Robust unit commitment using the parametric cost function approximation

FERC, Washington, D.C.

June 25-26, 2017

Raymond Perkins
Warren B. Powell
CASTLE Labs
Princeton University
http://www.castlelab.princeton.edu

Mission statement

- Main points:
  - » Classical view is that ISOs use a "deterministic model" for unit commitment.
  - » Considerable research has been done on the "stochastic unit commitment problem" that uses a stochastic lookahead model using scenario trees.
  - » Our thesis:
    - "stochastic programming" is a form of ***policy*** (a stochastic lookahead model) for solving a ***stochastic problem*** (the real world).
    - ISOs use a modified deterministic lookahead model, where the modifications enforce reserve requirements to ensure a robust solution.
    - We call this a ***parametric cost function approximation***, and argue that this is also a form of policy for solving ***stochastic unit commitment problems***.
  - » We will describe weaknesses in the use of scenario trees, and argue why the parametric CFA (which is current industry practice) is likely to be much more effective for handling uncertainty in this context.

Variability and uncertainty

Wind

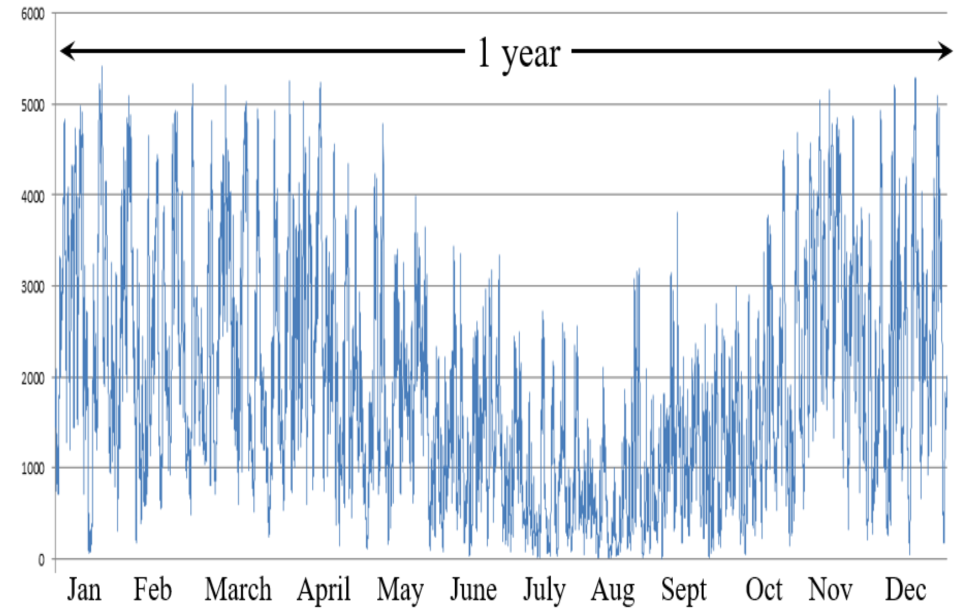

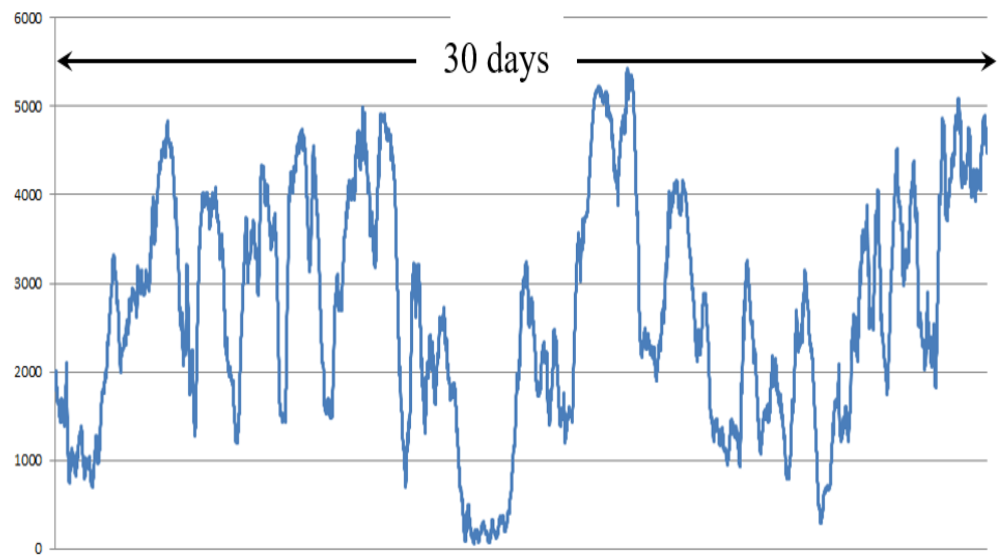

Solar

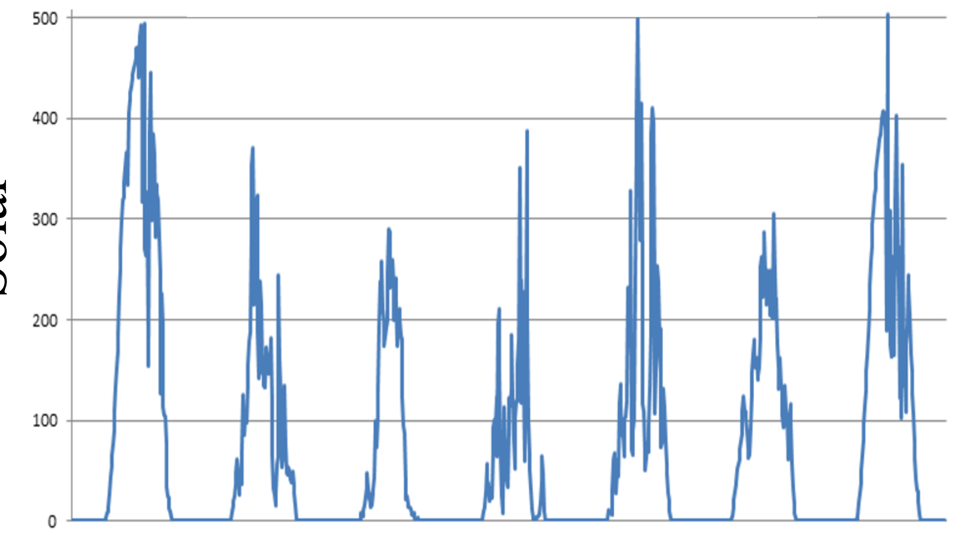

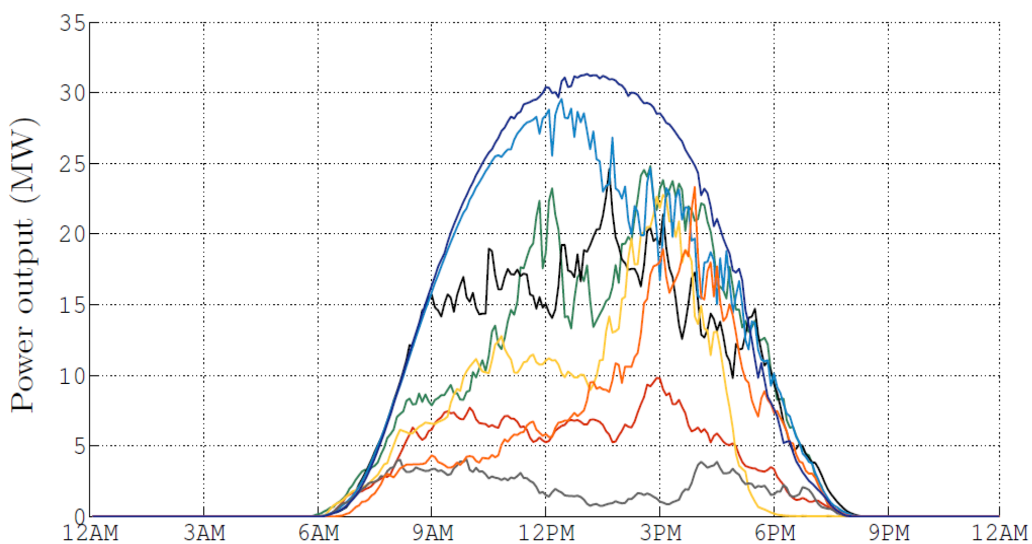

Variability and uncertainty

- Illustration of forecasted wind power and actual
  - » The forecast (black line) is deterministic (at time t, when the forecast was made). The actuals are stochastic.

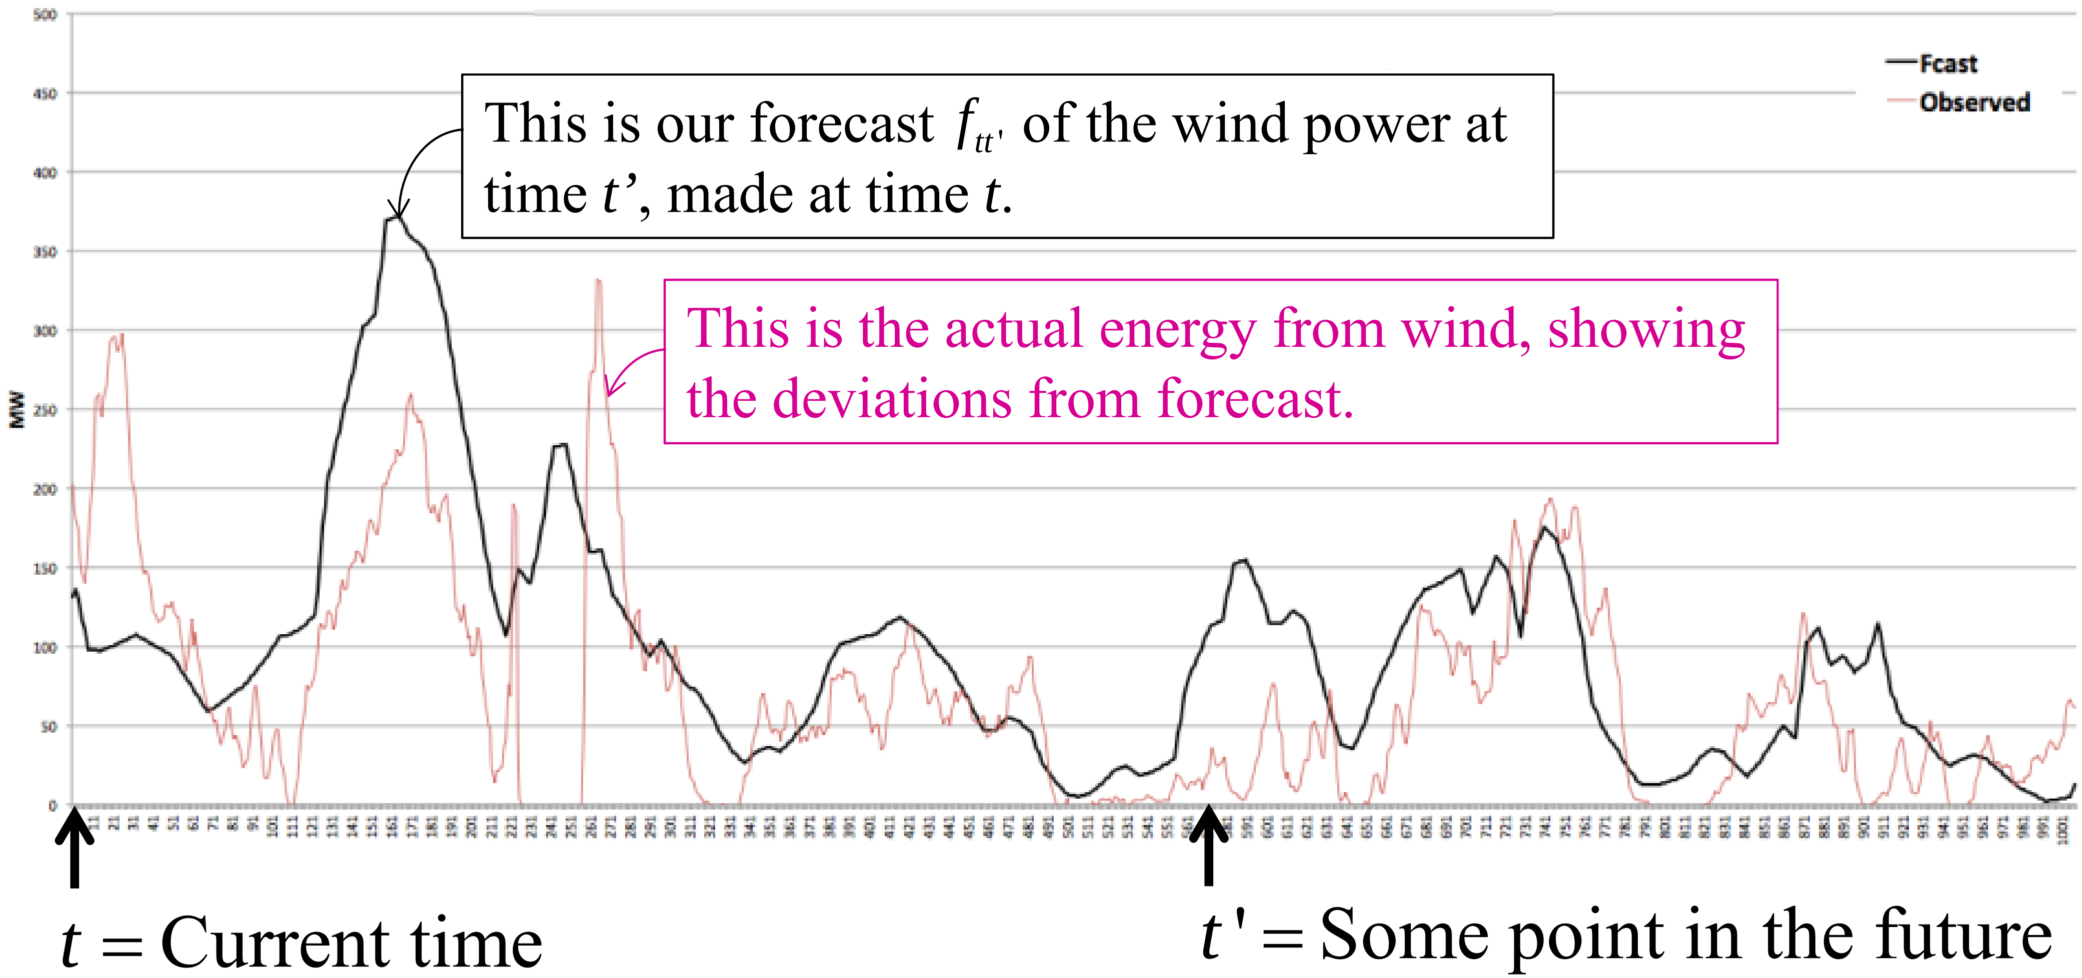

Variability and uncertainty

- Forecasts evolve over time as new information arrives:

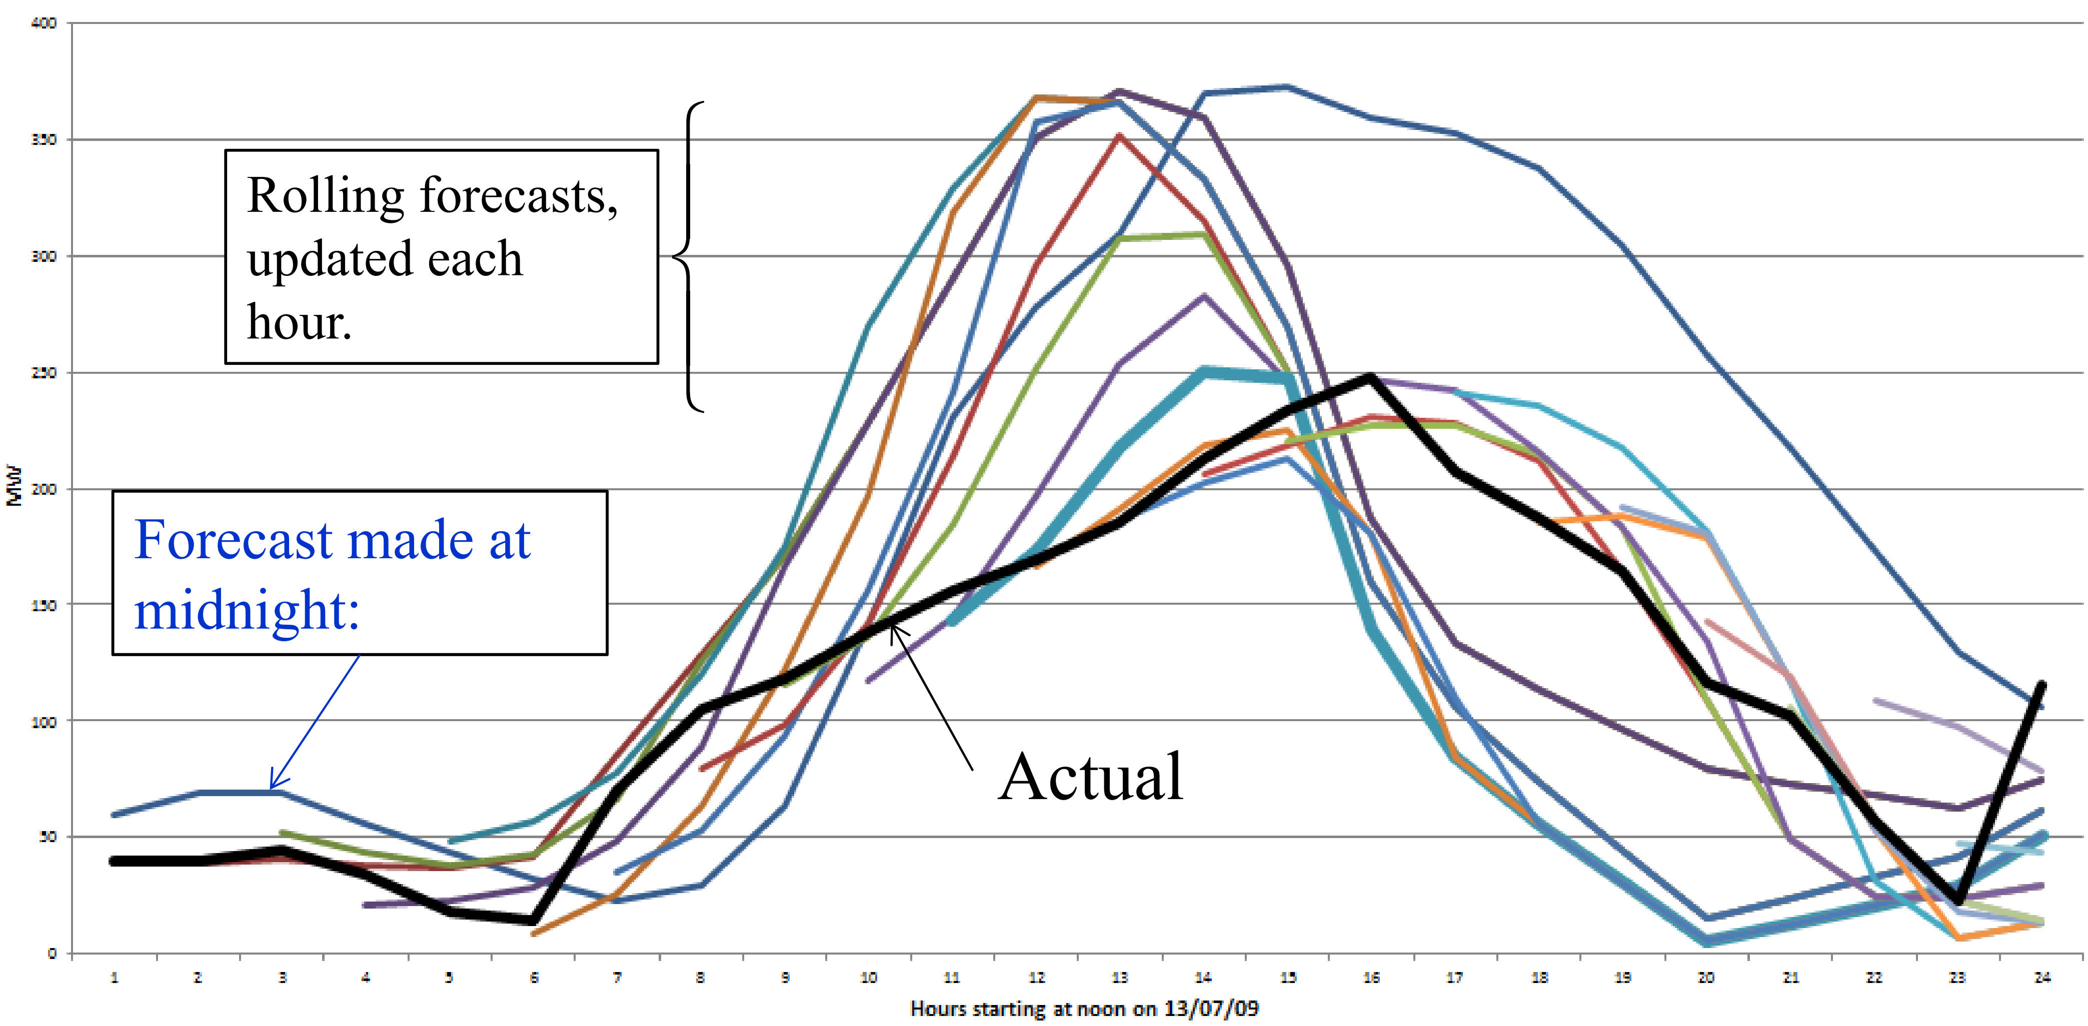

Variability and uncertainty

- ISOs handle uncertainty using a sequence of decisions
  - » Day-ahead, intermediate term and real-time planning each address different types of decisions

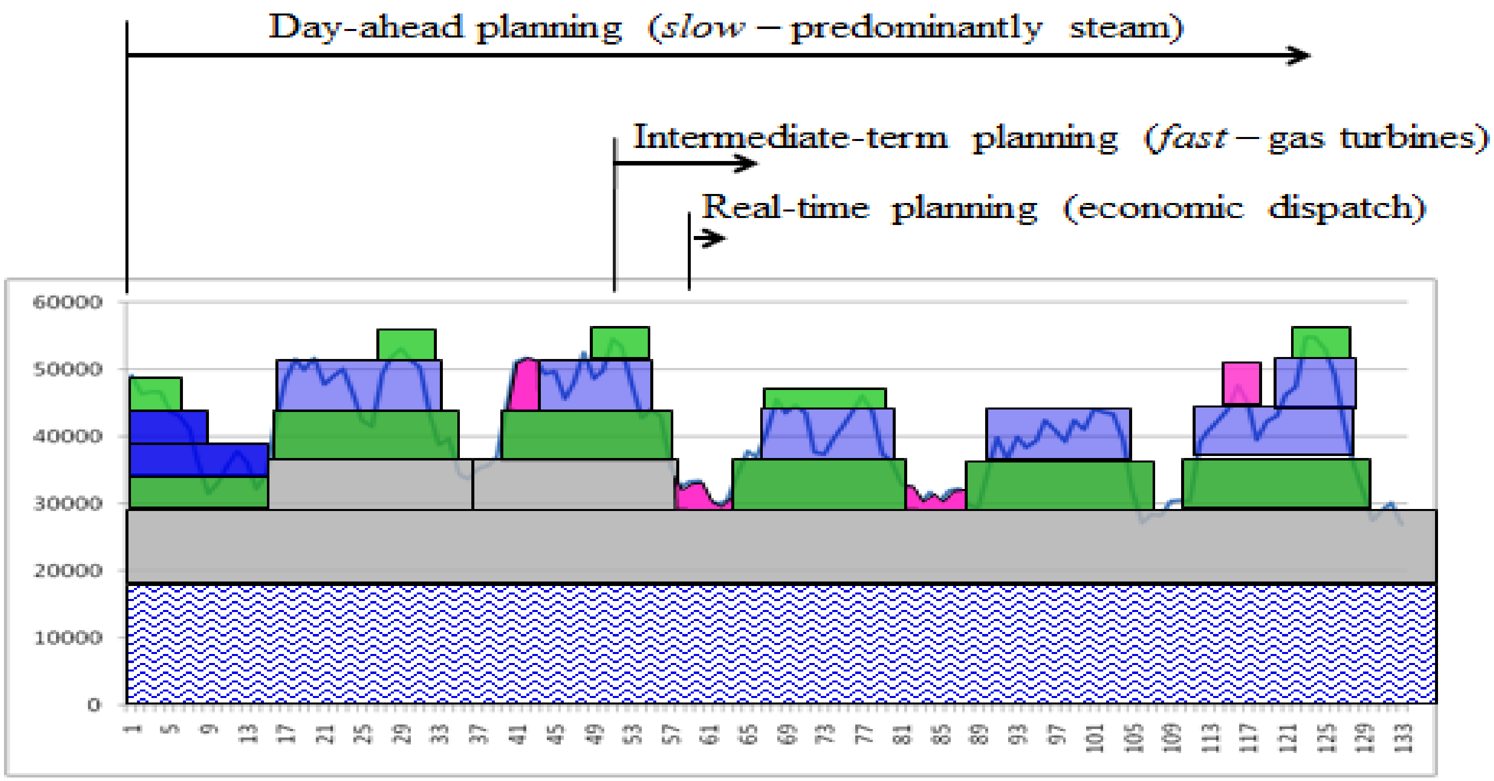

Lecture outline

- General modeling framework
- Stochastic lookahead policies
- The parametric cost function approximation
- Conclusions

Lecture outline

- General modeling framework
- Stochastic lookahead policies
- The parametric cost function approximation
- Conclusions

General modeling framework

- The objective function

$$\min_{\pi} \mathbb{E}^{\pi} \left\{ \sum_{t=0}^{T} C\left(S_t, X_t^{\pi}(S_t)\right) | S_0 \right\}$$

Expectation over all random outcomes

Cost function

State variable

Decision function (policy)

Finding the best policy

Given a *system model* (transition function)

$$S_{t+1} = S^M\left(S_t, x_t, W_{t+1}(\omega)\right)$$

We refer to this as the *base model* to avoid confusions with *lookahead models* we will introduce later.

General modeling framework

- There are two fundamental strategies for solving sequential decision problems:
  - » Policy search – Search over a parameterized class of functions for making decisions to optimize some metric.

$$\min_\theta E\left\{\sum_{t=0}^{T} C\left(S_t, X_t^\pi(S_t \mid \theta)\right) \mid S_0\right\}$$

  - » Lookahead approximations – Approximate the impact of a decision now on the future.

Policy search

- Policy search – Two types of policies:
  - » Analytical functions that directly map states to actions ("policy function approximations")
    - Lookup tables
      - "when in this state, take this action"
    - Parametric functions
      - Order-up-to policies: if inventory is less than s, order up to S.
    - Locally/semi/non parametric
      - Release rate from a reservoir as a function of reservoir level
  - » Minimizing analytical approximations of costs and/or constraints ("cost function approximations")
    - Optimizing a deterministic model modified to handle uncertainty (buffer stocks, schedule slack)

$$X^{CFA}(S_t \mid \theta) = \arg\min_{x_t \in \bar{X}_t^{\pi}(\theta)} \bar{C}^{\pi}(S_t, x_t \mid \theta)$$

Lookahead policies

- Lookahead approximations – Approximate the impact of a decision now on the future:
  - » Approximate lookahead models – Optimize over an approximate model of the future:
    - Replace uncertain future with a deterministic approximation
    - Model future with a small sample of uncertain outcomes

$$X_t^*(S_t) = \arg\min_{x_t} \left(C(S_t, x_t) + \mathbb{E}\left\{ \min_{\pi\in\Pi} \left\{ \mathbb{E} \sum_{t'=t+1}^{T} C(S_{t'}, X_{t'}^{\pi}(S_{t'})) \,|\, S_{t+1} \right\} |\, S_t, x_t \right\} \right)$$

  - » Approximating the value of being in a downstream state using machine learning ("value function approximations")

$$X_t^{VFA}(S_t) = \arg\min_{x_t} \left(C(S_t, x_t) + \bar{V}_t^x \left(S_t^x(S_t, x_t) \right) \right)$$

Lookahead policies

- The ultimate lookahead policy is optimal

$$X_t^*(S_t) = \arg\max_{x_t} \left(C(S_t, x_t) + \mathbb{E} \left\{ \max_{\pi \in \Pi} \left\{ \mathbb{E} \sum_{t'=t+1}^{T} C(S_{t'}, X_{t'}^{\pi}(S_{t'})) \mid S_{t+1} \right\} \mid S_t, x_t \right\} \right)$$

Maximization that we cannot compute

Expectations that we cannot compute

Designing policies

- The ultimate lookahead policy is optimal

$$X_t^*(S_t) = \arg\max_{x_t} \left(C(S_t, x_t) + \mathbb{E}\left\{ \max_{\pi \in \Pi} \left\{ \mathbb{E} \sum_{t'=t+1}^{T} C(S_{t'}, X_{t'}^{\pi}(S_{t'})) \mid S_{t+1} \right\} \mid S_t, x_t \right\} \right)$$

- Instead, we have to solve an approximation called the *lookahead model:*

$$X_t^*(S_t) = \arg\max_{x_t} \left(C(S_t, x_t) + \tilde{\mathbb{E}}\left\{ \max_{\tilde{\pi} \in \tilde{\Pi}} \left\{ \tilde{\mathbb{E}} \sum_{t'=t+1}^{t+H} C(\tilde{S}_{tt'}, \tilde{X}_{t'}^{\pi}(\tilde{S}_{tt'})) \mid \tilde{S}_{t,t+1} \right\} \mid \tilde{S}_{tt}, x_t \right\} \right)$$

  » *A lookahead policy* works by approximating the *lookahead model.*

Lookahead policies

- We use a series of approximations:
  - » Horizon truncation – Replacing a longer horizon problem with a shorter horizon
  - » Stage aggregation – Replacing multistage problems with two-stage approximation.
  - » Outcome aggregation/sampling – Simplifying the exogenous information process
  - » Discretization – Of time, states and decisions
  - » Dimensionality reduction – We may ignore some variables (such as forecasts) in the lookahead model that we capture in the base model (these become *latent* variables in the lookahead model).

Four (meta)classes of policies

Policy search

1) **Policy function approximations (PFAs)**
   - » Lookup tables, rules, parametric/nonparametric functions

2) **Cost function approximation (CFAs)**
   - » $X^{CFA}(S_t \mid \theta) = \arg\min_{x_t \in \bar{\mathcal{X}}_t^{\pi}(\theta)} \bar{C}^{\pi}(S_t, x_t \mid \theta)$

Lookahead approximations

3) **Policies based on value function approximations (VFAs)**
   - » $X_t^{VFA}(S_t) = \arg\min_{x_t} \left(C(S_t, x_t) + \bar{V}_t^x \left(S_t^x(S_t, x_t) \right) \right)$

4) **Direct lookahead policies (DLAs)**
   - » *Deterministic lookahead/rolling horizon proc./model predictive control*

$$X_t^{LA-D}(S_t) = \arg \min_{\tilde{x}_{tt}, \ldots, \tilde{x}_{t,t+H}} C(\tilde{S}_{tt}, \tilde{x}_{tt}) + \sum_{t'=t+1}^{T} C(\tilde{S}_{tt'}, \tilde{x}_{tt'})$$

   - » *Chance constrained programming*

$$P[A_t x_t \leq f(W)] \leq 1 - \delta$$

   - » *Stochastic lookahead /stochastic prog/Monte Carlo tree search*

$$X_t^{LA-S}(S_t) = \arg\min_{\tilde{x}_{tt}, \tilde{x}_{t,t+1}, \ldots, \tilde{x}_{t,t+T}} C(\tilde{S}_{tt}, \tilde{x}_{tt}) + \sum_{\tilde{\omega} \in \tilde{\Omega}_t} p(\tilde{\omega}) \sum_{t'=t+1}^{T} C(\tilde{S}_{tt'}(\tilde{\omega}), \tilde{x}_{tt'}(\tilde{\omega}))$$

   - » *"Robust optimization"*

$$X_t^{LA-RO}(S_t) = \arg \min_{\tilde{x}_{tt}, \ldots, \tilde{x}_{t,t+H}} \max_{w \in W_t(\theta)} C(\tilde{S}_{tt}, \tilde{x}_{tt}) + \sum_{t'=t+1}^{T} C(\tilde{S}_{tt'}(w), \tilde{x}_{tt'}(w))$$

Lookahead policies

- ❑ Lookahead policies peek into the future
  - » Optimize over deterministic lookahead model

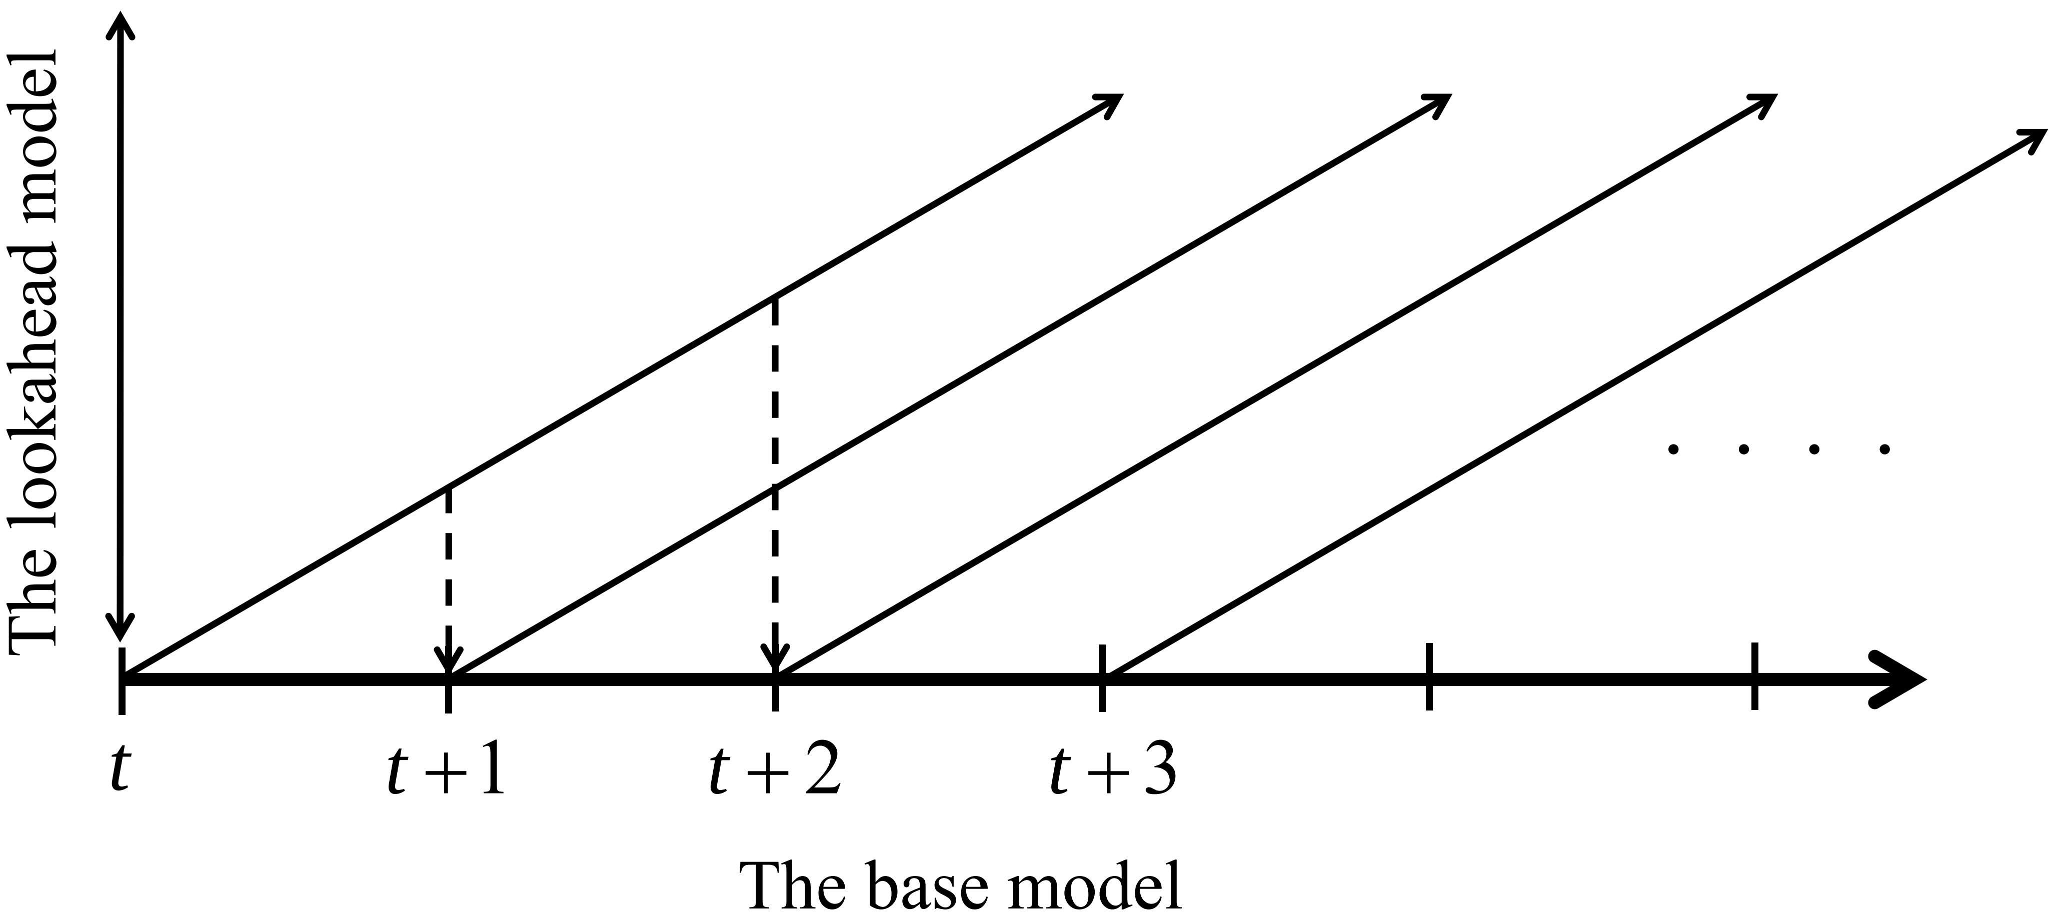

Lookahead policies

- ❑ Lookahead policies peek into the future
  - » Optimize over deterministic lookahead model

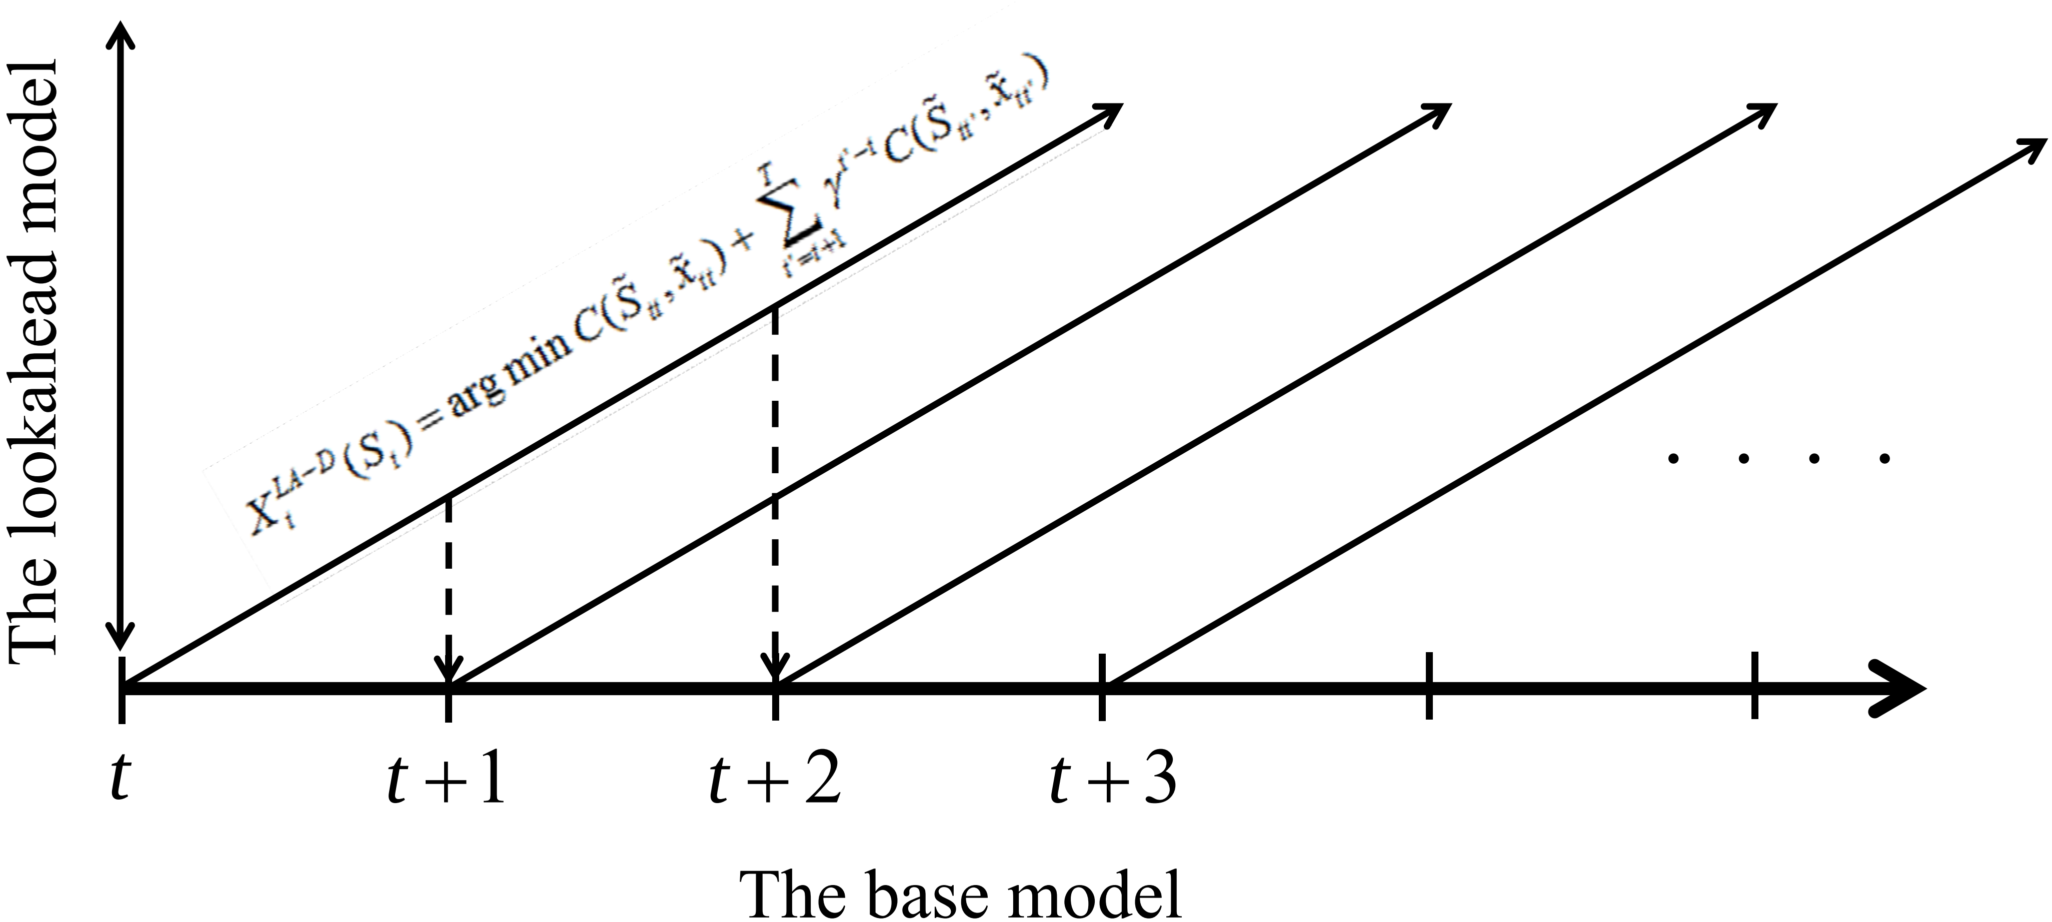

Lookahead policies

- Lookahead policies peek into the future
  - » Optimize over deterministic lookahead model

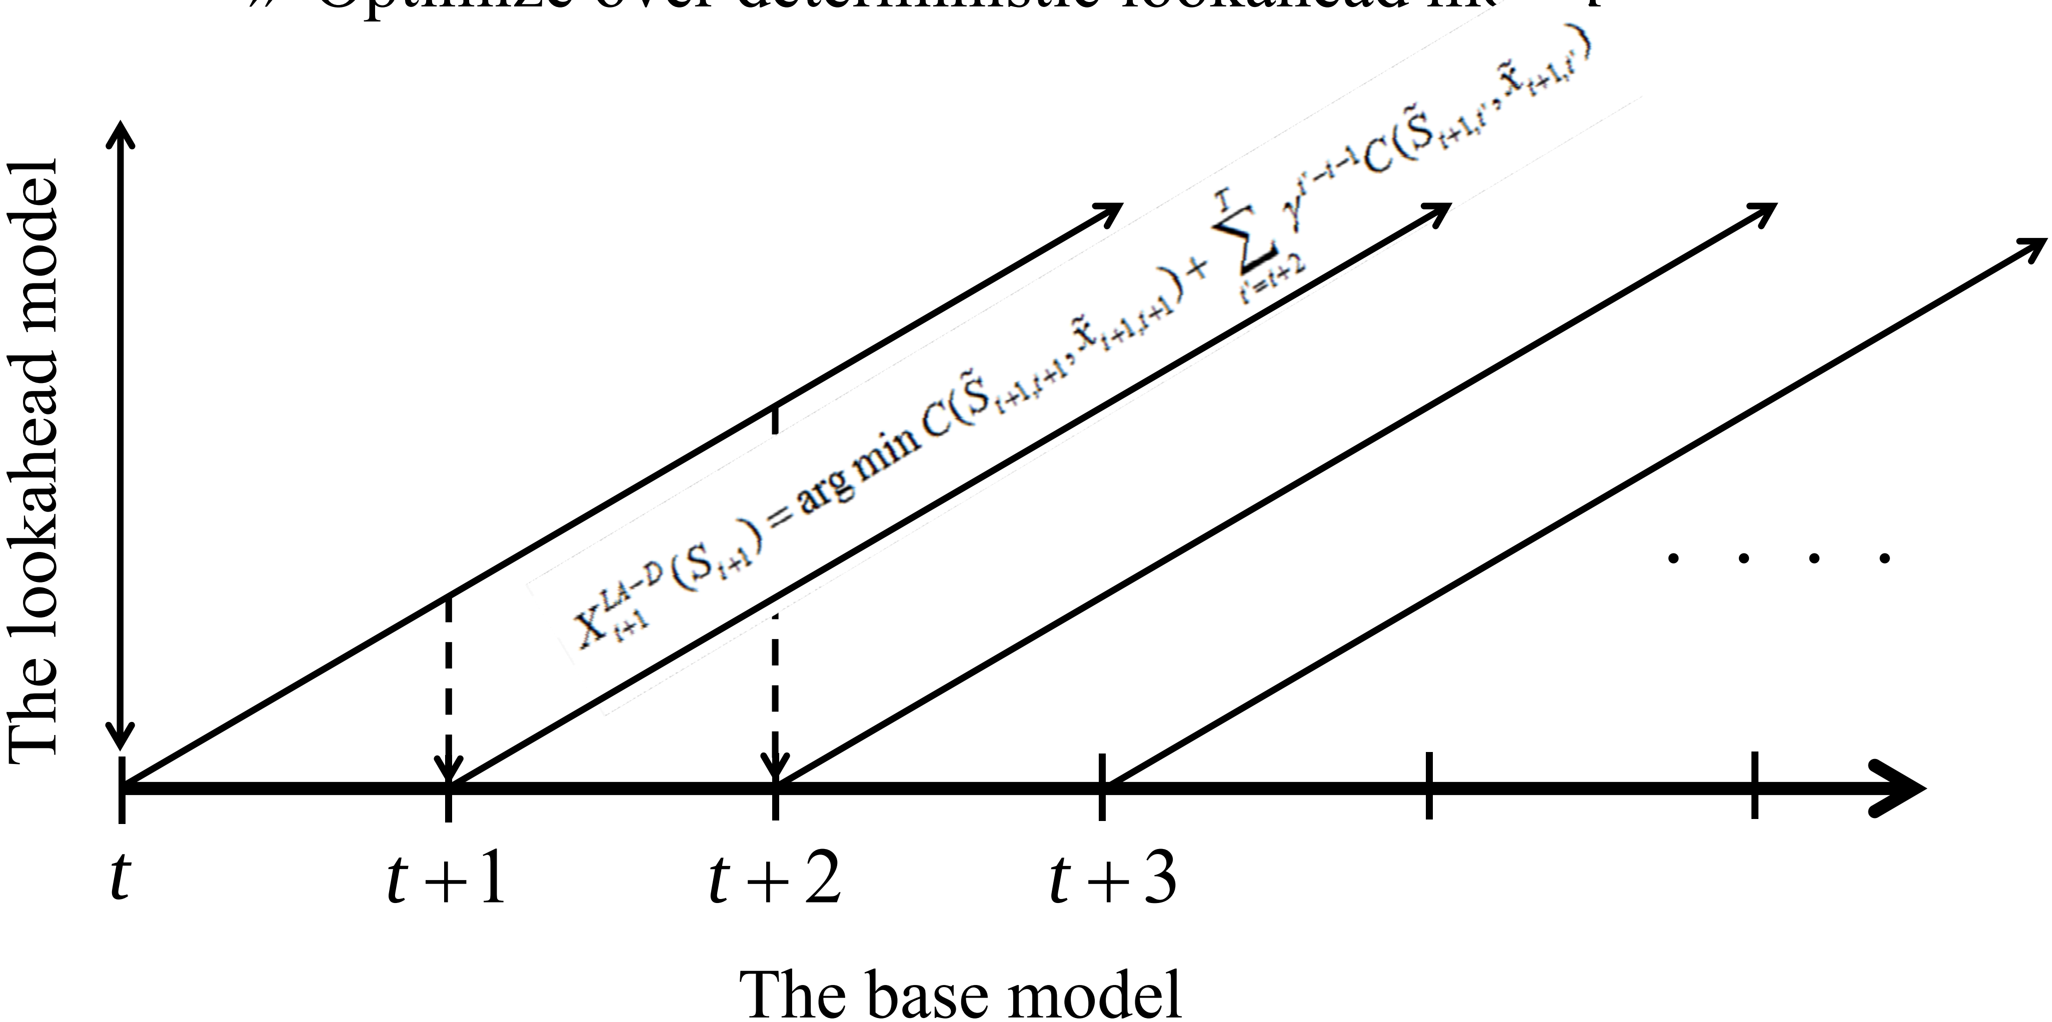

Lookahead policies

- ❑ Lookahead policies peek into the future
  - » Optimize over deterministic lookahead model

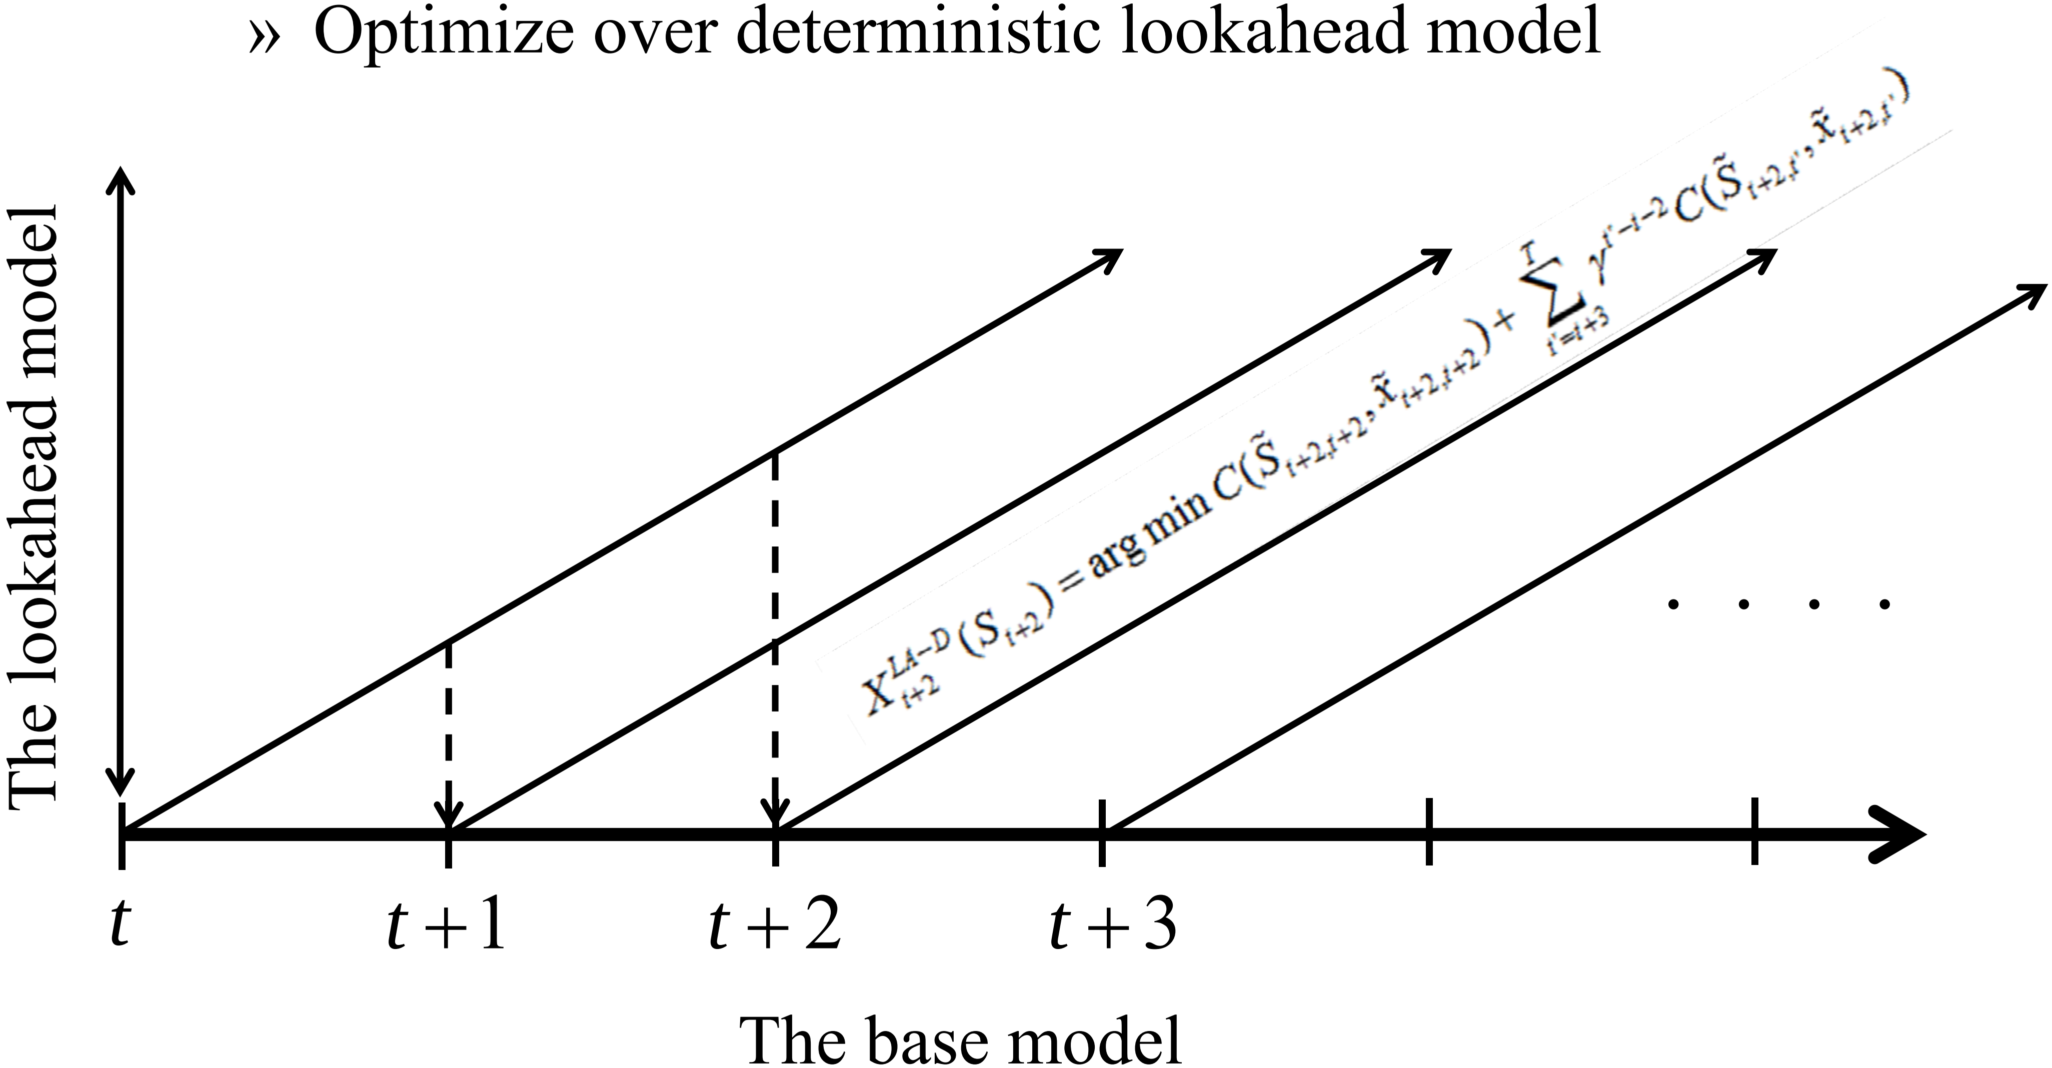

Lecture outline

- General modeling framework
- Stochastic lookahead policies
- The parametric cost function approximation
- Conclusions

Stochastic lookahead policies

- Stochastic lookahead
  - » Here, we approximate the information model by using a Monte Carlo sample to create a scenario tree:

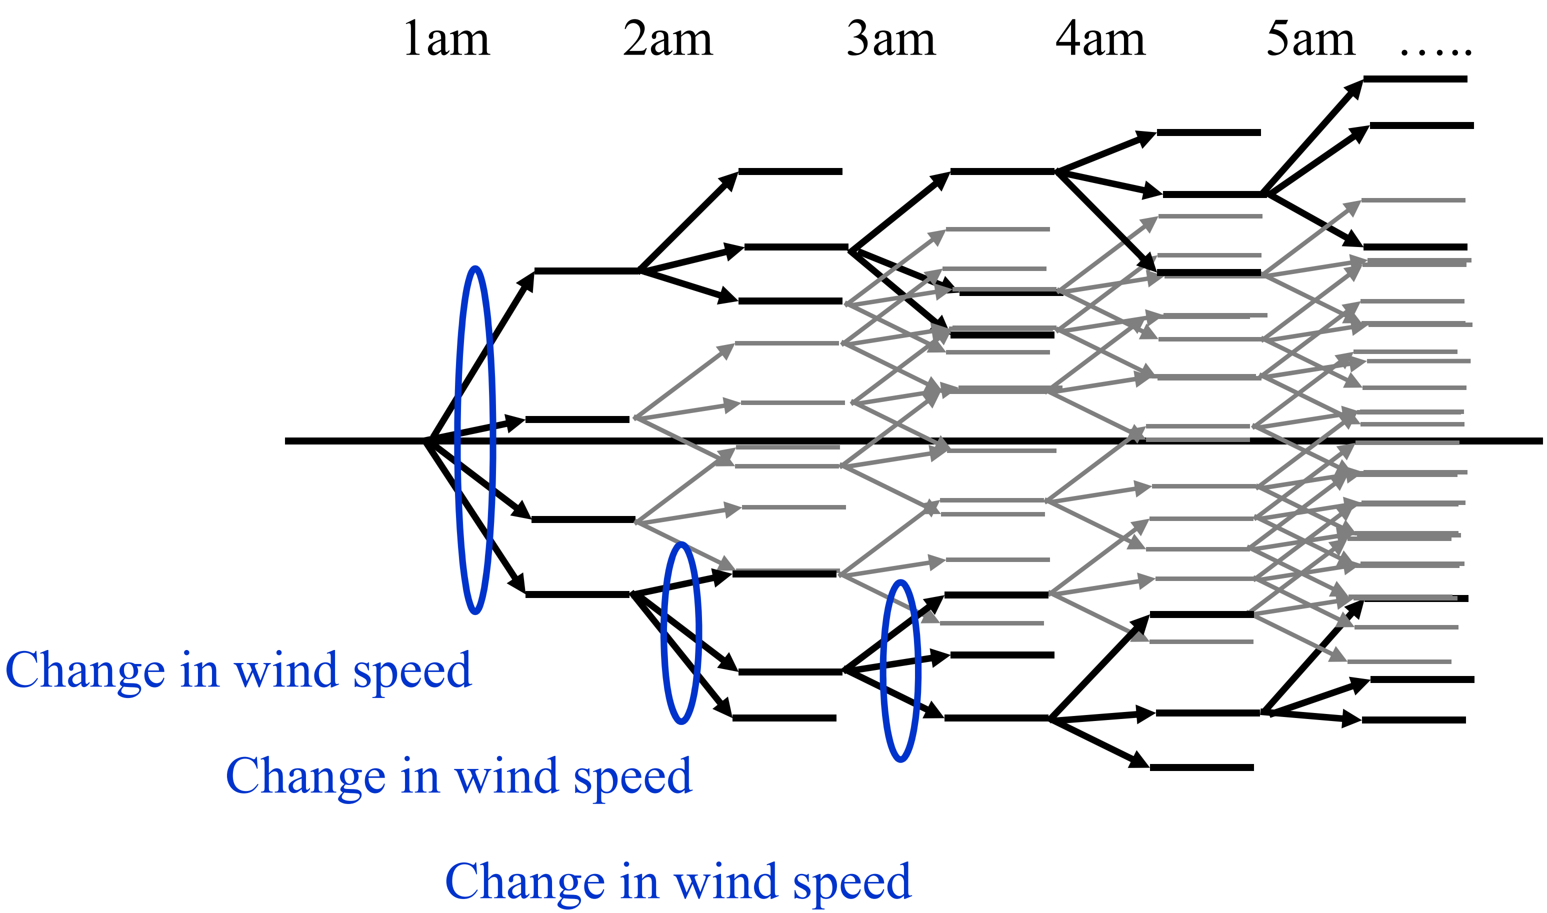

Stochastic lookahead policies

- We can then simulate this *lookahead policy* over time:

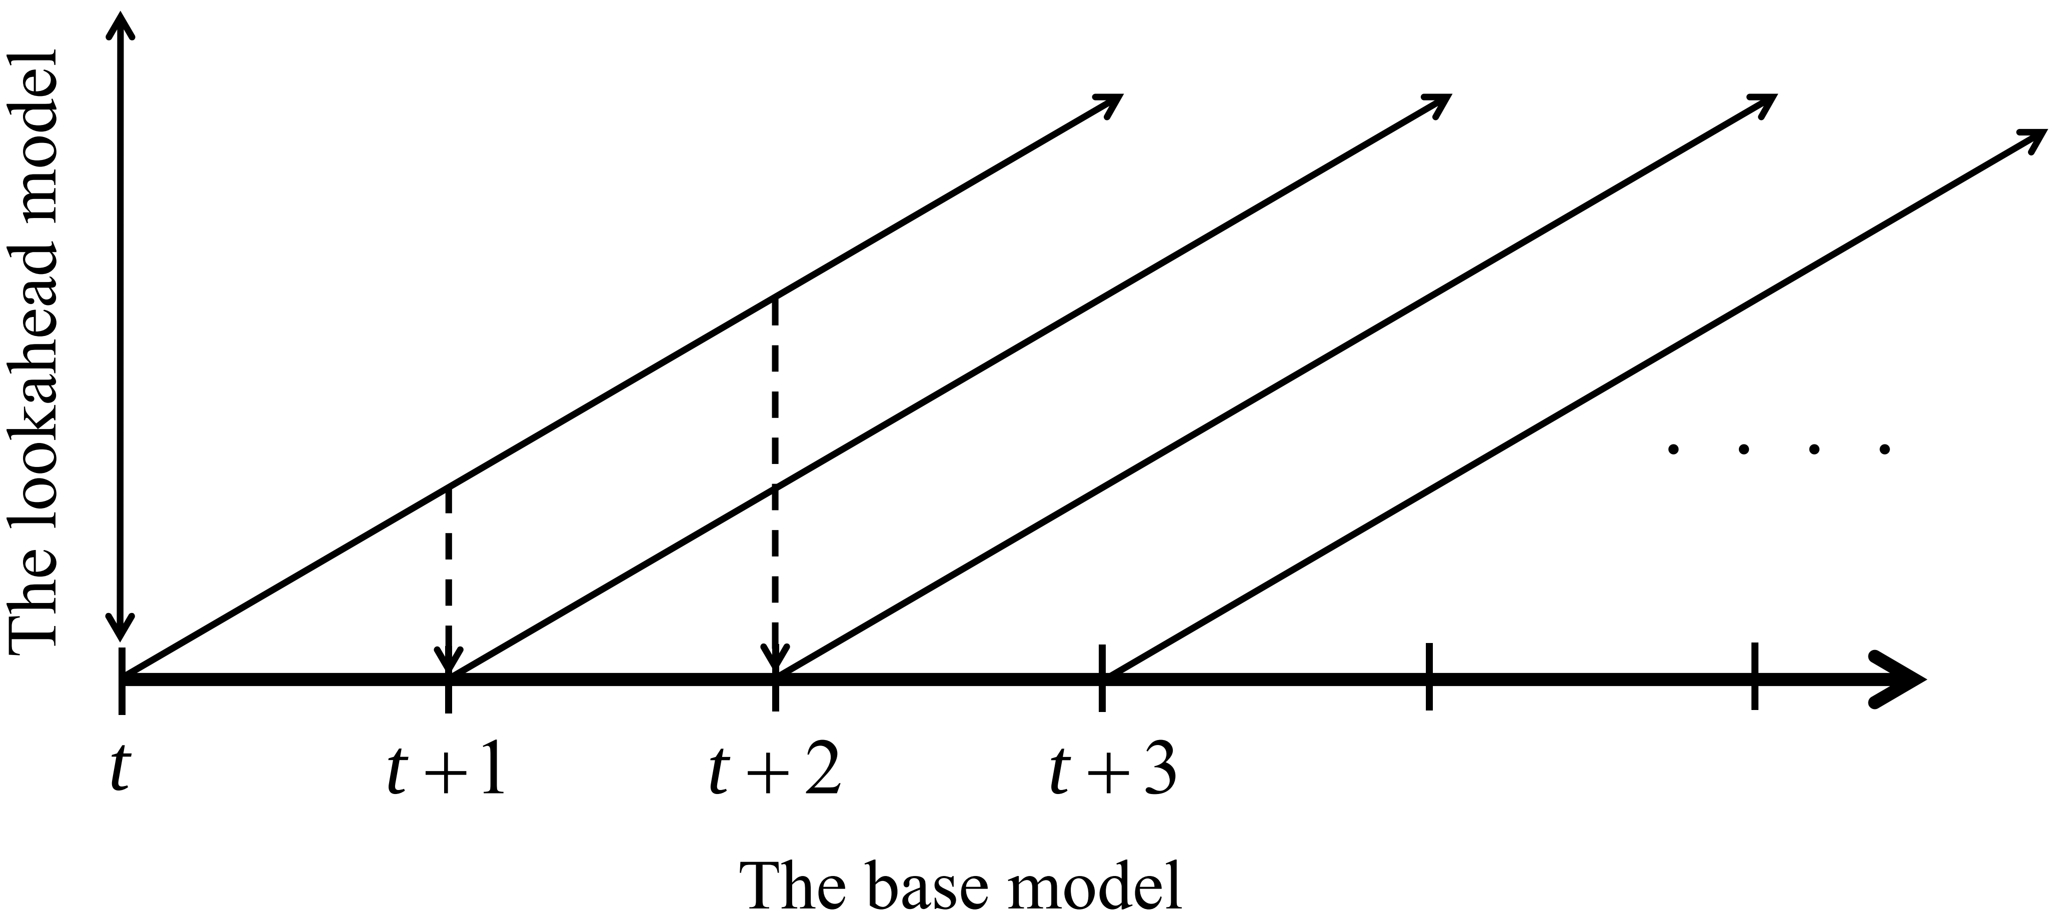

Stochastic lookahead policies

- We can then simulate this *lookahead policy* over time:

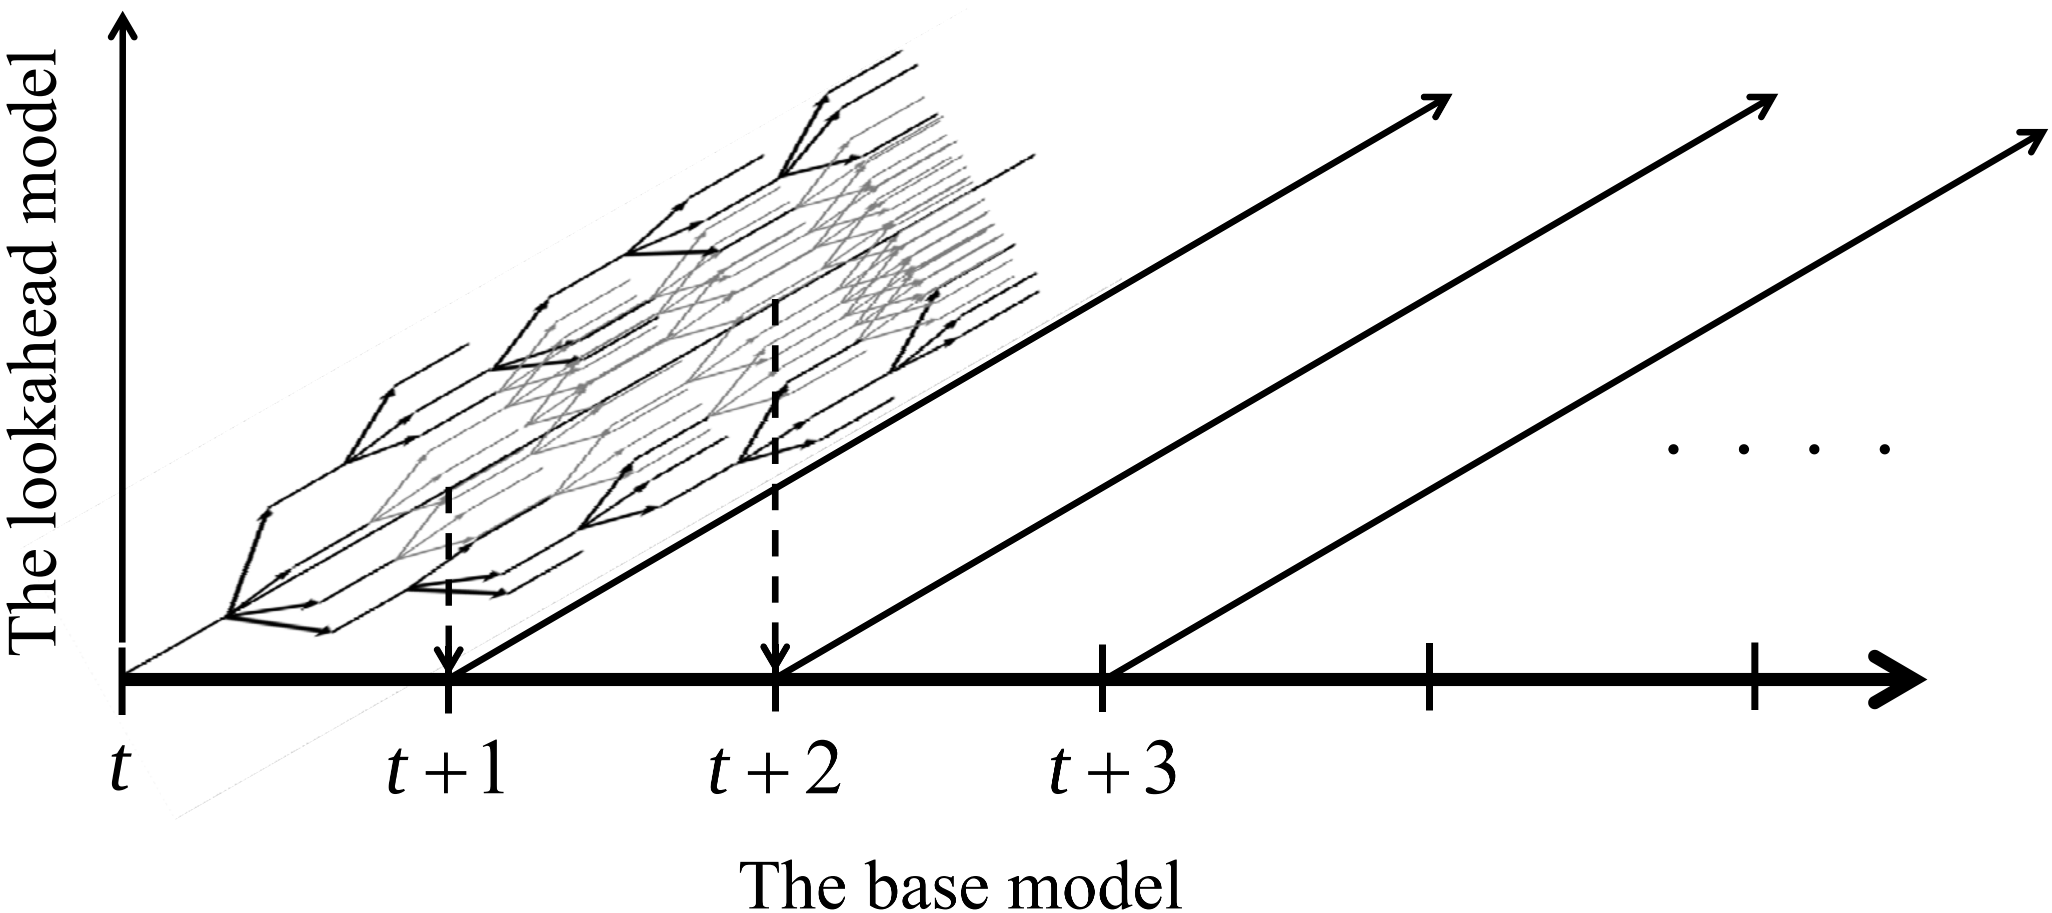

Stochastic lookahead policies

- We can then simulate this *lookahead policy* over time:

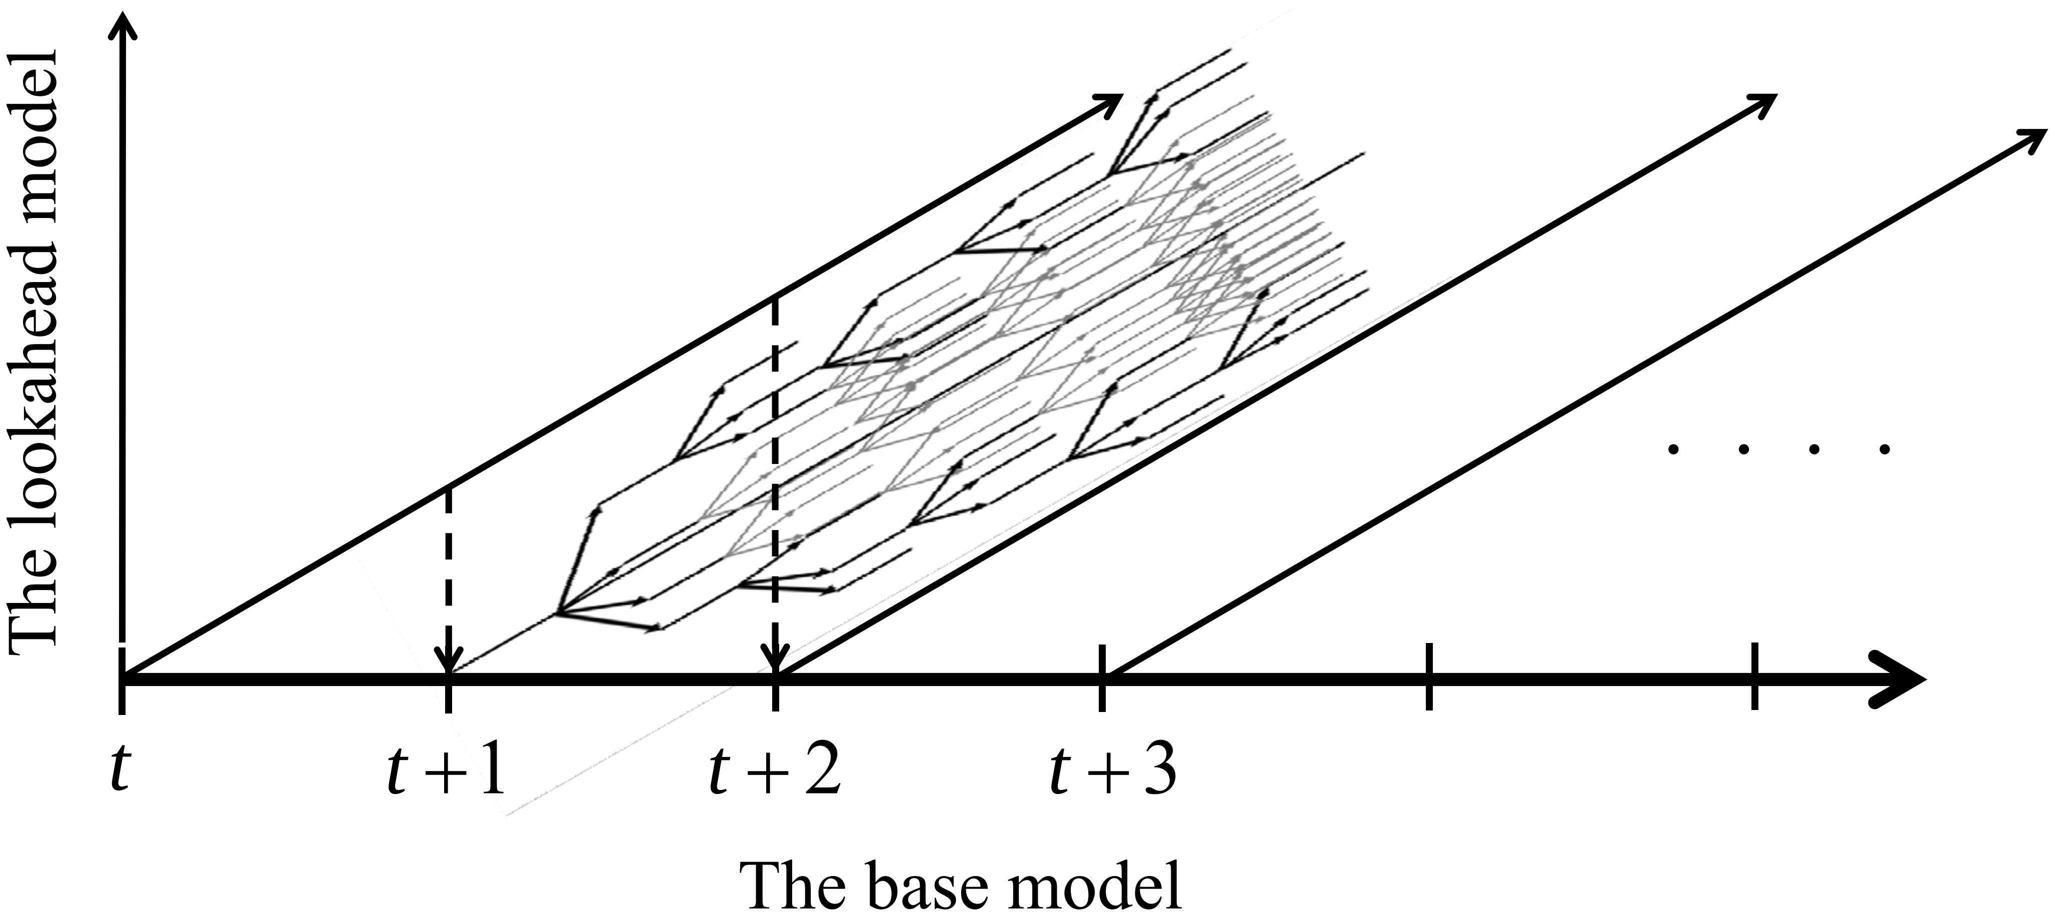

Stochastic lookahead policies

- We can then simulate this *lookahead policy* over time:

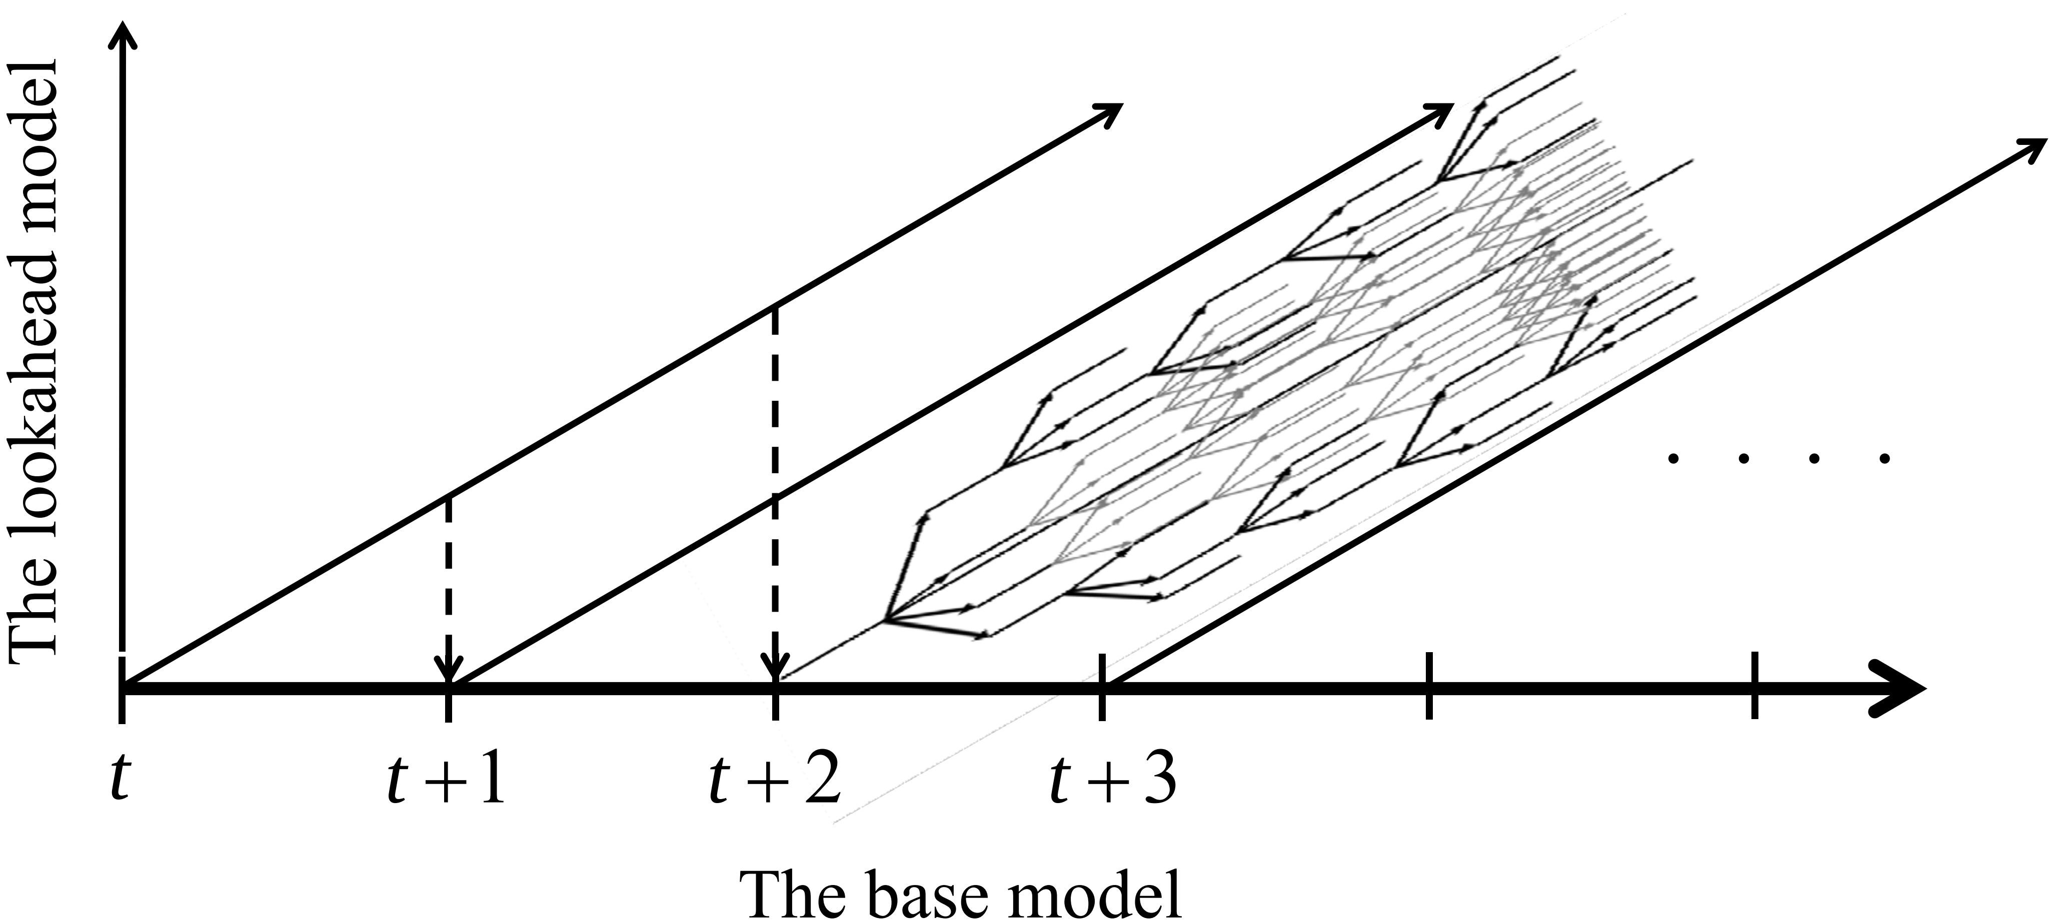

Stochastic lookahead policies

- Two stage lookahead approximation

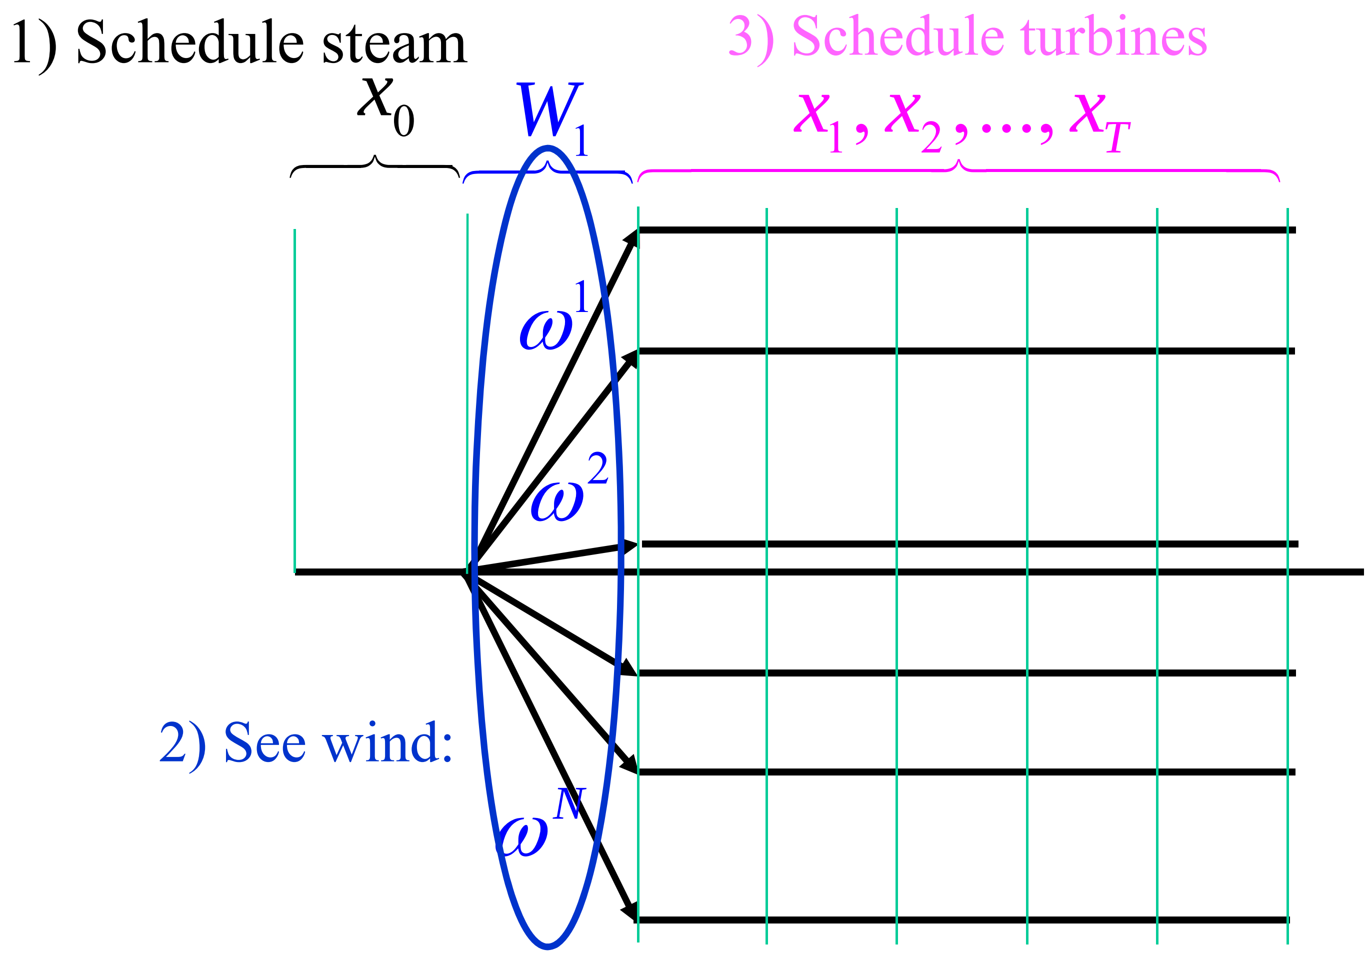

Stochastic lookahead policies

- Creating wind scenarios (Scenario #1)

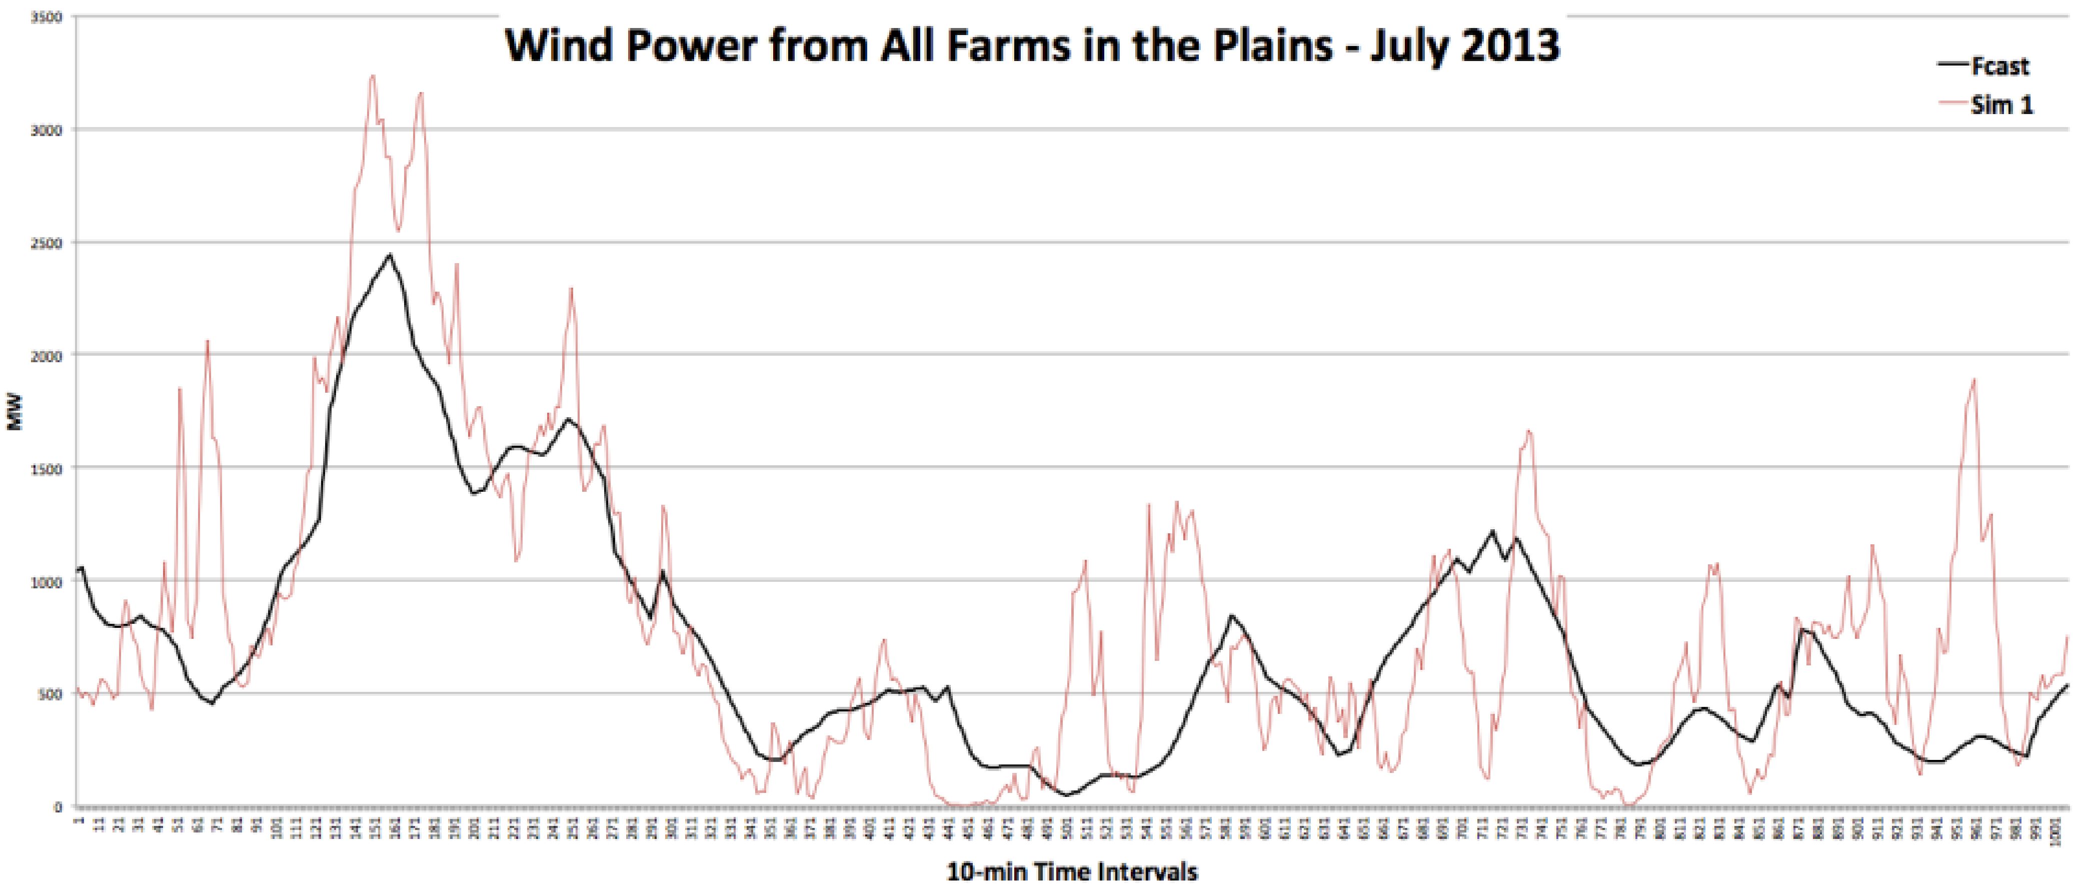

Stochastic lookahead policies

- Creating wind scenarios (Scenario #2)

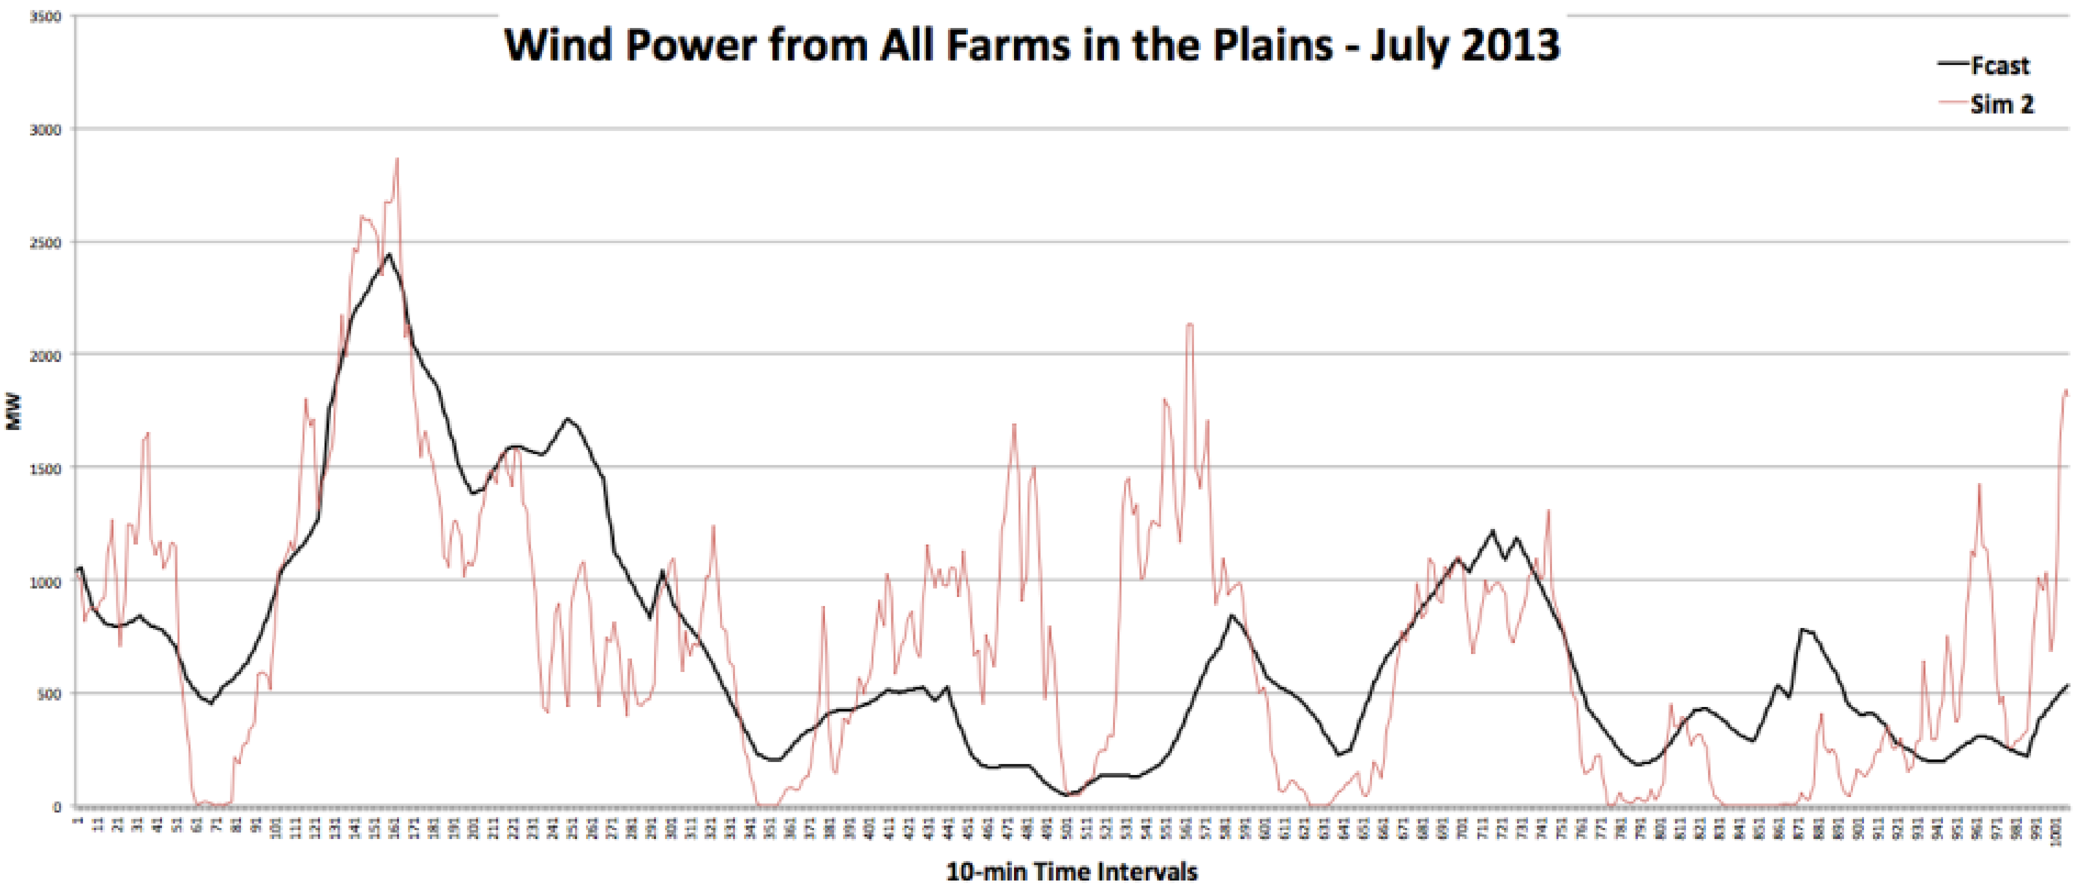

Stochastic lookahead policies

- Creating wind scenarios (Scenario #3)

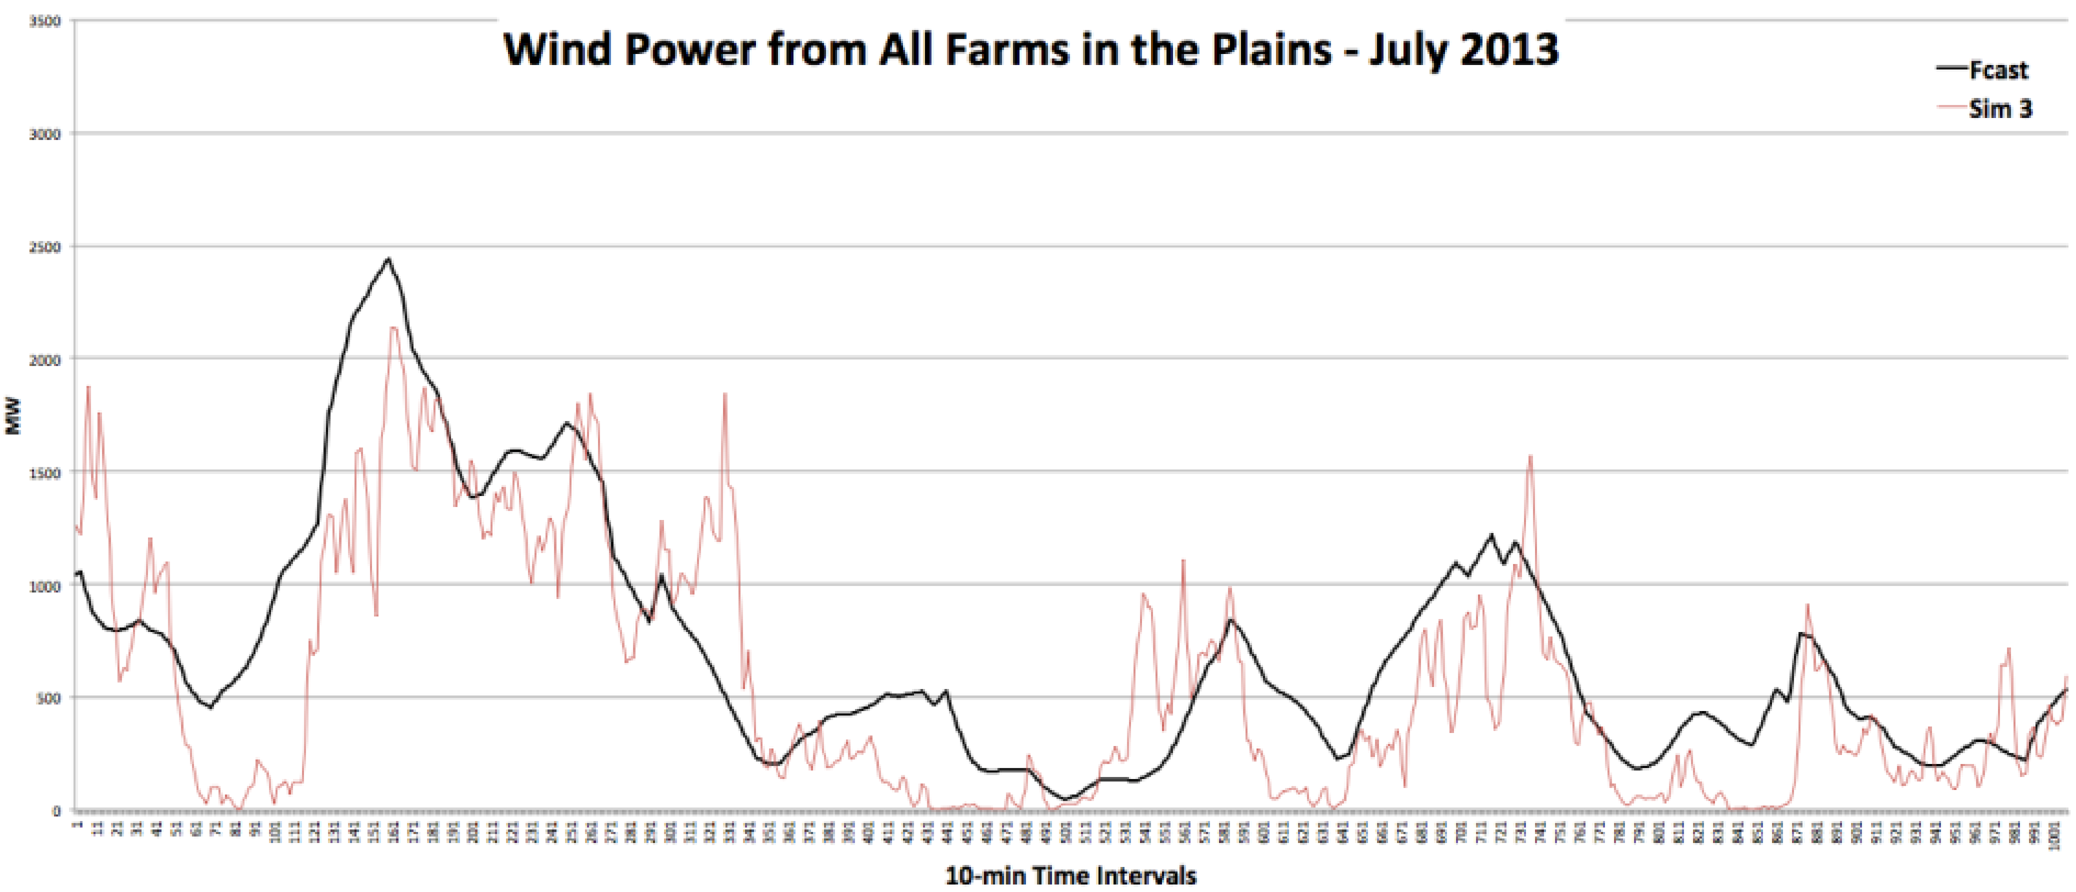

Stochastic lookahead policies

- Creating wind scenarios (Scenario #4)

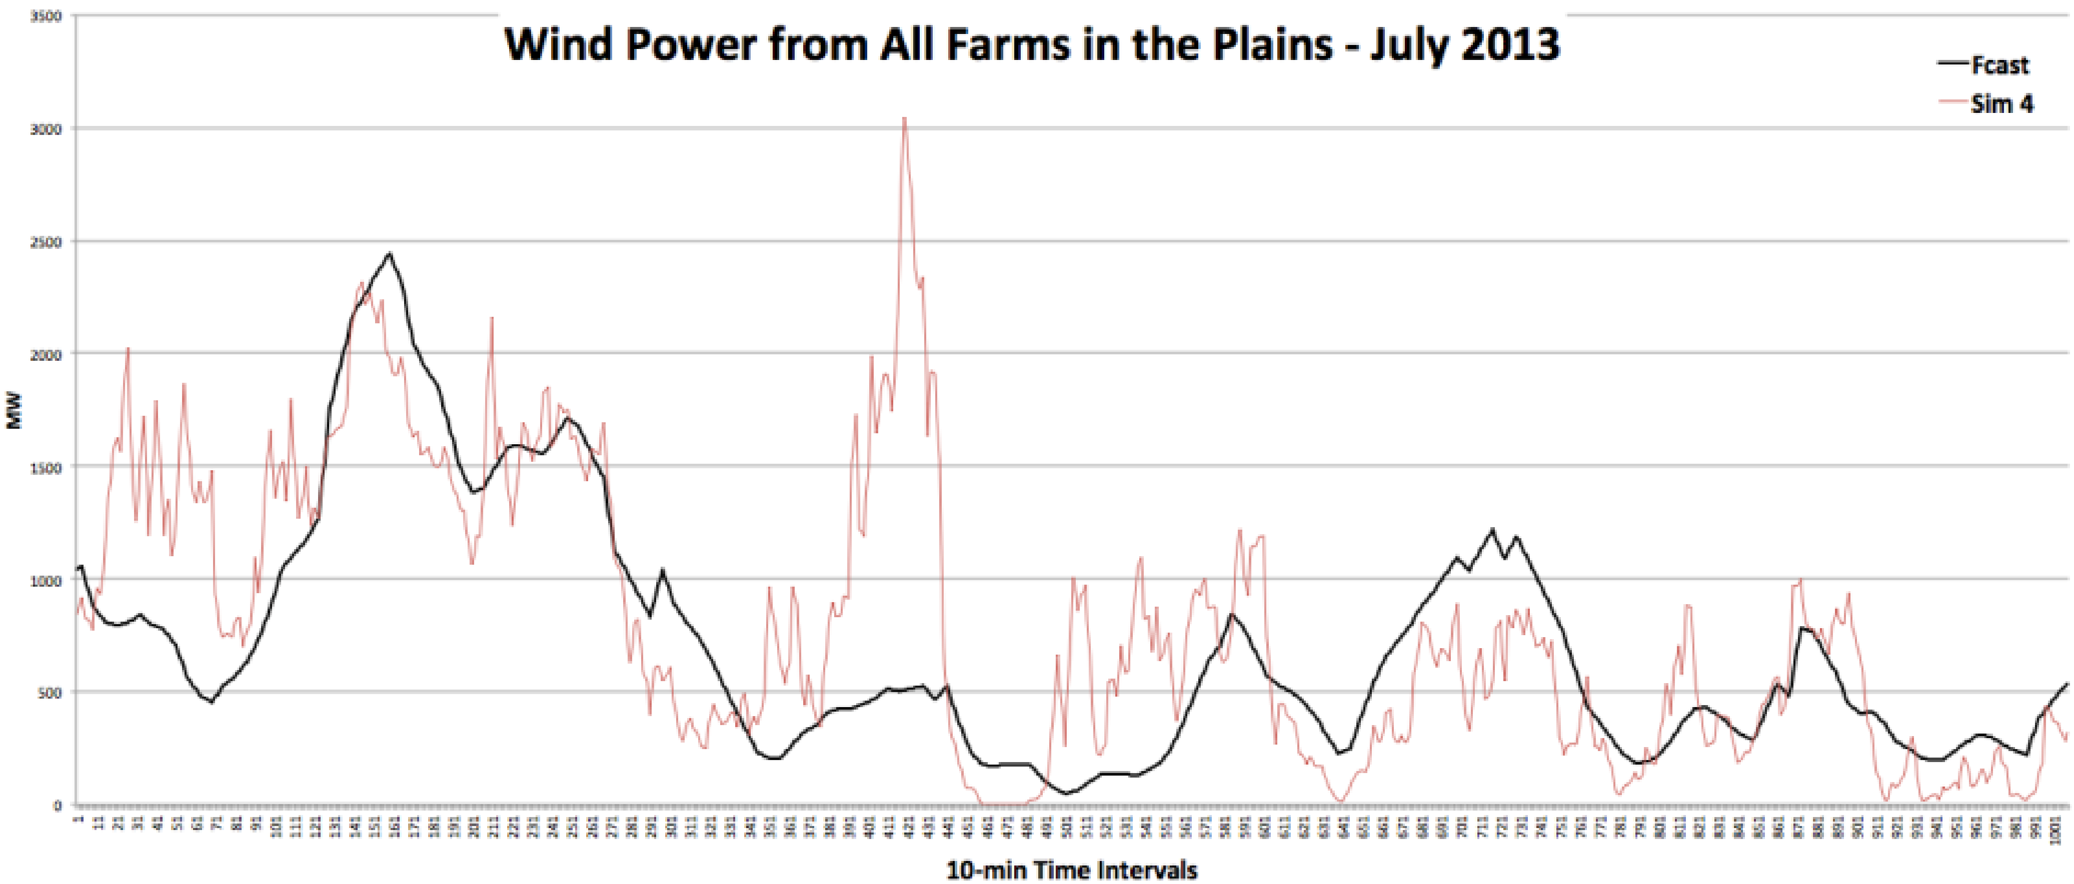

Stochastic lookahead policies

- Creating wind scenarios (Scenario #5)

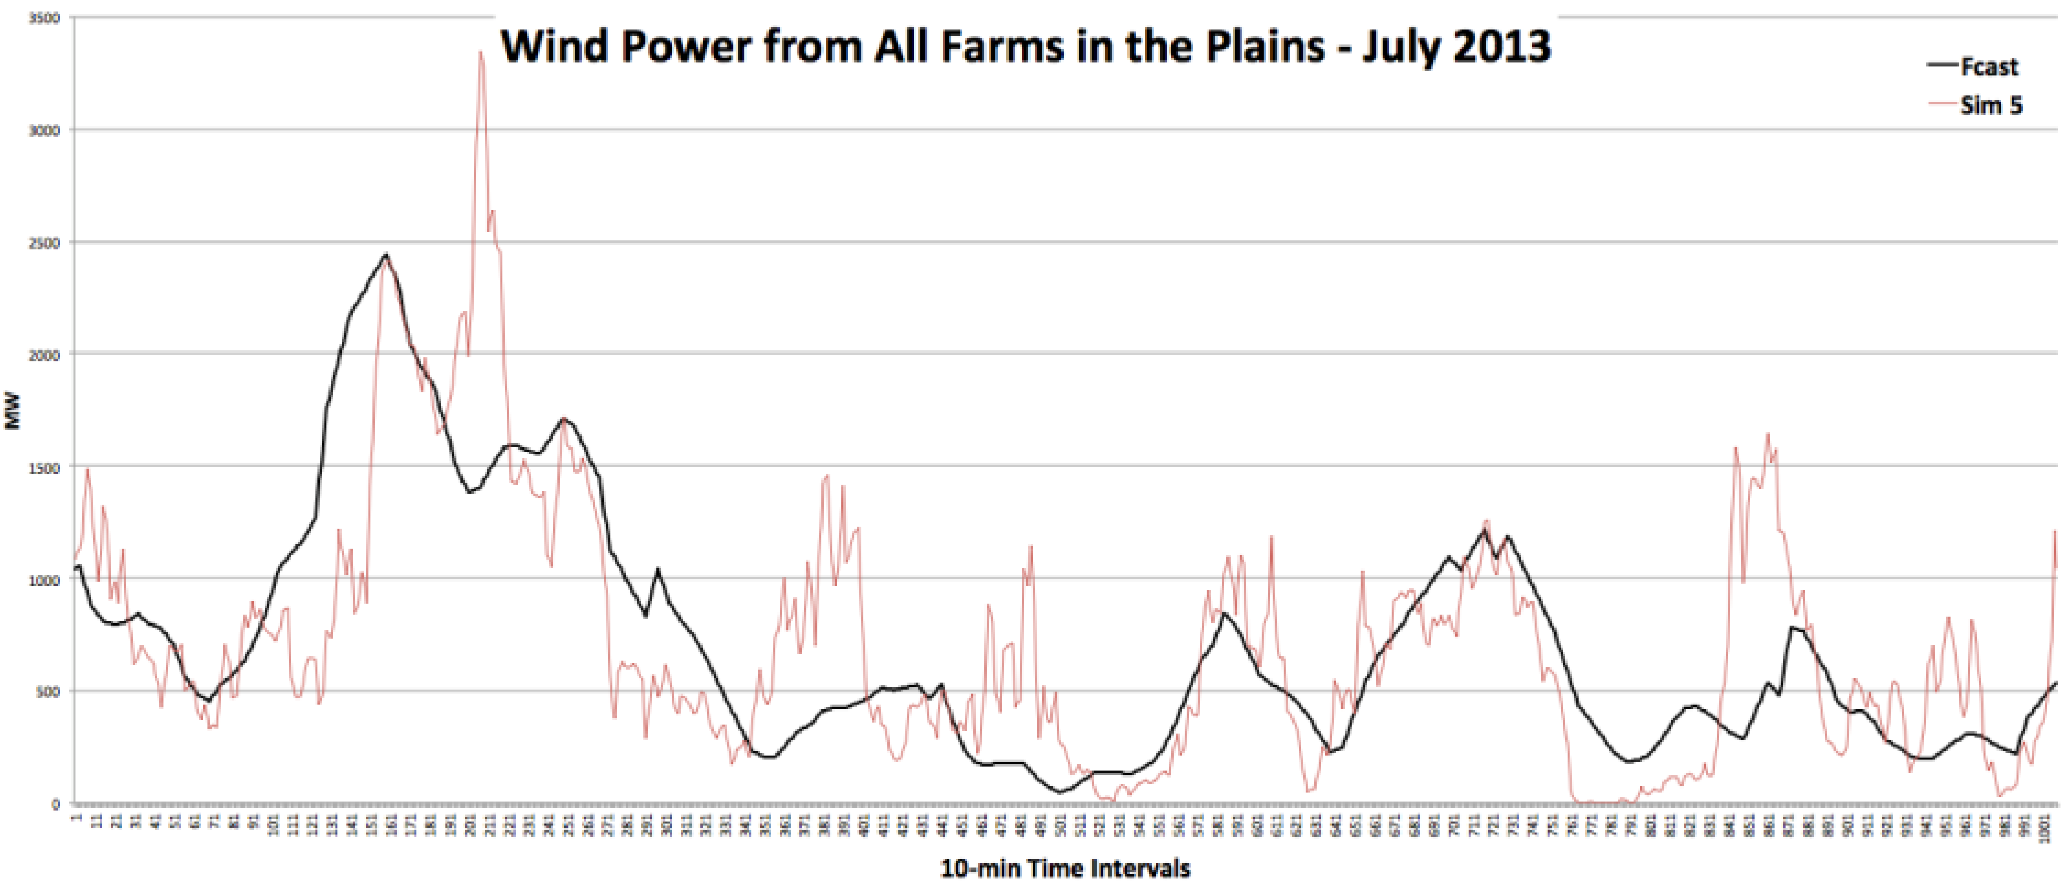

Stochastic lookahead policies

- The two-stage approximation

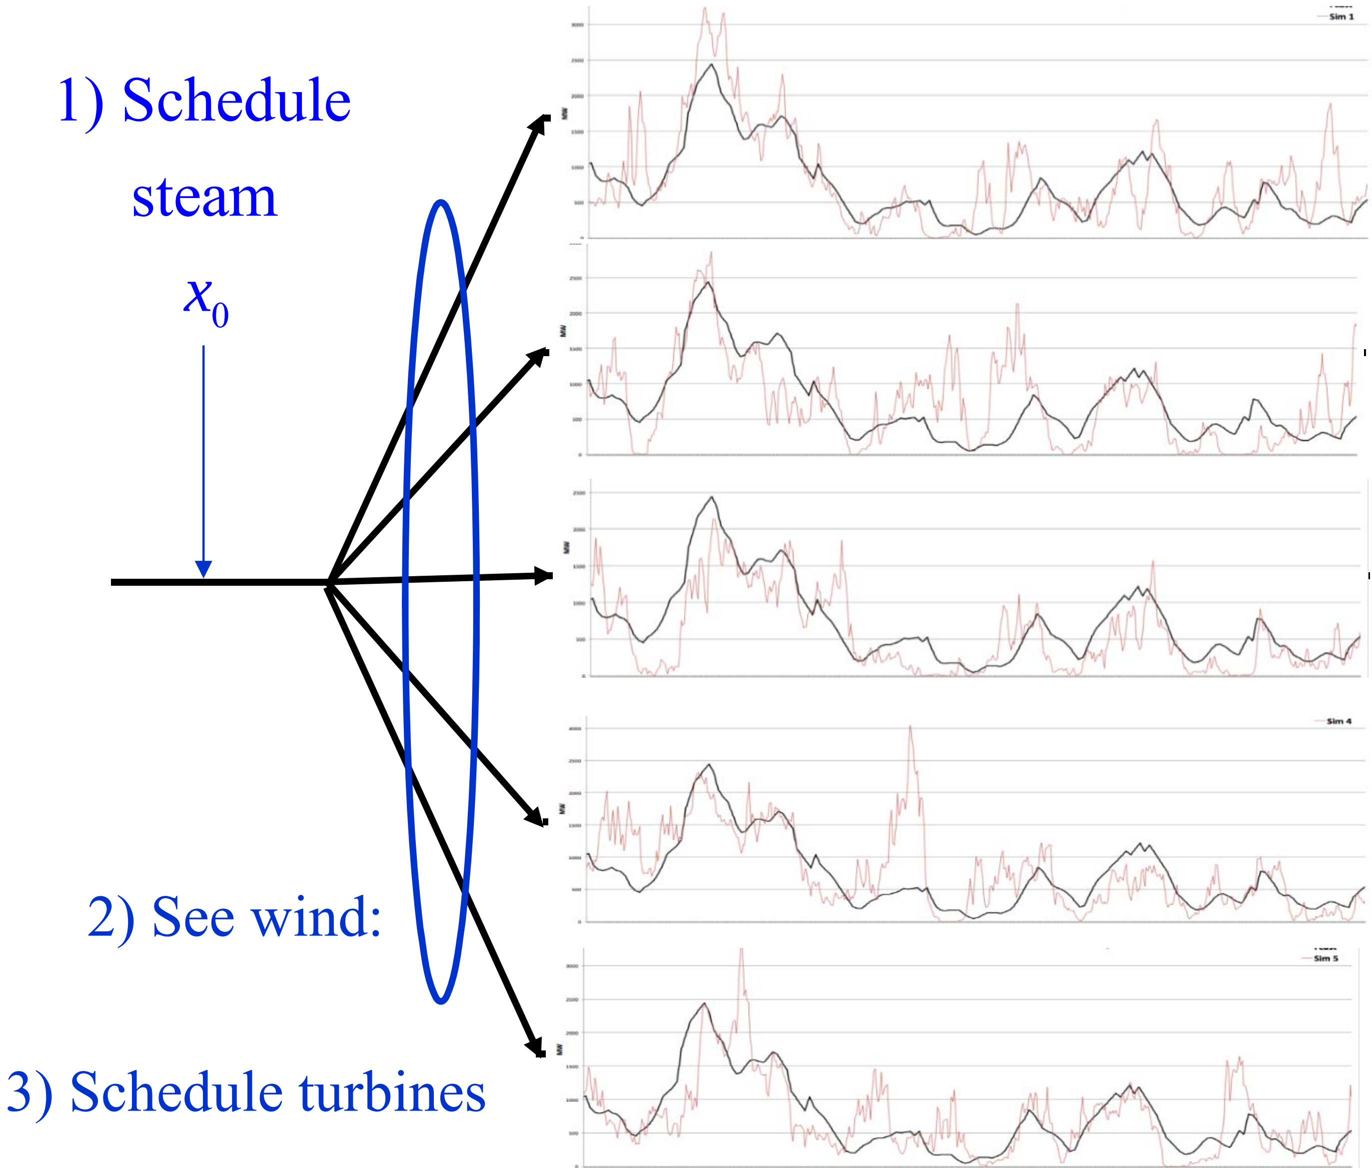

Stochastic lookahead policies

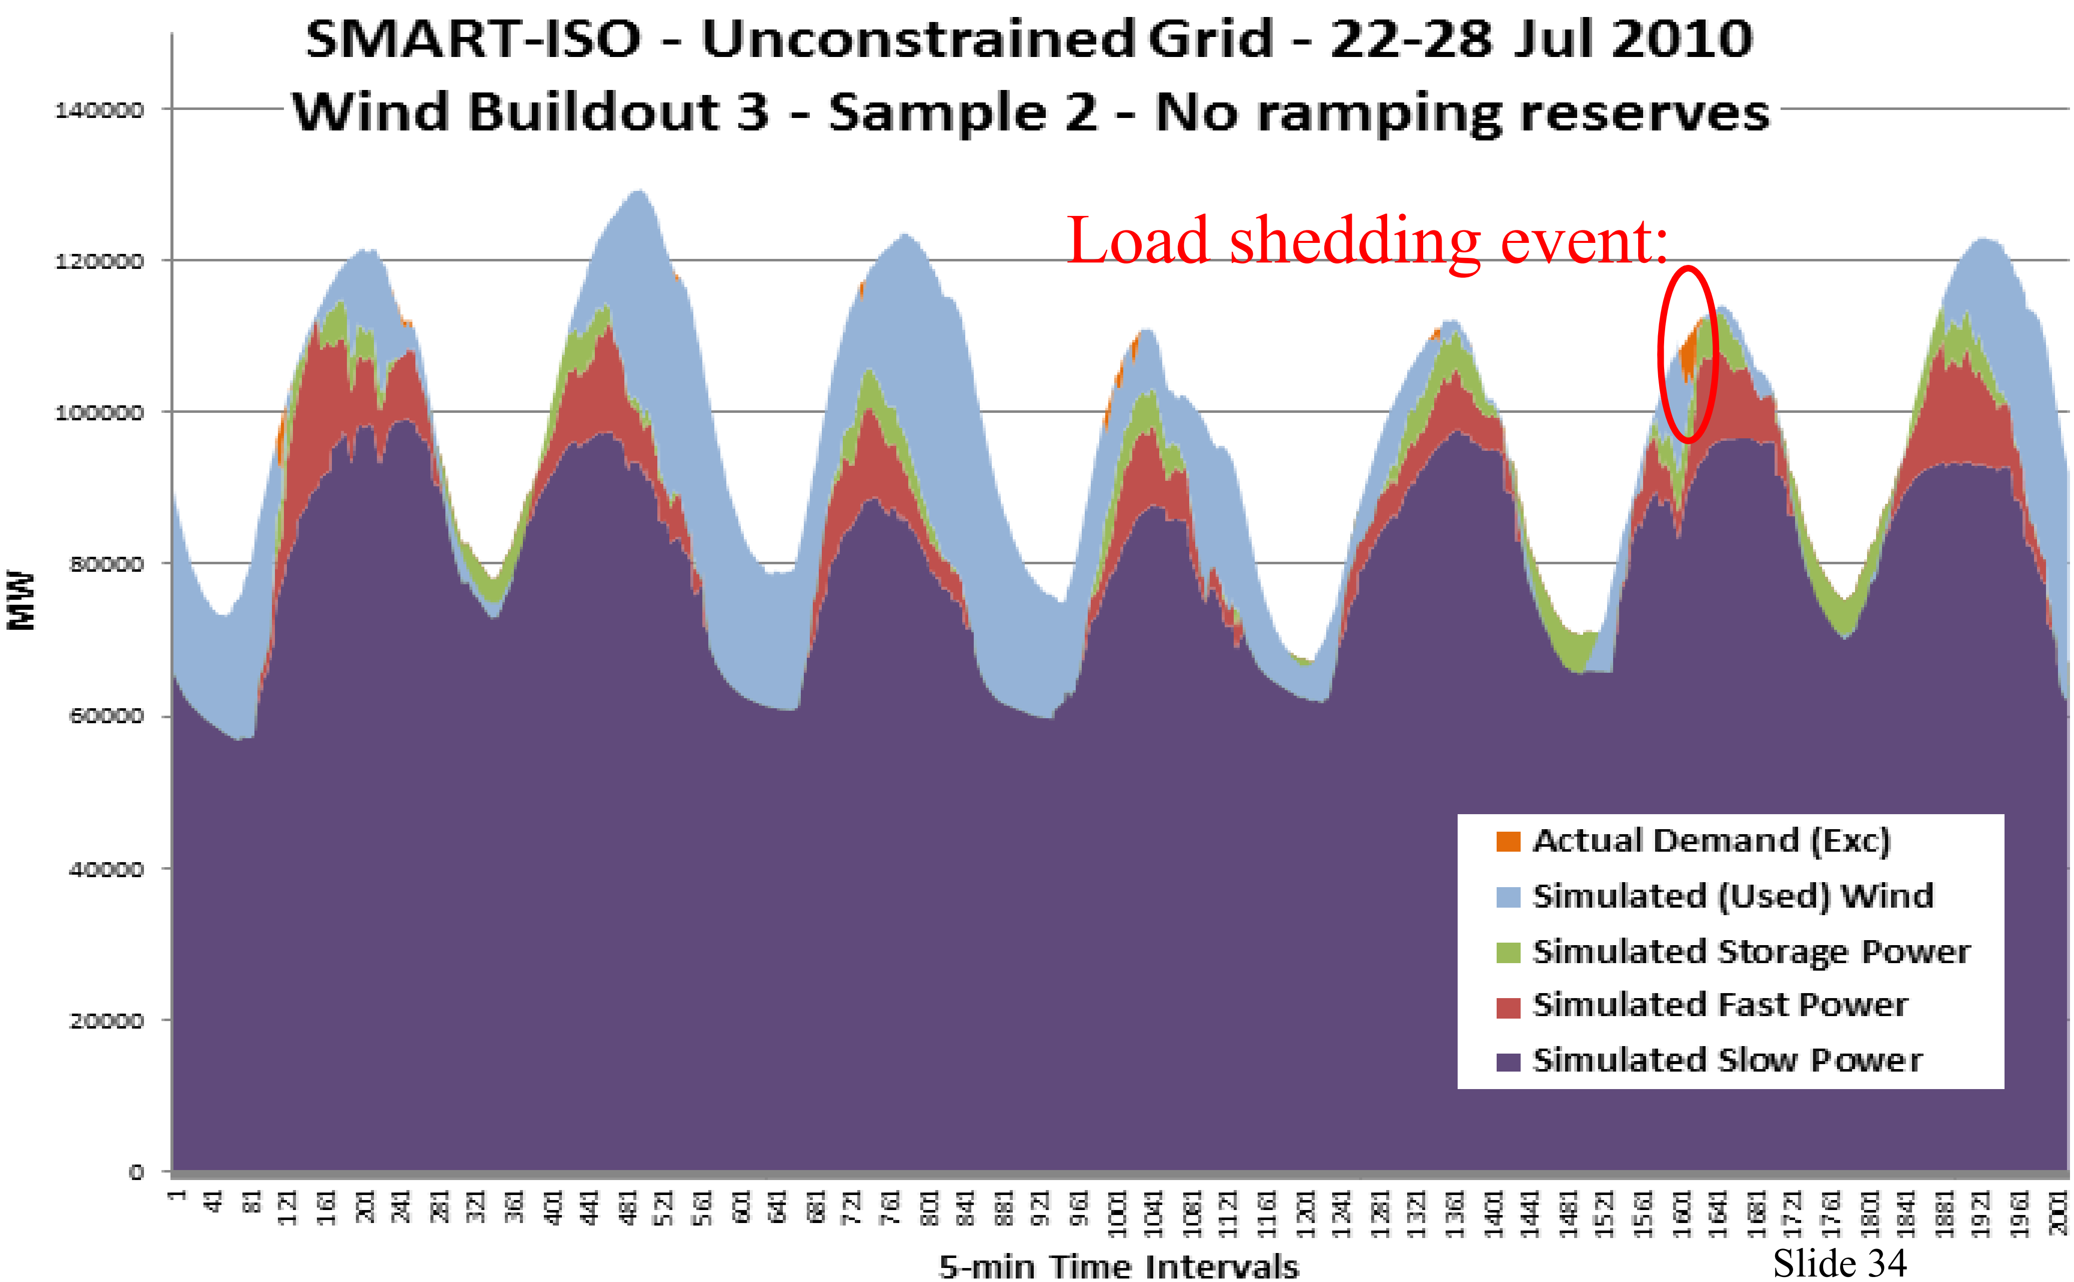

Stochastic lookahead policies

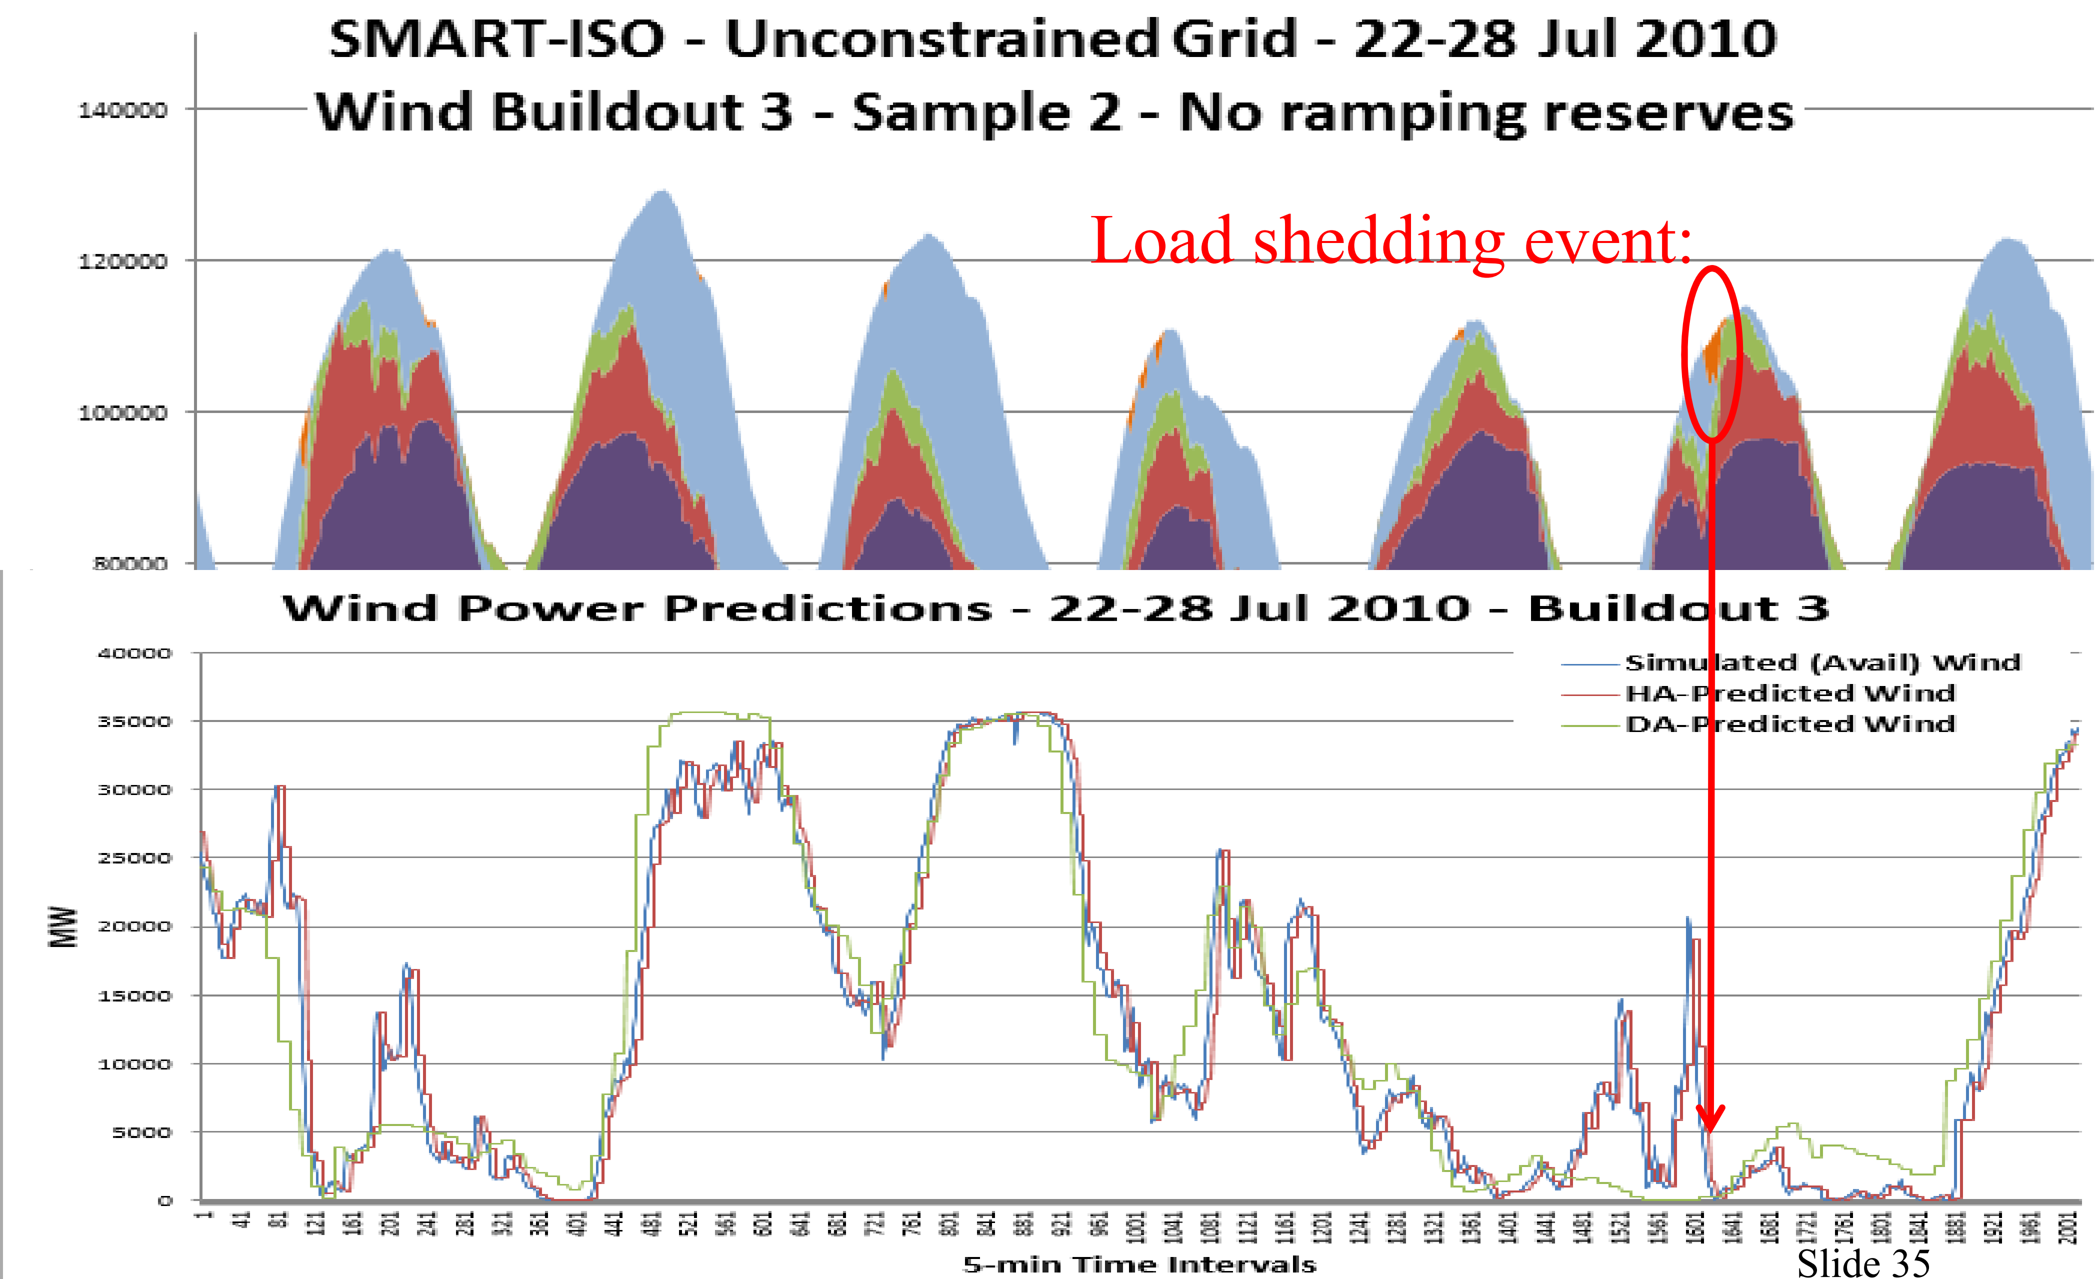

Stochastic lookahead policies

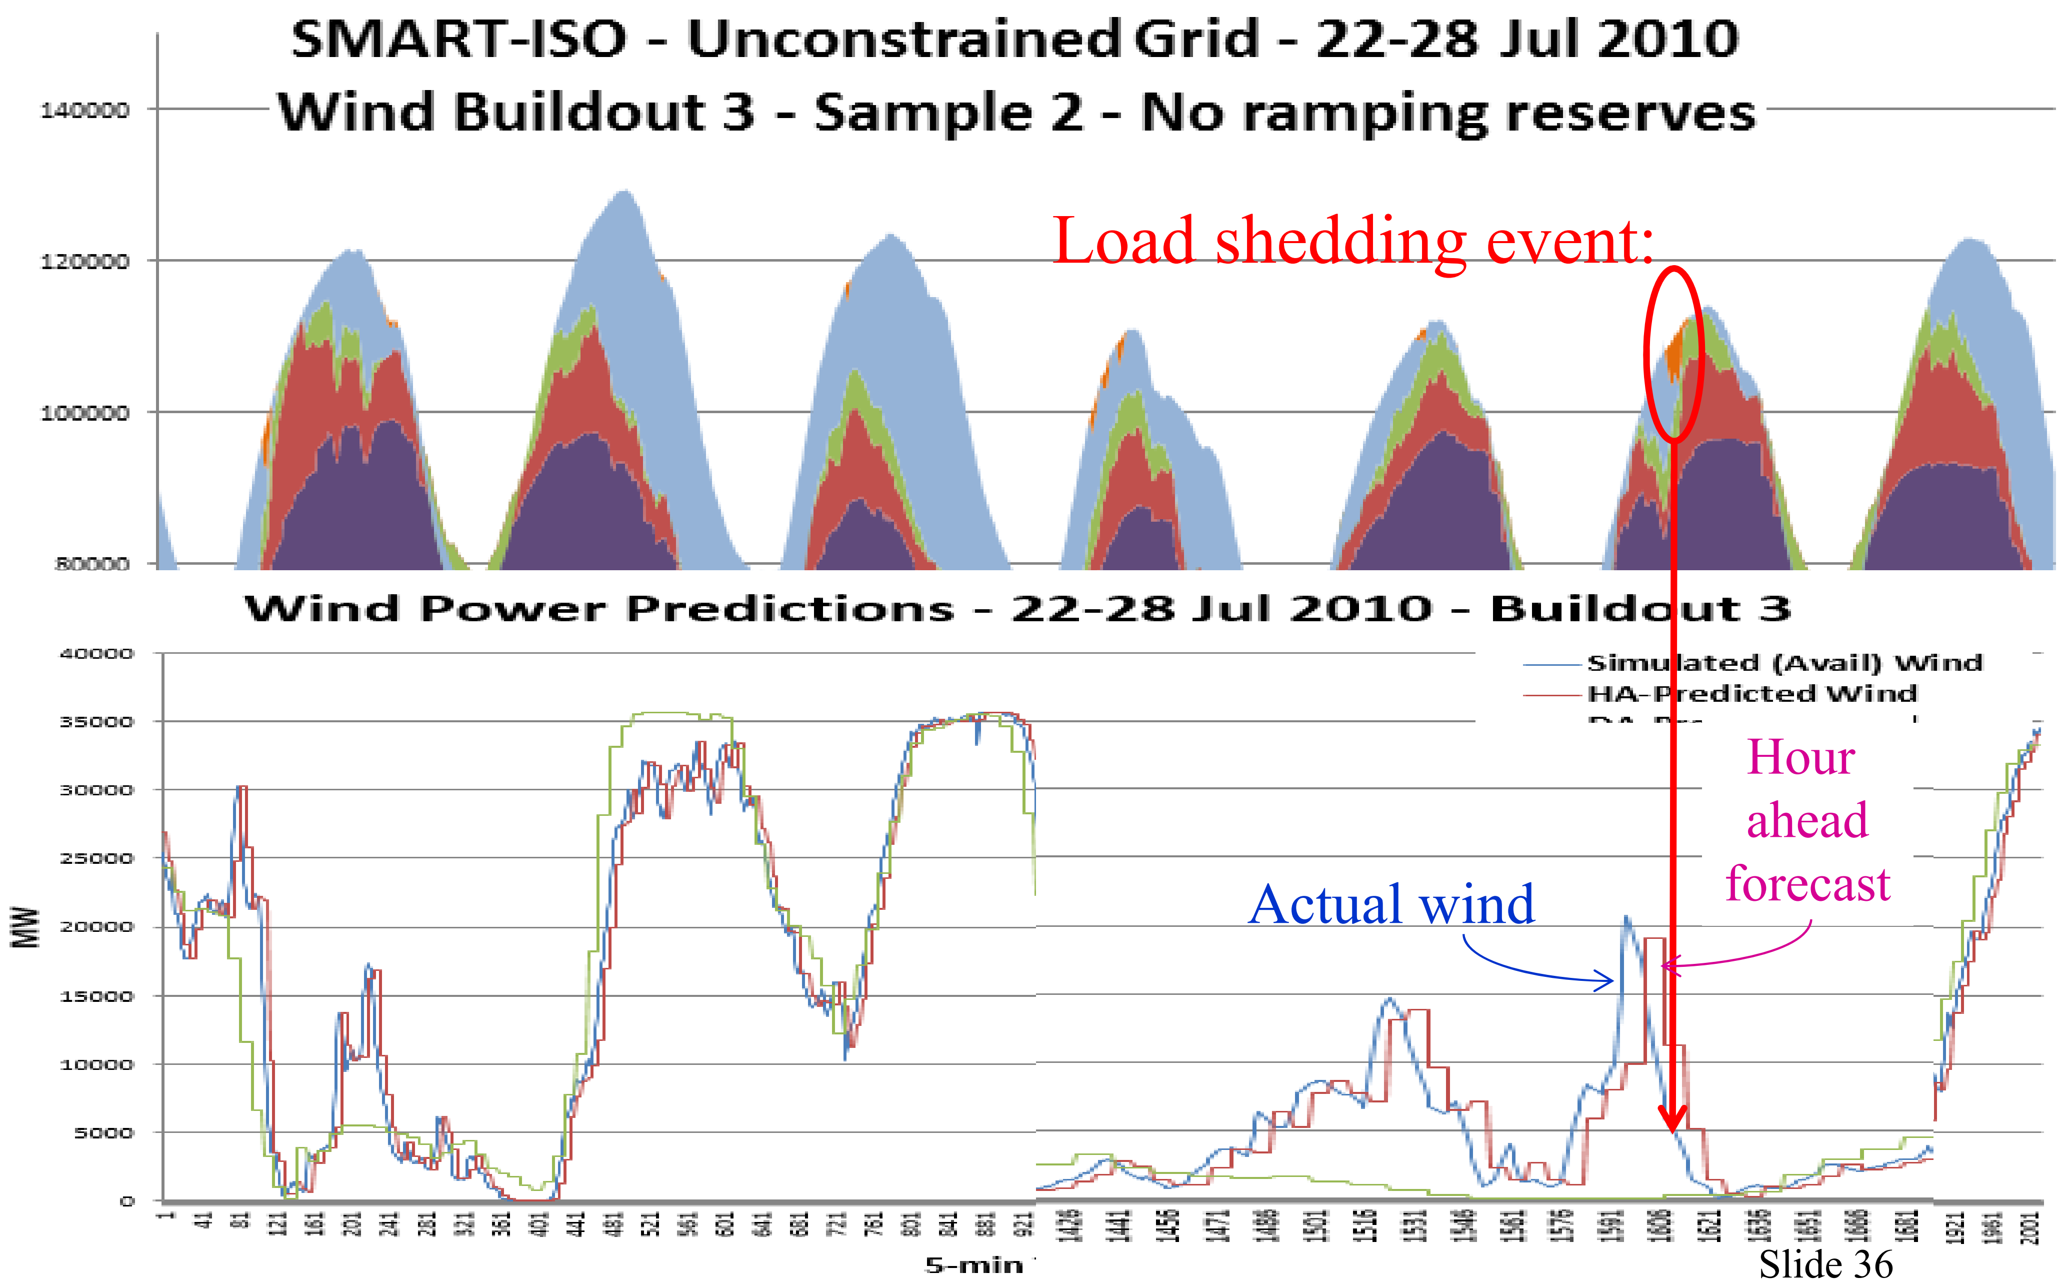

Stochastic lookahead policies

- Note the dip in steam just when the outage occurs – This is because we forecasted that wind would cover the load.

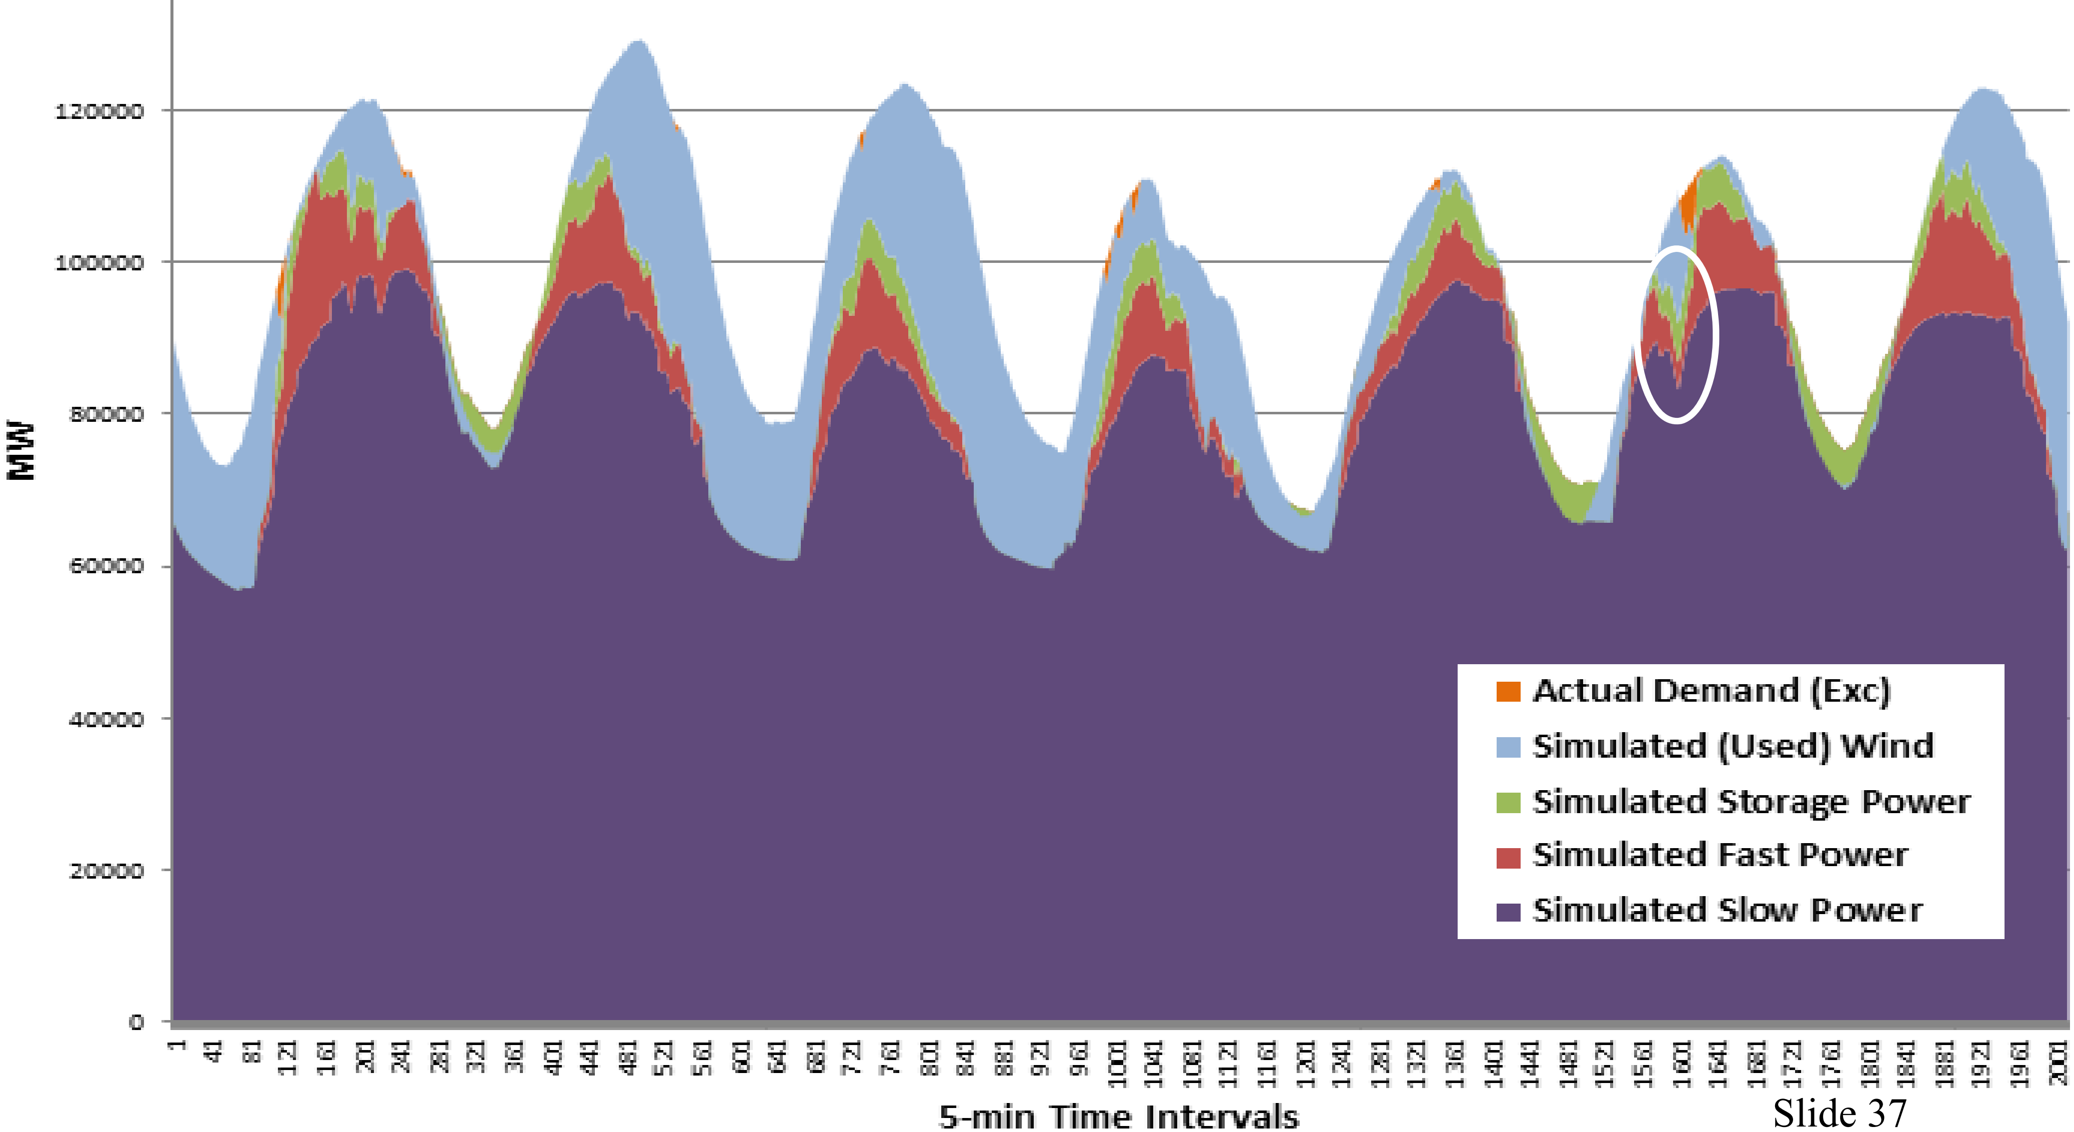

Stochastic lookahead policies

- When we allow the model to see into the future (the drop in wind), it schedules additional steam the day before. But only when we need it!

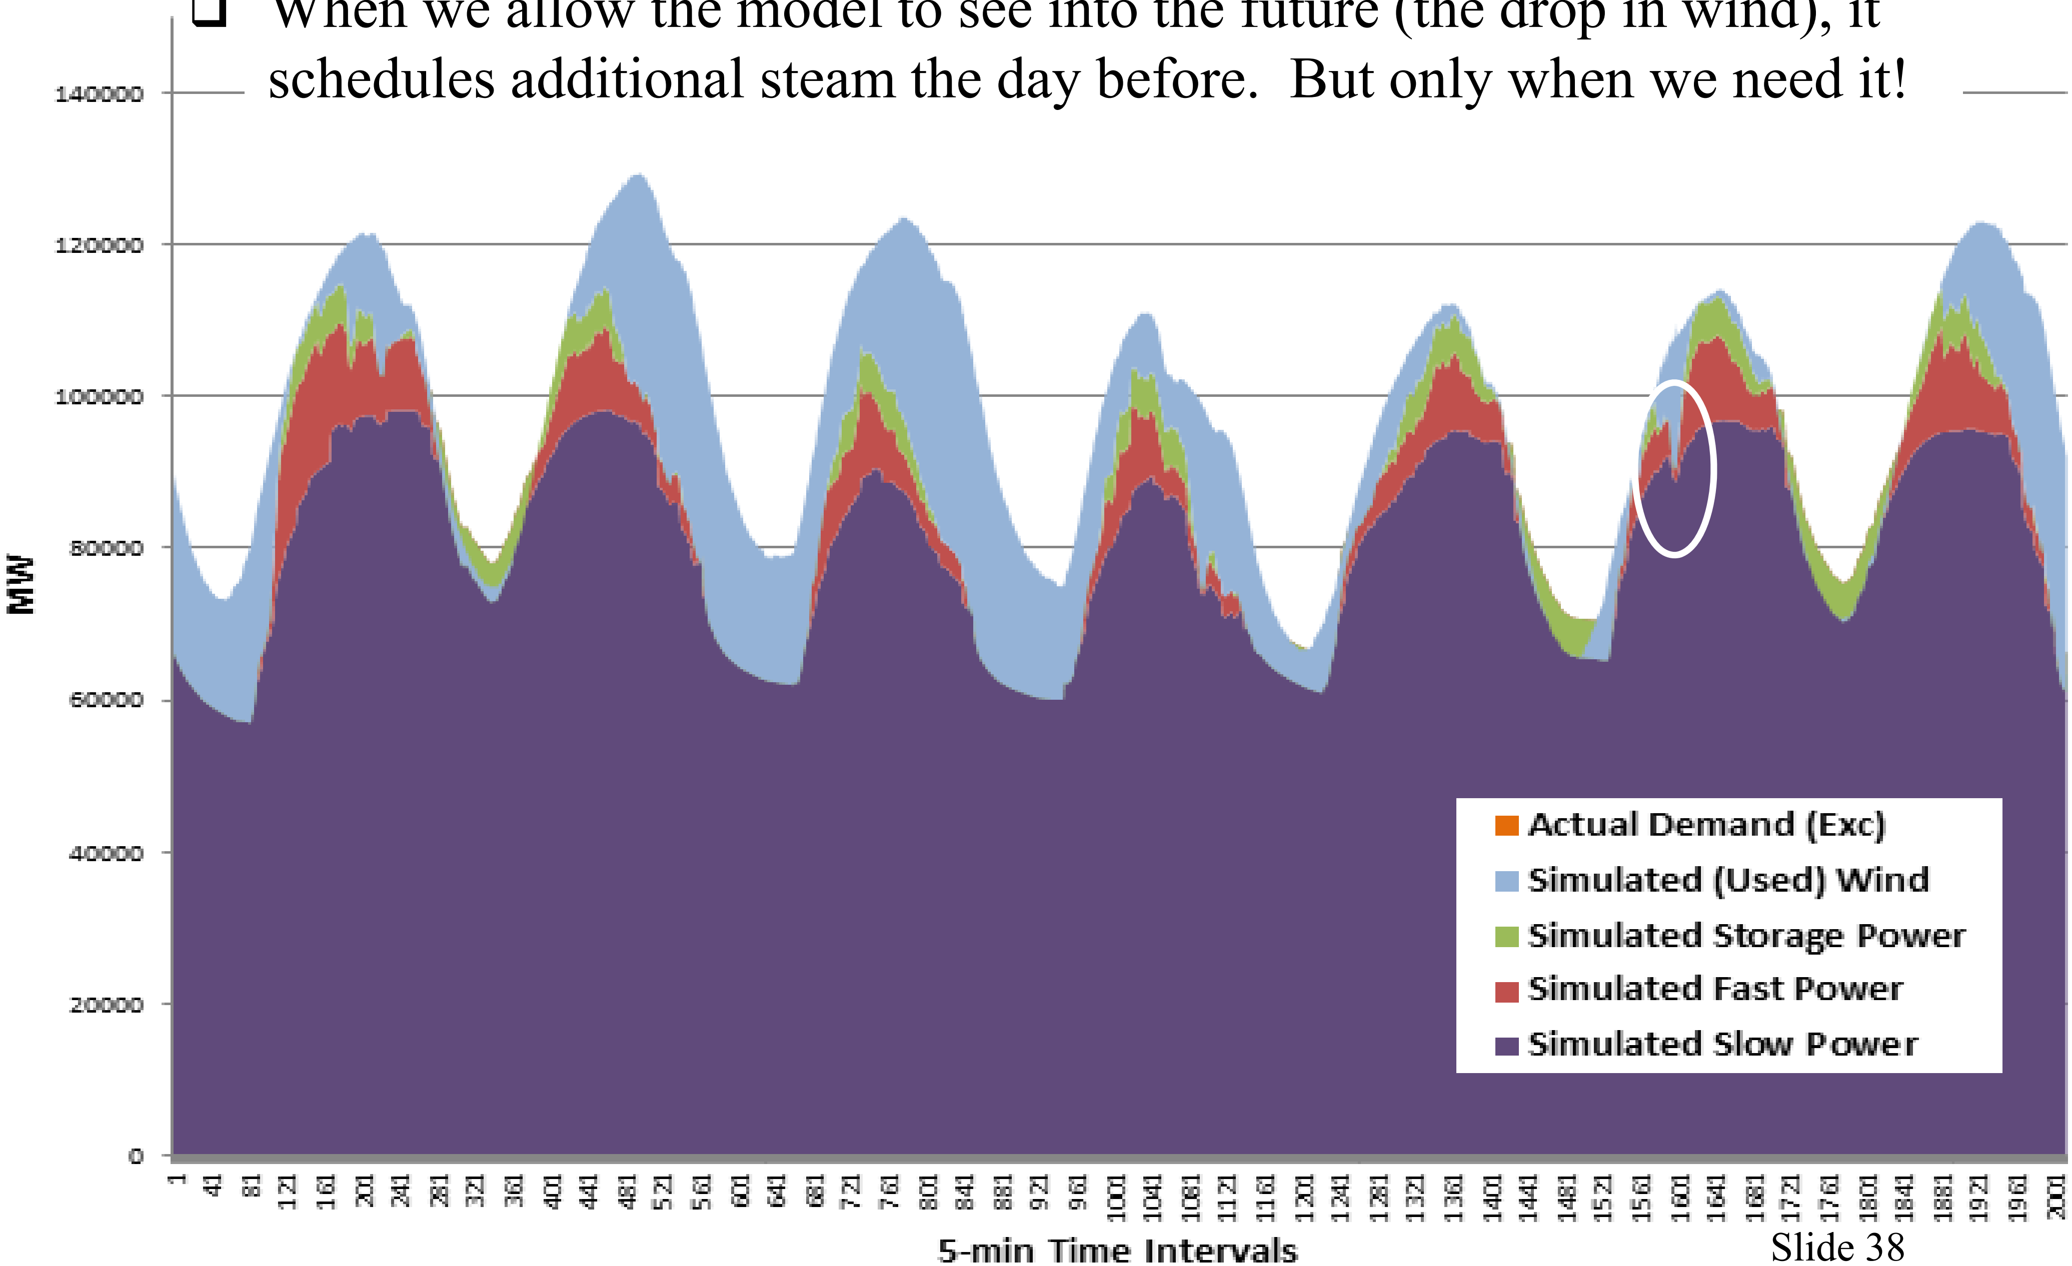

Stochastic lookahead policies

- The two-stage approximation

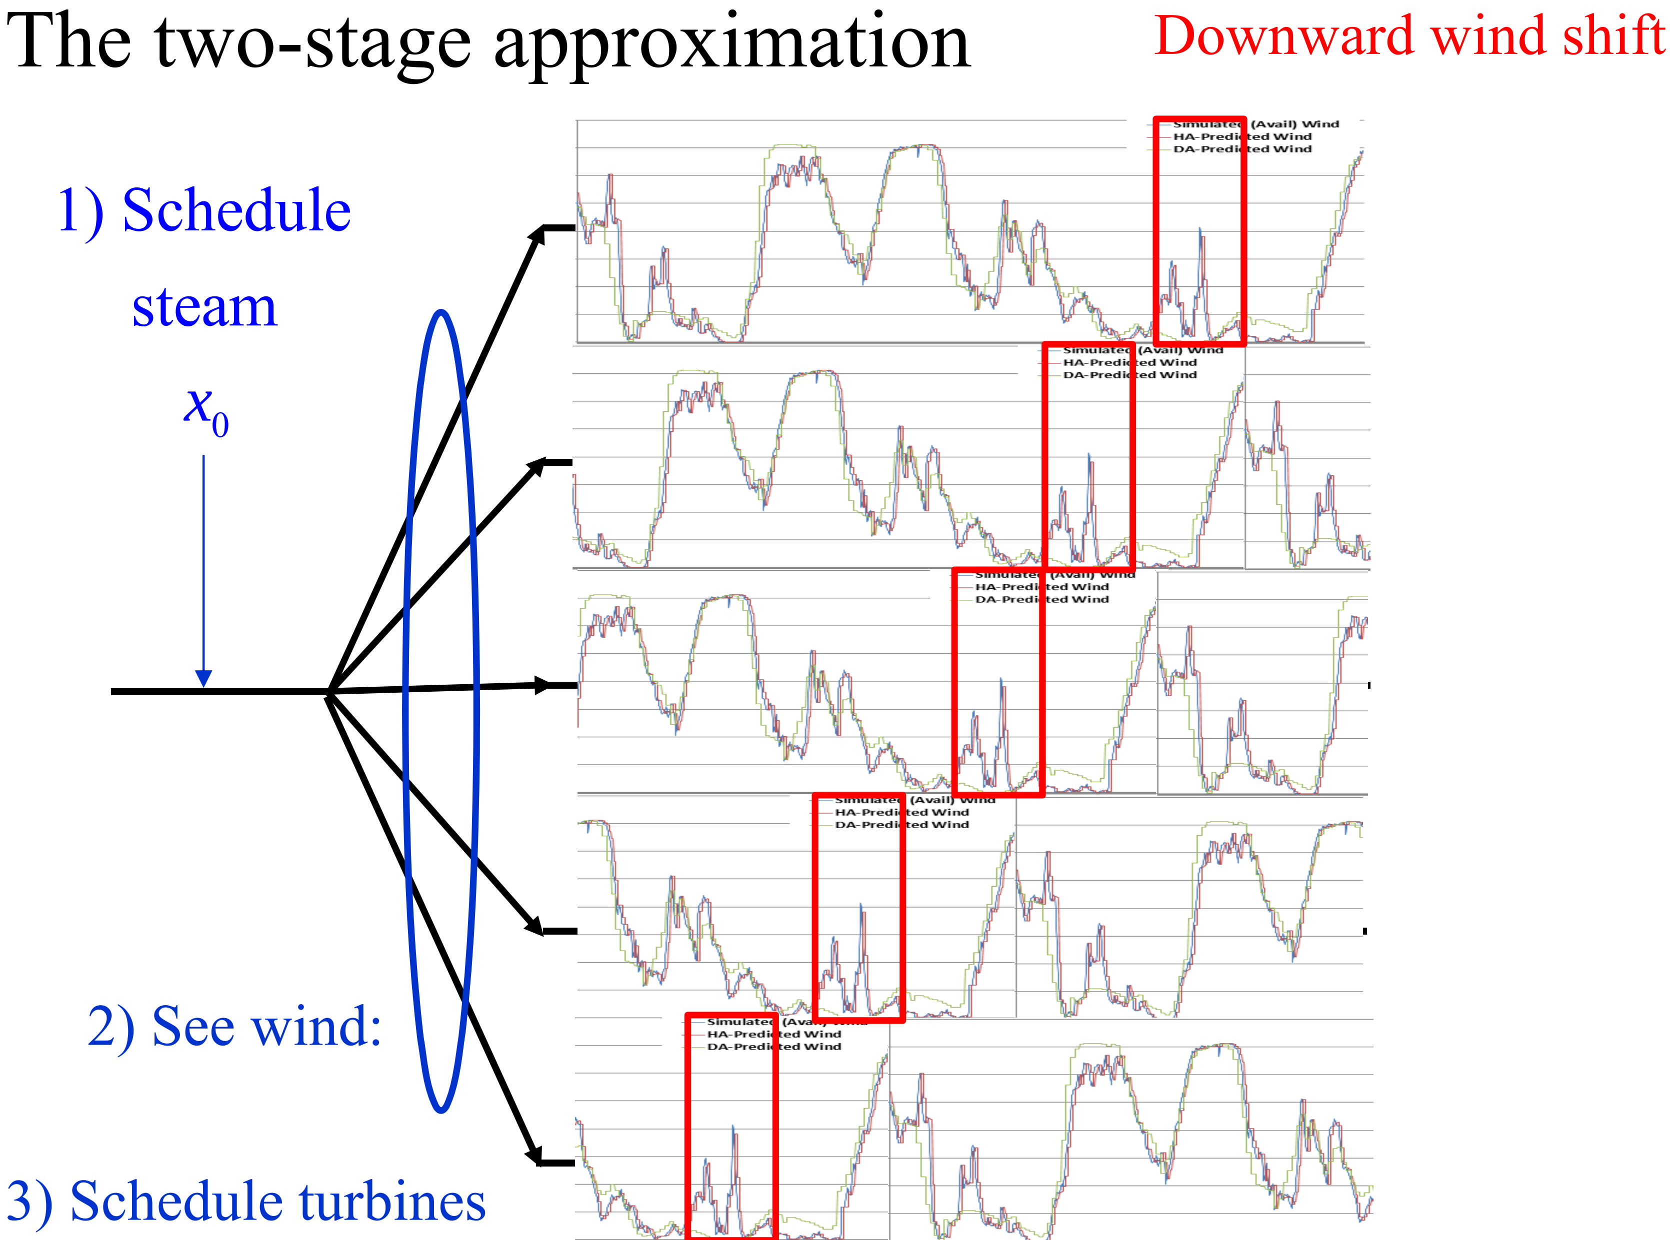

Lecture outline

- General modeling framework
- Stochastic lookahead policies
- The parametric cost function approximation
- Conclusions

Four (meta)classes of policies

1) **Policy function approximations (PFAs)**

» Lookup tables, rules, parametric/nonparametric functions

2) **Cost function approximation (CFAs)**

» $X^{CFA}(S_t \mid \theta) = \arg\min_{x_t \in \bar{\mathcal{X}}_t^{\pi}(\theta)} \bar{C}^{\pi}(S_t, x_t \mid \theta)$

3) **Policies based on value function approximations (VFAs)**

» $X_t^{VFA}(S_t) = \arg\min_{x_t} \left(C(S_t, x_t) + \bar{V}_t^{x}\left(S_t^{x}(S_t, x_t) \right) \right)$

4) **Lookahead policies**

» *Deterministic lookahead/rolling horizon proc./model predictive control*

$$X_t^{LA-D}(S_t) = \arg \min_{\tilde{x}_{tt}, \ldots, \tilde{x}_{t,t+H}} C(\tilde{S}_{tt}, \tilde{x}_{tt}) + \sum_{t'=t+1}^{T} C(\tilde{S}_{tt'}, \tilde{x}_{tt'})$$

» *Chance constrained programming*

$$P[A_t x_t \leq f(W)] \leq 1 - \delta$$

» *Stochastic lookahead /stochastic prog/Monte Carlo tree search*

$$X_t^{LA-S}(S_t) = \arg\min_{\tilde{x}_{tt}, \tilde{x}_{t,t+1}, \ldots, \tilde{x}_{t,t+T}} C(\tilde{S}_{tt}, \tilde{x}_{tt}) + \sum_{\tilde{\omega} \in \tilde{\Omega}_t} p(\tilde{\omega}) \sum_{t'=t+1}^{T} C(\tilde{S}_{tt'}(\tilde{\omega}), \tilde{x}_{tt'}(\tilde{\omega}))$$

» *"Robust optimization"*

$$X_t^{LA-RO}(S_t) = \arg \min_{\tilde{x}_{tt}, \ldots, \tilde{x}_{t,t+H}} \max_{w \in W_t(\theta)} C(\tilde{S}_{tt}, \tilde{x}_{tt}) + \sum_{t'=t+1}^{T} C(\tilde{S}_{tt'}(w), \tilde{x}_{tt'}(w))$$

Parametric cost function approximation

- A deterministic lookahead model
  - » Optimize over all decisions at the same time

$$\min_{\substack{(x_{tt'})_{t'=1,\ldots,24} \\ (y_{tt'})_{t'=1,\ldots,24}}} \sum_{t'=t}^{t+H} C(x_{tt'}, y_{tt'})$$

Steam generation

Gas turbines

  - » In a deterministic model, we mix generators with different notification times:
    - Steam generation is made day-ahead
    - Gas turbines can be planned an hour ahead or less

Parametric cost function approximation

- A deterministic lookahead policy
  - » This is the policy produced by solving a deterministic lookahead model

$$X_t^{\pi}(S_t) = \min_{\substack{(x_{tt'})_{t'=1,\ldots,24} \\ (y_{tt'})_{t'=1,\ldots,24}}} \sum_{t'=t}^{t+H} C(x_{tt'}, y_{tt'})$$

Steam generation

Gas turbines

  - » *No ISO uses a deterministic lookahead model. It would never work, and for this reason they have never used it. They always modify the model to produce a robust solution.*

Parametric cost function approximation

- A robust CFA policy
  - » The ISOs introduce reserves:

$$X_t^{\pi}(S_t \mid \theta) = \min_{\substack{(x_{tt'})_{t'=1,\dots,24} \\ (y_{tt'})_{t'=1,\dots,24}}} \sum_{t'=t}^{t+H} C(x_{tt'}, y_{tt'})$$

$$x_{t,t'}^{\max} - x_{t,t'} \geq \theta^{up} L_{tt'}$$ Up-ramping reserve

$$x_{t,t'} - x_{t,t'}^{\max} \geq \theta^{down} L_{tt'}$$ Down-ramping reserve

  - » This modification is a form of parametric function (a parametric cost function approximation). It has to be tuned to produce a robust policy.

Parametric cost function approximation

- An energy storage problem:

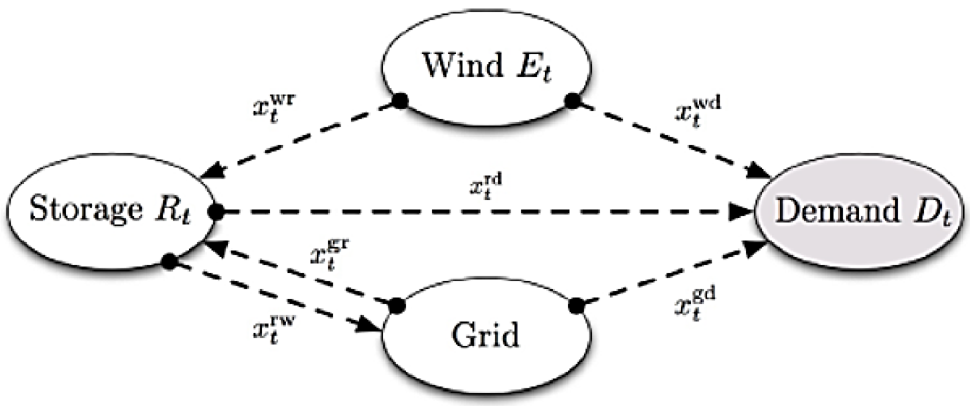

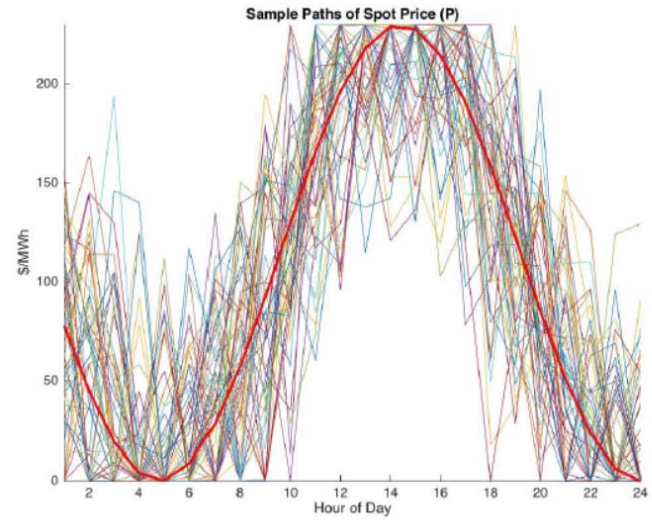

Figure: Sample paths of spot prices (P_t)

The state of the system can be represented by the following five dimensional vector,

$$S_t = (R_t, E_t, P_t, D_t, G_t)$$

where

- $R_t \in [0, R_{\max}]$ is the level of energy in storage at time t
- E_t is the amount of energy available from wind
- P_t is the spot price of electricity
- D_t is the power demand
- G_t is the energy available from the grid

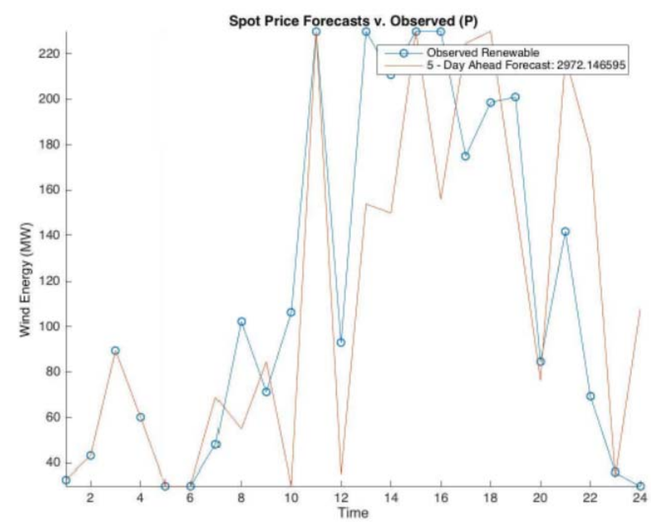

Parametric cost function approximation

- Benchmark policy – Deterministic lookahead

$$X_t^{\text{D-LA}}(S_t) = \underset{x_t,(\tilde{x}_{tt'},t'=t+1,\dots,t+H)}{\text{argmin}} \left(C(S_t, x_t) + \left[\sum_{t'=t+1}^{t+H} \tilde{c}_{tt'} \tilde{x}_{tt'} \right] \right)$$

$$\begin{aligned}
x_t^{wd} + \beta^d x_t^{rd} + x_t^{gd} &\le D_t, \\
x_t^{gd} + x_t^{gr} &\le G_t, \\
x_t^{rd} + x_t^{rg} &\le R_t, \\
x_t^{wr} + x_t^{gr} &\le R_{\max} - R_t, \\
x_t^{wr} + x_t^{wd} &\le E_t, \\
x_t^{wr} + x_t^{gr} &\le \gamma^c, \\
x_t^{rd} + x_t^{rg} &\le \gamma^d
\end{aligned}$$

Parametric cost function approximation

- Parametric cost function approximations
  - » Replace the constraint

$$x_t^{wr} + x_t^{wd} \le E_t,$$

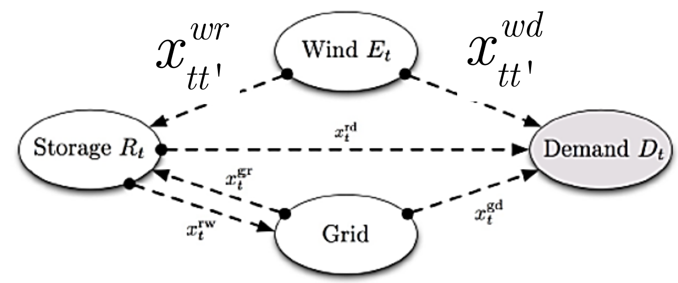

with:

- » Lookup table modified forecasts (one adjustment term for each time $\tau = t' - t$ in the future):

$$x_{tt'}^{wr} + x_{tt'}^{wd} \le \theta_{t'-t} F_{tt'}^{E}$$

- » Exponential function for adjustments (just two parameters)

$$x_{tt'}^{wr} + x_{tt'}^{wd} \le \theta_1 e^{\theta_2 (t'-t)} F_{tt'}^{E}$$

- » Constant adjustment (one parameter)

$$x_{tt'}^{wr} + x_{tt'}^{wd} \le \theta F_{tt'}^{E}$$

Parametric cost function approximation

- Optimizing the CFA:
  - Let $\bar{F}(\theta,\omega)$ be a simulation of our policy given by

$$\bar{F}(\theta,\omega) = \sum_{t=0}^{T} C\left(S_t(\omega), X_t^{\pi}(S_t(\omega) \mid \theta)\right)$$

  - We then compute the gradient with respect to θ

$$\nabla_{\theta} F(\theta) = \mathbb{E}\left\{\nabla_{\theta} \bar{F}(\theta,\omega)\right\}$$

  - The parameter θ is found using a classical stochastic gradient algorithm:

$$\theta^{n+1} = \theta^{n} + \alpha_n \nabla_{\theta} F(\theta^{n},\omega^{n+1})$$

We tested several stepsize formulas and found that ADAGRAD worked best:

$$\alpha_n = \frac{\eta}{\sqrt{G_t + \epsilon}}$$

Parametric cost function approximation

- Optimizing the CFA:
  - » We compute the gradient by applying the chain rule

$$\nabla_\theta \bar{F} = \left(\frac{\partial C_0}{\partial X_0}\cdot\frac{\partial X_0}{\partial \theta}\right)+\sum_{t'=1}^{T}\left[\left(\frac{\partial C_{t'}}{\partial S_{t'}}\cdot\frac{\partial S_{t'}}{\partial \theta}\right)+\left(\frac{\partial C_{t'}}{\partial X_{t'}(S_t|\theta)}\cdot\left(\frac{\partial X_{t'}(S_t|\theta)}{\partial S_{t'}}\cdot\frac{\partial S_{t'}}{\partial \theta}+\frac{\partial X_{t'}(S_t|\theta)}{\partial \theta}\right)\right)\right],$$

  - » Where the interaction from one time period to the next is captured using

$$\frac{\partial S_{t'}}{\partial \theta} = \frac{\partial S_{t'}}{\partial S_{t'-1}}\cdot\frac{\partial S_{t'-1}}{\partial \theta}+\frac{\partial S_{t'}}{\partial X_{t'-1}(S_{t-1}|\theta)}\cdot\left[\frac{\partial X_{t;-1}(S_{t-1}|\theta)}{\partial S_{t'-1}}\cdot\frac{\partial S_{t'-1}}{\partial \theta}+\frac{\partial X_{t'-1}(S_{t-1}|\theta)}{\partial \theta}\right].$$

  - » Assuming there are no integer variables, these equations are quite easy to compute.
  - » For real stochastic unit commitment problems, we are going to need to use a derivative-free algorithm.

Parametric cost function approximation

- Optimal adjustment parameters θ for each model

More accurate forecasts

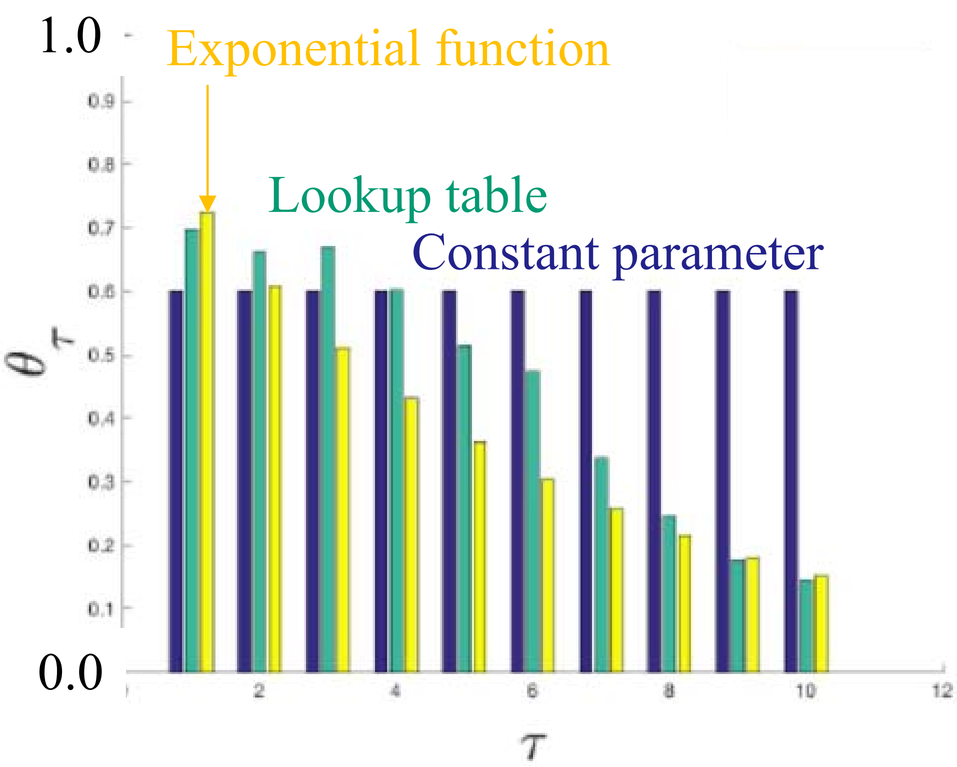

(a) θ_τ for $\sigma_f = 20$ and $\gamma = R_{\max}$

Less accurate forecasts

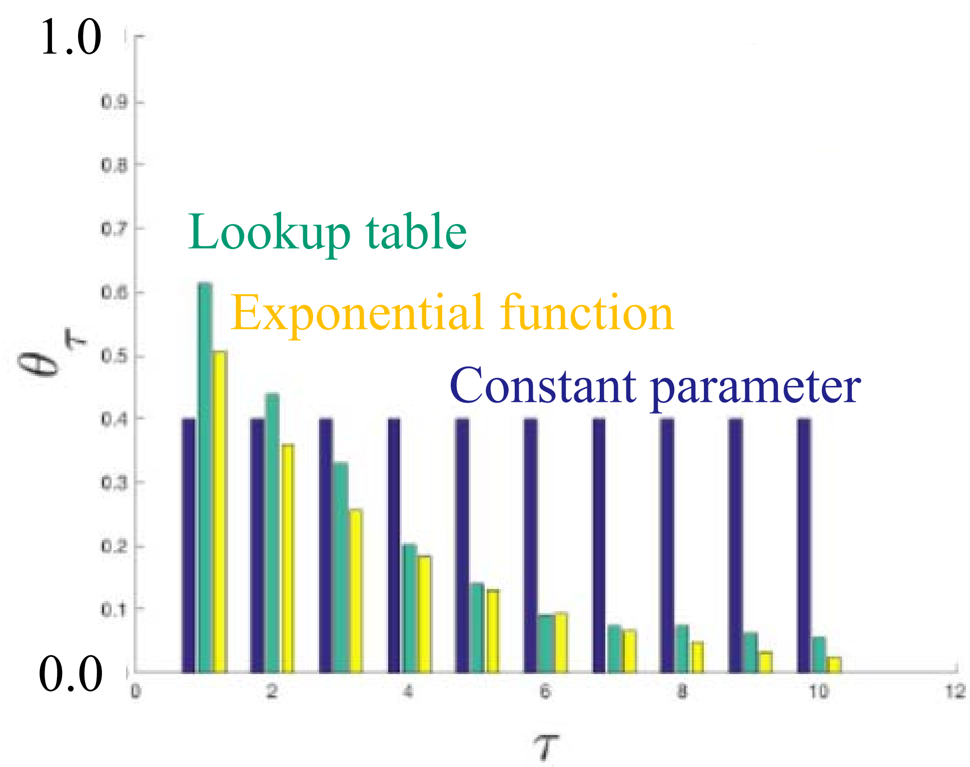

(a) θ_τ for $\sigma_f = 35$ and $\gamma = R_{\max}$

Parametric cost function approximation

- Improvement over deterministic benchmark:

Decreasing forecast accuracy: ⟶

	$\sigma_f = 20$	$\sigma_f = 25$	$\sigma_f = 30$	$\sigma_f = 35$
Constant	13%	13%	16%	17%
Lookup table	20%	22%	26%	25%
Exponential decay	14%	22%	26%	26%

» Significant improvement over deterministic benchmark (an untuned lookahead policy).

Parametric cost function approximation

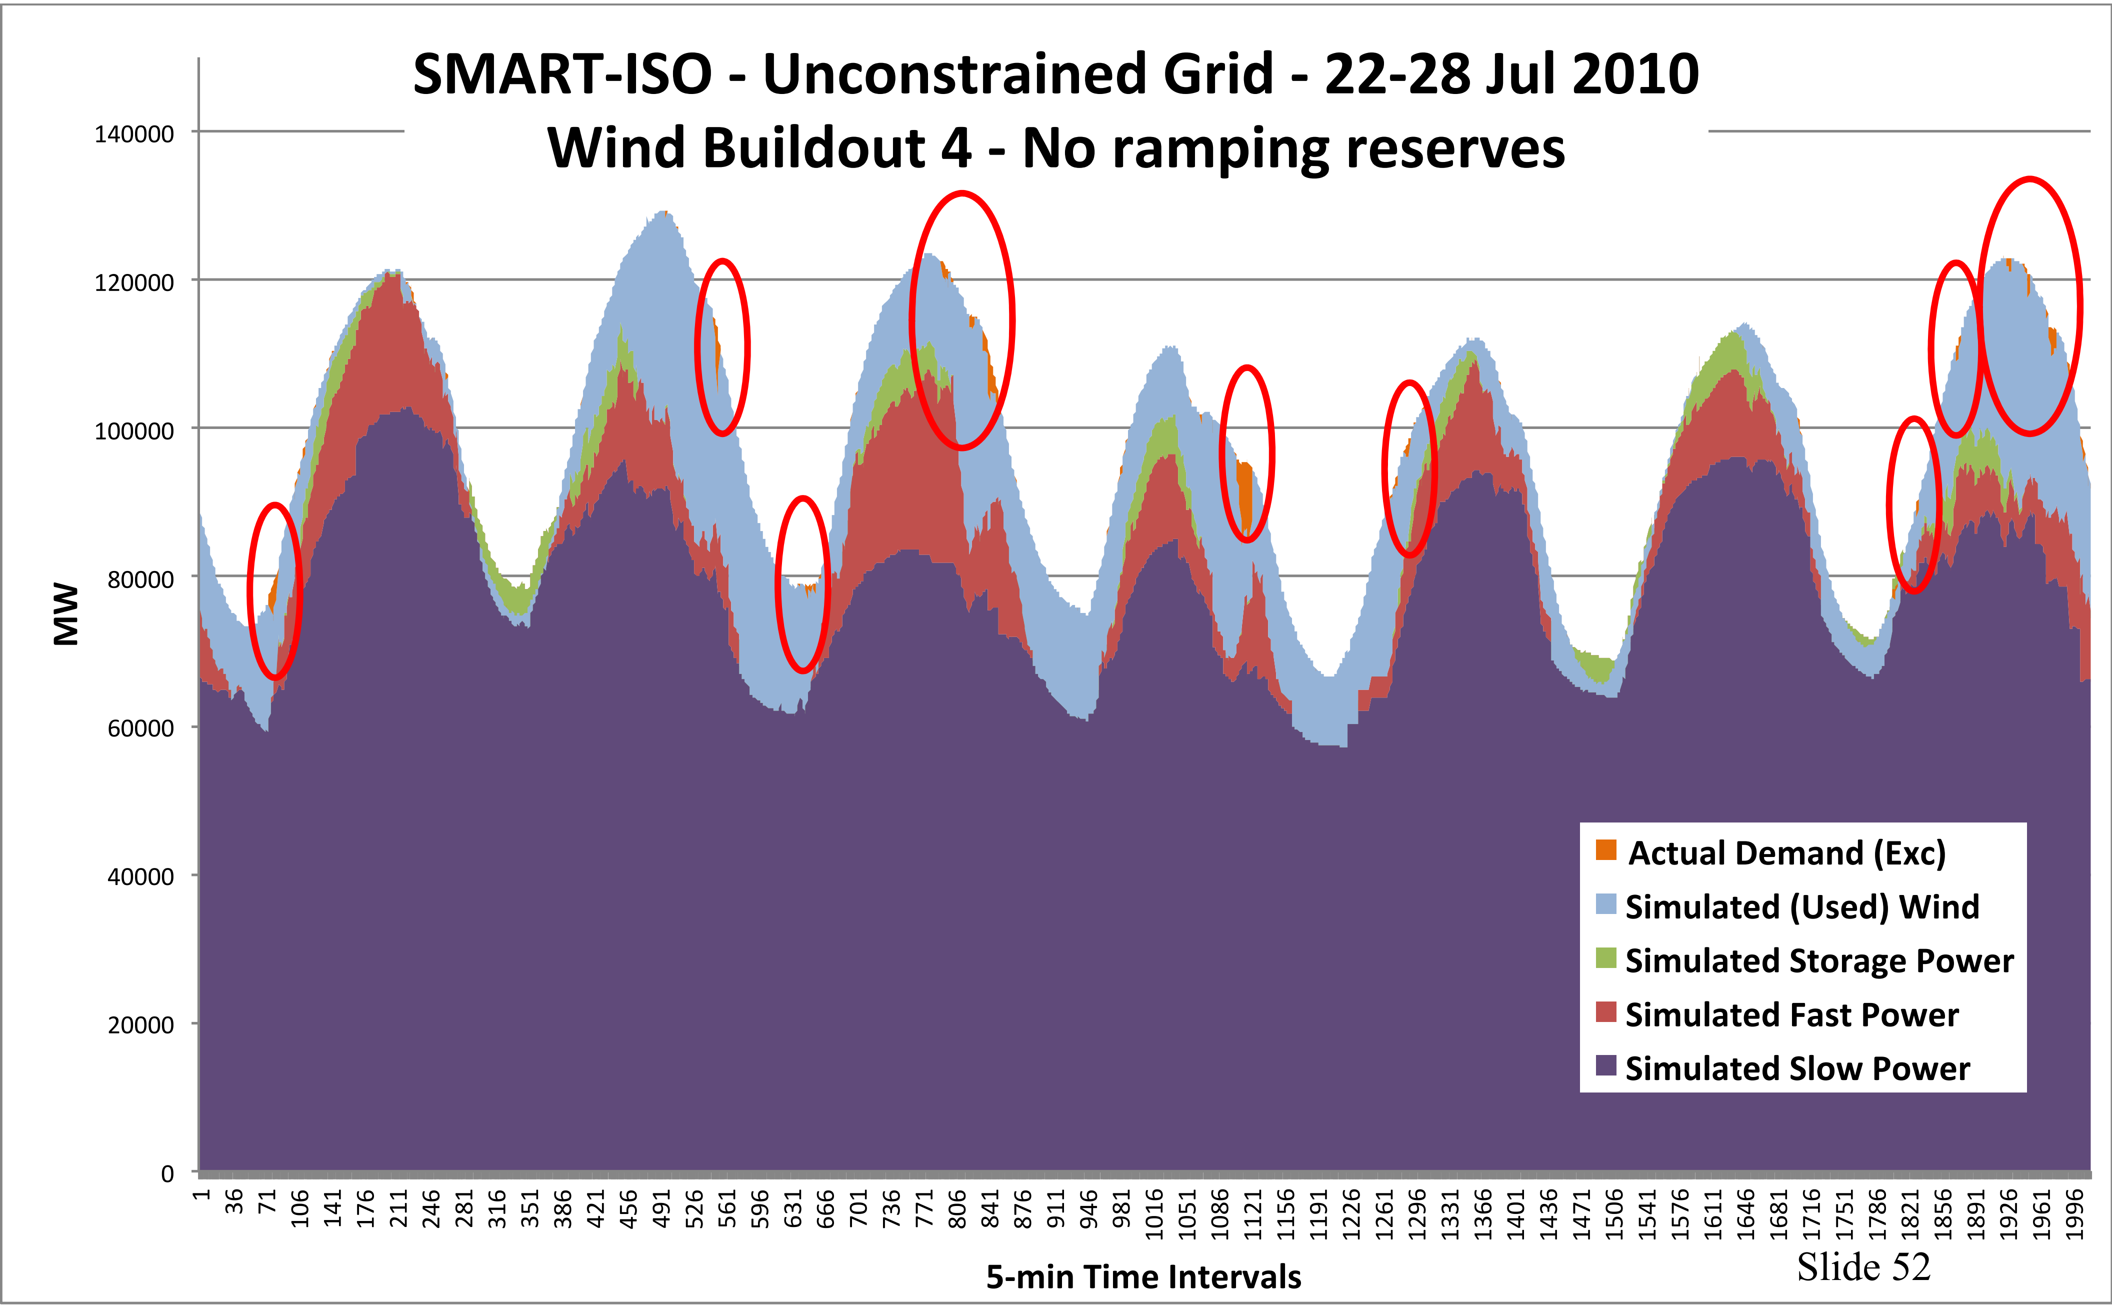

Parametric cost function approximation

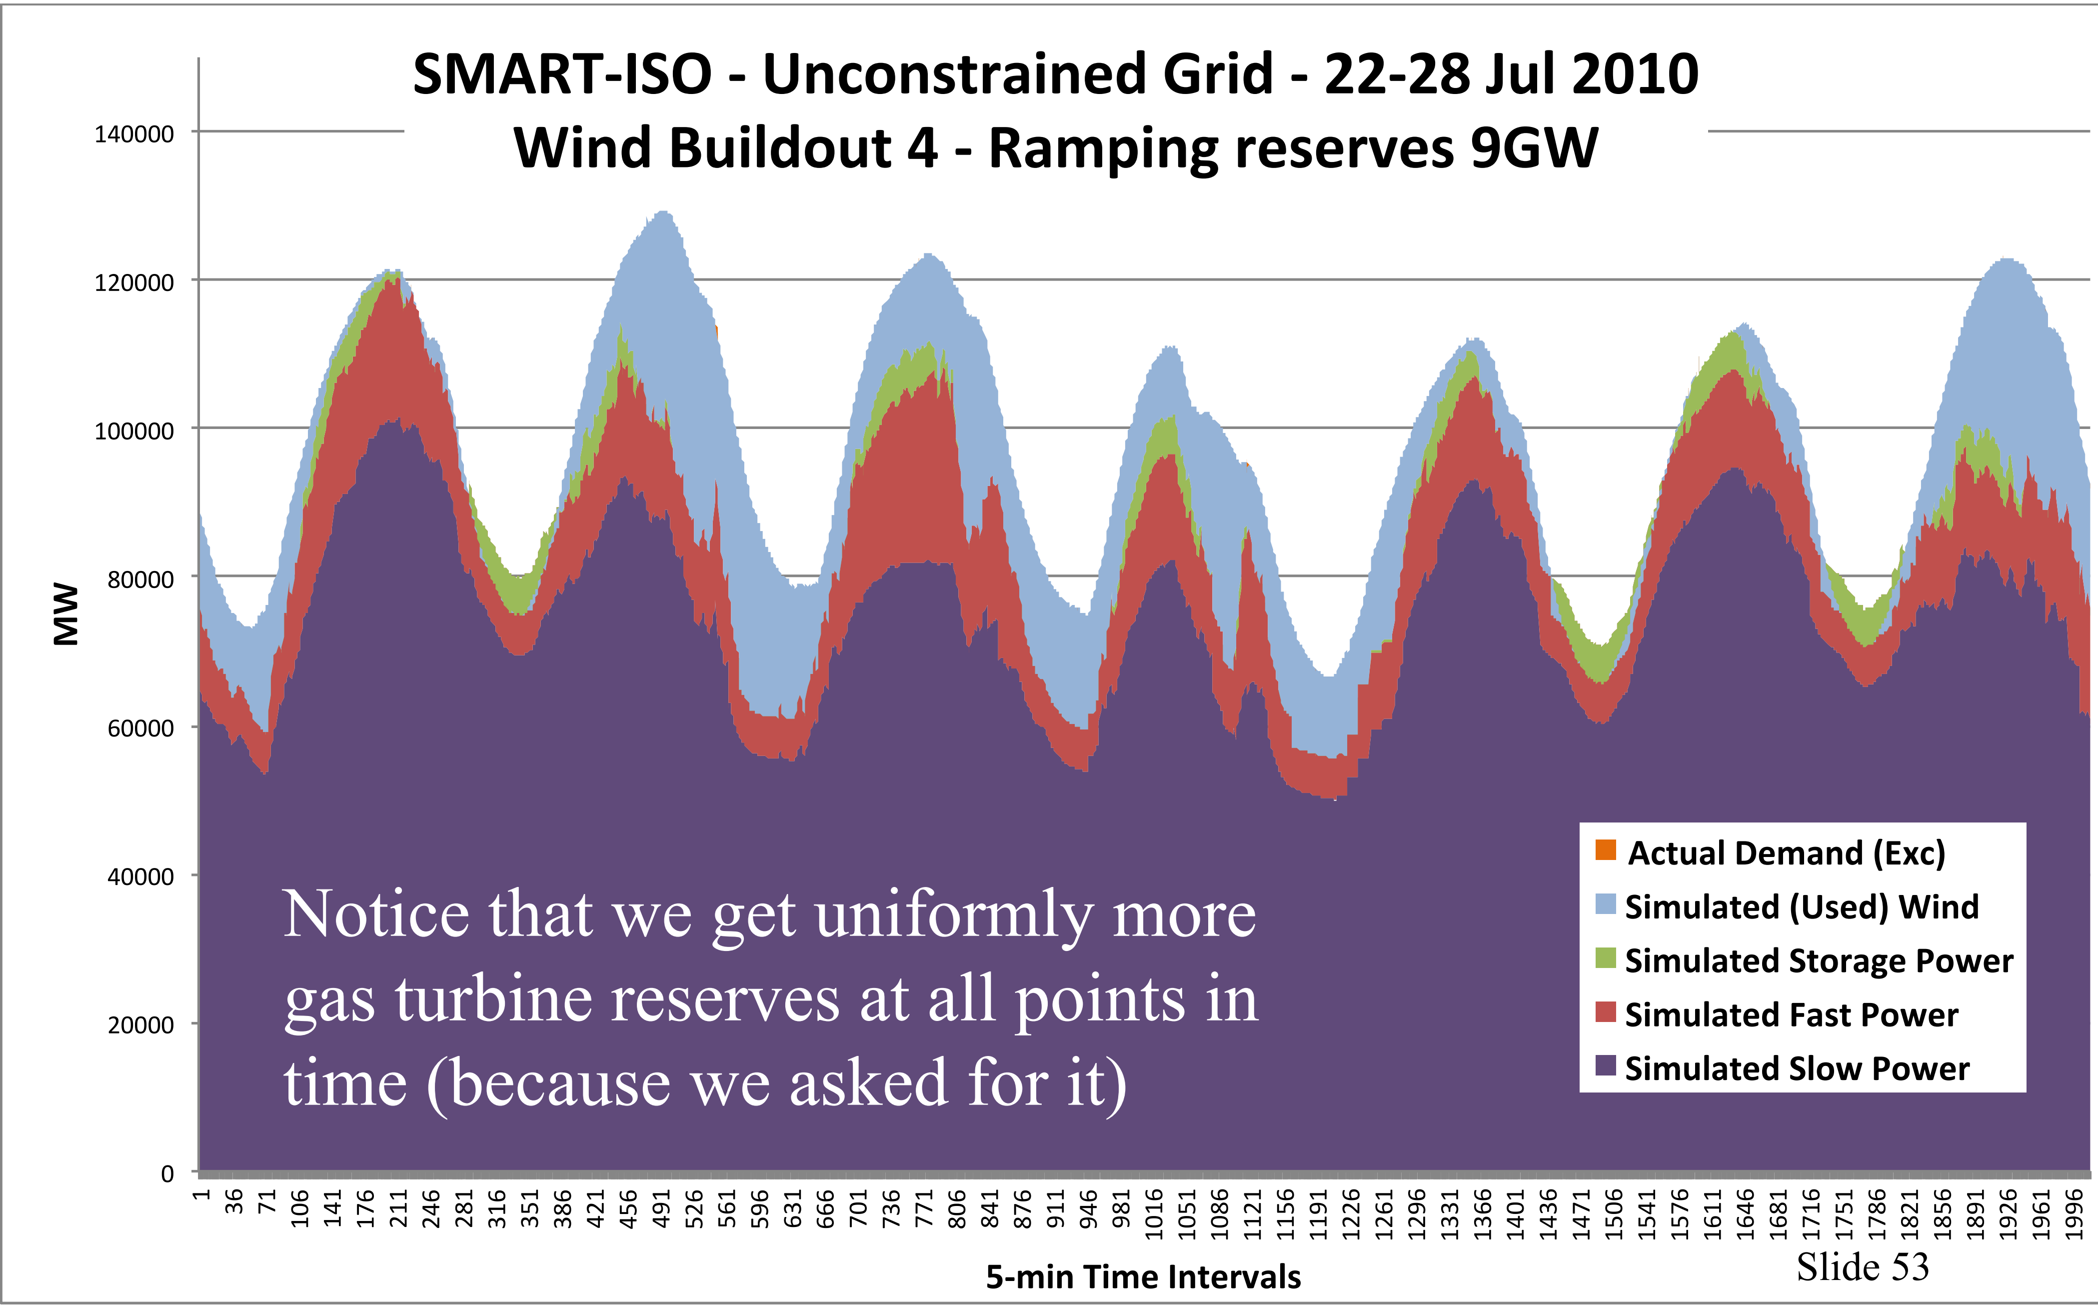

Notice that we get uniformly more gas turbine reserves at all points in time (because we asked for it)

Paper available on arXiv:

Stochastic Optimization with Parametric Cost Function Approximations

Raymond T. Perkins III[a], Warren B. Powell[b]

[a] raymondp@princeton.edu
[b] powell@princeton.edu

Abstract

A widely used heuristic for solving stochastic optimization problems is to use a deterministic rolling horizon procedure which has been modified to handle uncertainty (e.g. buffer stocks, schedule slack). This approach has been criticized for its use of a deterministic approximation of a stochastic problem, which is the major motivation for stochastic programming. We recast this debate by identifying both deterministic and stochastic approaches as policies for solving a stochastic base model, which may be a simulator or the real world. Stochastic lookahead models (stochastic programming) require a range of approximations to keep the problem tractable. By contrast, so-called deterministic models are actually parametrically modified cost function approximations which use parametric adjustments to the objective function and/or the constraints. These parameters are then optimized in a stochastic base model which does not require making any of the types of simplifications required by stochastic programming. We formalize this strategy and describe a gradient-based stochastic search strategy to optimize the parameters.

Keywords: Stochastic Optimization, Stochastic Programming, Decisions under uncertainty, Parametric Cost Function Approximation, Cost Function Approximation, Policy Search

Lecture outline

- General modeling framework
- Stochastic lookahead policies
- The parametric cost function approximation
- Conclusions

Conclusions:

» Weaknesses of stochastic programming (with scenario trees)
  - Computationally intensive
  - It does not ensure robust solutions (needs too many scenarios)
  - Many approximations are required (e.g. two stage).

» Parametric cost function approximations:
  - Generalizes standard industry practice
  - Uses domain knowledge to ensure robust policies across a much wider range of scenarios.
  - Resulting models can be implemented using existing commercial software.

» Parametric CFAs open up an entirely new research directions for stochastic unit commitment:
  - Propose new parameterizations to achieve robustness at lowest cost.
  - Design algorithms (either derivative-based or derivative-free) to optimize the parameterized policy.
  - Need to design algorithms for online learning.

Thank you!

http://energysystems.Princeton.edu