# Optimization of Intra-Day Battery Energy Storage

Workshop on Forecasting and Mathematical Modeling for Renewable Energy
PIMS, July 27, 2023

Mike Ludkovski

Dept of Statistics & Applied Probability UC Santa Barbara
Partially supported by ARPA-E PERFORM program

STATISTICS & APPLIED PROBABILITY CENTER FOR FINANCIAL MATHEMATICS & ACTUARIAL RESEARCH

UCSB

# Introduction

- Screenshot from **CAISO**: California grid operator
- Y-axis: power demand: 20-30 GW
- Net load = power demand minus renewable energy
- Note that load is as seen by the utility, i.e. net of behind-the-meter rooftop solar, etc.
- Observe the extreme ramp between 5 and 8pm: the Duck Curve

Typical day in 2023 in California (a Friday in May)

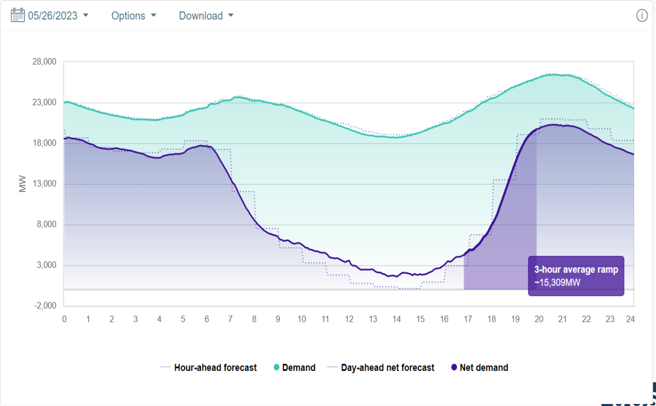


Courtesy of CAISO.com

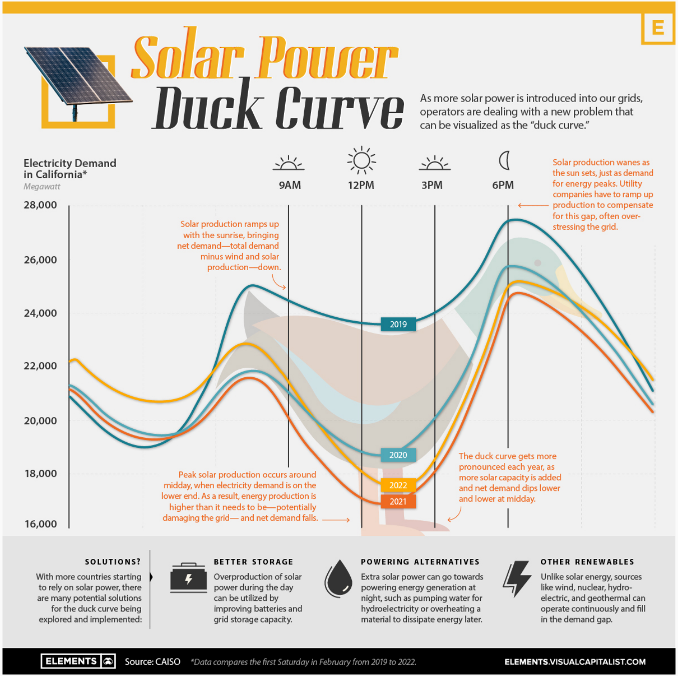

Solar Power
Duck Curve
As more solar power is introduced into our grids, operators are dealing with a new problem that can be visualized as the "duck curve."
Electricity Demand in California*
Megawatt
9AM
12PM
3PM
6PM
28,000
26,000
24,000
22,000
20,000
18,000
16,000
Solar production ramps up with the sunrise, bringing net demand—total demand minus wind and solar production—down.
Solar production wanes as the sun sets, just as demand for energy peaks. Utility companies have to ramp up production to compensate for this gap, often over-stressing the grid.
2019
2020
2022
2021
Peak solar production occurs around midday, when electricity demand is on the lower end. As a result, energy production is higher than it needs to be—potentially damaging the grid— and net demand falls.
The duck curve gets more pronounced each year, as more solar capacity is added and net demand dips lower and lower at midday.
SOLUTIONS?
With more countries starting to rely on solar power, there are many potential solutions for the duck curve being explored and implemented:
BETTER STORAGE
Overproduction of solar power during the day can be utilized by improving batteries and grid storage capacity.
POWERING ALTERNATIVES
Extra solar power can go towards powering energy generation at night, such as pumping water for hydroelectricity or overheating a material to dissipate energy later.
OTHER RENEWABLES
Unlike solar energy, sources like wind, nuclear, hydro-electric, and geothermal can operate continuously and fill in the demand gap.
ELEMENTS
Source: CAISO
*Data compares the first Saturday in February from 2019 to 2022.
ELEMENTS.VISUALCAPITALIST.COM


UCSB

# Another Day

Wednesday in April

- Net load was forecast to be negative on the day prior
- Realized net load between 10am-4pm was $\sim 6 - 7GW$ higher, and never below 5GW
- What's going on?

Courtesy of CAISO.com

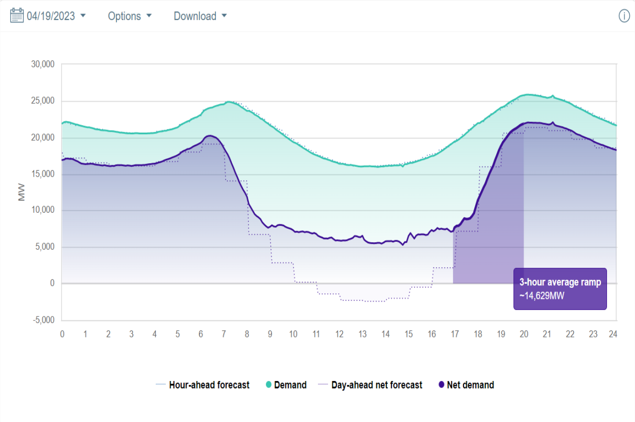


$\Longrightarrow$ Renewable Curtailment

UCSB

# Sunny California: Too much of a good thing

- Curtailment was 700 GWh in April 2023 (=23 GWh per day or about 5% of total demand)
- In 2022 total solar production averaged 3409 TWh/month, curtailed 672 TWh/month or 20%! (total demand is about 15000 TWh)
- Massive hit on the profitability of solar farms
- Throwing away incredible amount of free energy
- Too much power during the day, then a shortage in late evening

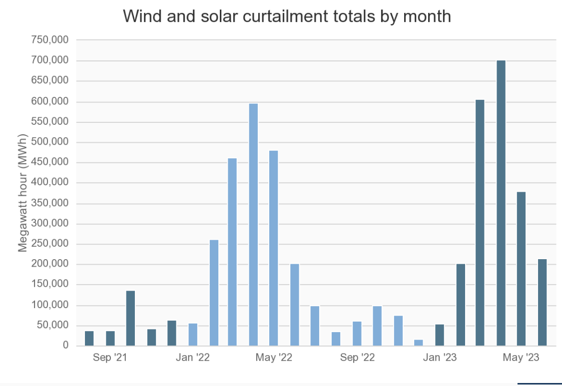


Renewable Curtailment

# BESS: Battery Energy Storage Systems

- Energy Storage is the centerpiece of solving these problems
- Now a meaningful segment of the power mix in many regions, experiencing exponential growth
- Creates many market design challenges
- Storage is the most interesting and practically relevant application of stochastic control for renewables
- This talk: short-duration storage (battery based)

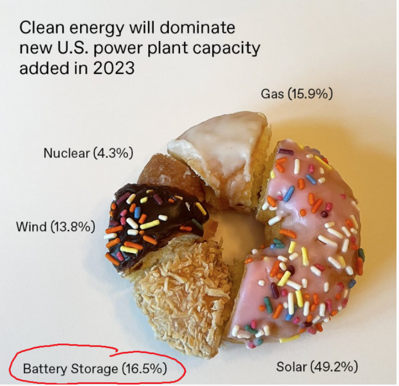


Image by Canary Media, June 1, 2023

UCSB

# Uses for Battery Storage

- **Energy arbitrage**: buy low during mid-day, discharge in the evening
- **Peak shaving**: provide power at the highest net-load hour(s)
- **Duck curve smoothing**: provide power to reduce the ramp rate
- **Supply firming**: complement random fluctuations in renewable generation
- **Load shaping**: better track preset load profile
- **Curtailment insurance**: avoid renewable curtailment
- **Congestion relief**: replace spatial transmission with a temporal one
- Ancillary services

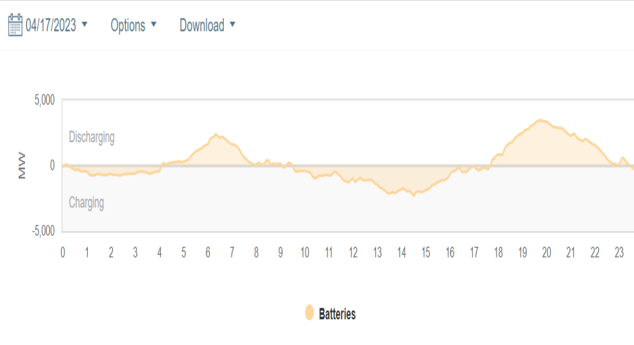


Image from caiso.com

UCSB

# Ancillary Services

- **Regulation-up/down**: resources to balance the grid due to short-term fluctuations: must respond to automatic control signals (on the scale of seconds) to increase or decrease their operating levels depending upon the need
- **(non)-Spinning reserve**: standby generation capacity that can be ramped to a specified load within 10 minutes when dispatched. Mostly for up –only triggered by contingency events (gen down)
- **Frequency regulation**: control system frequency close to 60Hz
- **Flexible Ramping Product**: flexible ramp capability to meet the potential net load movement in RT dispatch from load and renewables
- FRP targets net load changes between two dispatch intervals (e.g 95% interval based on historical forecast errors)
- ERCOT Contingency Reserve Service: version of FRP introduced in June ’23

UCSB

# Ancillary Service Markets

- Rules differ market-to-market
- Asset bid on AS day-ahead: acceptance probability
- If selected, receive a capacity payment
- If needed in RT (deployment probability), receive fees based on opportunity cost
- AS deployment is based on bid capacity, so could bid a partial fraction of the battery

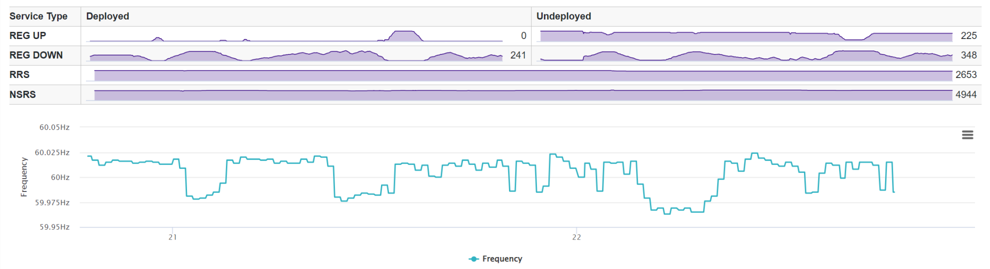

# Interlude presented by ChatGPT: "Battery Storage and Solar Energy"

Solar power, the radiant star,
Its beams of gold stretch near and far.
They dance on panels, silicon's braid,
Harvesting light, a celestial crusade.

But the sun's a fickle, wandering sprite,
At dusk she rests, in the veil of night.
Yet fear not, for a hero's here,
Battery storage, resolute and clear!

In vessels of power, they hold the might,
To store the sun's gifts, day or night.
When skies turn dark and stars align,
Battery energy surges, a grand design.

No more dependence on fossil fuel,
The future's bright, a cleaner rule.
With battery storage and solar rays,
We'll forge a path to greener days.

UCSB

# Optimizing Battery Storage Dispatch

# Storage Dispatch

Classical: buy low/sell high. Constrained by:

- Maximum (dis)charge rate (MW)
- Battery capacity (MWh)
- Roundtrip efficiency (%)
- Initial and terminal SOC conditions
- *Below:* synthetic example with OU prices;
- *Right:* real-life storage dispatch across AS markets

courtesy of CAISO.com

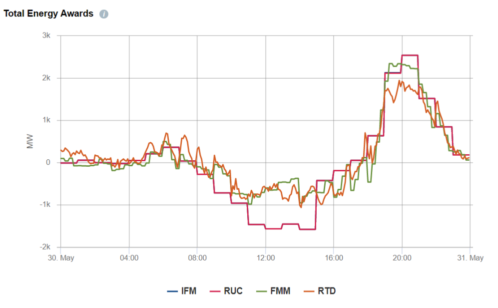


IFM=Integrated Forward Market, RUC=Residual Unit Commitment,

FFM=Fifteen Minute Market, RTD=Real Time Dispatch

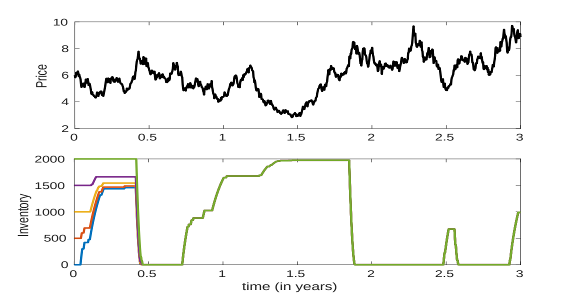


UCSB

# Battery Storage as Switching Control

- Inventory $I_t$ (endogenous)
- Price process/stochastic factors $\mathbf{P}_t$ (typically exogenous SDE)
- Terminal condition $W(\mathbf{P}_T, I_T)$
- Control $m_t$: $I_{t_{k+1}} = I_{t_k} + a(m_{t_k})$
- Total revenue on $[t_k, T]$: $v(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) := \sum_{s=k}^{K-1} e^{-r(t_s - t_k)}[\pi(\mathbf{P}_{t_s}, c_{t_s}(m_{t_{s+1}}))\Delta t - K(m_{t_s}, m_{t_{s+1}})] + e^{-r(T-t_k)} W(\mathbf{P}_T, I_T)$

**Solution Approaches:**

- PDE/HJB
- Rolling Intrinsic
- Regression Monte Carlo (Warin 2010, Bauerle-Riess 2016, L-Maheshwari 2019)
- Reinforcement Learning (Hure et al, 2020)

UCSB

# Regr Monte Carlo for Storage: Machine Learning Perspective

- Value function

$$V(t_k, \mathbf{P}, I, m) = \sup_{m_{t_k}} \mathbb{E}\left[ v(t_l, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) \middle| \mathbf{P}_{t_k} = \mathbf{P}, I_{t_k} = I, m_{t_k} = m \right]$$

- Continuation value: $q(t_k, \mathbf{P}, I, m) := \mathbb{E}\left[ e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I, m) \middle| \mathbf{P}_{t_k} = \mathbf{P} \right]$
- Dynamic Programming equation

$$V(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) = \max_{m \in \mathcal{J}} \mathbb{E}\left[ \pi^{\Delta}(\mathbf{P}_{t_k}, m_{t_k}, m) + e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I_{t_{k+1}}(m), m) \middle| \mathbf{P}_{t_k} \right]$$

- Optimal control $m^*_{t_{k+1}}(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k})$ is the arg max above.
- The pricing model is viewed as stochastic simulator(s) that produces noisy pathwise observations conditional on the control
- Learn the continuation or q-value $q(t, \mathbf{P}, I, m)$: cost-to-go conditional on next-step regime $m$; done recursively over $t$
- Sub-modules:

;

UCSB

# Regr Monte Carlo for Storage: Machine Learning Perspective

- Value function

$$V(t_k, \mathbf{P}, I, m) = \sup_{m_{t_k}} \mathbb{E}\left[ v(t_l, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) \middle| \mathbf{P}_{t_k} = \mathbf{P}, I_{t_k} = I, m_{t_k} = m \right]$$

- Continuation value: $q(t_k, \mathbf{P}, I, m) := \mathbb{E}\left[ e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I, m) \middle| \mathbf{P}_{t_k} = \mathbf{P} \right]$
- Dynamic Programming equation

$$V(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) = \max_{m \in \mathcal{J}} \mathbb{E}\left[ \pi^{\Delta}(\mathbf{P}_{t_k}, m_{t_k}, m) + e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I_{t_{k+1}}(m), m) \middle| \mathbf{P}_{t_k} \right]$$

- Optimal control $m^*_{t_{k+1}}(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k})$ is the arg max above.
- The pricing model is viewed as stochastic simulator(s) that produces noisy pathwise observations conditional on the control
- Learn the continuation or q-value $q(t, \mathbf{P}, I, m)$: cost-to-go conditional on next-step regime $m$; done recursively over $t$
- Sub-modules:
    - Approximation of the conditional expectation defining $q$;

UCSB

# Regr Monte Carlo for Storage: Machine Learning Perspective

- Value function

$$V(t_k, \mathbf{P}, I, m) = \sup_{m_{t_k}} \mathbb{E}\left[ v(t_l, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) \middle| \, \mathbf{P}_{t_k} = \mathbf{P}, I_{t_k} = I, m_{t_k} = m \right]$$

- Continuation value: $q(t_k, \mathbf{P}, I, m) := \mathbb{E}\left[ e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I, m) \middle| \, \mathbf{P}_{t_k} = \mathbf{P} \right]$
- Dynamic Programming equation

$$V(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) = \max_{m \in \mathcal{J}} \mathbb{E}\left[ \pi^{\Delta}(\mathbf{P}_{t_k}, m_{t_k}, m) + e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I_{t_{k+1}}(m), m) \middle| \mathbf{P}_{t_k} \right]$$

- Optimal control $m^*_{t_{k+1}}(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k})$ is the arg max above.
- The pricing model is viewed as stochastic simulator(s) that produces noisy pathwise observations conditional on the control
- Learn the continuation or q-value $q(t, \mathbf{P}, I, m)$: cost-to-go conditional on next-step regime $m$; done recursively over $t$
- Sub-modules:
  - Approximation of the conditional expectation defining $q$;
  - Evaluation of the optimal control [trivial if $m$ is discrete]

UCSB

# Regr Monte Carlo for Storage: Machine Learning Perspective

- Value function

$$V(t_k, \mathbf{P}, I, m) = \sup_{m_{t_k}} \mathbb{E}\left[ v(t_l, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) \middle| \mathbf{P}_{t_k} = \mathbf{P}, I_{t_k} = I, m_{t_k} = m \right]$$

- Continuation value: $q(t_k, \mathbf{P}, I, m) := \mathbb{E}\left[ e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I, m) \middle| \mathbf{P}_{t_k} = \mathbf{P} \right]$
- Dynamic Programming equation

$$V(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k}) = \max_{m \in \mathcal{J}} \mathbb{E}\left[ \pi^{\Delta}(\mathbf{P}_{t_k}, m_{t_k}, m) + e^{-r\Delta t} V(t_{k+1}, \mathbf{P}_{t_{k+1}}, I_{t_{k+1}}(m), m) \middle| \mathbf{P}_{t_k} \right]$$

- Optimal control $m^*_{t_{k+1}}(t_k, \mathbf{P}_{t_k}, I_{t_k}, m_{t_k})$ is the arg max above.
- The pricing model is viewed as stochastic simulator(s) that produces noisy pathwise observations conditional on the control
- Learn the continuation or q-value $q(t, \mathbf{P}, I, m)$: cost-to-go conditional on next-step regime $m$; done recursively over $t$
- Sub-modules:
  - Approximation of the conditional expectation defining $q$;
  - Evaluation of the optimal control [trivial if $m$ is discrete]
  - Evaluation of the pathwise continuation value

UCSB

## DEA: Dynamic Emulation Algorithm (L-Maheshwari 2020)

- A template for a simulation-based approach to stochastic storage
- Unifies Optimal Stopping RMC and Storage RMC
- A plug-and-play modular algorithm interpreted as a (machine/statistical) learning problem. Resembles Reinforcement Learning.
- Nests existing literature + offers MANY more choices
- Applied to stylized Gas Storage and Microgrid Control settings
- Solve the storage control problem $\Leftrightarrow$ Recursively identify the control map + value function given a pathwise reward simulator

UCSB

# DEA: Dynamic Emulation Algorithm (L-Maheshwari 2020)

- A template for a simulation-based approach to stochastic storage
- Unifies Optimal Stopping RMC and Storage RMC
- A plug-and-play modular algorithm interpreted as a (machine/statistical) learning problem. Resembles Reinforcement Learning.
- Nests existing literature + offers MANY more choices
- Applied to stylized Gas Storage and Microgrid Control settings
- Solve the storage control problem $\Leftrightarrow$ Recursively identify the control map + value function given a pathwise reward simulator
- Classic RMC: generate forward paths, use the resulting stochastic mesh to solve the DPE, employ cross-sectional regression for the conditional expectation.
- Key challenge for applying RMC is how to handle the endogenous $I_t$ – cannot do forward path simulation
    - Inventory path back-propagation and quasi-simulation
    - Treat $I_t$ as a *parameter*, solve a collection of 1-D problems in $P_t$
    - Control randomization

UCSB

# Training Design Choices

How to pick the training sites in $(P, I)$:

- Space Filling: explore continuation values throughout the input domain:
  - Quasi Monte Carlo sequences (Sobol)
  - Latin Hypercube Sampling
  - Gridded
- Probabilistic – reflects the distribution of $(P_t, \hat{I}_t)$
- Adaptive – target efficient learning of the action boundaries

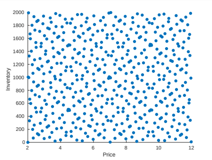


**(a)** 2D Sobol QMC sequence

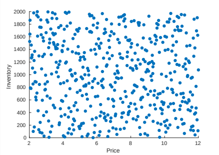


**(b)** 2D randomized space-filling

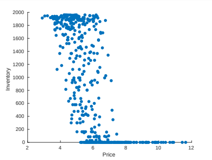


**(c)** Joint Probabilistic design

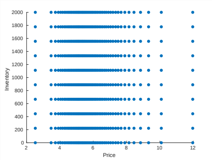


**(d)** Conventional product design

# Realistic Net Load Model

# Conditional Simulations

Daily ISO operations:

- Receive weather forecast for tomorrow
- Optimization solver for day-ahead Unit Commitment (hourly scale)
- Next day: Economic Dispatch based on realized generation/load (15-min scale)
- Intrinsic stochasticity of renewable production: forecast errors of 10% for 24-hours out are common and unavoidable
- Wrong-way correlation between renewable assets: when you need renewables they are less likely to be there

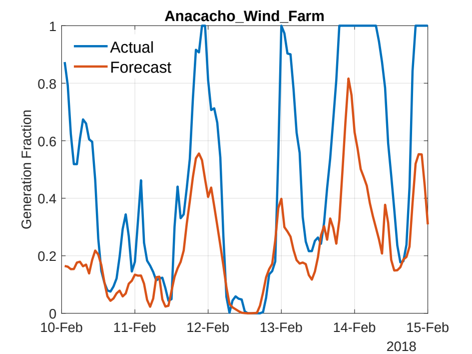


Hourly generation over 5 days at a wind farm in TX.

UCSB

# ARPA-E PERFORM Data

- Day-ahead forecast + realized generation at hourly frequency.
- Re-analysis of NWP provided by NREL.
- Directly work in MW (no GHI/wind speed)
- Solar: 226 assets; Wind: 264 assets
- Varying technologies, sizes and production behavior.
- High degree of locational correlation.
- Aim: generative model for scenarios of realized generation across **all** renewable assets

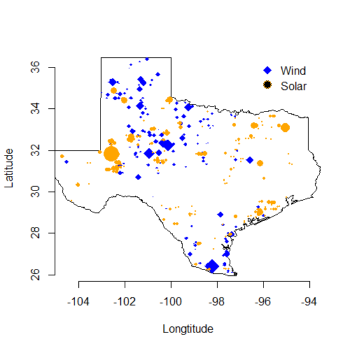


NREL ERCOT dataset

UCSB

# Building a Statistical Engine for Renewable Generation

- Model jointly forecast and actuals
- Data are vectors of dim 24 (no time-series)
- Multiple statistical challenges to resolve:
    - Data scarcity for correlation estimation
    - Point masses
    - Non-gaussian, highly heteroskedastic distributions
    - Layers of seasonality – raw data is fundamentally non-stationary
    - Spatial dependence

Blue Bell solar farm:

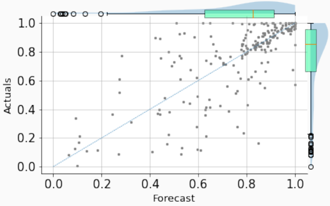


Baffin wind farm:

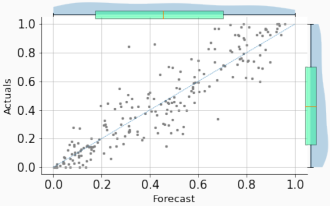


UCSB

# Calibration

- Convert raw MWh into production ratios
- Estimate mean- and max-generation for each hour/day of the year
- Standardize different days to make them comparable for inferring correlations
- Handle point masses through conditional inference
- Several calibration layers to account for seasonality, asset characteristics and non-Gaussianity
- To handle varying number of active hours across assets work in a transformed PCA factor space

Scalability to **hundreds** of assets is the secret sauce

UCSB

# Hierarchical Clustering

- Estimation of large covariance matrices is inherently unstable
- With only a few hundred data points, can only reliably estimate a small matrix
- Imposing structure is necessary (sparsity, regularization, hierarchy)
- Do this by recursively by clustering assets based on empirical correlations
- Cluster via simulated annealing (search in the high-dimensional space of potential clusters)
- Correlate wind vs solar vs load at the top level

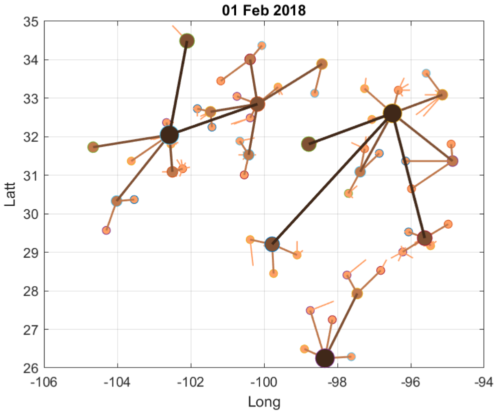


Overall wind clusters (4 levels total)

UCSB

# Simulated Annealing

- At each hierarchy level: employ simulated annealing to build clusters
- Repeatedly propose to (i) split an existing cluster into 2 or (ii) move one asset from its cluster to another
- Energy function $E(\mathcal{C}) := \sum_{\mathfrak{c} \in \mathcal{C}} \left[ 1 + \frac{\kappa}{|\mathfrak{c}|-1} \sum_{i,j \in \mathfrak{c}} \left( 1 - \rho_{i,j}^2 \right) \right]$
- Accept proposals based on the resulting change $\Delta E$ and current temperature schedule: Accept w.p.1 if $\Delta E_\ell < 0$, otherwise with prob. $\exp(-T_\ell \Delta E_\ell)$

More comments:

- Correlation degradation is inevitable with regularization. Ensure that preserve very high correlations (some assets are essentially adjacent to each other and have correlation of $> 99\%$)
- Current method does not make any explicit use of geographical patterns
- Clustering can be date-dependent

UCSB

# Simulations

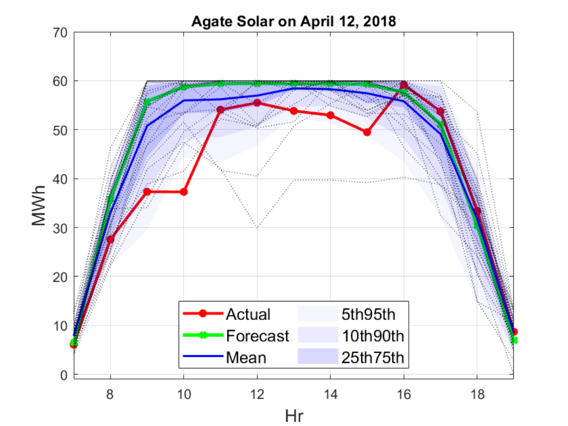


- Clusters yield the calibrated high-dim covariance matrix
- Populate with multivariate Gaussian samples
- Reverse all the steps to obtain simulations joint across hours + assets
- Runs on the sub-second scale per scenario (once clustering is done)

UCSB

# More Simulations

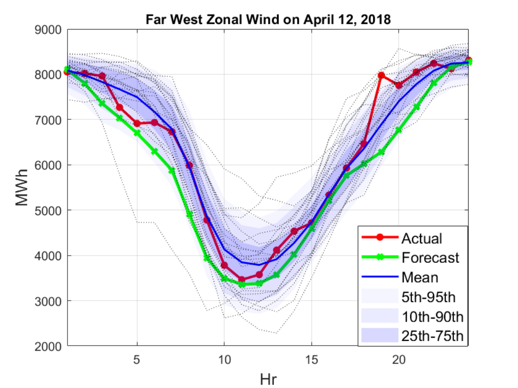


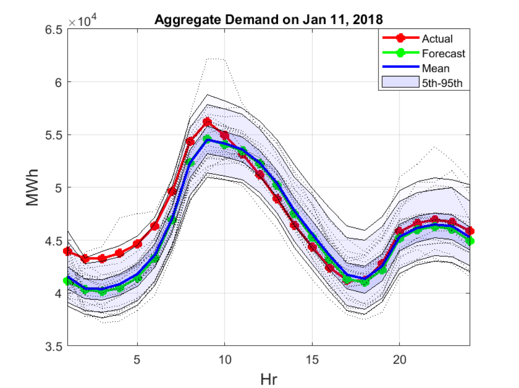


Scenarios are joint: can aggregate scenario-by-scenario

Far West has 48 wind farms. Total load is over 8 ERCOT zones.

UCSB

# Temporal Correlation

Empirical correlation matrix across the active hours of April 13, 2018

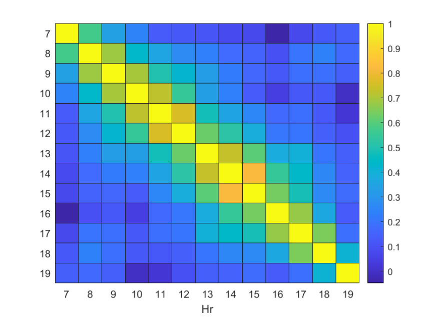


Adamstown solar (13 active hours)

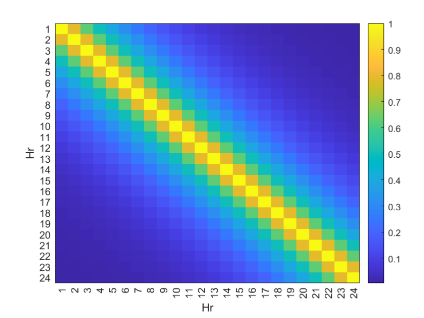


Amazon wind farm.

UCSB

# Putting it Together: First Pass

- Fixed database of 1000 net load scenarios from above engine
- Different distribution for each day (relies on 2018 data/day-ahead forecasts)
- Hourly frequency (soon 15-min dispatch intervals)
- 4MWh battery with 1 MW charge rate and 92% round-trip efficiency
- Participates in regulation control throughout the day (random effect added to inventory)
- Sole objective of energy arbitrage
- Zero initial and terminal SOC
- DEA solver with Gaussian Process surrogates on $(P_t, I_t)$

UCSB

# Sample Inventory Paths on 4 different days

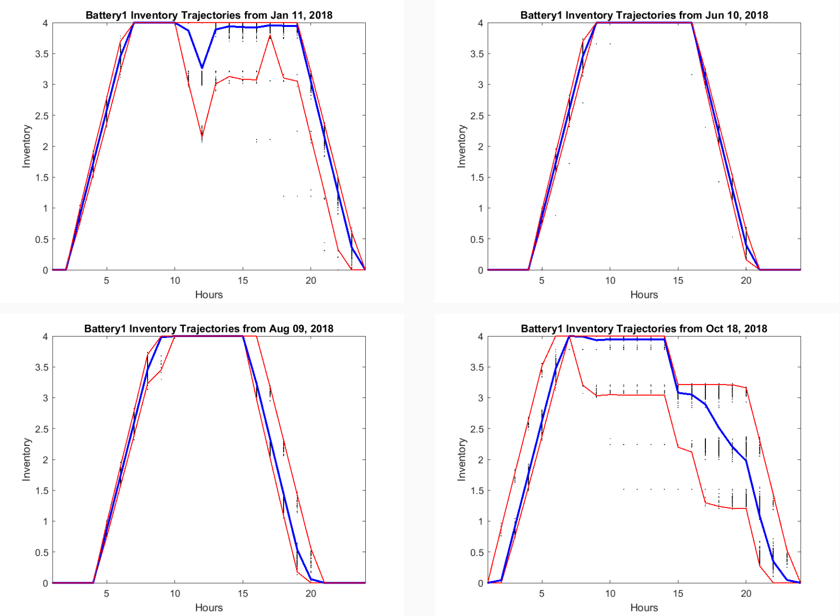


UCSB

# Bells and Whistles

- Multi-dimensional state:
  - Several factors (load, wind, solar, possibly zonal values, ...)
  - Cycling restriction (swing-like optionality)
- Payoff structure:
  - Function of current price (function of load)
  - Function of change in load (ramp)
  - Randomized due to deployment probability, regulation-up/down signal, etc.
  - Spatial structure: battery location vs the grid
- Need a **flexible, dim-agnostic** storage optimizer
- Run the solver hundreds of times on an intra-day basis

UCSB

# The Big Picture

- Level 1: optimizing battery operations by a single asset owner (price-taker)
- Level 2: optimal bidding strategy (how to allocate between AS and energy arbitrage) – economic value of storage across time and across market products
- Level 3: aggregate behavior of batteries at system level (mean field game)
- Level 4: market design + regulation + encouraging socially beneficial behavior (principal-agent)
- A nexus of complicated numerical tasks (with multi-layer engineering characteristics) and modeling tasks

UCSB

# References


Ludkovski, M., Maheshwari, A. (2020)
**Simulation Methods for Stochastic Storage Problems: A Statistical Learning Perspective**
*Energy Systems* 11, 377-415

Mike Ludkovski, Glen Swindle, Eric Grannan (2022)
**Large Scale Probabilistic Simulation of Renewables Production**
*Arxiv*, `arXiv:2205.04736`


THANK YOU!

UCSB

# Market Failure: Sep 7, 2022

Batteries began discharging in the middle of the afternoon, when there was still plenty of solar power and other supplies available to meet electricity demand. That depleted the cushion before it was more critically needed in the early evening, when the state was on the brink of rotating blackouts as demand hit an all-time record and solar supplies started dropping as the sun set.... The reason for Tuesday's earlier-than-expected deployment of batteries likely has to do with market signals. The way the California power market is set up now, energy storage systems are called upon to dispatch by the grid operator when the wholesale power price hits a cap of $1,000 a megawatt-hour...

## California's Battery Problems Heighten Threat of Power Outages

The way the state triggers batteries to supply the grid is opening the door to mismatches with demand.

Power inverters outside the battery building at the Moss Landing Energy Storage Facility in Moss Landing, California. *Photographer: David Paul Morris/Bloomberg*



By Mark Chediak
September 7, 2022 at 3:24 PM PDT *Updated on September 8, 2022 at 5:40 AM PDT*

The batteries that help fortify California's electric grid are kicking in at times when they're not really needed, draining the power source before more critical junctures and heightening the chances of blackouts as a blistering heat wave punishes the state.