can be stored and burned as a carbon-free fuel later on [11]. Another idea, known as the Concrete Tower, involves storing kinetic energy in blocks of recycled concrete stacked as high as small skyscrapers [44].

[18] Electricity *demand-response* refers to the integration of smart monitoring systems into an electric grid. To avoid being caught off guard by drastic changes in demand, utility operators try to predict demand in real time and ramp up or ramp down the power supply accordingly.

10 Future Work

There are still numerous questions that our model as is could be used to answer, many of which concern simply the input parameters. For instance, by adjusting the weighted average of FF and RES plant costs, we could model systems with different mixes of FF plants and of RES plants; this would be a simple way to approximate a geographic region. As an example, all FF plants could be coal plants and all renewable energy plants could be solar PV, which is a plausible set up for the energy transition in parts of India. Similarly, we could test different mixes of BSS and PHS in the sensitivity analysis for storage costs. We could also test different carbon tax starting prices and/or growth rates in the sensitivity analysis for renewable penetration. Furthermore, we could change cost inputs that depend on tech stage to reflect more conservative or more optimistic technological projections for 2030 and 2050, and see how this affects achievable renewable penetration levels for different carbon tax schedules, etc.

One tradeoff that our model could incorporate without major structural changes is that between storage and curtailment. Currently, our model requires some storage capacity to be built depending on the level of renewable penetration. However, since the (capital) costs of storage are so high, optimizing for the most cost-effective renewable power supply in the real world will likely entail building excess renewable energy and tolerating some degree of curtailment. In future iterations of the model, we could model this by providing two options when building a RES plant: 1) build the original required storage capacity, or 2) build a larger RES plant by some percentage, and accordingly less storage capacity. This would allow us to test how the optimal policy changes when different levels of curtailment are tolerated.

Another area that would be interesting to explore is making the carbon tax schedule more dynamic. By introducing stochasticity into the tax schedule, we could model the unpredictability of climate policy in the real world governments. For instance, if using a U.S.-based model, we might allow the carbon tax schedule to change every four or eight years to reflect changes in the federal administration and its influence on a nationwide carbon tax. We could also build the concept of carbon dividends into our model by allotting some percentage of carbon tax costs from the previous year towards building RES plants in the current year. It would be interesting to see how the conflict between an increasing and a decreasing amount of carbon emissions are reflected in these carbon dividends. In our current model, the carbon tax is deterministic given time, but it would also be interesting to make the carbon tax schedule dependent on annual or cumulative carbon emissions. This would allow us to test a wider array of carbon tax policies, and determine how rewarding a decrease in carbon emissions (in addition to penalizing existing emissions) might affect the optimal policy.

Finally, we recognize that our model is an extremely generalized picture of the energy transition and intentionally avoids modeling aspects of the energy system in a way that would tie it to a specific region geographically or economically. However, there is still room within this framework to include universal characteristics of renewable energy or annual load profiles, namely their stochastic nature. Although the results would represent averaged global observations of uncertainty, we could still incorporate stochasticity into how we model renewable power generation and annual power demand. This would allow us to draw more meaningful conclusions about the additional reliability costs of transitioning a fossil fuel to a renewable energy system.

11 Conclusion

As of December 2019, 73 countries present at the United Nations Framework Convention on Climate Change (UNFCCC) made formal commitments and drafted policy roadmaps towards net zero emissions by 2050 [50]. These targets are motivated by the rapidly shrinking global carbon budget, beyond which climate models suggest that constraining global temperature increase to 2 or even 3 degrees Celsius would be impossible. Throughout this paper we have used values for the starting price of carbon based on carbon taxes that exist in the real world, and using our model found that rendering even 50% renewable penetration cost-optimal would require a steep starting price of carbon. However, only 6 out of 25 countries with nationally implemented carbon taxes have set their price above 40 USD/ton [6]. Amongst the countries with high carbon prices it is unclear whether they will maintain an exponential increase in price, even though we saw that this growth rate was a key factor in incentivizing optimal policies to reach high levels of renewable penetration.

In our model, the need for a steep starting price of carbon is mainly due to energy storage costs that begin already high, and increase exponentially because storage requirements increase exponentially with renewable penetration. Although this is a difficult argument to make to policymakers, exponential increases in the price of carbon are thus required to achieve linear increases in renewable penetration. However, here is where the simplified nature of our model falls short. Between the actual price of carbon and the cost that a carbon tax places on society, there a great deal of room for innovative climate policies. For instance, the carbon dividends aspect of the CLC proposal is a highly effective way of reversing the potentially regressive nature of a universal carbon tax. In conclusion, we hope that this paper sheds a more realistic light on the cumulative costs of transitioning away from fossil fuel energy systems, and reveals how a critical factor in incentivizing this process is for countries to implement carbon taxes that are as ambitious as their decarbonization targets.

Acknowledgements

I would like to sincerely thank my advisors, Jackson Killian and Andrew Perrault, for their detailed guidance and mentorship over the course of this project.

Appendix

A Model Parameters

Table 2: Parameters used to run the MDP model and their baseline values (`p_v3_baseline` in the code repository). Any units in $ are adjusted to 2020-USD and tons refer to metric tons.

Variable	Description	Value	Units	Tech Stage
`n_years`	Number of years for which model is run	50		
`n_plants`	Total number of power plants	10		
`n_tech_stages`	Number of tech logical stages possible to advance through	3		
`ff_size`	Nameplate capacity of a FF plant	600	MW	
`ff_capacity`	Capacity factor of a FF plant	60	%	
`res_lifetime`	Number of years a FF power plant can run before needing renewal	32	%	
`res_size`	Nameplate capacity of a RES plant	*	MW	
`res_capacity`	Capacity factor of a RES plant	30	%	
`res_lifetime`	Number of years a RES power plant can run before needing renewal	20	%	
`c_co2_init`	Starting carbon tax price	41.84	$/ton	
`c_co2_inc`	Coefficient in carbon price function	5		
`co2_tax_type`	Whether carbon tax function is linear or exponential	EXP**		
`c_ff_cap`	Construction costs of a FF plant	3429	$/kW	
`c_ff_fix`	Fixed operation & maintenance costs of a FF plant	37	$/kW	
`c_ff_var`	Variable operation & maintenance costs of a FF plant	0.012	$/kWh	
`ff_emit`	Annual carbon emissions of a FF plant	1.65	kg/kWh	
`c_res_cap`	Construction costs of a RES plant	1354, 881, 574	$/kW	Dependent
`storage_coefs`	Proportion of storage system built using BSS and PHS, respectively	50, 50	%	
`storage_mix`	Coefficients for the exponential function that calculates storage required at some renewable penetration	9.99e-3, 02.31e-2, -9.99e-3	%	
`bss_hrs`	Number of hours that BS system can hold charge for	4	$/kWh	
`c_bss_cap`	Construction costs of a BS system	390, 213, 160	$/kWh	Dependent
`c_bss_fix`	Fixed operation & maintenance costs of a BS system	38	$/kW	
`c_bss_var`	Variable operation & maintenance costs of a BS system	1.75	$/kWh	
`c_phs_cap`	Construction costs of a PH system	170	$/kWh	
`c_phs_fix`	Fixed operation & maintenance costs of a PH system	15	$/kW	
`p_adv_tech`	Probability that tech logical stage advances to the next stage	11, 13, 0	%	Dependent
`disc_rate`	Discount rate ($1-$ discount factor) used to run MDP	3	%	

*Calculated using `fplant_size`, `fplant_capacity`, `rplant_capacity` so that total available power in the system is constant.

**Either "LIN" or "EXP" are accepted values.

B Baseline Value Estimates

Table 3: Estimates and data sources used to calculate parameter values for baseline and other scenarios. Any units in $ are adjusted to 2020-USD and tons refer to metric tons.

Category	**Description**	**Value**	**Units**	**Source**
–	Discount rate	3	%	[49]
FF: coal	Capital cost	4409	$/kW	[5]
	Fixed O&M cost	51.24	$/kW	[5]
	Variable O&M cost	0.0052	$/kWh	[5]
	Capacity factor	81	%	[15]
	Rate of carbon emissions	2.21	kg/kWh	[26]
FF: natural gas	Capital cost	1094	$/kW	[5]
	Fixed O&M cost	11.27	$/kW	[5]
	Variable O&M cost	0.0063	$/kWh	[5]
	Capacity factor	22	%	[15]
	Rate of carbon emissions	0.92	kg/kWh	[26]
RES: solar PV	Capital cost	1210	$/kW	[21]
	Capital cost (2030, 2050 projections)	687, 323	$/kW	[21]
	Fixed O&M cost	15.19	$/kW	[2]
	Variable O&M cost	0	$/kWh	[2]
	Capacity factor	25	%	[21]
RES: onshore wind	Capital cost	1497	$/kW	[22]
	Capital cost (2030, 2050 projections)	1075, 825	$/kW	[22]
	Fixed O&M cost	26.22	$/kW	[2]
	Variable O&M cost	0	$/kWh	[2]
	Capacity factor	35	%	[21]
Storage: BSS	Capital cost	390	$/kW	[8]
	Capital cost (2030, 2050 projections)	213, 160	$/kW	[8]
	Fixed O&M cost	38	$/kW	[8]
	Variable O&M cost	1.75	$/kWh	[8]
Storage: PHS	Capital cost	170	$/kW	[17]
	Fixed O&M cost	15	$/kW	[17]
	Variable O&M cost	0.0055	$/kWh	[17]
CO2 tax	CLC proposal	41.84	$/ton	[48]
	Stern Review	446.04	$/ton	[45]
	Japan	2.65	$/ton	[6]
	South Africa	8.52	$/ton	[6]
	United Kingdom	23.63	$/ton	[6]
	Finland	68.52	$/ton	[6]
	Switzerland	99.11	$/ton	[6]
	Sweden	123.18	$/ton	[6]

References

[1] Frank Ackerman and Elizabeth A. Stanton. *The Cost of Climate Change.* National Resources Defense Council, at: https://www.nrdc.org/sites/default/files/cost.pdf. 2008.

[2] *Annual Energy Outlook 2020: Cost and Performance Characteristics of New Generating Technologies.* U.S. Energy Information Administration, at: https://www.eia.gov/outlooks/aeo/assumptions/pdf/table_8.2.pdf. 2019.

[3] Marshall Burke, Solomon M Hsiang, and Edward Miguel. "Global non-linear effect of temperature on economic production". In: *Nature* 527.7577 (2015), pp. 235–239.

[4] *California Imports the Most Electricity from Other States; Pennsylvania Exports the Most.* U.S. Energy Information Administration, at: https://www.eia.gov/todayinenergy/detail.php?id=38912. 2019.

[5] *Capital Cost Estimates for Utility Scale Electricity Generating Plants.* U.S. Energy Information Administration, at: https://www.eia.gov/analysis/studies/powerplants/capitalcost/pdf/capcost_assumption.pdf. 2016.

[6] *Carbon Pricing Dashboard: Up-to-Date Overview of Carbon Pricing Initiatives.* World Bank, at: https://carbonpricingda 2020.

[7] Michael Child et al. "Flexible electricity generation, grid exchange and storage for the transition to a 100% renewable energy system in Europe". In: *Renewable energy* 139 (2019), pp. 80–101.

[8] Wesley Cole and A. Will Frazier. *Cost Projections for Utility-Scale Battery Storage.* National Renewable Energy Laboratory, at: https://www.nrel.gov/docs/fy19osti/73222.pdf. 2019.

[9] *Cost of Transitioning to 100-Percent Renewable Energy.* Institute for Energy Research, at: https://www.instituteforenerg of-transitioning-to-100-percent-renewable-energy/. 2019.

[10] *Data & Statistics.* International Energy Agency, at: https://www.iea.org/data-and-statistics. 2020.

[11] Jason Deign. *Siemens Tests Ammonia as a Form of Energy Storage for Renewables.* Greentech Media, at: https://www.greentechmedia.com/articles/read/siemens-ammonia-hydrogen-energy-storage. 2020.

[12] Paul Denholm and Maureen Hand. "Grid flexibility and storage required to achieve very high penetration of variable renewable electricity". In: *Energy Policy* 39.3 (2011), pp. 1817–1830.

[13] Paul Denholm and Robert M Margolis. "Evaluating the limits of solar photovoltaics (PV) in electric power systems utilizing energy storage and other enabling technologies". In: *Energy Policy* 35.9 (2007), pp. 4424–4433.

[14] *Digest of UK Energy Statistics (DUKES): electricity.* Gov.UK, at: https://www.gov.uk/government/statistics/electricity chapter-5-digest-of-united-kingdom-energy-statistics-dukes. 2019.

[15] *Electric Power Monthly with Data for February 2020.* U.S. Energy Information Administration, at: https://www.eia.gov/electricity/monthly/current_month/epm.pdf. 2020.

[16] *Energy and The Environment Explained: Outlook for Future Emissions.* U.S. Energy Information Administration, at: https://www.eia.gov/energyexplained/energy-and-the-environment/outlook-for-future-emissions.php. 2020.

[17] *Energy Storage Technology and Cost Characterization Report.* Pacific Northwest National Laboratory, at: https://www.energy.gov/sites/prod/files/2019/07/f65/Storage%20Cost%20and%20Performance%20Characterizatio 2019.

[18] Annette Evans, Vladimir Strezov, and Tim J Evans. "Assessment of utility energy storage options for increased renewable energy penetration". In: *Renewable and Sustainable Energy Reviews* 16.6 (2012), pp. 4141–4147.

[19] *Fact Sheet: Energy Storage: electricity.* Environmental and Energy Study Institute, at: https://www.eesi.org/papers/view storage-2019. 2019.

[20] Abdolvahhab Fetanat and Ehsan Khorasaninejad. "Size optimization for hybrid photovoltaic–wind energy system using ant colony optimization for continuous domains based integer programming". In: *Applied Soft Computing* 31 (2015), pp. 196–209.

[21] *Future of Solar Photovoltaic: Deployment, investment, technology, grid integration and socio-economic aspects.* International Renewable Energy Agency, at: https://www.irena.org/-/media/Files/IRENA/Agency/Publicatior 2019.

[22] *Future of Wind: Deployment, investment, technology, grid integration and socio-economic aspects.* International Renewable Energy Agency, at: https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Oc 2019.

[23] Lawrence H Goulder and Andrew R Schein. "Carbon taxes versus cap and trade: a critical review". In: *Climate Change Economics* 4.03 (2013), p. 1350010.

[24] Lawrence H. Goulder Goulder and Marc Hafstead. *A Numerical General Equilibrium Model for Evaluating U.S. Energy and Environmental Policies.* Resources for the Future, at: https://media.rff.org/archive/files/sharepoint Numerical-General-Equilibrium-Model.pdf. 2013.

[25] Marc Hafstead. *Analysis of Alternative Carbon Tax Price Paths for the Climate Leadership Council (CLC) Carbon Dividends Plan.* Climate Leadership Council, at: https://clcouncil.org/reports/RFF-IB-18-07-rev_mar19.pdf. 2019.

[26] *How much carbon dioxide is produced per kilowatthour of U.S. electricity generation?* U.S. Energy Information Administration, at: https://www.eia.gov/tools/faqs/faq.php?id=74&t=11. 2018.

[27] Becca Jones-Albertus. "Confronting the Duck Curve: How to Address Over-Generation of Solar Energy". In: *Office of Energy Efficiency & Renewable Energy* (2017).

[28] Azadeh Kamjoo et al. "Multi-objective design under uncertainties of hybrid renewable energy system using NSGA-II and chance constrained programming". In: *International Journal of Electrical Power & Energy Systems* 74 (2016), pp. 187–194.

[29] Goksin Kavlak, James McNerney, and Jessika E Trancik. "Evaluating the causes of cost reduction in photovoltaic modules". In: *Energy policy* 123 (2018), pp. 700–710.

[30] Anil Kilickaplan et al. "An energy transition pathway for Turkey to achieve 100% renewable energy powered electricity, desalination and non-energetic industrial gas demand sectors by 2050". In: *Solar Energy* 158 (2017), pp. 218–235.

[31] Vincent Krakowski et al. "Feasible path toward 40–100% renewable energy shares for power supply in France by 2050: A prospective analysis". In: *Applied energy* 171 (2016), pp. 501–522.

[32] *Learn More About Interconnections.* Office of Electricity, at: https://www.energy.gov/oe/services/electricity-policy-coordination-and-implementation/transmission-planning/recovery-act-0.

[33] Richard Loulou. *Documentation for the TIMES Model Part I.* International Energy Agency, at: https://iea-etsap.org/docs/Documentation_for_the_TIMES_Model-Part-I_July-2016.pdf. 2016.

[34] Akbar Maleki and Alireza Askarzadeh. "Artificial bee swarm optimization for optimum sizing of a stand-alone PV/WT/FC hybrid system considering LPSP concept". In: *Solar Energy* 107 (2014), pp. 227–235.

[35] Sonam Norbu and Santanu Bandyopadhyay. "Power Pinch Analysis for optimal sizing of renewable-based isolated system with uncertainties". In: *Energy* 135 (2017), pp. 466–475.

[36] Filip Patocka. *Environmental Impacts of Pumped Storage Hydro Power Plants.* Norwegian University of Science and Technology, at: https://pdfs.semanticscholar.org/b3e7/c716775cba240aec39f6521c3af28c821a16.pdf. 2014.

[37] Hannah Ritchie and Max Roser. *Energy.* Our World in Data, at: https://ourworldindata.org/energy#all-charts-preview. 2014.

[38] Ayesha Sadiqa, Ashish Gulagi, and Christian Breyer. "Energy transition roadmap towards 100% renewable energy and role of storage technologies for Pakistan by 2050". In: *Energy* 147 (2018), pp. 518–533.

[39] Matthew R Shaner et al. "Geophysical constraints on the reliability of solar and wind power in the United States". In: *Energy & Environmental Science* 11.4 (2018), pp. 914–925.

[40] Masoud Sharafi and Tarek Y ELMekkawy. "Multi-objective optimal design of hybrid renewable energy systems using PSO-simulation based approach". In: *Renewable Energy* 68 (2014), pp. 67–79.

[41] Nahum Shimkin. *Lecture Notes: Dynamic Programming – Finite Horizon.* The Andrew & Erna Viterbi Faculty of Electrical Engineering, at: https://webee.technion.ac.il/shimkin/LCS11/ch2_DP_finite.pdf. 2011.

[42] Abebe Asfaw Solomon, Daniel M Kammen, and D Callaway. "The role of large-scale energy storage design and dispatch in the power grid: a study of very high grid penetration of variable renewable resources". In: *Applied Energy* 134 (2014), pp. 75–89.

[43] Abebe Asfaw Solomon et al. "How much energy storage is needed to incorporate very large intermittent renewables?" In: *Energy Procedia* 135 (2017), pp. 283–293.

[44] Julian Spector. *Can Newcomer Energy Vault Break the Curse of Mechanical Grid Storage?* Greentech Media, at: https://www.greentechmedia.com/articles/read/energy-vault-stacks-concrete-blocks-to-store-energy. 2018.

[45] Nicholas Stern and Nicholas Herbert Stern. *The economics of climate change: the Stern review.* cambridge University press, 2007.

[46] Richard Stuebi. *Power Plant Flexibility: Increasingly Valuable in a More Challenging Operating Environment.* General Electric, at: https://www.ge.com/power/transform/article.transform.articles.2017.oct.power-plant-flexibility-increa. 2017.

[47] James Temple. *The $2.5 trillion reason we can't rely on batteries to clean up the grid.* 2018.

[48] *The Four Pillars of Our Carbon Dividends Plan.* Climate Leadership Council, at: https://clcouncil.org/our-plan/. 2019.

[49] Hal Varian. "Recalculating the costs of global climate change". In: *The New York Times* 14.12 (2006), p. 2006.

[50] Catherine Benson Wahlen. *73 Countries Commit to Net Zero CO2 Emissions by 2050.* International Institute for Sustainable Development, at: https://sdg.iisd.org/news/73-countries-commit-to-net-zero-co2-emissions-by-2050/. 2019.

[51] Henrik Zsiborács et al. "Intermittent renewable energy sources: The role of energy storage in the european power system of 2040". In: *Electronics* 8.7 (2019), p. 729.