Testing the Stability of PJM's "Resource Carve Out" Capacity Market Proposal Using an Agent Based Model

Andrew M Mowry

Topic:	This paper follows the open FERC docket EL16-49, which pertains to PJM's prospective capacity market redesign.
Purpose:	To investigate the stability of PJM's preferred design outcome, the Extended Resource Carve Out with Expanded Minimum Offer Price Rule (Ex-RCO w/ Ex-MOPR).
Methodology:	Simulating the time evolution of an abstracted Energy+Capacity market by coding an Agent Based Model with PJM's specific proposal.
Findings:	PJM's Ex-RCO w/ Ex-MOPR proposal succeeds in preventing capacity market price suppression due to state-sponsored subsidies, but fails to control over-procurement of capacity without a cap on carve-outs.

Contents

Introduction

Electricity markets that have restructured to allow wholesale generation competition, such as those operated by PJM or ISO-NE, are struggling with opposing mandates to preserve market-based competition and to allow states to effect policy outcomes through market intervention. In the present case, subsidized generation – often favored for carbon-reduction or RPS ends – threatens to marginalize and bankrupt unsubsidized generation by undercutting market prices in the energy and capacity markets. Under pressure from competing political forces, the system operators and their regulators are attempting to redesign aspects of their markets to reach a compromise solution.

PJM is embroiled in a long-running process of this kind at FERC, focusing on a proposed redesign of its capacity market construct, the Reliability Pricing Model (RPM). At the center of the debate is the requirement, due to the legislated Federal Power Act (FPA), that wholesale rates not be "unjust and unreasonable". Under many interpretations of this rule, state subsidies for certain power generation resources act through the current capacity market structure to create unjust and unreasonable wholesale rates.

This paper explores the solution that PJM has proposed to reform the RPM. Specifically, I apply an "agent based model" (ABM), which generates aggregate market behavior that arises from the interaction of many individual market participants. By simulating the key attributes of the energy and capacity markets, the model provides insight into the circumstances that might lead to price suppression, investment cycles, and under- or over-procurement of capacity.

The conclusions from my simulations generally agree with the representations that PJM presents in its submissions to FERC: under the proposal, capacity prices are not uncompetitively suppressed, and subsidized resources are able to attain capacity commitments. However, the simulations show clearly that the capacity construct does not adequately deal with high penetrations of subsidized resources and may require revisiting in the near-term future.

The paper is structured as follows: first, an overview of PJM's capacity market restructuring is given; second, the ABM as a methodology is introduced; third, my specific implementation of the ABM to PJM's proposed RPM rules is described; fourth, the main simulation results are presented and discussed; and fifth, conclusions are made and possible extensions to the paper are proposed.

Background: Overview of PJM's Capacity Market Restructuring

Capacity markets are usually motivated by the need to ensure resource adequacy. The "missing money problem" in energy markets, especially those such as PJM without robust real-time scarcity and reserves pricing, necessitates an alternative way of procuring generation for extreme times. The capacity market is a common solution. Other methods of handling this problem, such as administrative payments, exist, but are not handled here and are not seriously entertained by PJM.

For several reasons – including price volatility inherent from the coupling with the energy market, shaky notions of "reliability" and "adequacy", the ever-present specter of market power (both from large buyers and sellers), and the psychological role of the capacity market as a "make-whole" market – the capacity market in PJM, called the Reliability Pricing Model (RPM), is under constant review and redesign. For the past several years, the RPM has operated with a Minimum Offer Price Rule (MOPR), specifying a floor price below which a new resource may not offer capacity into the market.

The purpose of the MOPR is to protect against buyer-side market power: a large buyer may benefit economically from propping up a generator below cost if it has the effect of suppressing capacity prices.

In 2016, Calpine Corporation launched a complaint against PJM at FERC alleging that the MOPR as implemented produced "unjust and unreasonable" rates by allowing too many opportunities for state subsidized resources to impact capacity prices. Calpine, as a merchant generation owner, was itself threatened by the potential for capacity clearing prices to be suppressed in this way. FERC agreed with Calpine, ordering PJM to come up with a better solution. PJM's first attempt to do so was rejected by FERC in June 2018 for failing to go far enough in protecting competitive rates.[1] The docket (EL16-49) has since brimmed, not only with PJM's updated proposal and subsequent defense, but with the submissions of dozens of intervenors representing different constituencies, showing how contentious this issue is. As mentioned in the introduction, this debate is occurring in the context of positive state policies designed to affect those states' generation mixes, for example Renewable Portfolio Standards.

PJM's updated proposal, which we now focus on, was originally presented a la carte to FERC, but PJM later clarified in its Nov 6th filing that it advocates for the adoption of two components together: the expanded MOPR (ex-MOPR) and the extended Resource Carve-Out (ex-RCO).[2] The ex-MOPR would force all new and existing generators not yet having cleared an RPM auction (with certain narrower exceptions than those currently in place) to offer in at or above their net CONE equivalent. This is a strong check on buyer-side market power, but arguably disables states from achieving policy goals. Net CONE is a high barrier to entry, and a state would then not be able to subsidize its preferred technology into a capacity award. Consequently the state's native load would end up purchasing capacity products from non-preferred technologies, frustrating the state's policy goals.

To address this concern, PJM introduced the RCO construct to provide an (immediately unsatisfying) mechanism to clear these preferred and subsidized resources: by electing RCO status for a delivery year the resources are able to ignore the ex-MOPR and secure capacity obligations using below-cost offers, but gain no corresponding monetary awards. Instead, an amount of load commensurate to the "carved-out resource" capacity is rebated from RPM auction revenue; this rebate can be directed to the carved-out resource or not, depending on state policies. This is an unsatisfying solution because, although it re-empowers states to manage their own capacity resources, it does nothing to address potential capacity price suppression for unsubsidized resources.

As a foil to *this* concern, PJM proposed a second round in the RPM auction, which, together with the first, is known as the extended RCO. Mechanically, PJM in the second round removes the carved-out resources from the offer supply curve altogether and resolves the auction to determine what should have been the competitive clearing price. This competitive/unsuppressed price is awarded to all those who cleared in the first round. Those resources that are inframarginal in the second round but not in the first round only receive the difference between the two clearing prices, the so-called "inframarginal rent". (See PJM's October 2 Submittal for full details on the ex-MOPR and ex-RCO proposals.[3])

The intervenors in the docket have raised many concerns to this proposal. LS Power, for instance, argues out that there might be a need for a "cap" on total RCO MWs allowed in the auction to

[1] Calpine Corp. v. PJM Interconnection, L.L.C., 163 FERC ¶ 61,236 (Jun. 29, 2018) ("June 29 Order").
[2] Reply Submission of PJM Interconnection, L.L.C., Docket Nos. EL16-49-000, et al., (Nov. 6, 2018) ("PJM Reply")
[3] Initial Submission of PJM Interconnection, L.L.C., Docket Nos. EL16-49-000, et al., (Oct. 2, 2018) ("October 2 Submittal").

prevent price suppression despite the novel market structures.[4] (No analysis was provided.) Load interests worried that clearing prices may trend in the other direction and rise. What can be expected to happen to the intermarginal units between the first and second rounds, and do their auction incentives change? As this paper focuses on "unjust and unreasonable" rates, these questions around price formation are of clear interest.

Background: Overview of Agent Based Modelling (ABM)

To investigate the possible behavior of PJM's ex-MOPR and ex-RCO proposal introduced above, I use a stochastic, agent-based model (ABM). The core intuition for the ABM is that it models each generator (agent) as a discrete decision-maker that is able to make decisions over the course of a simulation according to market rules and any information it has access to. By creating many such agents with varied attributes and preferences and allowing them to interact, it is possible to generate believable and intuitive market-level patterns. The ABM has been used before to study structural changes in energy and capacity markets, recently by Hach and Spinler, who applied one to the UK energy market.[5]

As Hach and Spinler introduce, other types of models are often employed to study capacity markets. "Capacity expansion models" based on linear programming obviously excel at finding optimal solutions, but complex decision processes and uncertainties can't be accounted for. ABMs (or "dynamic investment models" as the pair calls them) have their relative strength in "market feedback, technology competition, … and strategic behavior". In an earlier paper, Hach et al. mention the use of real options models for single project valuations, but does not expand their use to a marketplace.[6]

It is important to make a note here about the type of modelling that this paper attempts: it is not to predict the future but rather to consider what could occur in the future and how those situations might arise. To do anything else, as Herbert Simon reminds, when considering a dynamic problem like this, would be foolish.[7] Accordingly this paper treats the modelling much like a system dynamics problem, full of feedback loops and control theory and not good for much more than identifying modes of behavior. Indeed, Mauricio Cepeda recently published a study of European capacity markets explicitly framed in terms of system dynamics.[8]

Methodology: Application of ABM to PJM's Proposed Capacity Market Reform

This paper's application of an ABM, which I now describe, follows the broad structure used by Hach and Spinler.[9] It comprises a core loop that represents one year's energy market followed by a capacity auction for the next year. In each iteration of the loop, (1) a set of agents (power plants) is determined, (2) the agents compete to clear the energy market, (3) they compete, along with a set of

---

[4] Initial Brief of LS Power Associates, L.P., Docket Nos. EL16-49-000, et al., at 13 (Oct. 2, 2018)

[5] Hach, Daniel and Spinler, Stefan, Robustness of Capacity and Energy-Only Markets - A Stochastic Dynamic Capacity Investment Model (October 11, 2014). Available at SSRN: https://ssrn.com/abstract=2508696

[6] Daniel Hach, et al., Capacity market design options: A dynamic capacity investment model and a GB case study, European Journal of Operational Research, Volume 249, Issue 2, 2016, Pages 691-705, ISSN 0377-2217, https://doi.org/10.1016/j.ejor.2015.08.034.

[7] Simon, Herbert A. "Prediction and Prescription in Systems Modeling." *Operations Research*, vol. 38, no. 1, 1990, pp. 7–14. *JSTOR*, JSTOR, www.jstor.org/stable/171293.

[8] *Cepeda*, *Mauricio.* Assessing cross-border integration of capacity mechanisms in coupled electricity markets. *Energy Policy*, 2018, vol. 119, issue C, 28-40. https://econpapers.repec.org/RePEc:eee:enepol:v:119:y:2018:i:c:p:28-40

[9] Hach and Spinler, supra Note 5.

new agents (proposed new power plants), to clear the capacity market, and (4) each agent is allowed to make a decision to enter or exit next year's market according to their economics and the rules of the market. I now detail each of these steps in more detail.

The Agents

The hallmark of an agent-based model is the use of numerous, heterogeneous agents. In this model each agent represents a power plant of a particular technology type, although I have chosen to abstract the agents from real generation technologies. There are two types of technology types I allow to compete, Type A and Type B, with attributes summarized below.

Type	Build Cost	Fixed Cost	VOM	Heat Rate	Life	Capacity	Subsidized
	\$/MW	\$/MWy	\$/MWh	Btu/Wh	Years	MW	
Type A	\$ 500,000	\$ 20,000	\$ 20.00	-	20	500	No
Type B	\$ 800,000	\$ 20,000	\$ 20.00	-	20	500	Yes

These attributes together define the short-run and long-run economics of the power plant. To take each in turn: "Build Cost" is the amount that must be paid as a lump sum upon entering the market based upon the plant's installed capacity, and we see that Type B is 160% the cost Type A; "Fixed Cost" is the amount that must be recovered each year for a plant to break even, including such expenses as major maintenance, and is based on the plant's installed capacity; "VOM" or "Variable O&M" is the cost to produce 1 MWh of electricity independent of fuel; "Heat Rate" describes the fuel-to-power conversion efficiency and is assumed zero for each technology;[10] "Life" is the operating lifetime of a plant, crucially over which it must recoup its Build Cost to be a sensible investment; "Capacity" is the plant's maximum MW output over an hour; "Subsidized" connotes whether subsidies affect this plant.

Two dynamics are important to discuss: subsidization and heterogeneity. Subsidies, as the conduit for state policy, are crucial to the question that this model investigates; they are modelled in a very specific way that elucidates their dynamics over time. Both ITC-style[11] and PTC-style[12] subsidies are considered. The ITC-style subsidy is redeemable in the first year of operation, and acts as a Build Cost rebate; the PTC-style subsidy is redeemable on every MWh of energy produced and acts as a VOM rebate. Each of these subsidies intensifies on a linear schedule over the first 20 years of the simulation, at which point Type B ends up with an effective subsidized Build Cost of \$400,000/MW and an effective subsidized VOM of \$10/MWh. Thus, the Type B technology begins the simulation relatively expensive, and ends the simulation relatively cheap (on a subsidized basis). There is no technology learning in this model, i.e. the agent attribute table does not change.

Heterogeneity between agents, which causes a dispersion in their economics and the decisions they make, enters the model upon agent creation. The attributes above are taken as average values for each technology, but an agent upon instantiation has values greater than or less than these averages by a random amount up to 20%. Thus, in each iteration some power plants are more or less competitive than others. Over time, successful power plants will tend to have more efficient characteristics, and the population average attributes would be expected to cluster around the better values. At the beginning

---

[10] This is a major simplifying assumption for the model, as it need not account for fuel prices. Short run variable costs are thus driven solely by VOM.

[11] Investment Tax Credit, as it is implemented by the IRS in the USA.

[12] Production Tax Credit, as it is implemented by the IRS in the USA.

of the simulation, however, when a random assortment of power plants is created, there is a wide diversity of plant characteristics.

The major simplification I make in this stage of the model is, however, in the structural heterogeneity of the agents. Only two types of agents are considered, and they are mostly similar. This is an intentional design choice that more clearly isolates the dynamics I care about. Although the model provides useful conclusions in this simplified state – and so I argue it is complicated enough – there are undoubtedly other interesting dynamics that are excluded from the model due to the lack of, for example, peaker plants, intermittent technologies, fuel price variability, reliability, and the like.

The Energy Market

The energy market structure I model is based on hourly least-cost dispatch, with a number of simplifications to make the simulation more tractable. These include:

- No ramping constraints; each generator is assumed to have an infinite ramp capability;
- No partial commitments; each generator is either at base load or no load each hour;
- Vertical demand curve; each hour the demand is constant;
- No “strategic” offers; a generator is assumed to offer only at short-run marginal cost.

Thus the energy price each hour is easily determinable by crossing a cost-ordered dispatch stack against the vertical demand curve and finding the marginal supplier’s variable cost of production. In cases where the fixed demand exceeds all available supply, a Value of Lost Load (VOLL), set to an arbitrary large number, is used as the clearing price.

The supply stack in each year is composed of all the agents participating in the market that year. Without reliability considerations or fuel variability, the supply stack is identical from one hour to the next within the same year.

The hourly demand quantities are derived from PJM’s 2018 RTO load. The first 1000 hours of the year are taken as the full year for the simulation, and each hour is normalized such that the maximum demand in the year is equal to 10 GW. This scale combined with a 15% reserve requirement and an average generator size of 500 MW leads to an equilibrium number of ~23 generators to service the market. From year to year the exact same hourly demand series is repeated. The load distribution and a sample hour’s supply and demand curves are shown below.

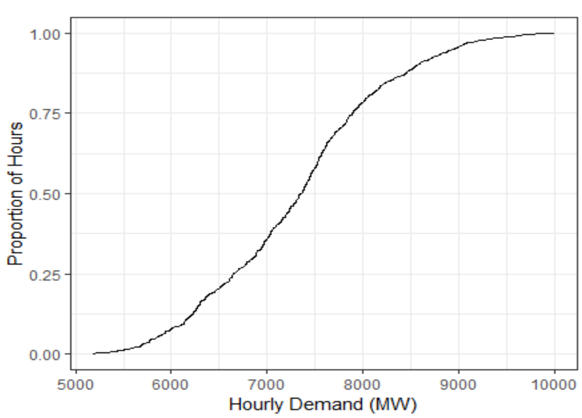

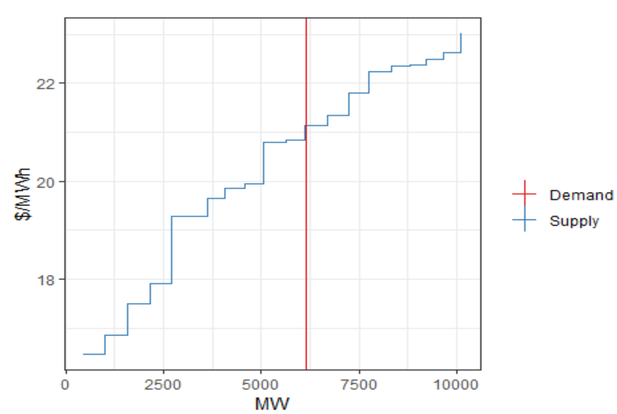

The Capacity Market

The competition between heterogeneous agents in the energy market above leads to revenue differences between them, and almost inevitably some of the agents will not be able to cover their Fixed Costs in a given year. A capacity market mechanism, in a simple sense, adds revenue to the market sufficient to keep total installed capacity above a certain level. As in the energy market I described above, the modelled capacity market is simplified. Because the point of this paper is to analyze specifically PJM's proposed RPM construct, I need to justify the simplifications more fully.

The first simplification is in the creation of the administrative "demand curve", or what PJM terms the Variable Resource Requirement curve. While in reality PJM arduously creates a piecewise function of Procured MW incorporating measures of Net CONE, energy market expectations, and target reserve levels, my model sets a simple vertical demand quantity at 115% of the maximum observed demand. I argue that this simplification is not meaningful to any conclusions I draw about the Ex-RCO and Ex-MOPR constructs; it is important only to have a consistent demand-side process that is able to successfully procure capacity.

The second simplification is the use of a 1-year forward market instead of the 3-year forward market that PJM uses to procure capacity. The 3-year window contributes to reduced revenue volatility and a longer planning horizon, both useful qualities when planning large capital investments, sourcing funding, managing politics, etc. My model does not suffer from a shortage of investable funds or lengthy construction project timetables, and so the benefits are irrelevant. Thinking of installed capacity as a dynamic system, a longer capacity windows dampens the system. As we will see, the simulation does not exhibit underdamped behavior, and so I argue that this simplification, too, is not meaningful.

The third simplification, as in the energy market, is the lack of creative generator offers. This is perhaps the most consequential shortcoming of my model, as the interaction of agents via their decision strategies is how the market-level dynamics unfold. My model assumes that agents offer into the capacity market "at cost", again as in the energy market, by solving for the recurring payment streams they need to break even over the long run. While this is a logical strategy – and one that displays intuitive dynamics over time, as I will show – it is not how all market participants behave. Drawing from my own experience in the merchant power sector, it is common to calculate a theoretical "cost" of the obligations that a capacity award confers, which cost is then used as an offer price. This cost can differ markedly from the break-even cost methodology my agents use; no doubt there exist other strategies as well. Further discussion of the agents' calculations is left to the next sub-section, and further discussion of the gaps in my model is left for the "Extensions" section.

Outside of these simplifications, the modelled market structure is true to PJM's proposed ex-RCO and ex-MOPR proposals. Subsidized resources are subjected to the MOPR, resources electing the RCO can reverse that decision in later years, the two round auction is faithfully represented, capacity awards entail must-offer obligations[13], and inframarginal rent is appropriately allocated. Despite the above simplifications, these structures constrain and shape agent behavior appropriately.

The Decision-Making Processes of Agents

The strategies of agents were described briefly above; here I will lay them out in more detail. Generally, an agent behaves rationally and individually, using an NPV-style approach calculated over the

[13] Although without any reliability concerns for the agents, there is no "bite" to these obligations.

remaining lifetime of the agent. Expectations about future conditions are derived from past years' occurrences weighted according to a recency bias.

At the beginning of a year, the agents who participate in the energy market are simply those who decided at the end of the last year to remain or to enter. In the special case of the first year, the agents who participate are randomly instantiated in sufficient number to satisfy the capacity reserve requirement. Thus, there is no special logic or decision-making at the beginning of an iteration.

During the energy market phase, each agent offers at its subsidized short-run marginal cost. Generally the subsidized short-run marginal cost of agent *i* at time *t* is,

$$SRMC_i(t) = VOM_i + Heat\ Rate_i * Fuel\ Cost(t) - PTC_i.$$

The fuel cost, though, is always zero, so each agent simply offers at its VOM less the PTC-style subsidies to which it is entitled.

During the capacity market phase, each agent generically offers at its expected long-run revenue shortfall, aiming to achieve breakeven at a minimum. This logic is realized differently for existing agents and for prospective agents (those offering unconstructed capacity into the market). A prospective agent conducts an NPV calculation, backsolving for the constant stream of capacity payments necessary, in conjunction with expected energy market revenues, to balance the plant construction cost over its operating lifetime. Thus a prospective agent *i* would solve for *CapRev* in the equation,

$$NPV_i = 0 = \frac{(Build_i - ITC_i + EEM_i + CapRev)}{(1+r)^1} + \sum_{t=2}^{Life_i} \left[\frac{(EEM_i + CapRev)}{(1+r)^t}\right],$$

where *CapRev* is the constant annual capacity payment, *EEM* is the *Expected Energy Margin* in forward years, and *r* is a discount rate[14]. The *EEM* for each agent is calculated based on the weighted average of its hypothetical dispatch in previous years' energy markets.[15] Earlier years are discounted increasingly as they retreat into the past. So for agent *i* calculating its *EEM* in year *k* for year *k+1*,

$$EEM_i = \frac{\sum_{t=k}^{k-LAG}\left[\frac{1}{(k-t+1)^{BIAS}} * (Rev_{i,t} - VOM_i + PTC_i)\right]}{\sum_{t=k}^{k-LAG} \frac{1}{(k-t+1)^{BIAS}}},$$

where LAG is a constant limiting the lookback period for calculating this weighted average, and BIAS is a constant indicating how much early years should be discounted compared to recent ones.[16]

An existing agent uses the same logic, but has a lower *CapRev* threshold since it suffers no build costs. MOPR-style considerations are easily handled in this framework by toggling on or off the ITC and PTC contributions as the situations warrant. (These Boolean switches are not represented in the above equations.)

The option to elect RCO status introduces a small intricacy to the capacity market logic, since that status must be elected prior to the capacity market clearing. In order to make this election, each

[14] The discount rate is held constant across all agents at 8% in my simulations.
[15] When calculating the hypothetical dispatch, an agent assumes it is a price taker. It does not attempt to "re-dispatch" previous years.
[16] Like the discount rate, the LAG and BIAS constants are assumed the same for all agents, at 5 years and 2 respectively.

agent calculates an expected capacity clearing price *ECC*, akin to EEM_i, for the forthcoming auction.[17] If the agent expects not to be able to clear the coming auction (perhaps it is subsidized and cannot compete under the modelled MOPR rules), it will consider electing RCO status.

This decision to elect RCO hinges on an important variable, what I am calling the "RCO Passback". Under the proposed ex-RCO construct, PJM rebates to states the capacity payments for load that is associated with a carved-out resource. The carved-out resource does not receive any capacity revenue from PJM, then, but relies on its state to arrange, presumably, a bilateral capacity contract. The capacity payments from this contract are the RCO Passback: the state is "passing back" capacity revenues to the subsidized, carved-out resource. At issue for the economic decision to elect RCO is how this RCO Passback compares to the PJM RPM capacity clear. If, for example, the RCO Passback = 100% of the clear, i.e. the state makes a resource whole to the unsubsidized capacity clearing price, there will be small incentive *not* to elect RCO. This variable is explored in this paper's simulations.

Finally, after the capacity market is cleared, existing agents must decide whether to remain in or to exit the market, and potential agents must decide whether or not to build. For agents who cleared the capacity market or elected RCO the choice is simple: they remain or enter, both because it is economical for them and because they are obligated. Agents who did not clear the capacity market, however, may still choose to remain or enter the market if they believe they can subsist on the energy market alone. For several reasons, that does not occur in this model.[18] Agents who were "intermarginal", i.e. they cleared the second round but not the first round of the ex-RCO RPM, have a similar choice but benefit from the inframarginal rent to defray some costs.

Results: Three Simulations and Their Dynamics

I now present the results of three simulations showing possible dynamics of the resource composition in PJM as a result of the ex-RCO and ex-MOPR proposal, if implemented. Each simulation represents a fully implemented market with competition between unsubsidized and subsidized resources as described in the preceding section. The difference between simulations is only the RCO Passback rate, which varies from 50% to 90% of the full capacity clearing price and which I have identified as the key tuning parameter. (Changing other parameters, e.g. the ITC/PTC schedules or introducing technology learning, would change the dynamics, but in less relevant ways. These will be discussed in the final section.)

Full Simulation with RCO Passback Set Low, at 50%

There are three principal charts I use to show the results of the simulation: Figure 1, which shows the supply stack evolution over time with a focus on technology type; Figure 2, which shows the

[17] The *ECC* uses the same weighting system as the EEM_i, but it is identical for all agents since capacity clears are identical for all agents. Thus the formula for *ECC* is more simple:

$$ECC = \frac{\sum_{t=k}^{k-LAG}\left[\frac{1}{(k-t+1)^{BIAS}}*(CC_t)\right]}{\sum_{t=k}^{k-LAG}\frac{1}{(k-t+1)^{BIAS}}},$$

where CC_t is an actual historical capacity clear.

[18] Some structural reasons that might make this choice attractive, e.g. extremely volatile energy prices due to power congestion, are not considered at all in this model, and others, e.g. magnificent subsidies to just a few resources promising outlandish energy margins, are not instantiated.

entering and exiting of generators on a per-plant basis; Figure 3, which shows the prices that drive the investment decisions; and Figure 4, which shows how the agent attributes evolve over time under competition. (The captions to the charts provide more explanation for what is pictured.) Together, these charts tell the story of the simulation.

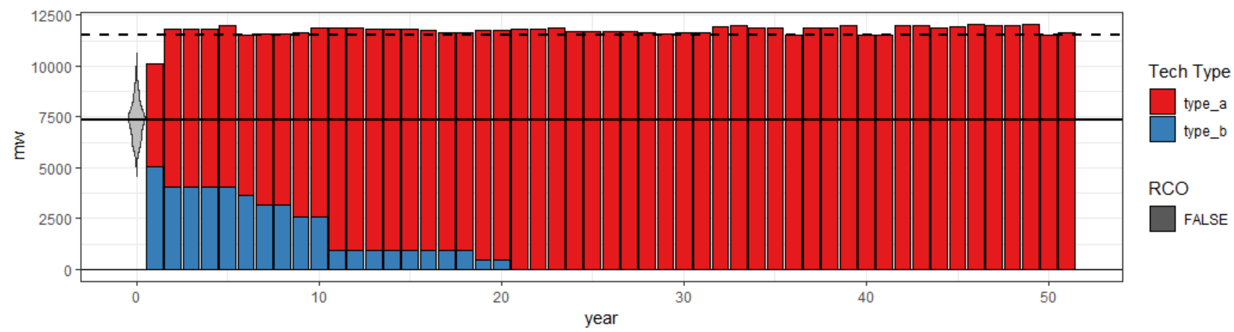

Figure 1: Supply & Demand. The grey violin at Year 0 indicates the hourly demand distribution, which is the same for every year. The solid horizontal line indicates the mean of this distribution, and the dashed represents the maximum demand multiplied by the reserve requirement. The stacked bars represent the total installed capacity each year, colored by technology type. Any resources that have elected RCO treatment are greyed out slightly.

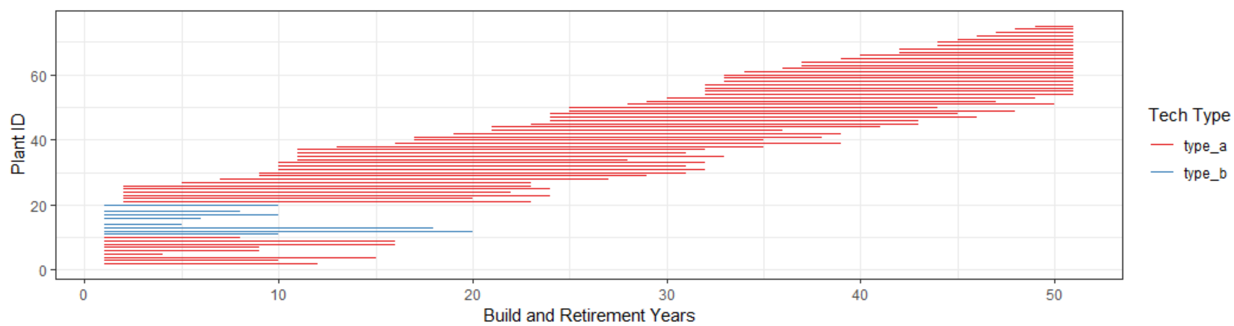

Figure 2: Agent Lifecycle. Each horizontal line indicates the lifespan of one agent, beginning at its first year in the energy market, and ending on its last year. Each line is colored according to its technology type. The line will end either due to the end of the operating lifetime, or due to an investment decision to avoid continued losses.

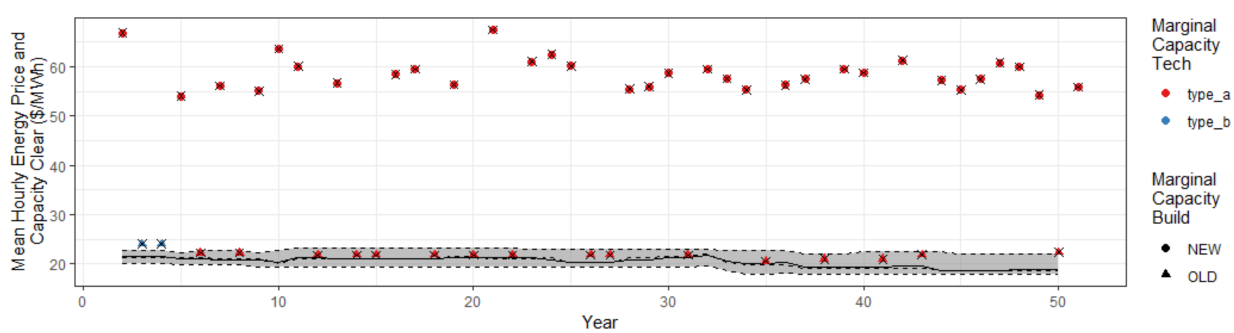

Figure 3: Prices. The shaded band represents the range of hourly energy prices that emerge each year, and the solid line running through its middle represents the median hourly energy price each year. The dots represent capacity clearing prices: they are colored according to the marginal technology and shaped according to whether the marginal agent was existing or prospective. Grey Xs are overlaid to show the Round 1 clear price, although in this first example it never separates from the Round 2 clear.

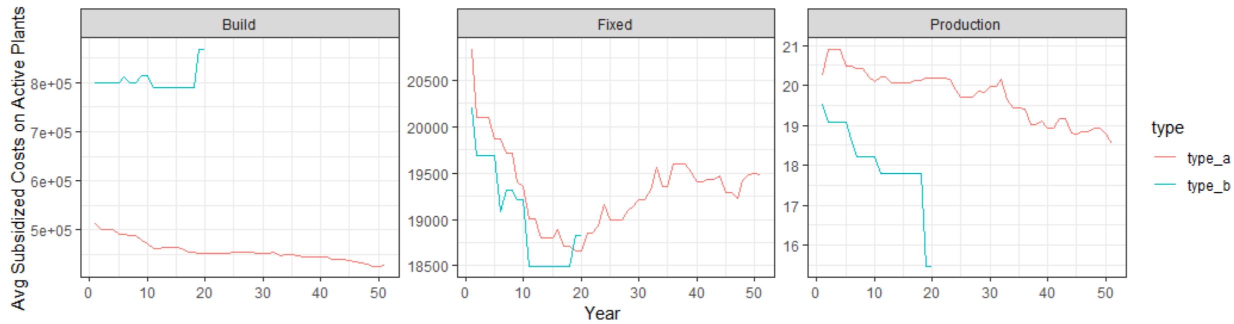

Figure 4: Attribute Evolution. Each line indicates how the attributes for each technology change over time by averaging those agents existing each year. The Build Costs and Production Costs (VOM) include the effects of ITC and PTC subsidies, respectively.

What is immediately apparent from Figure 1 is that the Type A technology (which is unsubsidized) comes to completely dominate the Type B technology by year 20, despite the impressive subsidies available to Type B technologies. Figure 2 shows, in fact, that the Type B generators exist only because they were randomly seeded at the start of the simulation; after the first year, no more Type B plants are built. Figure 3 shows that energy and capacity markets were well behaved over this simulation: energy prices were steady and decreasing over 50 years, and the capacity market oscillated between a high clear sufficient to incentivize Type A new builds (in years when plants retired) and a low clear sufficient to keep existing Type A plants in the market (in years when nothing retired). The Round 1 and Round 2 clearing prices never separated. Figure 4 gives the reason for the steady, declining energy prices: as time passed, only those Type A generators with the lowest build and production costs were selected for construction.

What occurred here is simple enough to see: the Expanded MOPR made the expensive but subsidized Type B technology uncompetitive in the normal capacity market, so it wouldn't clear as a normal resource. The low RCO Passback rate offered by the state, though, made even RCO status for potential Type B technologies uneconomic. Stepping into the NPV calculation, the expected energy margin plus 50% of the expected capacity clear was insufficient to incentivize Type B technology construction.

We can evaluate this situation from the perspective of three stakeholders: the competitive merchant generation, the state policymaker with an RPS agenda, and the market regulator. Everyone here is happy except the state policymaker: the merchant generators have, basically, a market undistorted by Type B technologies, and the market regulator has a steady market with functioning investment signals and well matched supply and demand. The state policymaker, however, has no Type B generators to show for the impressive subsidies already offered. Of course, the state should just raise the RCO Passback a bit to fix this situation, so it is perhaps, for the policymaker, an unrealistic equilibrium. We should not forget, however, that each subsidy offered (which includes, as a sort of capacity subsidy, the RCO Passback) entails a political battle.

Full Simulation with RCO Passback Set High, at 90%

We now turn to a situation in which the state offers an extremely generous RCO Passback rate, at 90% of the normal clearing price. This turns out to be a very disruptive policy:

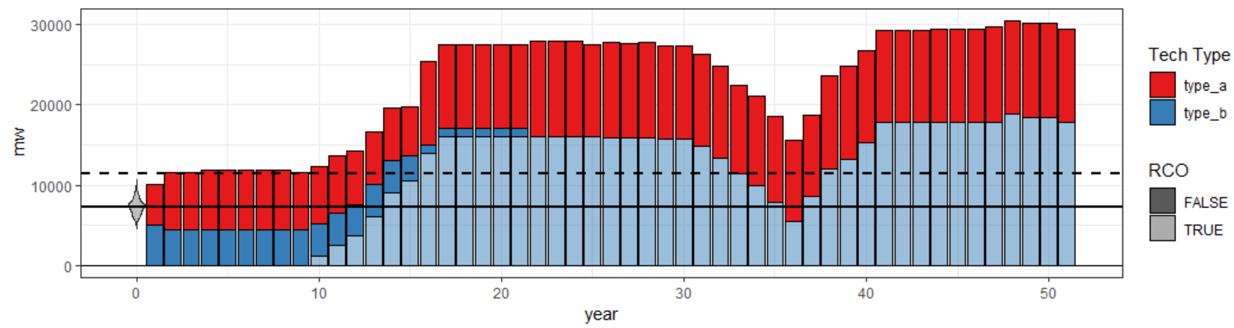

Figure 1: Supply & Demand.

Figure 2: Agent Lifecycle.

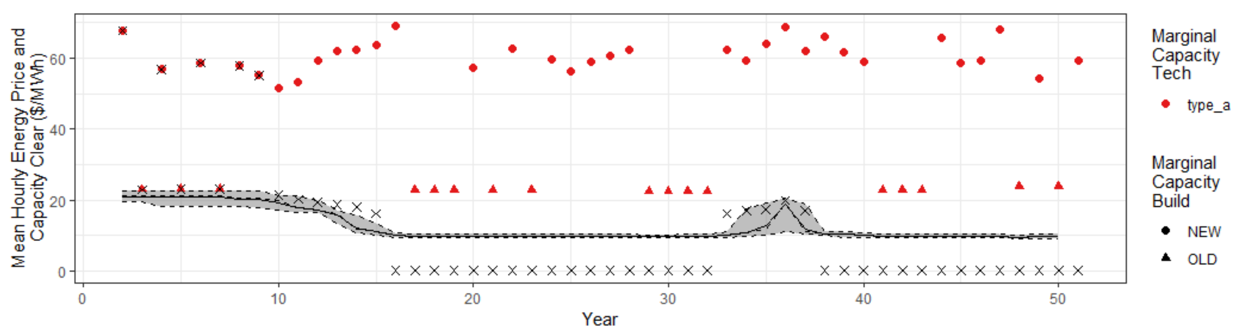

Figure 3: Prices.

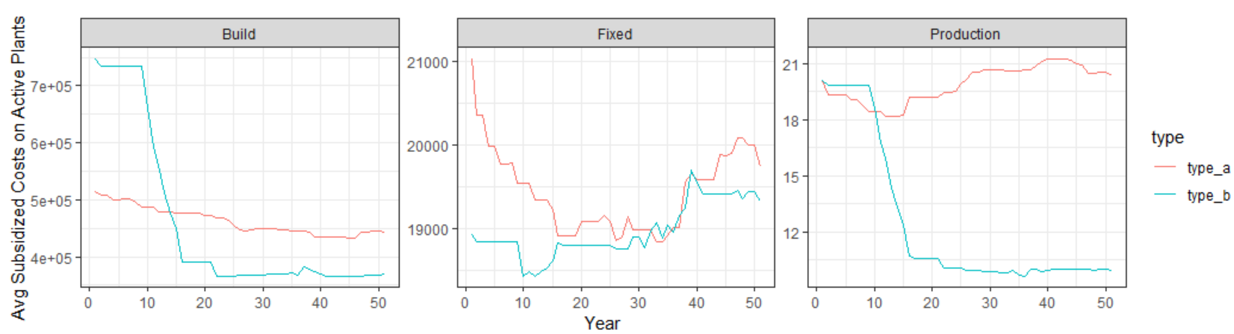

Figure 4: Attribute Evolution.

Figure 1 tells a strikingly different story than before. Although Type B initially loses market share, at year 10 Type B installed capacity balloons up towards the full reserve requirement as RCO-elected capacity and stays there throughout the rest of the simulation. Equally as interesting, Type A never loses its market share and instead also remains at the full reserve requirement: the market is seemingly procuring double capacity, and the total installed capacity exhibits a strange dislocation around year 35. Figure 3 shows the effect of Type B's entrance in the market: energy prices drop quickly towards the $10/MWh subsidized VOM cost, and Round 1 of the capacity auction separates drastically from Round 2 after year 15: Round 1 prices actually set at $0/MW for the majority of the simulation. The Round 2 capacity prices are higher than in the previous simulation, since the Type A resources are no longer getting any revenue from the energy market. Figure 4 shows similar competitive pressures as before, but the advantage of Type B technology is clearer.

The story here is of a completely broken market. By electing RCO, subsidized resources have access to 90% of the full capacity clear and are a no-brainer to build up until they completely saturate the energy market. (Though the energy market is saturated by Type B resources by year 15, Type B resources continue to enter the market as RCO for another few years given their lookback logic.) At that level of construction, Round 1 of the capacity market prices at $0/MW, since the entire demand can be satisfied with RCO resources. This means that the inframarginal rent created in Round 2 is actually sufficient to promote entry of new unsubsidized resources to maintain a duplicate set of capacity resources. The perturbation around year 35 occurs as Type B resources begin to retire (see Figure 2): there is no build signal until the energy market desaturates somewhat, at which point more Type B RCO resources are constructed.

This scenario of an overly-generous RCO Passback rate makes only the state policymakers happy: they finally get their preferred resource composition after years of hefty subsidies. The unsubsidized power plants, however, are almost completely blocked out of the energy market and now make only a break-even amount through the capacity market. The regulators are faced with a clearly inefficient market that procures entirely too much supply, and with higher headline capacity clears as well. This situation, unfortunately, seems more likely than the previous given that it is entirely within the power of policymakers to create and it is aligned with their interests to do so. (One imagines there would be a political emergency between years 10 and 15, however.)

Full Simulation with RCO Passback Set Medium, at 70%

The two extremes having been explored (happy policymakers and unhappy regulators on one side, vice versa on the other), it is possible to find a happy medium. By setting the RCO Passback rate to 70% we see an intermediate scenario:

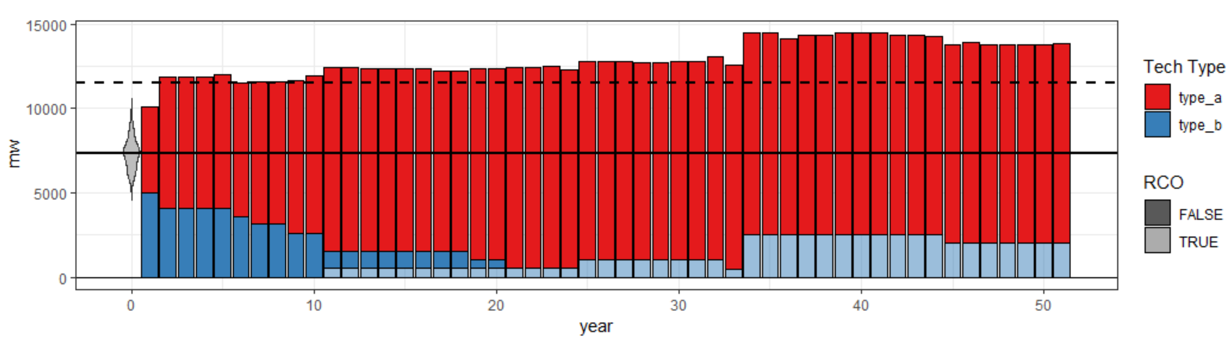

Figure 1: Supply & Demand.

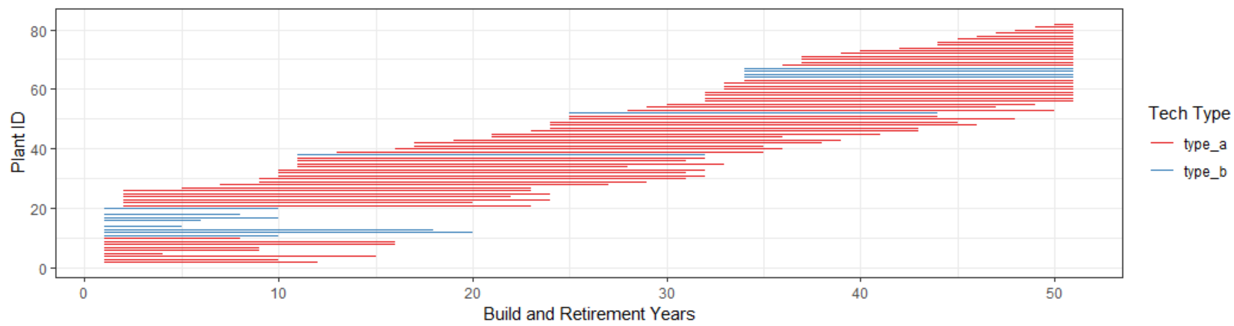

Figure 2: Agent Lifecycle.

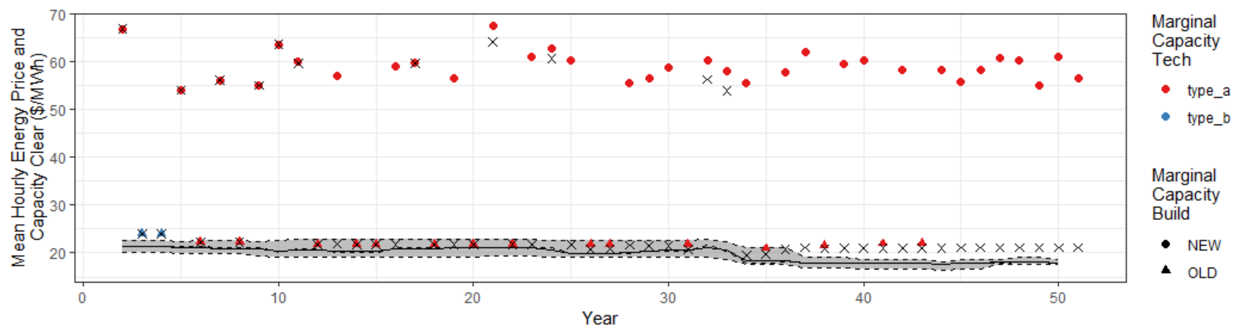

Figure 3: Prices.

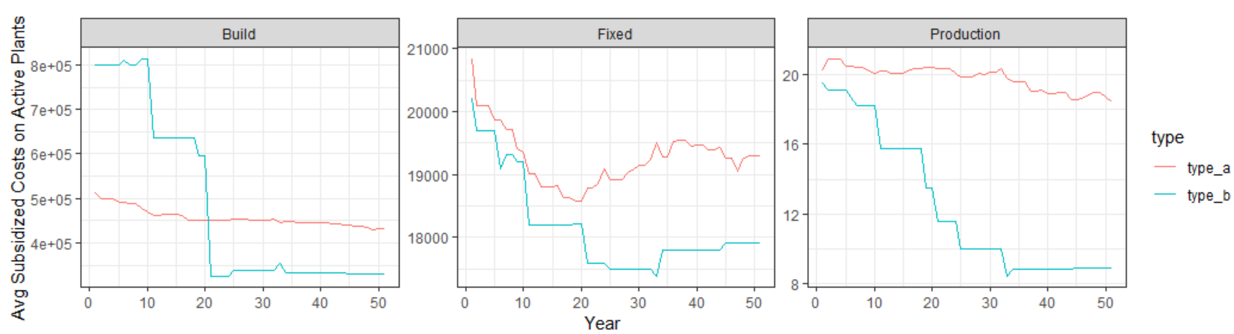

Figure 4: Attribute Evolution.

This situation looks reasonable, and it could be described as a successful compromise between state policy and competitive market interests. The outcome, however, is unlikely to be stable. The choice of RCO Passback rate to make this situation feasible is entirely dependent upon the cost structure and variability of the subsidized resources. The RCO Passback rate has to incentivize just a few eligible resources to select the RCO option, but not too many. In reality, it is unlikely this balance could be held for long, before technology learning, subsidy schedules, or something else upset the equilibrium.

Conclusions and Extensions

This paper shows that the market conditions sought by PJM – competitive prices with the participation of state-sponsored resources – exist in an unstable equilibrium. States will tend to enact policies with subsidies sufficient to materialize their preferred generation portfolios. This means increasing ITC-style, PTC-style, or "RCO Passback"-style subsidies until the modelled regime moves from Simulation 1 to Simulation 3. It is easy to overshoot the mark and land in Simulation 2, however, with vast overbuild, especially if other dynamics such as technology learning are taken into account. Though the world will exist for a time, it seems, around Simulation 3 if the ex-RCO and ex-MOPR are implemented, it will not be for long.

Easy fixes are not easy to come by and are essentially political. The “cap” previously mentioned due to LS Power would prevent the sort of cascade seen in Simulation 2 above, but it raises again the question of a market structure that directly impedes state policy interests. I suggest an avenue for exploration may be a recalculation of PJM’s VRR curve from Round 1 to Round 2 of the RPM auction. The proximate cause of the capacity overbuild seen in Simulation 2 is a Round 2 demand that does not account for previously procured RCO capacity. If this demand quantity were reduced, there would be less overbuild, though price suppression would crop up. Another avenue may be discounts to the inframarginal rent paid out to intermarginal agents, since this was the wealth transfer means by which new, unsubsidized capacity entered in Simulation 2. This solution, however, may be unacceptable as discriminatory. Although any solution is really just a tradeoff between interests, it is important to prioritize the overall sensibility of the eventual outcome. The outcome in Simulation 2 is not sensible and should be protected against.

There are other sources of instability that this paper did not investigate because they are not exclusively linked to PJM’s reform proposals. For example, resources so subsidized by an ITC or PTC that they can subsist on energy market revenues alone would cause a similar over-capacity problem but for entirely different reasons. (There is no MOPR style rule in the energy markets.)

The model presented in this paper is necessarily and intentionally limited in many ways. Still, there are specific ways the analysis could be usefully extended. These include:

- Expanding the scope to a two-region model. This model only considers a single capacity zone, but considering multiple zones linked by limited transmission elements would take a step closer to reality and perhaps show other instabilities.
- Comparing PJM’s recommendations with ISO-NE’s own proposed two-part capacity model.
- Expanding upon the technical components of the ABM. Many ABMs allow learning by agents between iterations, often through a genetic algorithm. Parameters like the discount rate and LAG period could be tuned in this way. More ambitiously and more consequentially, alternative strategies for agents could be explored and implemented, so that not everyone is “playing the same game”.
- Including reliability considerations. As alluded to above, performance requirements and reliability risks are integral to the calculus of real offer strategies in capacity markets like PJM’s. Specifically as it pertains to the “intermarginal” agents, who receive some payment but no capacity obligation, this feature could meaningful change the market-level dynamics.

Appendix A: Simulation Code

The code, written in the R language, to run this simulation is copied below, along with the settings used to generate the “medium RCO” simulation shown in this paper. To generate the others, only the “CARVED.OUT.CAP.FACTOR” variable need be changed.

```
# Some functions
generatePortfolio <- function(dt.tech, num.seed, spread.factor) {
  if(!is.numeric(num.seed)) stop()
  if(!is.numeric(spread.factor)) stop()
  if(!num.seed > 0) stop()
  if(!spread.factor %between% c(0, .99)) stop()

  dt.agents <- rbindlist(lapply(X = dt.tech$type, FUN = function(this.type) {
    dt.tech[type == this.type, list(type,
                                    mw_max = capacity*runif(num.seed, 1-spread.factor, 1+spread.factor),
                                    build = `build_cost`*runif(num.seed, 1-spread.factor, 1+spread.factor),
                                    hr = `heat_rate`*runif(num.seed, 1-spread.factor, 1+spread.factor),
                                    vom = `vom_cost`*runif(num.seed, 1-spread.factor, 1+spread.factor),
                                    fixed = `fixed_cost`*runif(num.seed, 1-spread.factor, 1+spread.factor),
                                    life = round(`life`*runif(num.seed, 1-spread.factor, 1+spread.factor)),
                                    age = 0, itc = 0, ptc = 0, rco = NA)]
  }))
  dt.agents[, `:=`(agent = 1:.N)]
  return(dt.agents)
}
calcEnergyMargin <- function(dt.pricing, dt.fuel, lst.agent) {
  if(lst.agent[["hr"]] == 0) {
    return(dt.pricing[, list(hour, gm = round(price - lst.agent[["vom"]], 2))])
  } else {
    dt.pricing <- merge(dt.pricing, dt.fuel, by = "hour")
    dt.pricing[, fuel := ifelse(lst.agent[["type"]] %in% c("ocgt", "ccgt"), hh,
                                ifelse(lst.agent[["type"]] == "nuclear", u,
                                       ifelse(lst.agent[["type"]] == "coal", coal, 0)))]
    dt.pricing[, gm := round(price -  fuel*lst.agent[["hr"]] - lst.agent[["vom"]], 2)]
    return(dt.pricing[, list(hour, gm)])
  }
}
calcPriceTakerHistoricalGm <- function(dt.agents, dt.hist.pricing, dt.hist.fuel.pricing, max.year, max.lag) {
  these.years <- seq(from = max.year, to = max(max.year-max.lag, 1), by = -1)
  lst.gm <- mapply(SIMPLIFY = F,  FUN = function(i, num.year) {
    dt.gm <- calcEnergyMargin(dt.pricing = dt.hist.pricing[year == num.year],
                                dt.fuel = dt.hist.fuel.pricing,
                                lst.agent = as.list(dt.agents[agent == i]))
    dt.gm <- dt.gm[, list(agent = i, year = num.year, gm, h = ifelse(gm + dt.agents[agent == i, ptc] >= 0, 1, 0))]
    dt.gm[, `:=`(gm = h*gm, credit = h*dt.agents[agent == i, ptc])]
    dt.gm[, list(gm = sum(gm), h = sum(h), credit = sum(credit)), by = list(agent, year)]
  },
  i = rep(dt.agents$agent, each = length(these.years)),
  num.year = rep(these.years, times = length(dt.agents$agent)))
  dt.gm <- rbindlist(lst.gm)
  return(dt.gm)
}
calcWavgExpNm <- function(dt.agents, dt.gm, recency.bias, this.year) {
  dt.agents <- merge(dt.agents, dt.gm, by = "agent")
  dt.agents[, weighting := (1/(this.year-year+1))^recency.bias]
  dt.wavg.gm <- dt.agents[, list(exp_gm = weighted.mean(gm, weighting),
                                 exp_h = weighted.mean(h, weighting)), by = list(agent, fixed)]
  dt.wavg.gm[, exp_nm := exp_gm-fixed]
  dt.wavg.gm[, fixed := NULL]
  return(dt.wavg.gm[])
}
calcNpv <- function(dt.agents, discount.rate, include.build, exclude.credits) {
  stopifnot(c("agent", "exp_nm", "build", "life", "age") %in% names(dt.agents),
            is.numeric(discount.rate))

  lst.npv <- lapply(X = dt.agents$agent, FUN = function(i) {
    exp.nm <- sapply(1:dt.agents[agent == i, life-age], FUN = function(n) {
      if(exclude.credits) {
        dt.agents[agent == i, exp_nm]/(1+discount.rate)^n
      } else {
        dt.agents[agent == i, exp_nm + exp_credit]/(1+discount.rate)^n
      }
    })
    if(include.build) {
      if(exclude.credits) {
        npv <- sum(exp.nm) - dt.agents[agent == i, build]/(1+discount.rate)^1
      } else {
        npv <- sum(exp.nm) - dt.agents[agent == i, build - itc]/(1+discount.rate)^1
      }
    } else {
      npv <- sum(exp.nm)
    }
    data.table(agent = i, npv = npv)
  })

  dt.npv <- rbindlist(lst.npv)
  return(dt.npv)
}
```

```
calcMmAnnuity <- function(i, discount.rate, dt.info) {
  if(dt.info[i, npv] > 0) {
    return(0)
  } else {
    return(uniroot(interval = c(0, 1e6), tol = 1, f = function(cap.stream, nrg.npv, lifetime) {
      sum(sapply(1:lifetime, FUN = function(n) cap.stream/(1+discount.rate)^n)) + nrg.npv
    }, nrg.npv = dt.info[i, npv], lifetime = dt.info[i, life-age])$root[1])
  }
}
setItcPtc <- function(dt.agents, itc.rate, ptc.rate, this.year, include.sub.tech, itc.max = .5, ptc.max = .5) {
  if(include.sub.tech) {
    dt.agents[type == WHICH.SUB, c("itc", "ptc") := list(pmin(build*itc.max, itc.rate*this.year),
                                                          pmin(vom*ptc.max, ptc.rate*this.year))]
    return(dt.agents)
  } else {
    return(dt.agents)
  }
}
calcMmAndStats <- function(dt.agents, dt.gm, recency.bias, discount.rate, enable.mopr, dt.past.agent.clears, this.year) {
  # Reduce to an expected Net Margin for future years (per MW), then NPV it
  dt.agents <- merge(dt.agents, calcWavgExpNm(dt.agents, dt.gm, recency.bias, this.year), by = "agent")
  dt.agents[, exp_credit := exp_h*ptc]
  # For MOPR, plants never having cleared before must continue to offer in with build costs
  # (New plants must always offer in with build costs)
  # (if it's year 1, assume everyone has already cleared)
  if(enable.mopr & this.year != 1) {
    never.cleared <- dt.past.agent.clears[, list(total_clear = sum(ifelse(is.na(clear), 0, 1))), by = agent]
    never.cleared <- never.cleared[total_clear == 0, agent]
  } else {
    never.cleared <- 0
  }
  dt.npv1 <- calcNpv(dt.agents[!agent %in% never.cleared & age != 0], discount.rate,
                     include.build = F, exclude.credits = enable.mopr)
  dt.npv2 <- calcNpv(dt.agents[ agent %in% never.cleared | age == 0], discount.rate,
                     include.build = T, exclude.credits = enable.mopr)
  dt.npv <- rbind(dt.npv1, dt.npv2)
  dt.agents <- merge(dt.agents, dt.npv, by = "agent", all.x = T)
  vec.mm <- sapply(1:nrow(dt.npv),
                   FUN = calcMmAnnuity,
                   discount.rate = discount.rate,
                   dt.info = dt.agents[, list(agent, life, age, npv)])
  dt.agents[, mm := vec.mm]

  return(dt.agents[])
}

# Main function
runSim <- function(lst.prev = list(), start.year = 1, through.year = THROUGH.YEAR) {

  #####
  # Intialize sim
  if(length(lst.prev) == 0) {
    dt.saved.agent.bio <- data.table()
    dt.saved.prices.nrg <- data.table()
    dt.saved.prices.cap <- data.table()
    dt.saved.mkt.nrg <- data.table()
    dt.saved.mkt.cap <- data.table()
    dt.saved.invest <- data.table()
    dt.agents.active <- data.table()
    cap.clear <- 0
  } else {
    dt.saved.agent.bio <- lst.prev[[1]]
    dt.saved.prices.nrg <- lst.prev[[2]]
    dt.saved.prices.cap <- lst.prev[[3]]
    dt.saved.mkt.nrg <- lst.prev[[4]]
    dt.saved.mkt.cap <- lst.prev[[5]]
    dt.saved.invest <- lst.prev[[6]]
    dt.agents.active <- lst.prev[[7]]
    cap.clear <- dt.saved.mkt.cap[year == max(year), clear][1]
  }

  for(this.year in start.year:through.year) {
    cat("\nYear", this.year)
    # this.year <- 1

    if(nrow(dt.agents.active) == 0) {
      # In first iteration, or if everyone went bust, start with a large portfolio
      if(this.year == 1) {
        dt.agents.active <- generatePortfolio(dt.tech, NUM.SEED, SPREAD.FACTOR)
        # Randomize the age of plants we just created, bounded between one and full life less one
        dt.agents.active[, age := pmin(pmax(1, round(runif(n = .N)*life)), life-1)]
        dt.agents.active <- setItcPtc(copy(dt.agents.active), ITC.RATE, PTC.RATE, this.year, INCLUDE.SUB.TECH)

        dt.saved.mkt.cap <- data.table(year = this.year, agent = dt.agents.active$agent, mm = NA, clear = NA, rco = NA)
        dt.saved.agent.bio <- unique(rbind(dt.saved.agent.bio, dt.agents.active[
          , list(agent, type, mw_max, build, hr, vom, fixed, life, itc, ptc)]))
      } else {
        dt.agents.active <- generatePortfolio(dt.tech, ceiling(NUM.SEED/3), SPREAD.FACTOR)
        dt.agents.active <- setItcPtc(copy(dt.agents.active), ITC.RATE, PTC.RATE, this.year, INCLUDE.SUB.TECH)
        dt.agents.active[, agent := (max(dt.saved.agent.bio$agent)+1):(max(dt.saved.agent.bio$agent)+.N)]
        dt.saved.agent.bio <- unique(rbind(dt.saved.agent.bio, dt.agents.active[
          , list(agent, type, mw_max, build, hr, vom, fixed, life, itc, ptc)]))
      }
    }
```

```
##########
# START  : Operational year sim
##########

#####
# Create energy supply curve (for each hour if using gas plants)
dt.supply <- dt.agents.active[, list(agent, type, mw = mw_max, hr, vom, ptc)]
if(INCLUDE.INT.TECH) dt.supply[type == WHICH.INT, mw := ifelse(runif(.N) <= INTERMITTENCY, mw, 1)]
if(USE.ABSTRACT.TECH) {
  dt.supply[, cost := vom - ptc]
  dt.supply <- dt.supply[order(cost)][, mw_cum := cumsum(mw)][]
  dt.supply <- as.data.table(merge.data.frame(data.table(hour = 1:HOUR.CUTOFF), dt.supply, all = T))
} else {
  dt.supply <- merge(dt.fuel,
                     as.data.table(merge.data.frame(data.table(hour = 1:HOUR.CUTOFF), dt.supply, all = T)),
                     by = "hour")
  dt.supply[, cost := hr*ifelse(type %in% c("ocgt", "ccgt"), hh,
                                ifelse(type == "nuclear", u,
                                       ifelse(type == "coal", coal, 0))) + vom - ptc]
  dt.supply <- dt.supply[order(hour, cost)][, mw_cum := cumsum(mw), by = hour][]
}
# ggplot() + theme_bw() + labs(x = "MW", y = "$/MWh") +
#   geom_step(data = dt.supply[hour == 1], mapping = aes(x = mw_cum, y = cost, color = "Supply")) +
#   geom_vline(data = dt.demand[hour == 1], mapping = aes(xintercept = mw, color = "Demand")) +
#   scale_color_brewer("", palette = 6, type = "qual")

#####
# Solve for hourly energy pricing, and record the marginal agent
dt.ds <- merge(dt.demand, dt.supply, by = "hour", suffixes = c("_d", "_s"), all = T)
dt.ds <- dt.ds[, list(hour, agent, excess = mw_cum - mw_d, cost)]
dt.ds <- dt.ds[excess > 0, .SD[excess == min(excess), list(cost, agent)], by = hour]
dt.prices <- merge(dt.ds, data.table(hour = 1:HOUR.CUTOFF), by = "hour", all.y = T)[is.na(cost), cost := VOLL][
  , list(year = this.year, hour, price = cost, ma = agent)]
# ggplot(dt.prices) + geom_line(mapping = aes(x = hour, y = price))

#####
# Calc GM for each agent using prices we just found
dt.gm <- calcPriceTakerHistoricalGm(dt.agents = dt.agents.active,
                                    dt.hist.pricing = dt.prices,
                                    dt.hist.fuel.pricing = dt.fuel,
                                    this.year, max.lag = 0)
if(INCLUDE.INT.TECH) dt.gm[agent %in% dt.saved.agent.bio[type == WHICH.INT, agent],
                           `:=`(gm = gm*INTERMITTENCY, h = h*INTERMITTENCY, credit = credit*INTERMITTENCY)]

###
dt.saved.prices.nrg <- rbind(dt.saved.prices.nrg, dt.prices)
dt.saved.mkt.nrg <- rbind(dt.saved.mkt.nrg, dt.gm)

###
dt.agents.active[, age := age + 1] # 1 fewer operational year left for the facilities
# Remove end-of-life gens; they don't participate in the auction for next year
dt.agents.future <- dt.agents.active[age != life]

##########
# FINISH : Operational year sim
# START  : Capacity/ICAP sim (still run when w/o cap mkt, just assumes clear is $0)
##########

# Identify potential new builds. These participate alongside existing plants in cap auction
dt.potential.builds <- generatePortfolio(dt.tech, NUM.POSSIBLE.INVEST, SPREAD.FACTOR)
dt.potential.builds <- setItcPtc(copy(dt.potential.builds), ITC.RATE, PTC.RATE, this.year, INCLUDE.SUB.TECH)
# relabel so unique from existing agents
potential.agents <- (max(dt.saved.agent.bio$agent)+1):(max(dt.saved.agent.bio$agent)+nrow(dt.potential.builds))
dt.potential.builds[, agent := potential.agents]

###
# Determine the "missing money" for old/new plants
# MM def: constant annual revenue needed to set NPV to zero, given Net Margin expectations from energy market
dt.players <- rbind(dt.potential.builds, dt.agents.future[, names(dt.potential.builds), with = F])
# Calc GM for each potential plant as price taker, over all available years (max 5 years)
dt.gm <- calcPriceTakerHistoricalGm(dt.agents = dt.players,
                                    dt.hist.pricing = dt.saved.prices.nrg,
                                    dt.hist.fuel.pricing = dt.fuel,
                                    max.year = this.year, max.lag = MAX.LAG)
if(INCLUDE.INT.TECH) dt.gm[agent %in% dt.players[type == WHICH.INT, agent],
                           `:=`(gm = gm*INTERMITTENCY, h = h*INTERMITTENCY, credit = credit*INTERMITTENCY)]
# Calc the MM - GM gets a WAvg inside this func
# The MOPR setting means that plants that have never cleared before have to offer in unsubsidized and w/ build costs
dt.players <- calcMmAndStats(dt.agents = dt.players, dt.gm = dt.gm,
                             RECENCY.BIAS, DISCOUNT.RATE, ENABLE.MOPR,
                             dt.past.agent.clears = dt.saved.mkt.cap[year == max(year), list(agent, clear)],
                             this.year)

if(NUM.NEW.INVEST != 0) {
  # Choose top N new investments to move forward into capacity market with existing
  allowed.builds <- dt.players[agent %in% potential.agents][order(mm/mw_max, -npv/mw_max, decreasing = F)][1:NUM.NEW.INVEST,
agent]
  allowed.builds <- allowed.builds[!is.na(allowed.builds)]
  dt.players <- dt.players[(agent %in% allowed.builds & agent %in% potential.agents) | !(agent %in% potential.agents)]
}

###
# Run the capacity market
```

```
    # Calc an exponentially smoothed capacity expectation
    if(nrow(dt.saved.prices.cap) == 0) {
      exp.cap.clear <- cap.clear
    } else {
      dt.contemplated.clears <- dt.saved.prices.cap[year >= (this.year - MAX.LAG)]
      exp.cap.clear <- dt.saved.prices.cap[, weighted.mean(r2_clear, (1/(this.year-year+1))^RECENCY.BIAS)]
    }

    if(ENABLE.CAP.MKT) {

      # Create demand curve
      quantity.demanded <- max(dt.demand$mw)*RESERVE.RATIO

      if(ENABLE.EX.RCO) {
        ### Round 1

        # New resources will choose to RCO if:
        # # They have subsidies
        # # They would NOT have cleared this year's auction (so expect not to clear next one)
        # # They will have positive NPV with last year's clear * ADJ
        # Existing RCO resources may stay so
        dt.potential.rco <- dt.players[(age == 0 & ((itc > 0 | ptc > 0 ) & mm > exp.cap.clear)) | (age > 0 & !is.na(rco))]

        if(nrow(dt.potential.rco) > 0) {
          # Recalc NPV and MM with ITC and PTC credit
          dt.npv <- rbind(calcNpv(dt.potential.rco[age == 0], DISCOUNT.RATE, include.build = T, exclude.credit = F),
                          calcNpv(dt.potential.rco[age != 0], DISCOUNT.RATE, include.build = F, exclude.credit = F))
          dt.potential.rco <- merge(dt.potential.rco[, npv := NULL], dt.npv, by = "agent", all.x = T)
          vec.mm <- sapply(1:nrow(dt.npv),
                           FUN = calcMmAnnuity,
                           discount.rate = DISCOUNT.RATE,
                           dt.info = dt.potential.rco[, list(agent, life, age, npv)])
          dt.potential.rco[, mm := vec.mm]
          # See if the RCO rebate from state will incentivize subsidized MMo to enter/stay
          # Previously RCO'd resources electing NOT to continue to RCO will take a shot in the normal auction (with MOPR)
          dt.rco <- dt.potential.rco[mm <= exp.cap.clear*CARVED.OUT.CAP.FACTOR]
          # Thos who choose RCO enter the first stage of auction with $0 offers
          dt.rco[, mm := 0]

          dt.players <- dt.players[!agent %in% dt.rco$agent]

          dt.supply <- rbind(dt.rco[, list(agent, type, mw = mw_max, mm)],
                             dt.players[, list(agent, type, mw = mw_max, mm)])
        } else {
          dt.rco <- copy(dt.potential.rco)
          dt.supply <- dt.players[, list(agent, type, mw = mw_max, mm)]
        }
        if(INCLUDE.INT.TECH) dt.supply[type == WHICH.INT, mw := mw*CAP.VALUE.INT]
        dt.supply <- dt.supply[order(mm)][, mw_cum := cumsum(mw)][]

        # Solve for marginal price
        dt.ds <- dt.supply[, list(agent, mm, mw_cum, mw_d = quantity.demanded)]
        dt.ds <- dt.ds[, excess := mw_cum - mw_d]
        dt.rco.results <- dt.ds[excess > 0][excess == min(excess), list(r1_clear = mm, r1_agent = agent, rco_mw =
sum(dt.rco$mw_max))]

        if(nrow(dt.rco.results) == 0) {
          dt.rco.results <- data.table(r1_clear = CAPACITY.RESERVE, r1_agent = NA, rco_mw = 0)
        }

      } else {
        dt.rco <- data.table(agent = vector("integer", 0L))
        dt.rco.results <- data.table(r1_clear = NA, r1_agent = NA, rco_mw = 0)
      }

      ### Round 2
      dt.supply <- dt.players[!agent %in% dt.rco$agent, list(agent, type, mw = mw_max, mm)]
      if(INCLUDE.INT.TECH) dt.supply[type == WHICH.INT, mw := mw*CAP.VALUE.INT]
      dt.supply <- dt.supply[order(mm)][, mw_cum := cumsum(mw)][]

      dt.ds <- dt.supply[, list(agent, mm, mw_cum, mw_d = quantity.demanded)]
      dt.ds <- dt.ds[, excess := mw_cum - mw_d]
      dt.r2.results <- dt.ds[excess > 0][excess == min(excess), list(r2_clear = mm, r2_agent = agent)]
      if(nrow(dt.r2.results) > 0) {
        dt.rco.results <- cbind(dt.rco.results, dt.r2.results)
      } else {
        dt.rco.results <- dt.rco.results[, list(r1_clear, r1_agent, rco_mw, r2_clear = CAPACITY.RESERVE, r2_agent = NA)]
      }

      if(!ENABLE.EX.RCO) dt.rco.results[, `:=`(r1_clear = r2_clear, r1_agent = r2_agent)]

      if(nrow(dt.r2.results) == 0) {
        cap.clear <- CAPACITY.RESERVE
        str.desc <- "RESERVE"
        str.type <- "RESERVE"
      } else {
        cap.clear <- dt.rco.results$r2_clear
        if(dt.rco.results$r2_agent %in% dt.players[age == 0, agent]) {
          str.desc <- paste0("NEW")
          str.type <- dt.players[agent == dt.rco.results$r2_agent, type]
        } else if(dt.rco.results$r2_agent %in% dt.players[age != 0, agent]) {
          str.desc <- paste0("OLD")
          str.type <- dt.players[agent == dt.rco.results$r2_agent, type]
        } else {
          stop()
```

```
        }
      }

    } else {
      cap.clear <- 0
      str.desc <- "NA"
      str.type <- "NA"
      dt.rco.results <- data.table(r1_clear = NA, r1_agent = NA, rco_mw = NA, r2_clear = NA, r2_agent = NA)
    }

    # Who cleared and did not
    if(ENABLE.CAP.MKT & ENABLE.EX.RCO) {
      inframarginal.rent <- cap.clear - dt.rco.results$r1_clear
      dt.rco[, `:=`(rco = cap.clear*CARVED.OUT.CAP.FACTOR,
                    exp_clear = exp.cap.clear,
                    clear = NA)]
      dt.players <- dt.players[!agent %in% dt.rco$agent]
      dt.players[, `:=`(rco = NA,
                        exp_clear = exp.cap.clear,
                        clear = ifelse(mm > cap.clear, NA,
                                       ifelse(mm > dt.rco.results$r1_clear, inframarginal.rent,
                                              cap.clear)))]
    } else {
      dt.rco[, `:=`(rco = NA,
                    exp_clear = NA,
                    clear = NA)]
      dt.players[, `:=`(rco = NA,
                        exp_clear = exp.cap.clear,
                        clear = ifelse(mm > cap.clear, NA, cap.clear))]
    }


    ##########
    # FINISH : Capacity/ICAP sim
    # START  : Retirement/Construction decisions
    ##########

    # Those without an award may still choose to stay/enter the market if they think they can get by on energy alone (fully
subsidized)
    # (This shouldn't happen with an RCO construct, since RCO status would be preferable)
    stay.no.clear <- dt.players[age != 0 & is.na(clear) & exp_nm + exp_credit > 0, agent]
    enter.no.clear <- calcNpv(dt.players[age == 0], DISCOUNT.RATE, include.build = T, exclude.credit = F)[npv > 0, agent]
    dt.no.clear <- dt.players[is.na(clear) & agent %in% c(stay.no.clear, enter.no.clear)]

    # Everyone who received a capacity award is obligated to enter or stay in the market
    # THIS IS WRONG. INFRAMARGINAL RENT RECEIVERS DONT HAVE TO STAY!!!
    dt.agents.active <- rbind(dt.players[!is.na(clear)], dt.no.clear)
    if(nrow(dt.rco) == 0) dt.rco <- dt.agents.active[0]
    dt.agents.active <- rbind(dt.agents.active, dt.rco[, names(dt.agents.active), with = F])
    # Re-index the new builds, creating a map
    dt.agents.active[age == 0, agent2 := (max(dt.saved.agent.bio$agent)+1):(max(dt.saved.agent.bio$agent)+.N)]
    dt.map <- dt.agents.active[, list(agent, agent2)][is.na(agent2), agent2 := agent]
    # Remap agents from capacity clear results
    dt.rco.results <- merge(dt.rco.results, dt.map, by.x = "r1_agent", by.y = "agent", all.x = T)[, list(r1_clear, r1_agent =
agent2, rco_mw, r2_clear, r2_agent)]
    dt.rco.results <- merge(dt.rco.results, dt.map, by.x = "r2_agent", by.y = "agent", all.x = T)[, list(r1_clear, r1_agent,
rco_mw, r2_clear, r2_agent = agent2)]
    dt.agents.active[, agent := ifelse(is.na(agent2), agent, agent2)]
    dt.agents.active[, agent2 := NULL]

    # Save results
    dt.saved.prices.cap <- rbind(dt.saved.prices.cap,
                                 cbind(data.table(year = this.year+1, price = cap.clear, exp_price = exp.cap.clear,
                                                  desc = str.desc, type = str.type),
                                       dt.rco.results))
    dt.saved.agent.bio <- unique(rbind(dt.saved.agent.bio, dt.agents.active[
      , list(agent, type, mw_max, build, hr, vom, fixed, life, itc, ptc)]))
    if(nrow(dt.agents.active) > 0) {
      dt.saved.mkt.cap <- rbind(dt.saved.mkt.cap,
                                dt.agents.active[, list(year = this.year+1, agent, mm, clear, rco)])
      dt.saved.invest <- rbind(dt.saved.invest,
                               dt.agents.active[, list(year = this.year, agent, exp_gm, exp_h, exp_nm, exp_credit, npv)])
    }

    ##########
    # FINISH : Retirement/Construction decisions
    # START  : Wrap up
    ##########

    # Increment year
    this.year <- this.year + 1
  }

  return(
    list(dt.saved.agent.bio,
      dt.saved.prices.nrg, dt.saved.prices.cap, dt.saved.mkt.nrg,
      dt.saved.mkt.cap, dt.saved.invest,
      dt.agents.active)
    )
}

#####
# KICKOFF THE SCRIPT, START BELOW
#####
# Read in 2017 PJM RTO demand
```

```
chr.username <- "mowry"
dt.demand.raw <- fread(paste0("C:/Users/", chr.username,
                              "/OneDrive/Documents/School/MIT/HKS API-166/Paper/Capacity Markets/pjm_rto_load_metered.csv"))

# read in Henry Hub from https://www.eia.gov/dnav/ng/hist/rngwhhdd.htm
dt.gas <- fread(paste0("C:/Users/", chr.username,
                "/OneDrive/Documents/School/MIT/HKS API-166/Paper/Capacity Markets/hh2017.csv"))
dt.gas <- merge(dt.gas[, list(hh, date = format(as.Date(paste0(year, "-", month, "-", day))))],
                data.table(date = format(seq(from = as.Date("2017-01-01"), to = as.Date("2017-12-31"), by = "1 day"))),
                by = "date", all.y = T)
if(which(!is.na(dt.gas$hh))[1] != 1) dt.gas[1:.N < which(!is.na(dt.gas$hh))[1], hh := dt.gas[!is.na(hh)][1, hh]]
for(i in 1:nrow(dt.gas)) if(is.na(dt.gas[i, hh])) dt.gas[i, hh := dt.gas[i-1, hh]]
dt.gas <- data.table(hour = 1:8760, hh = rep(dt.gas$hh, each = 24))

# Read in technology table
dt.tech.raw <- fread(paste0("C:/Users/", chr.username,
                            "/OneDrive/Documents/School/MIT/HKS API-166/Paper/Capacity Markets/tech_characteristics.csv"))
dt.tech.raw <- dt.tech.raw[2:.N]
for(i in 1:length(dt.tech.raw)) if(!any(is.na(as.numeric(dt.tech.raw[[i]])))) dt.tech.raw[[i]] <- as.numeric(dt.tech.raw[[i]])

#####
# Constants
set.seed(1) # set.seed(NULL)
#
COAL.PRICE <- 3; UR.PRICE <- round(30/(180E6), 2); OIL.PRICE <- 14; # fuel prices per MMBTU
(https://www.uranium.info/unit_conversion_table.php)
# if(USE.ABSTRACT.TECH) dt.tech <- dt.tech.raw[grepl("sub", type)] else dt.tech <- dt.tech.raw[!grepl("sub", type)]
USE.ABSTRACT.TECH <- T # use nuke/coal/gas or abstract?
dt.tech <- copy(dt.tech.raw)[enable == 1]
#
MAX.DEMAND <- 10000 # Scales actual demand MW to hit this max (limits agents, which speeds up sim)
VOLL <- 5000 # Energy price when supply insufficient to meet demand
HOUR.CUTOFF <- 1000 # Number of hours in year (out of 8760)
NUM.SEED <- round(MAX.DEMAND/dt.tech[, sum(capacity)]) # Number of new agents per tech seeded at start
#
NUM.POSSIBLE.INVEST <- 10 # Number of new agents per tech considered each year
NUM.NEW.INVEST <- 0 # Ultimate MW of new agents allowed to be built per year; dampens over building
SPREAD.FACTOR <- .2 # How much spread on tech characteristics; (0, .99)
DISCOUNT.RATE <- .08 # Investor discount rate
RECENCY.BIAS <- 2 # How much are historical years taken into account for financial projections. (Higher = less)
MAX.LAG <- 5 # How many years considered when forming gm/cap expectations
#
ENABLE.CAP.MKT <- F
ENABLE.MOPR <- T
ENABLE.EX.RCO <- T
CARVED.OUT.CAP.FACTOR <- 0.9 # % of RPM clear that a carve-out gets from its state
RESERVE.RATIO <- 1.15 # Administrative demand set at X% of last year's max load
CAPACITY.RESERVE <- dt.tech[, max(build_cost*1.25)] # If insufficient supply in cap mkt, price set to this
#
INCLUDE.INT.TECH <- F # Whether to include an intermittent resource
WHICH.INT <- "type_b" # Which resource becomes intermittent
INTERMITTENCY <- .35 # % of time resource is at full load (versus ~0)
# IMPORTANT: all intermittent gens are NOT CORRELATED in energy market (destabilizing)
CAP.VALUE.INT <- INTERMITTENCY # What % of rated capacity an intermittent resource is counted as for cap mkt purposes
#
INCLUDE.SUB.TECH <- F # Whether to include a subsidized tech
WHICH.SUB <- "type_b" # Which resource to subsidize (with ITC and PTC below)
ITC.RATE <- 40000 # $ of build costs credited each year by an ITC style subsidy
PTC.RATE <- 1 # $/MWh of PTC added every year (materializes in VOM reduction)

#####
# Run
dt.demand <- dt.demand.raw[1:HOUR.CUTOFF][, list(hour = 1:.N, mw = mw*(MAX.DEMAND/max(mw)))]
dt.fuel <- dt.gas[1:HOUR.CUTOFF, list(hour, hh, coal = COAL.PRICE, u = UR.PRICE, oil = OIL.PRICE)]

source(file = "./iterative model.R")
lst.res <- runSim(list(), 1, 50)
# lst.res <- runSim(lst.res, 31, 50)

#####
# Chart

# Merge agent information
dt.saved.agent.bio <- lst.res[[1]]
dt.saved.prices.nrg <- lst.res[[2]]
dt.saved.prices.cap <- lst.res[[3]]
dt.saved.mkt.nrg <- lst.res[[4]]
dt.saved.mkt.cap <- lst.res[[5]]
dt.saved.invest <- lst.res[[6]]
dt.agents.active <- lst.res[[7]]
dt.story <- Reduce(x = list(dt.saved.agent.bio, dt.saved.mkt.cap, dt.saved.mkt.nrg, dt.saved.invest),
                   f = function(a,b) merge(a,b,by=intersect(names(a),names(b)), all = T))
dt.story <- dt.story[!is.na(type)] # get rid of first-year invest rows

# Portfolio mix
ggplot(dt.story[, list(mw = sum(mw_max)), by = list(year, type, RCO = !is.na(rco))][order(type)]) +
  theme_bw() + geom_hline(yintercept = 0) +
  geom_violin(data = cbind(year = 0, dt.demand), mapping = aes(x = year, y = mw), fill = "grey", trim = F) +
  geom_bar(mapping = aes(x = year, y = mw, fill = type, alpha = RCO),
           stat = "identity", position = "stack", color = "black") +
  geom_hline(yintercept = mean(dt.demand$mw), size = 1) +
  geom_hline(yintercept = max(dt.demand$mw)*RESERVE.RATIO, linetype = "dashed", size = 1) +
  scale_fill_brewer("Tech Type", type = "qual", palette = 6) + scale_alpha_manual(values = c(1, .5))
# Agent longevity
ggplot(dt.story[, list(birth = min(year), death = max(year)), by = list(agent, type)]) +
```

```
  geom_segment(mapping = aes(x = birth, xend = death, y = agent, yend = agent, color = type)) +
  scale_color_brewer("Tech Type", type = "qual", palette = 6) +
  theme_bw() + labs(x = "Build and Retirement Years", y = "Plant ID")
# Prices
fyear <- dt.saved.prices.nrg[, list(test = sum(price==VOLL)), by = year][test == 0][1, year]
fyear <- max(fyear, dt.saved.prices.cap[price < CAPACITY.RESERVE][1, year])
ggplot(dt.saved.prices.nrg[year >= fyear, list(mi = min(price), me = median(price), mn = mean(price), ma = max(price)), by =
list(year)]) +
  geom_ribbon(mapping = aes(x = year, ymin = mi, ymax = ma), fill = "grey50", alpha = .5) +
  geom_line(mapping = aes(x = year, y = mi), linetype = "dashed") +
  geom_line(mapping = aes(x = year, y = ma), linetype = "dashed") +
  geom_line(mapping = aes(x = year, y = me), linetype = "dashed") +
  geom_line(mapping = aes(x = year, y = mn)) +
  geom_point(data = dt.saved.prices.cap[year >= fyear],
             mapping = aes(x = year, y = price/HOUR.CUTOFF, color = type, shape = desc), size = 2) +
  geom_point(data = dt.saved.prices.cap[year >= fyear],
             mapping = aes(x = year, y = r1_clear/HOUR.CUTOFF), size = 2, shape = "cross") +
  scale_color_brewer("Marginal \nCapacity \nTech", type = "qual", palette = 6) +
  scale_shape_discrete("Marginal \nCapacity \nBuild") +
  theme_bw() + labs(x = "Year", y = "Mean Hourly Energy Price and\nCapacity Clear ($/MWh)")
# Capacity factors
ggplot(dt.story[, list(cf = round(cf = sum(h/HOUR.CUTOFF*mw_max*ifelse(type == WHICH.INT & INCLUDE.INT.TECH == T, CAP.VALUE.INT,
1))/sum(mw_max)*100)),
                by = list(type, year)]) +
  geom_line(mapping = aes(x = year, y = cf, color = type)) +
  geom_point(mapping = aes(x = year, y = cf, color = type), alpha = .5) +
  scale_color_brewer("Tech Type", type = "qual", palette = 6) +
  theme_bw() + labs(x = "Year", y = "Capacity Factor %")
# RCO stats
fyear <- dt.saved.prices.cap[price < CAPACITY.RESERVE][1, year]
ggplot(melt(copy(dt.saved.prices.cap)[year >= fyear][, infra := r2_clear - r1_clear],
            id.vars = c("year"), measure.vars = c("r1_clear", "r2_clear", "infra"))) +
  geom_hline(yintercept = 0) + theme_bw() + labs(x = "Year", y = "RCO RPM Round Clears $/MW") +
  geom_line(mapping = aes(x = year, y = value, color = variable), size = 1)
# Tech improvement
ggplot(melt(dt.story[, list(year, type, Build = build-itc, Fixed = fixed, Production = vom-ptc)],
            id.vars = c("year", "type"))) +
  geom_boxplot(mapping = aes(x = year, y = value, fill = type, group = paste0(year, type)), outlier.size = 0) +
  facet_wrap(~variable, nrow = 1, scales = "free_y") + theme_bw() +
  scale_fill_brewer("Tech Type", type = "qual", palette = 6) + labs(x = "Year", y = "Subsidized Costs on Active Plants")
ggplot(melt(dt.story[, list(Build = mean(build-itc), Fixed = mean(fixed), Production = mean(vom-ptc)), by = list(year, type)],
            id.vars = c("year", "type"))) +
  geom_line(mapping = aes(x = year, y = value, color = type, group = paste0(variable, type))) +
  facet_wrap(~variable, nrow = 1, scales = "free_y") + theme_bw() +
  scale_fill_brewer("Tech Type", type = "qual", palette = 6) + labs(x = "Year", y = "Avg Subsidized Costs on Active Plants")
```