Robust Policies for Unit Commitment

FERC Technical Conference on Software
Washington, D.C.

June 23-25, 2014

Warren B. Powell
Hugo P. Simao
PENSA Laboratory
Princeton University
http://energysystems.princeton.edu

Designing a policy

- Dealing with uncertainty
  - » We have to design policies to manage the different forms of uncertainty.
  - » We do this by looking for *robust policies*, which are rules for making decisions.
  - » We write our optimization problem in the form:

$$\min_\pi E^\pi \left\{ \sum_{t=0}^{T} \gamma^t C\left(S_t, X^\pi(S_t)\right) \right\}$$

Search for the best policy

Averaging over multiple samples

Day-ahead, hour-ahead and real-time decisions

Objective function where $S_{t+1} = S^M\left(S_t, X^\pi(S_t), W_{t+1}\right)$

Designing a policy

- Dealing with uncertainty
  - We have to design policies to manage the different forms of uncertainty.
  - We do this by looking for *robust policies*, which are rules for making decisions.
  - We write our optimization problem in the form:

$$\min_{\pi} E^{\pi}\left\{\sum_{t=0}^{T}\gamma^{t}C\left(S_t, X^{\pi}(S_t)\right)\right\}$$

  - We refer to this as the *base model* which is typically calculated as a simulation:

$$\min_{\pi} \bar{F}^{\pi} = \frac{1}{N}\sum_{n=1}^{N}\sum_{t=0}^{T}\gamma^{t}C\left(S_t(\omega^n), X_t^{\pi}(S_t(\omega^n))\right)$$

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

Four classes of policies

1) Policy function approximations (PFAs)

» Lookup tables, rules, parametric functions

2) Cost function approximation (CFAs)

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

3) Policies based on value function approximations (VFAs)

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

4) Lookahead policies

» ***Deterministic lookahead:***

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

» ***Stochastic lookahead (e.g. stochastic trees)***

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

Lookahead policies

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

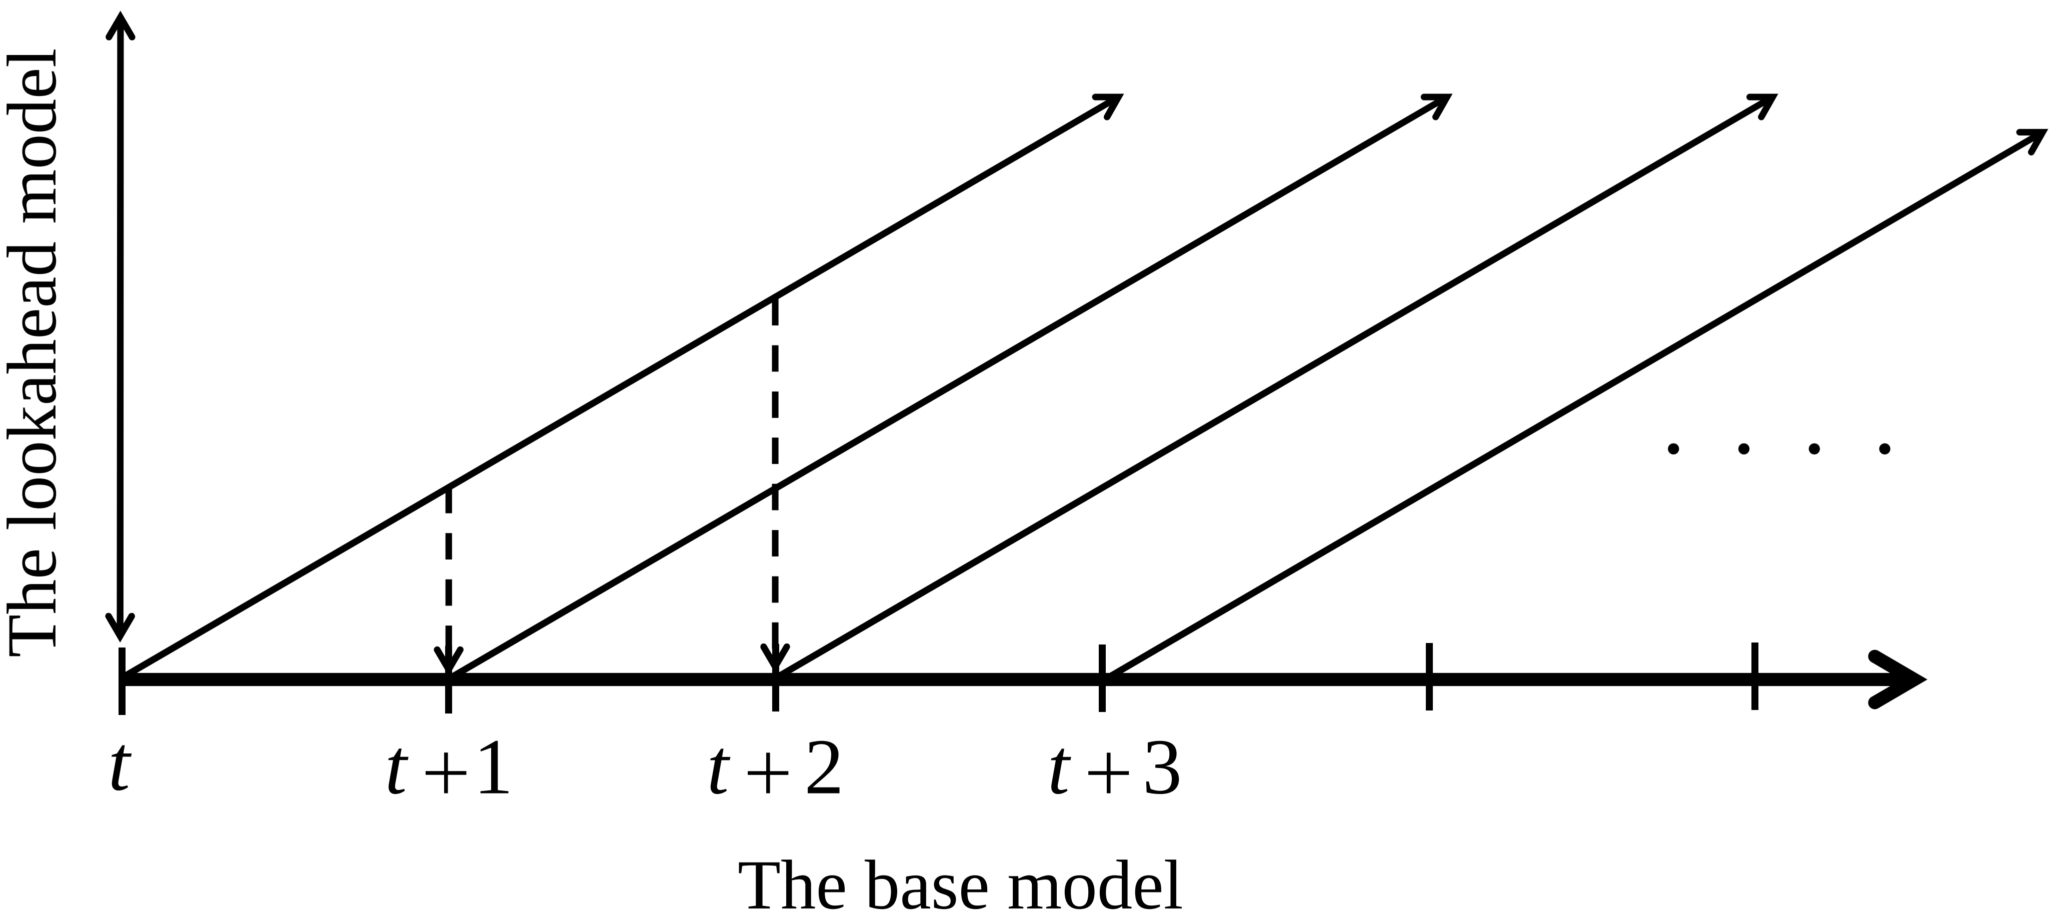

Lookahead policies

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

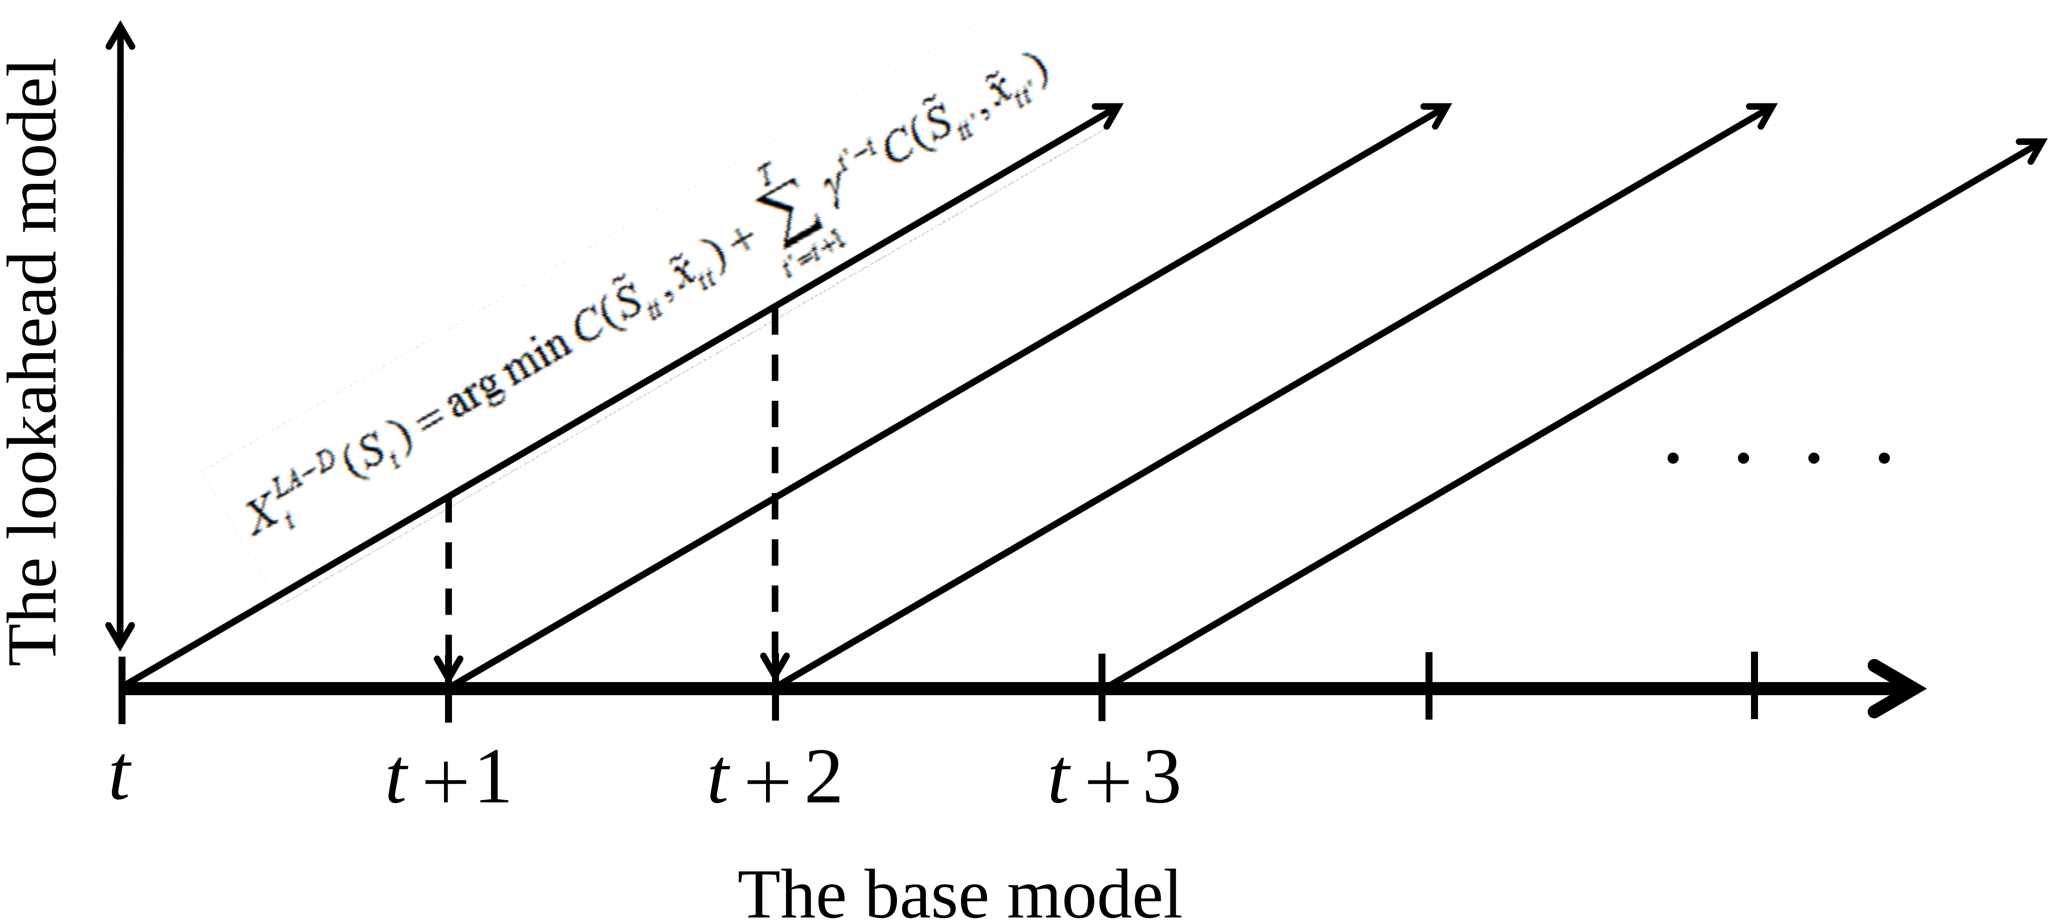

Lookahead policies

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

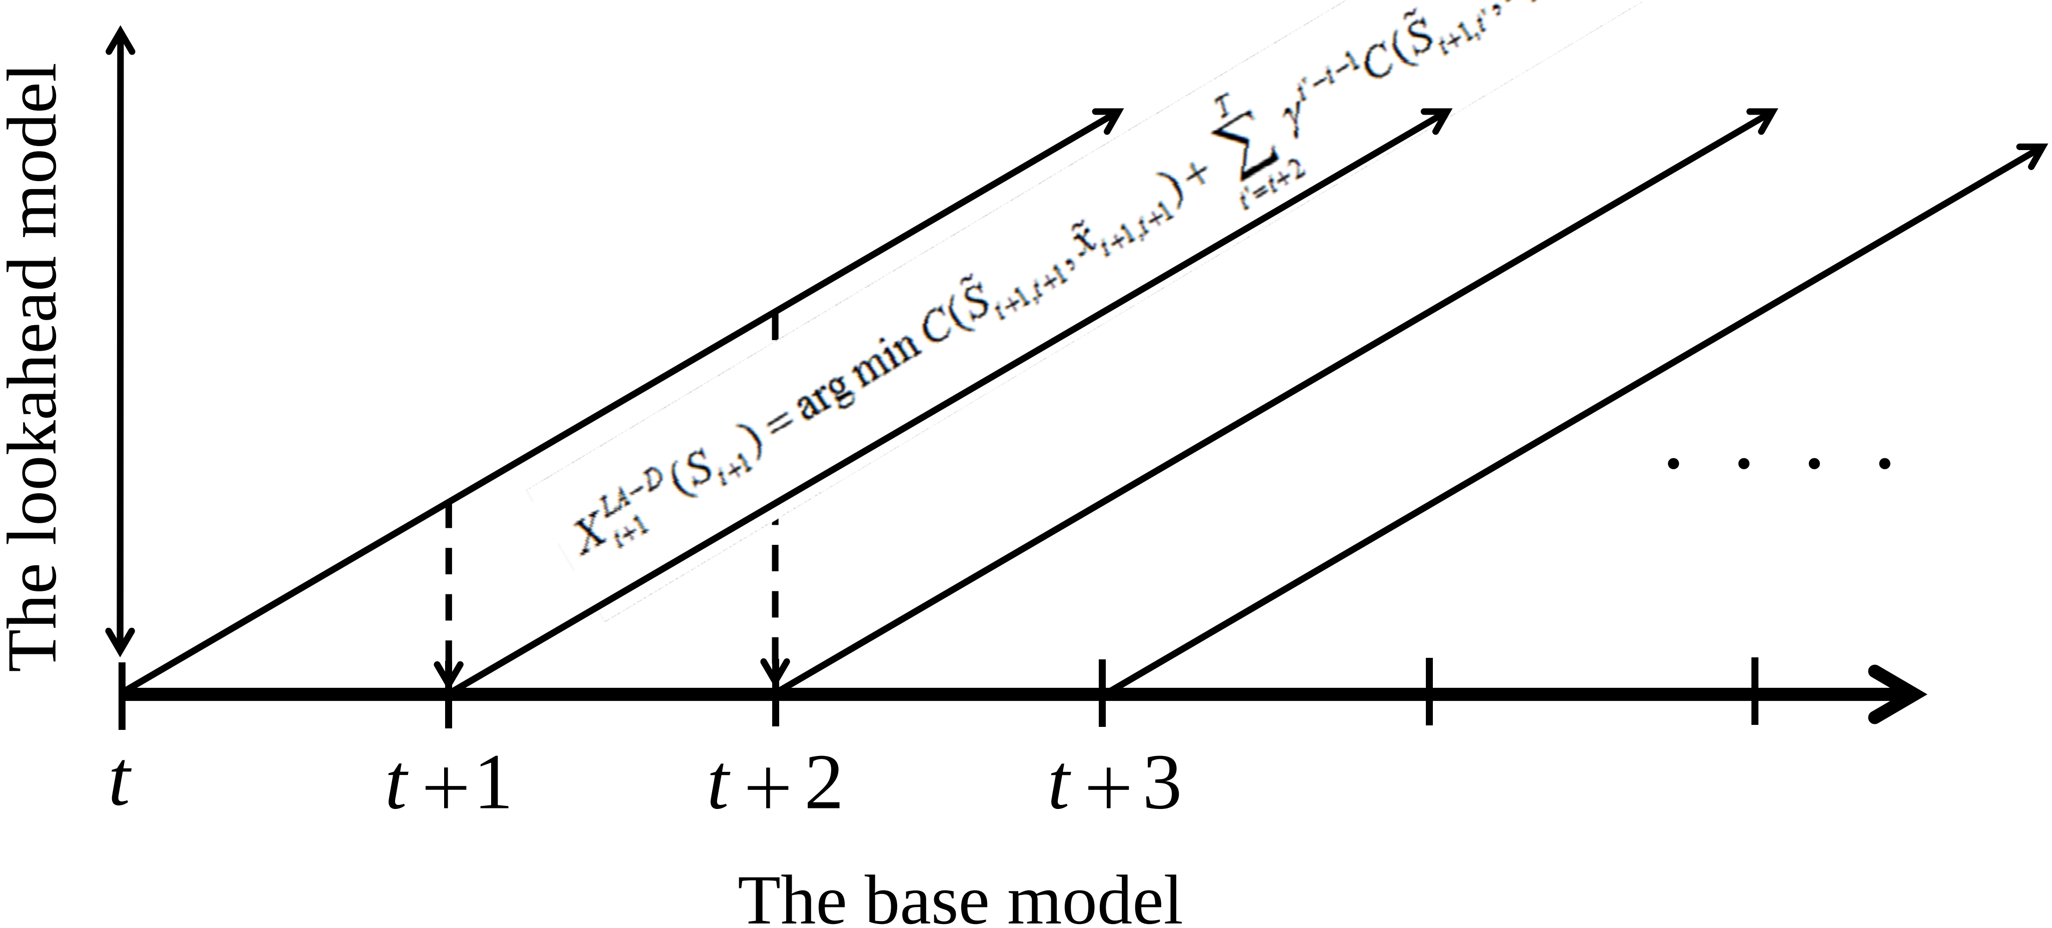

Lookahead policies

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

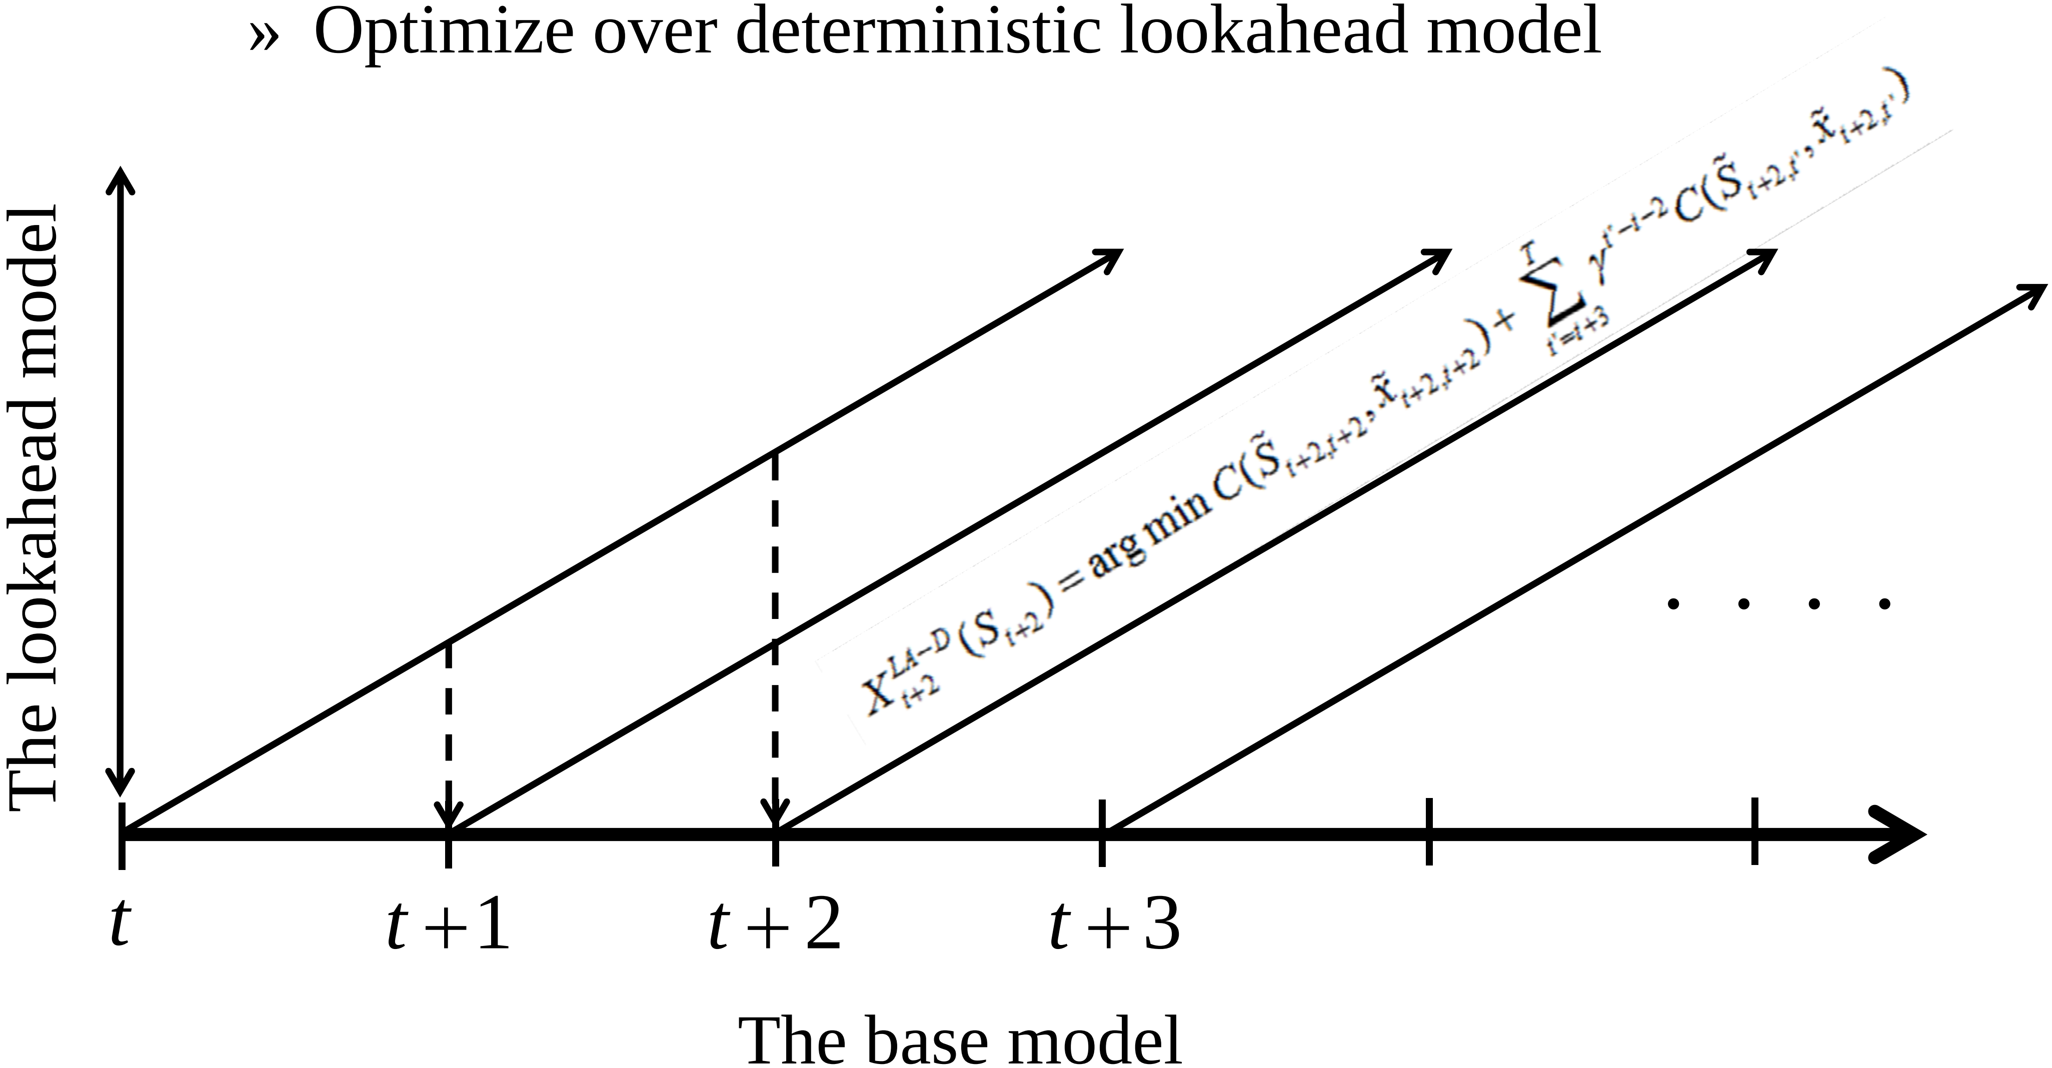

Four classes of policies

1) Policy function approximations (PFAs)

» Lookup tables, rules, parametric functions

2) Cost function approximation (CFAs)

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

3) Policies based on value function approximations (VFAs)

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

4) Lookahead policies

» ***Deterministic lookahead:***

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

» ***Stochastic lookahead (e.g. stochastic trees)***

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

Stochastic lookahead policies

- The optimal policy requires solving

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

Expectation that we cannot compute

Minimization that we cannot compute

Stochastic lookahead policies

- We use a series of approximations:
  - » Stage aggregation – Replacing multistage problems with two-stage approximations.
  - » Outcome aggregation/sampling – Simplifying the exogenous information process
  - » Discretization – Of time, states and decisions
  - » Horizon truncation – Replacing a longer horizon problem with a shorter horizon
  - » 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).

Stochastic lookahead policies

- We use a series of approximations:
  - » Stage aggregation – Replacing multistage problems with two-stage approximations.
  - » Outcome aggregation/sampling – Simplifying the exogenous information process
  - » Discretization – Of time, states and decisions
  - » Horizon truncation – Replacing a longer horizon problem with a shorter horizon
  - » 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).

Stochastic lookahead policies

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

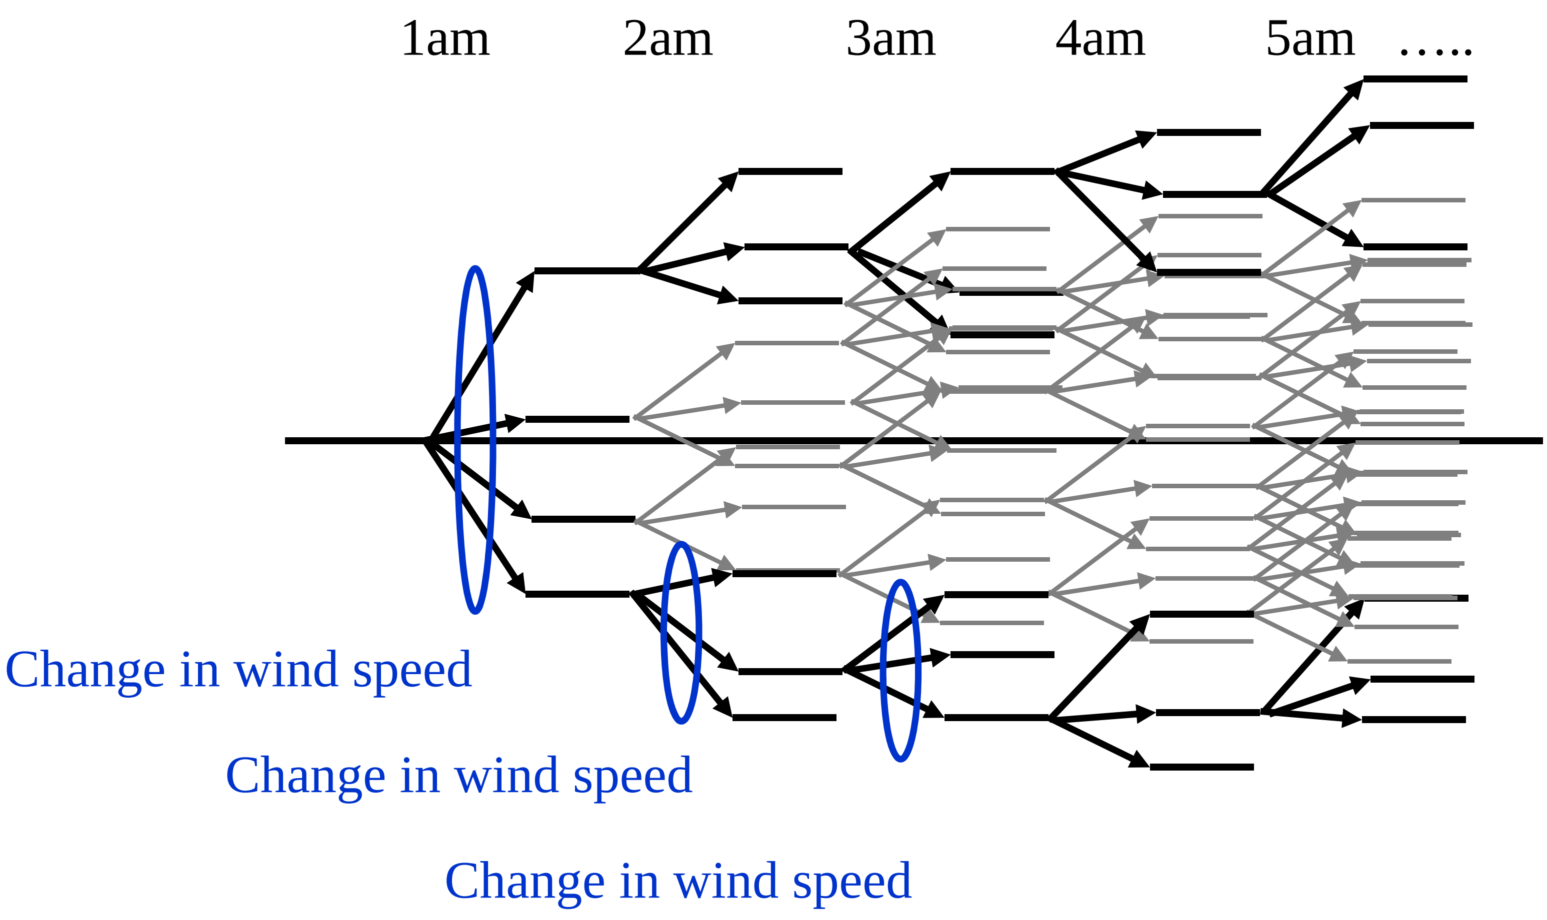

Stochastic lookahead policies

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

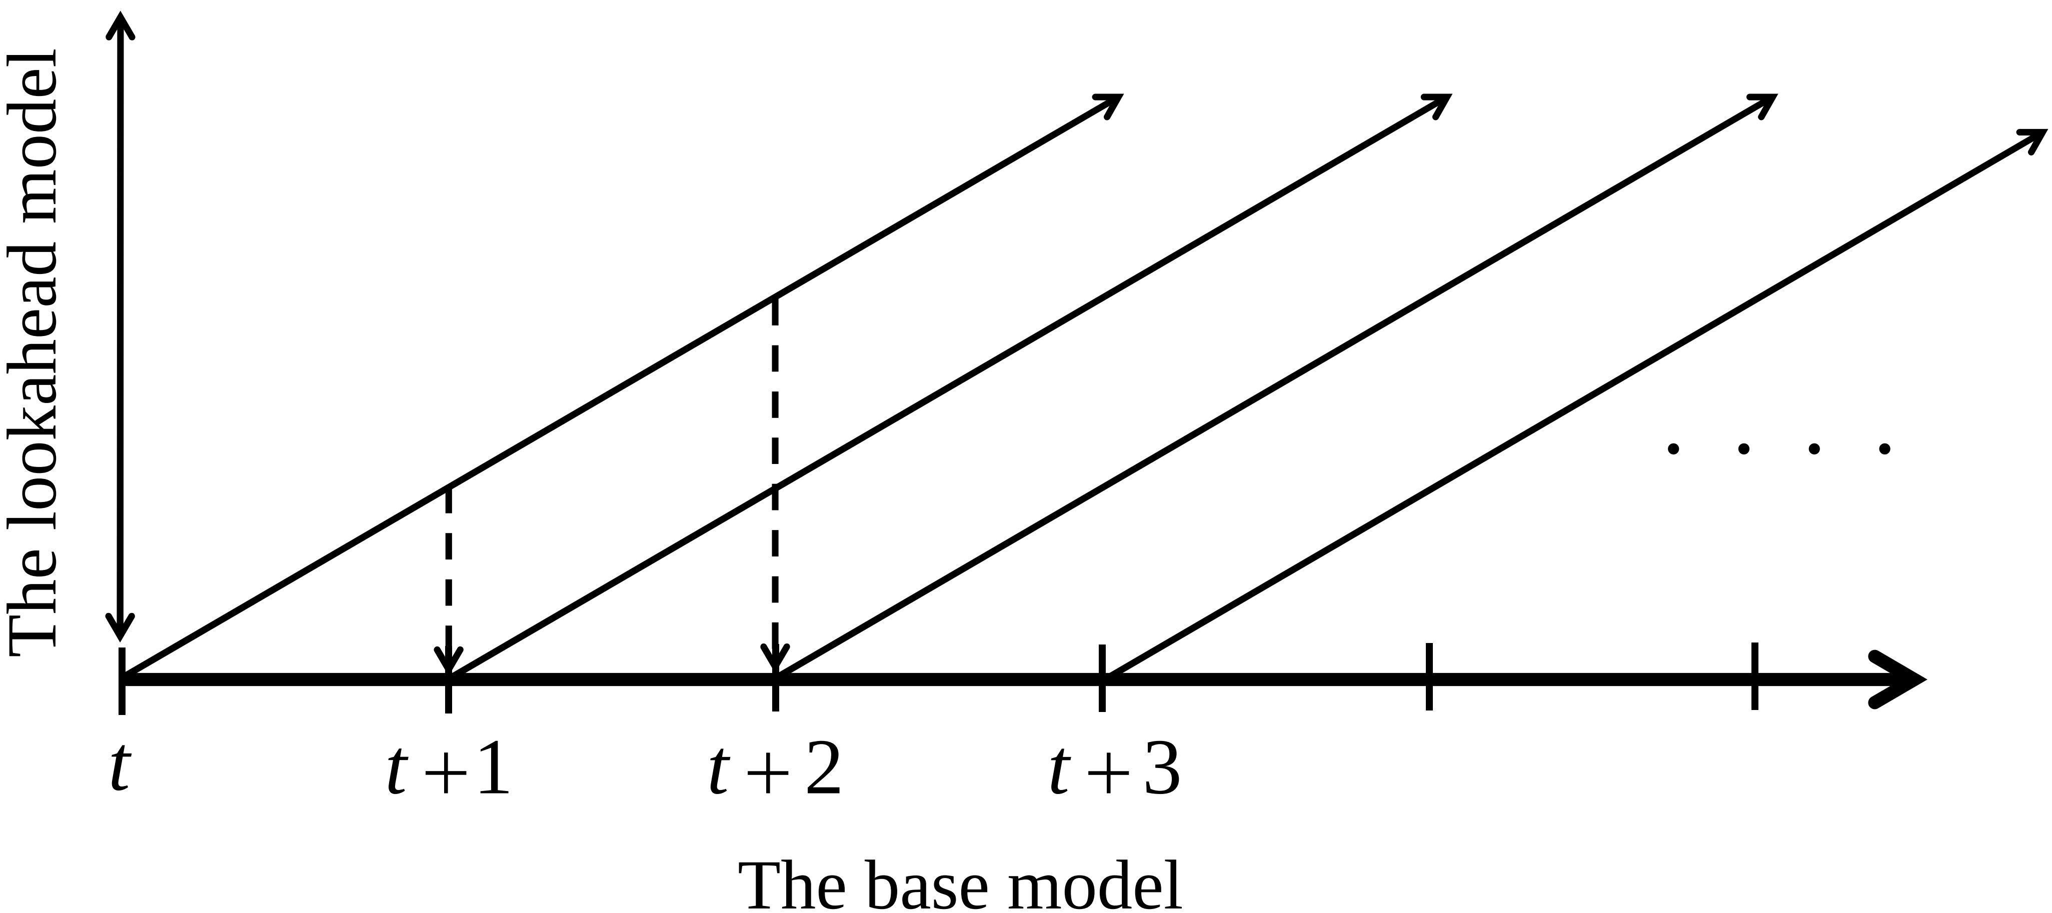

Stochastic lookahead policies

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

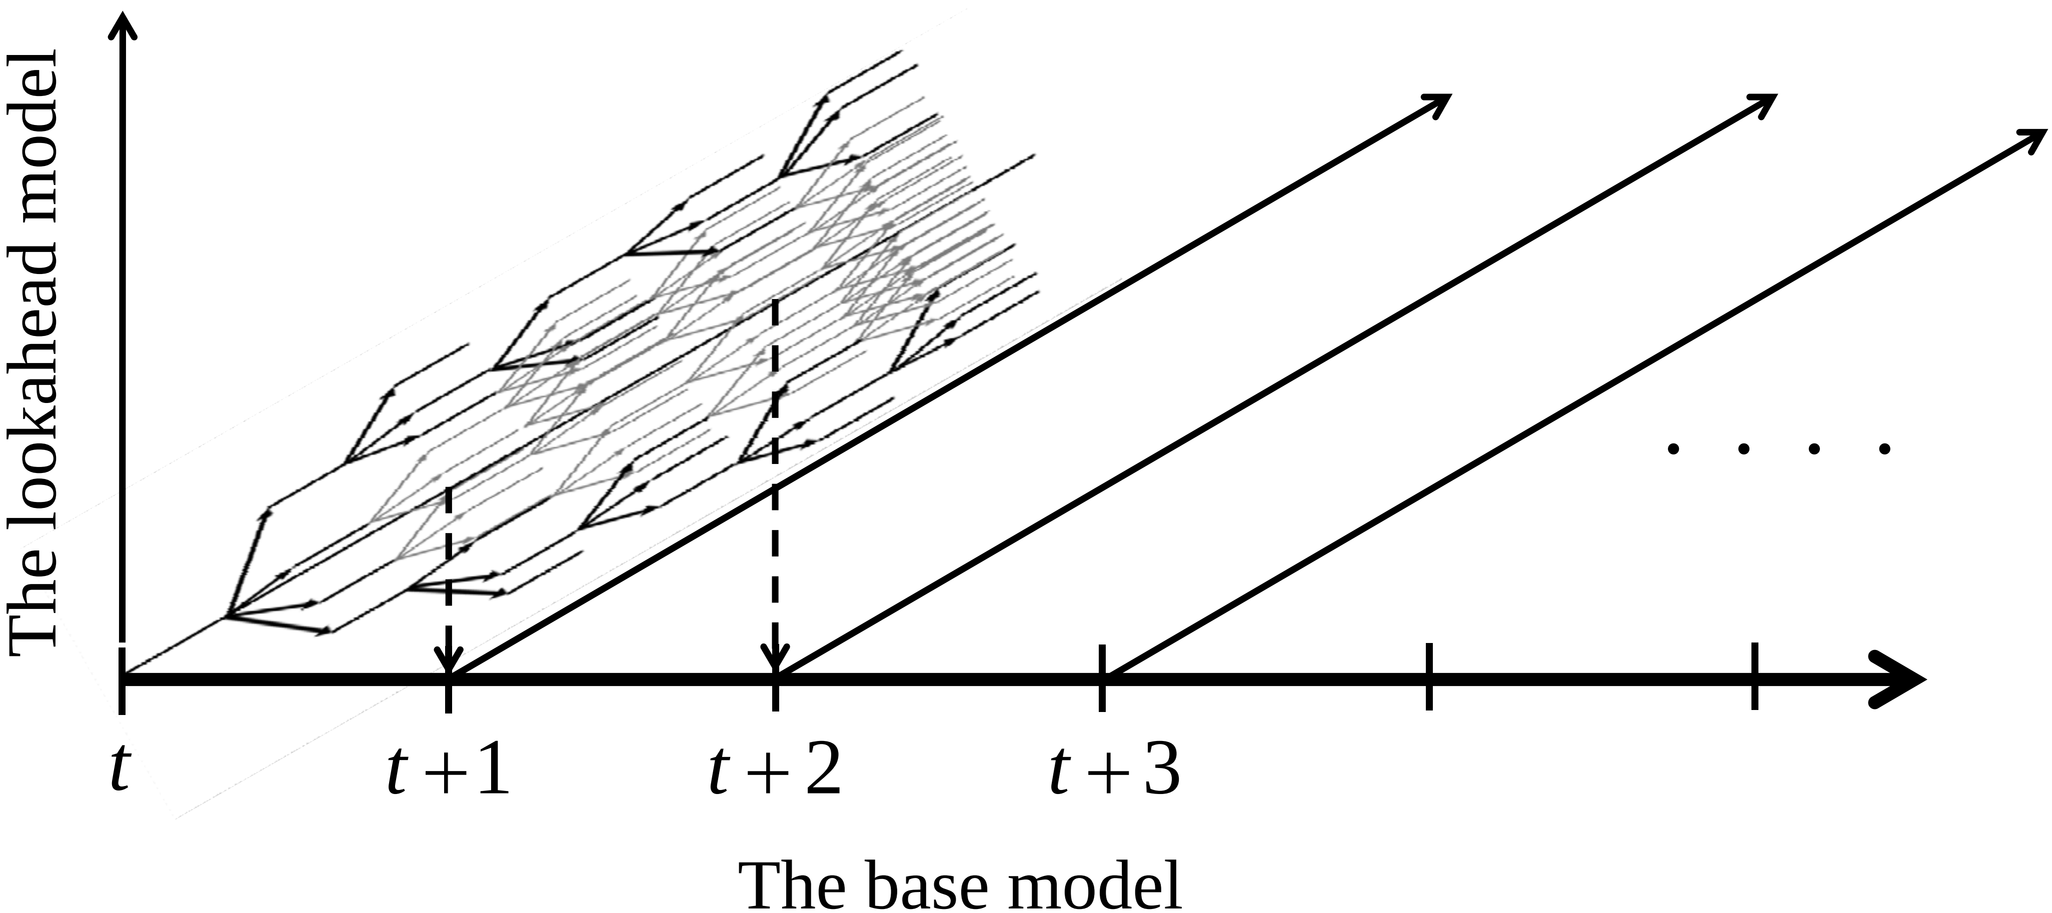

Stochastic lookahead policies

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

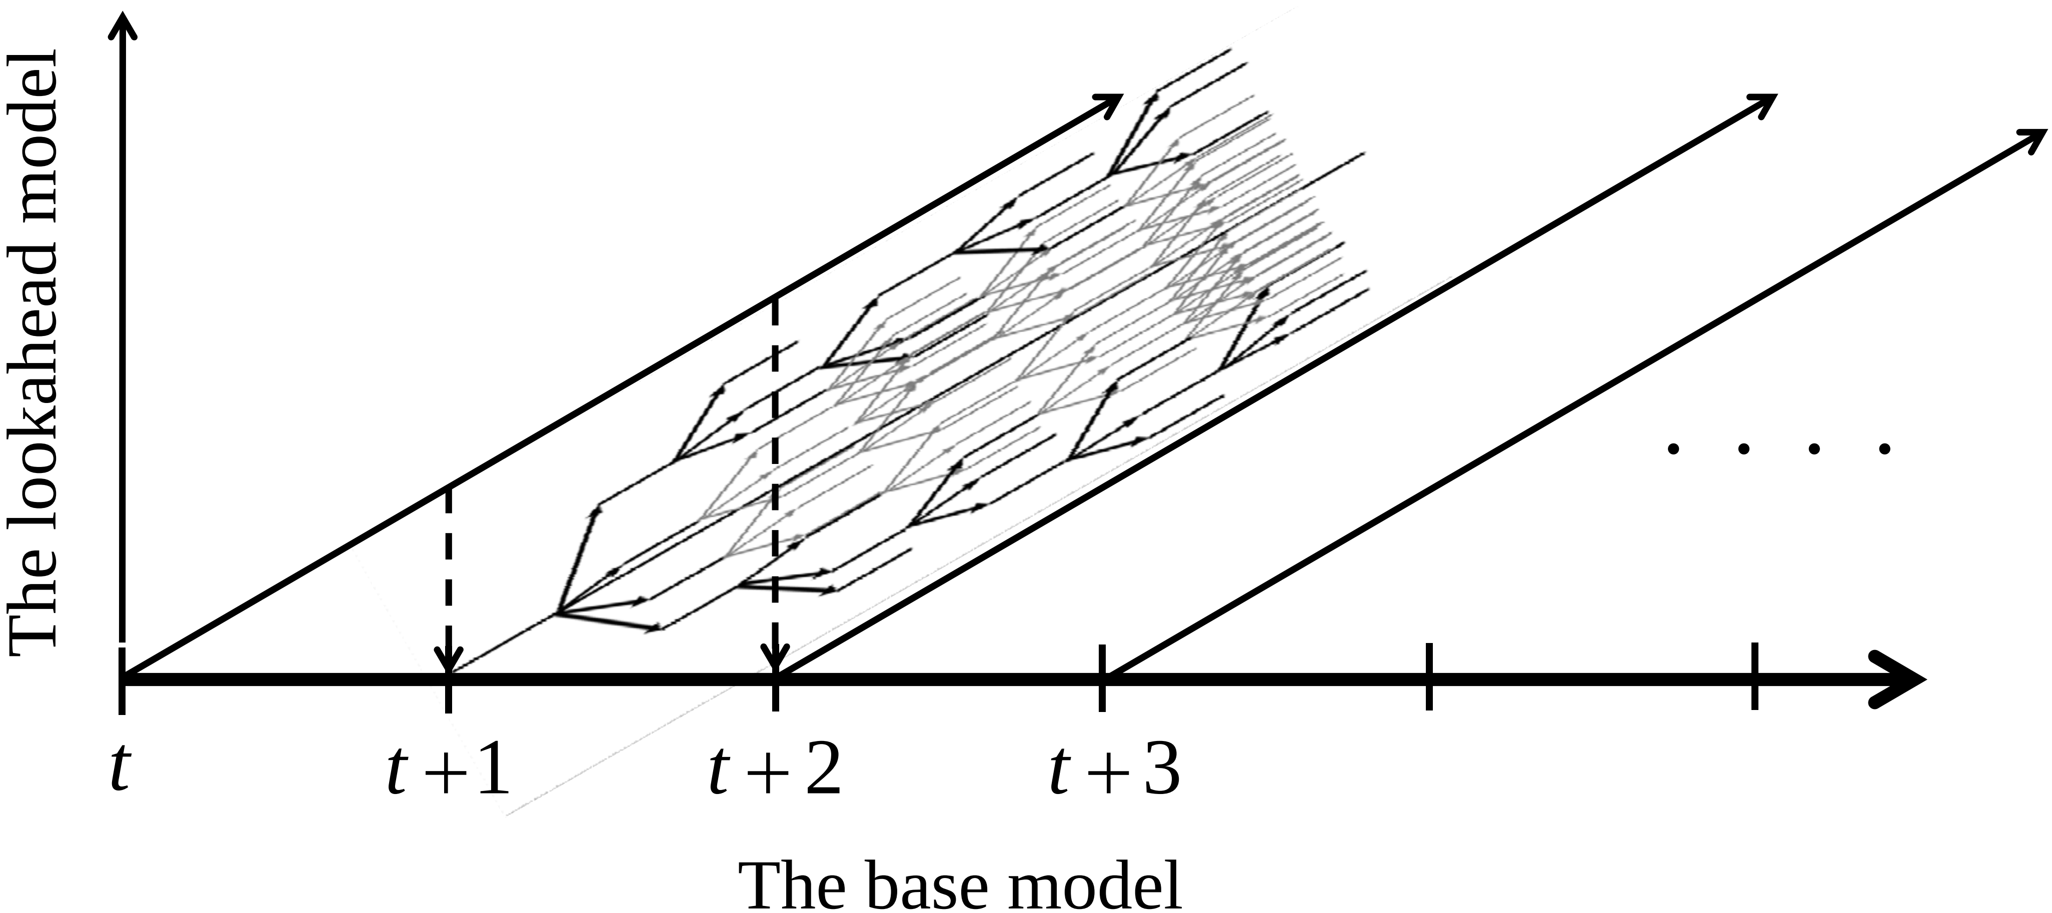

Stochastic lookahead policies

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

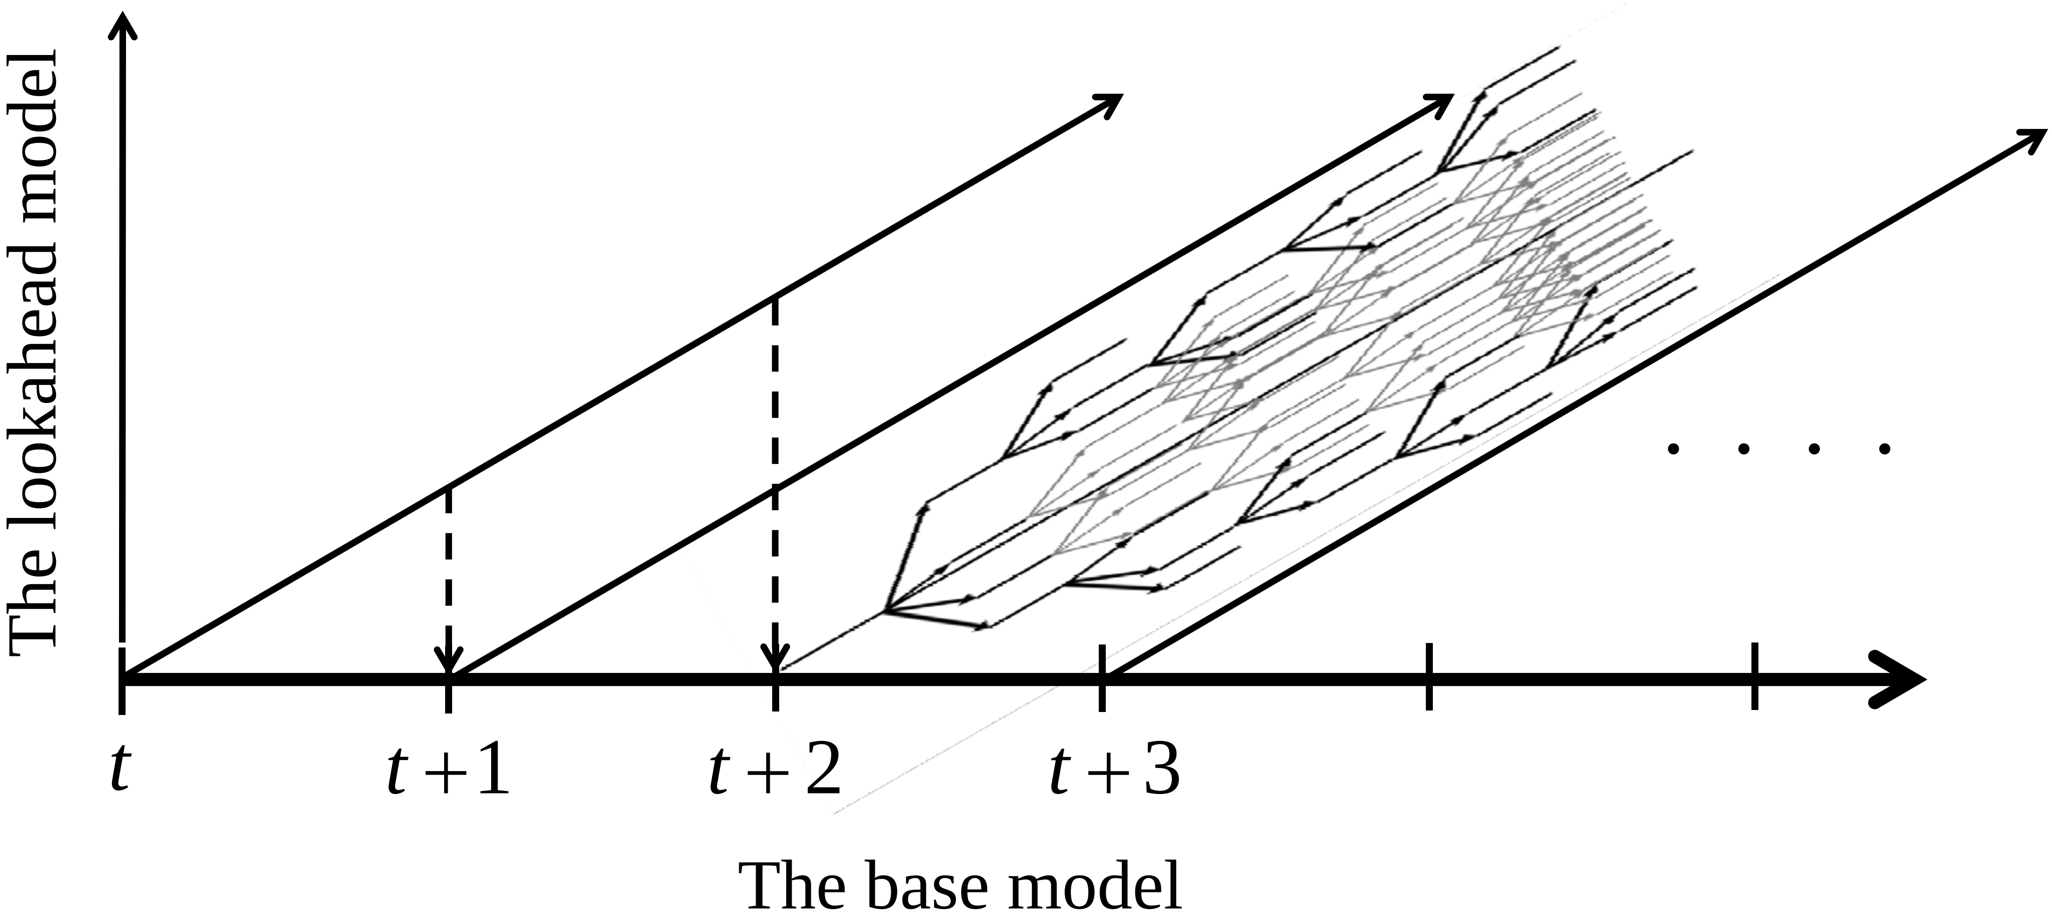

Stochastic lookahead policies

- Multistage lookahead approximation
  - The problem with multistage trees is that even with sparse sampling, they are too expensive to compute:

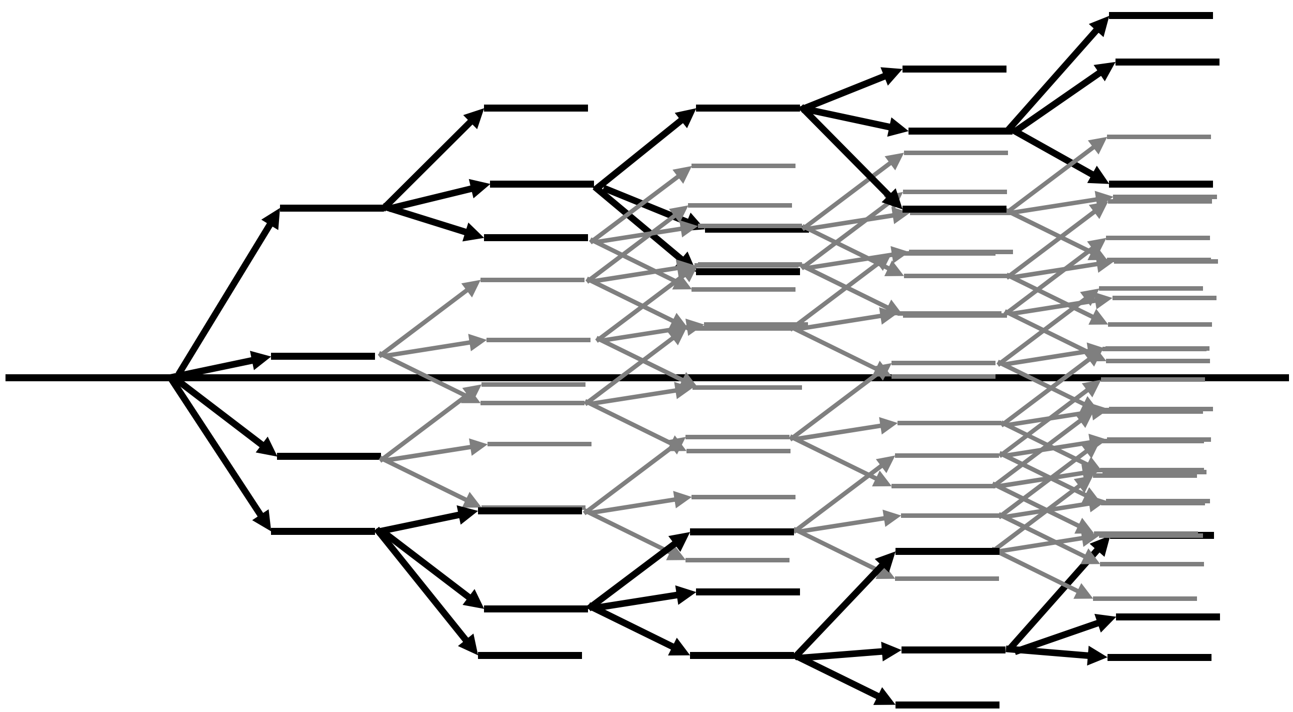

Stochastic lookahead policies

- Two stage lookahead approximation

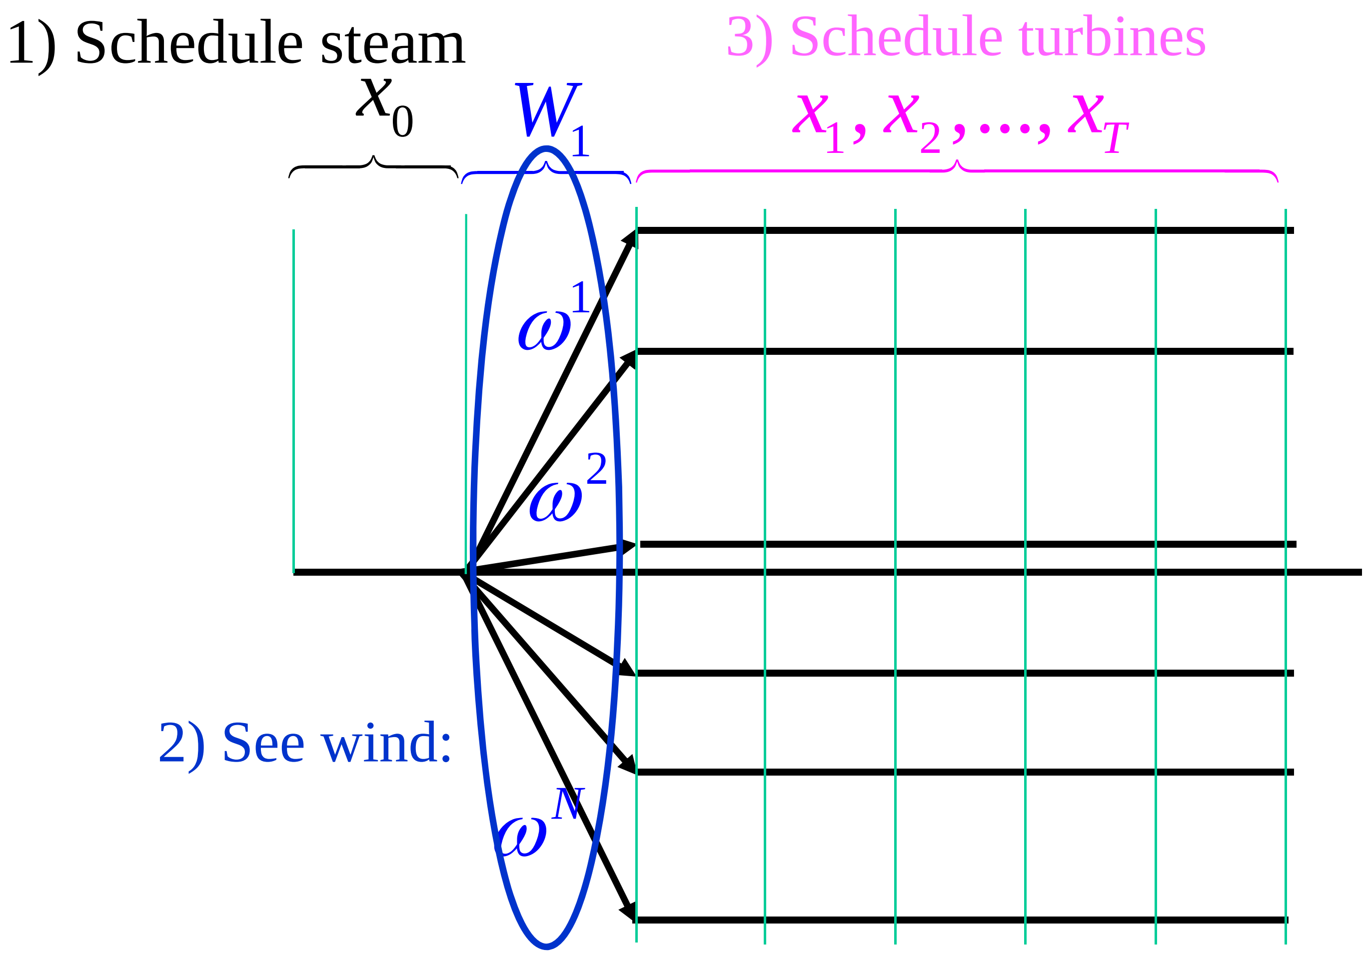

Stochastic lookahead policies

- Creating wind scenarios (Scenario #1)

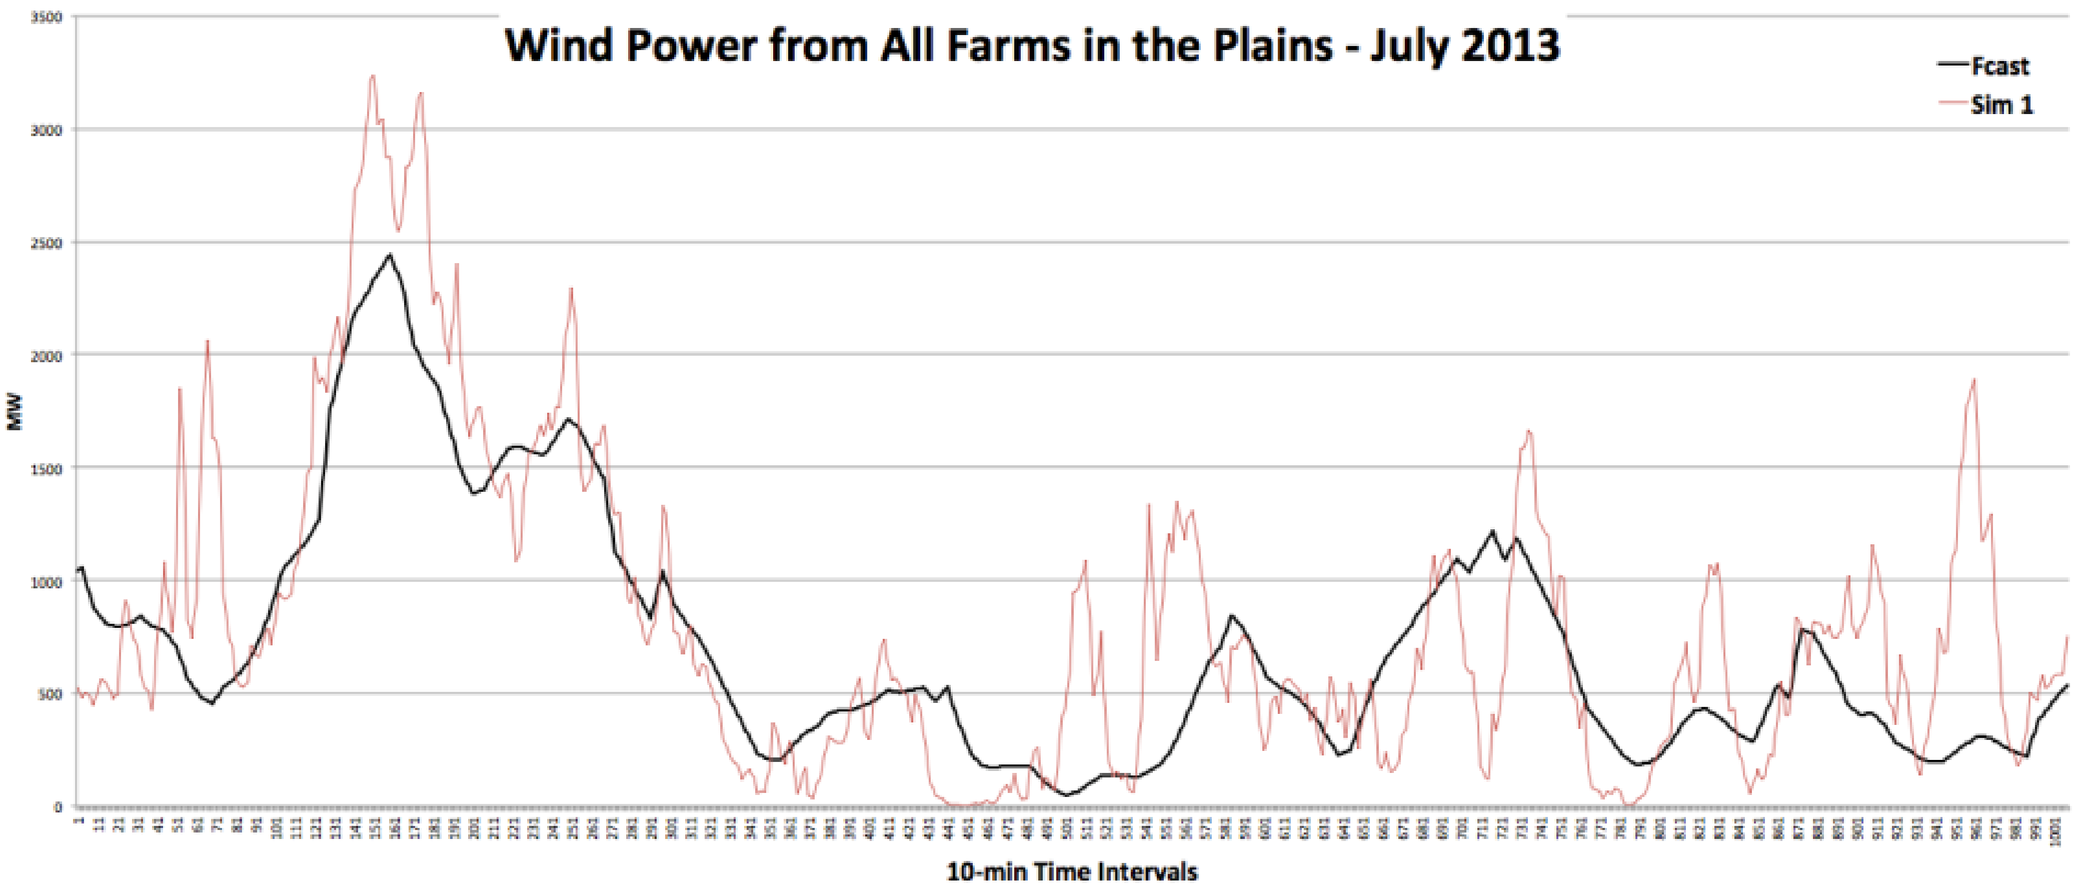

Stochastic lookahead policies

- Creating wind scenarios (Scenario #2)

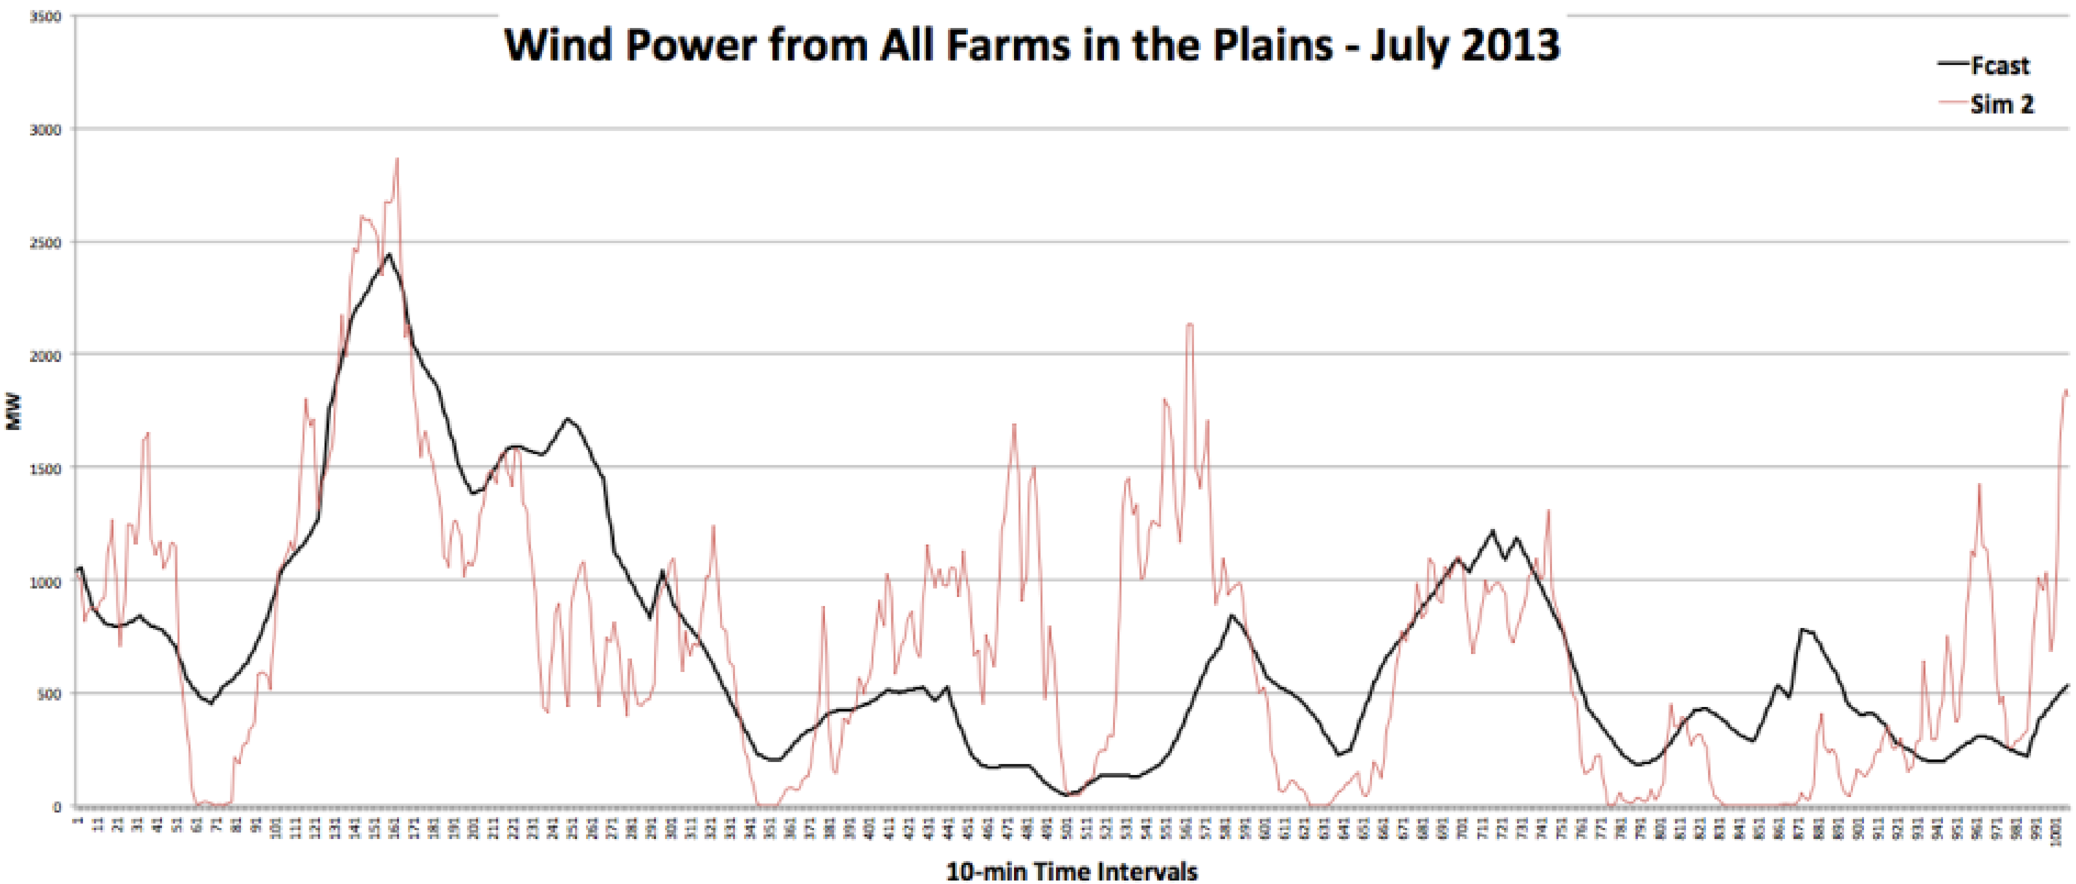

Stochastic lookahead policies

- Creating wind scenarios (Scenario #3)

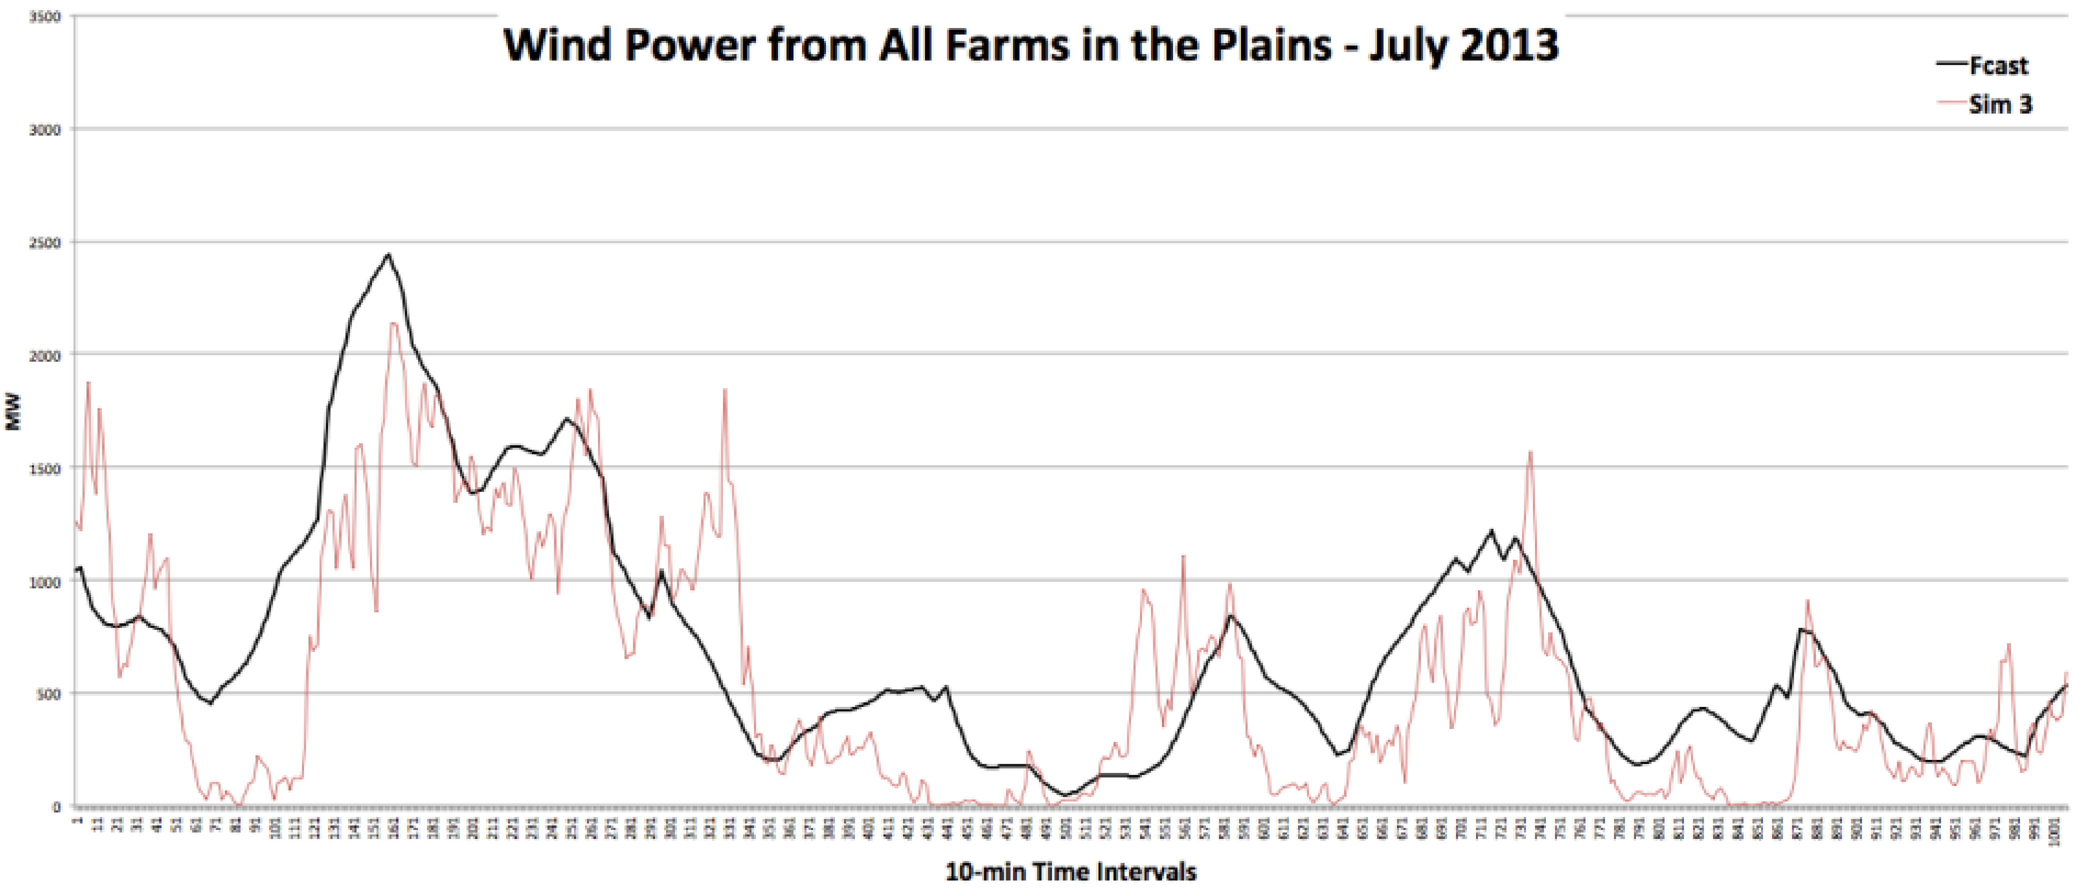

Stochastic lookahead policies

- Creating wind scenarios (Scenario #4)

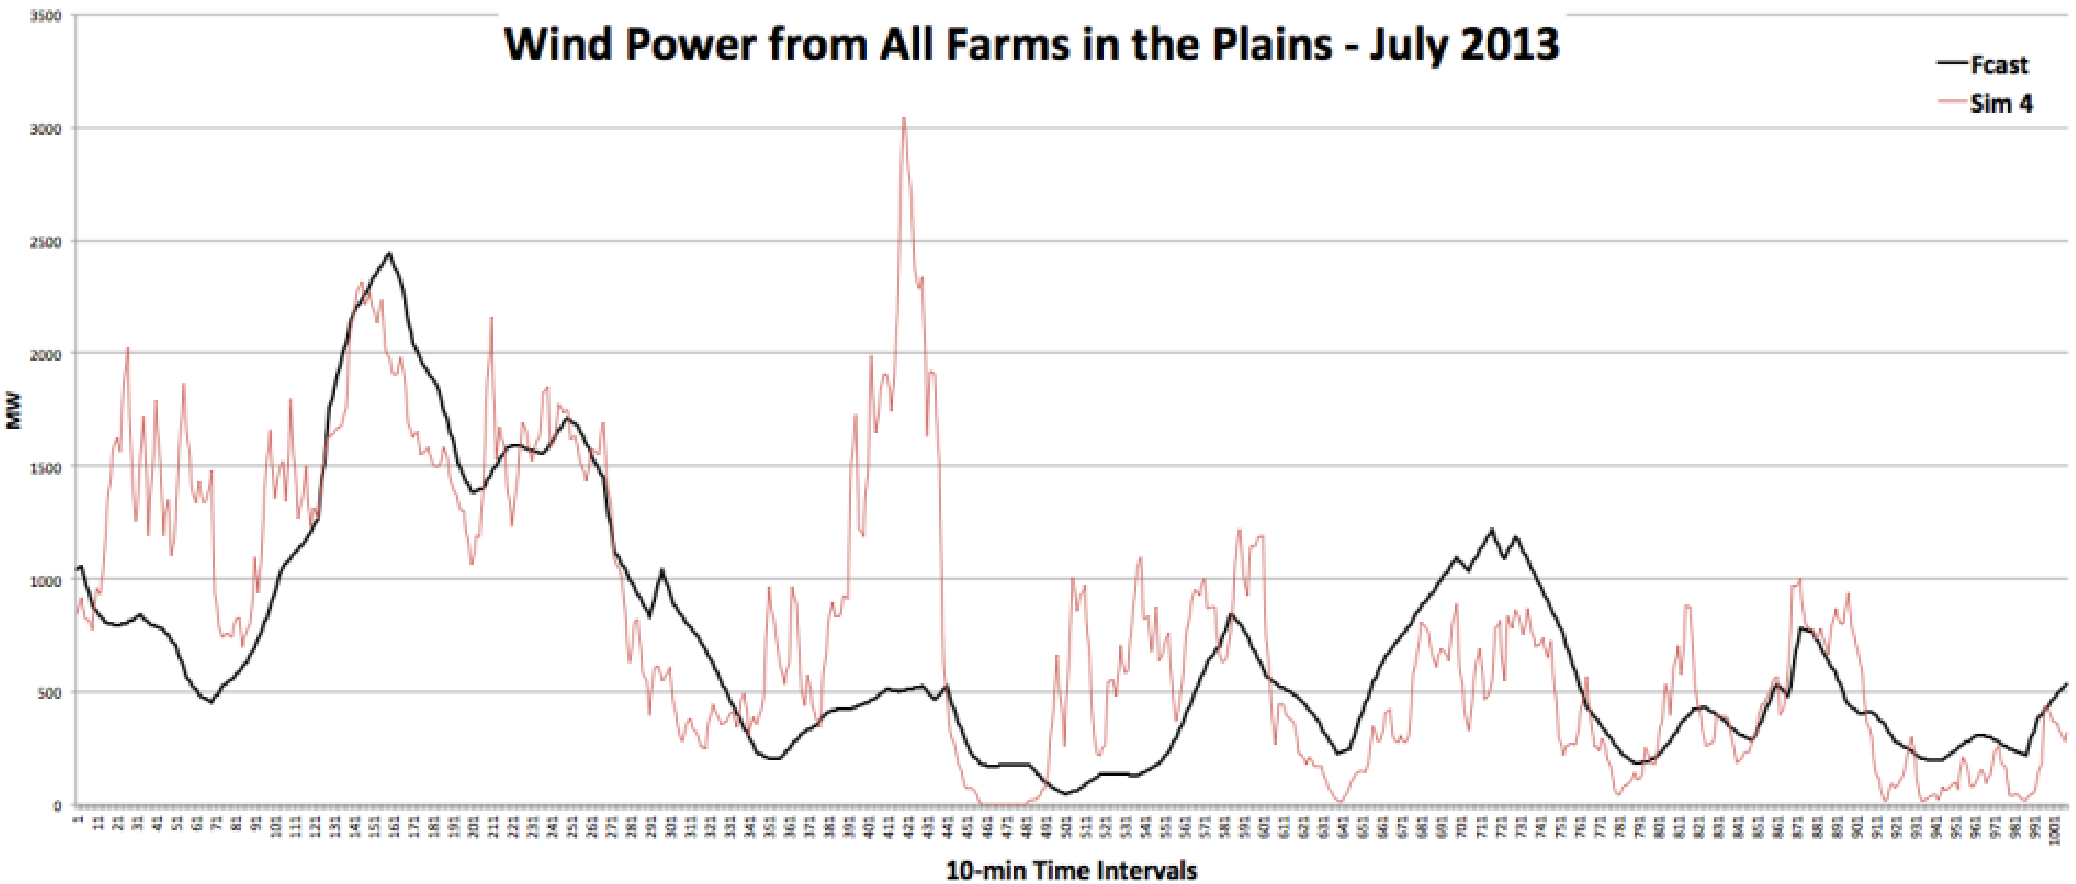

Stochastic lookahead policies

- Creating wind scenarios (Scenario #5)

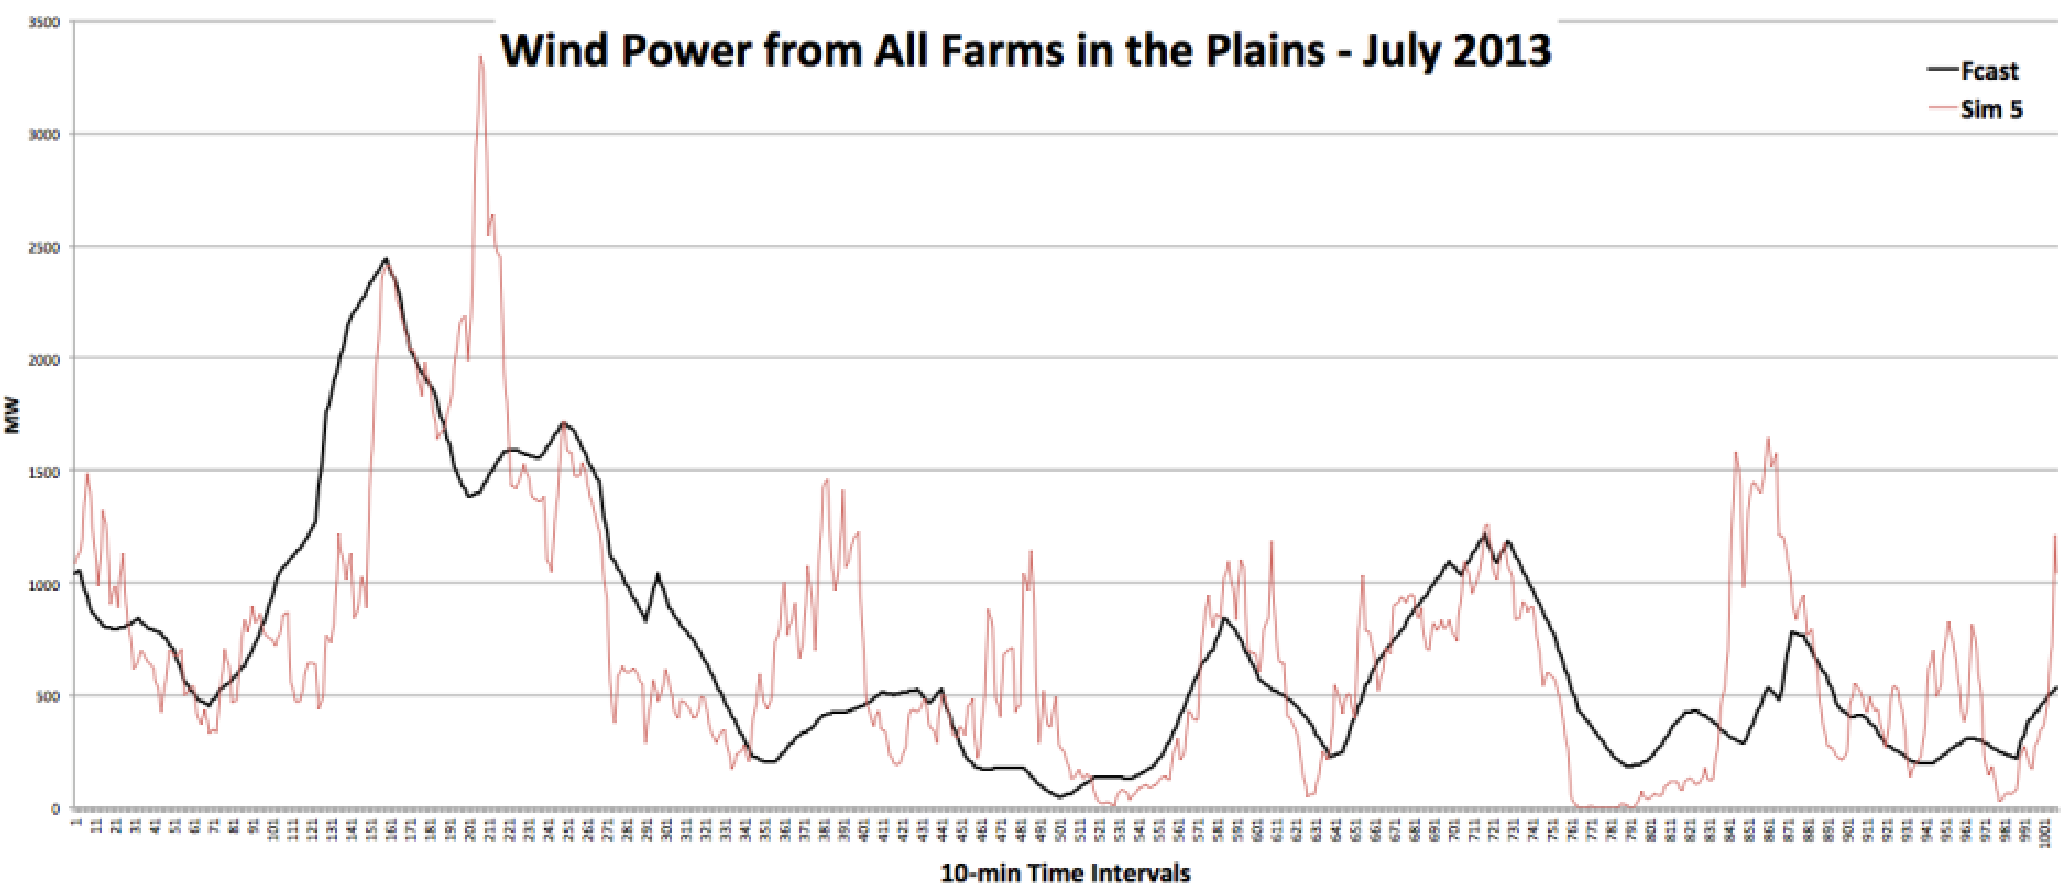

Stochastic lookahead policies

- The two-stage approximation

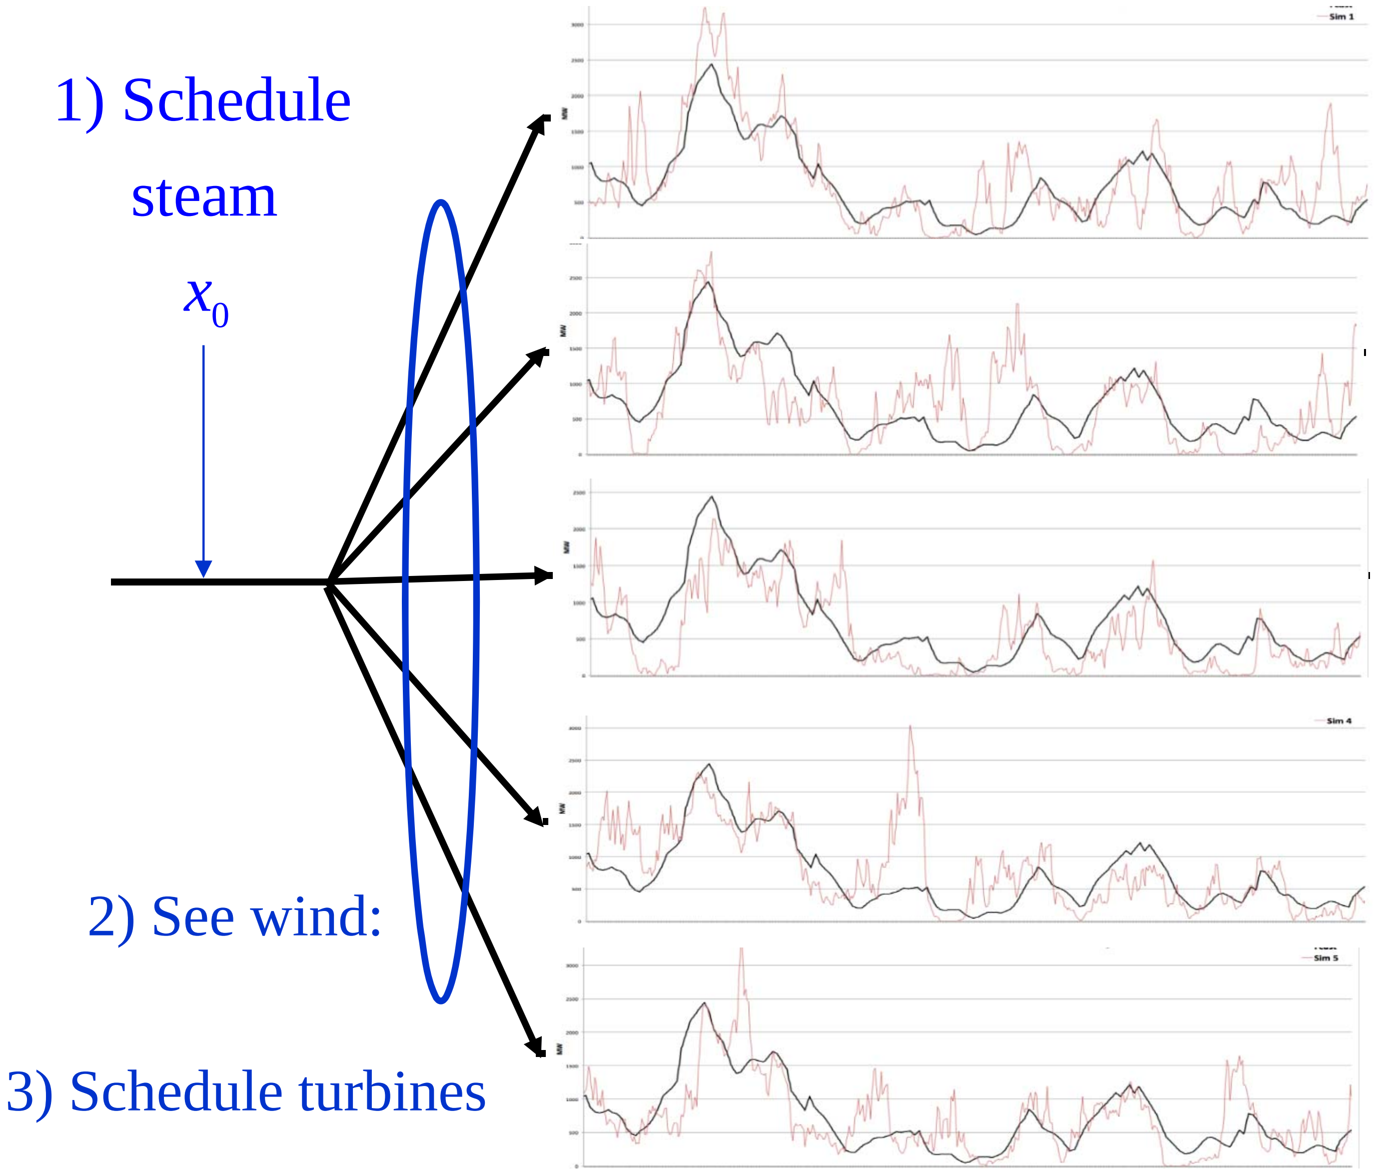

Stochastic lookahead policies

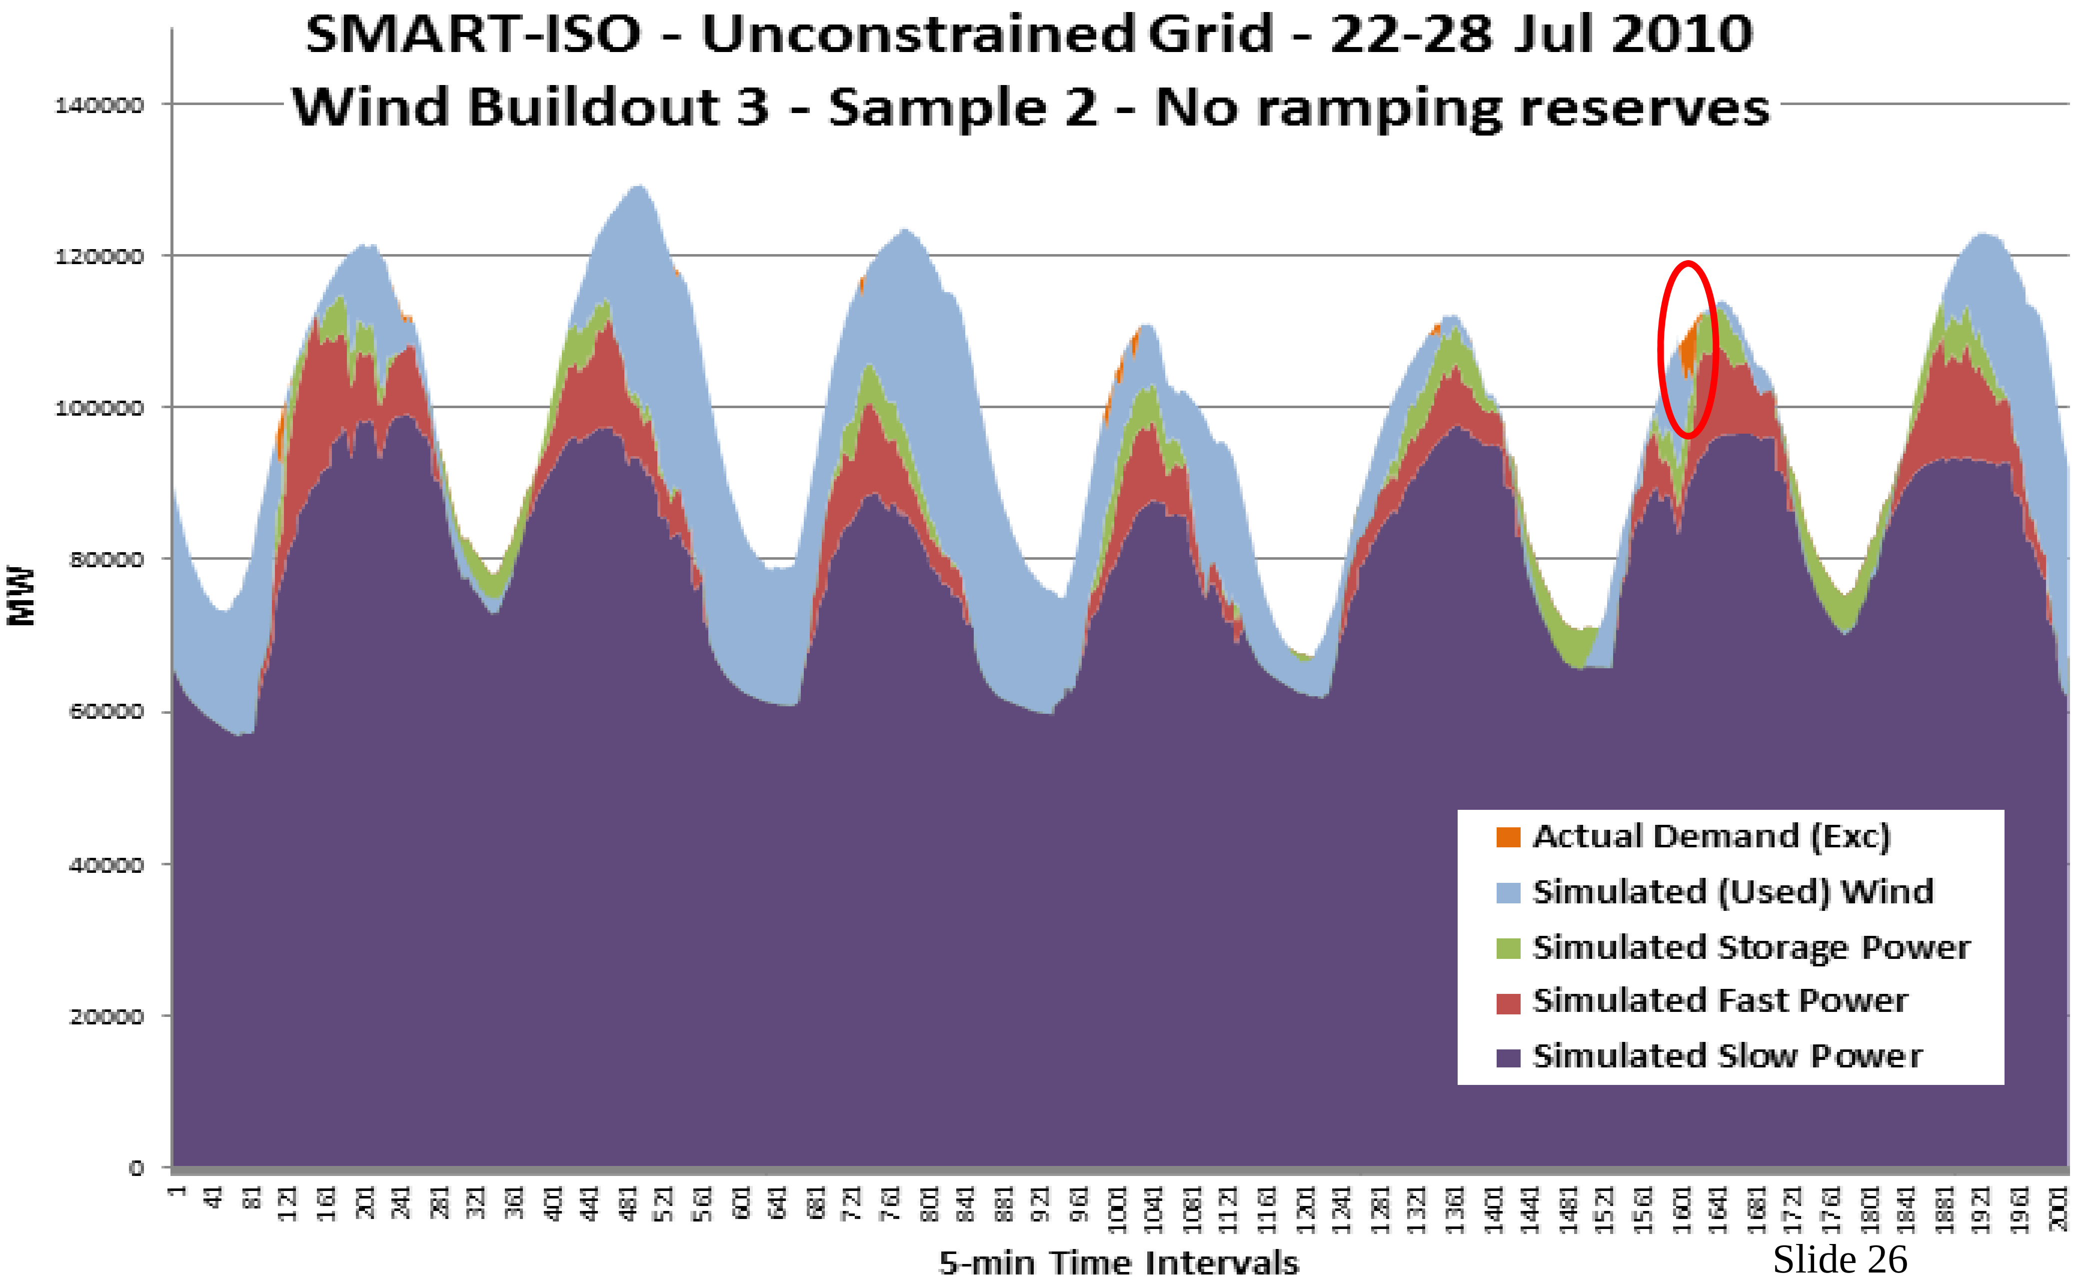

Stochastic lookahead policies

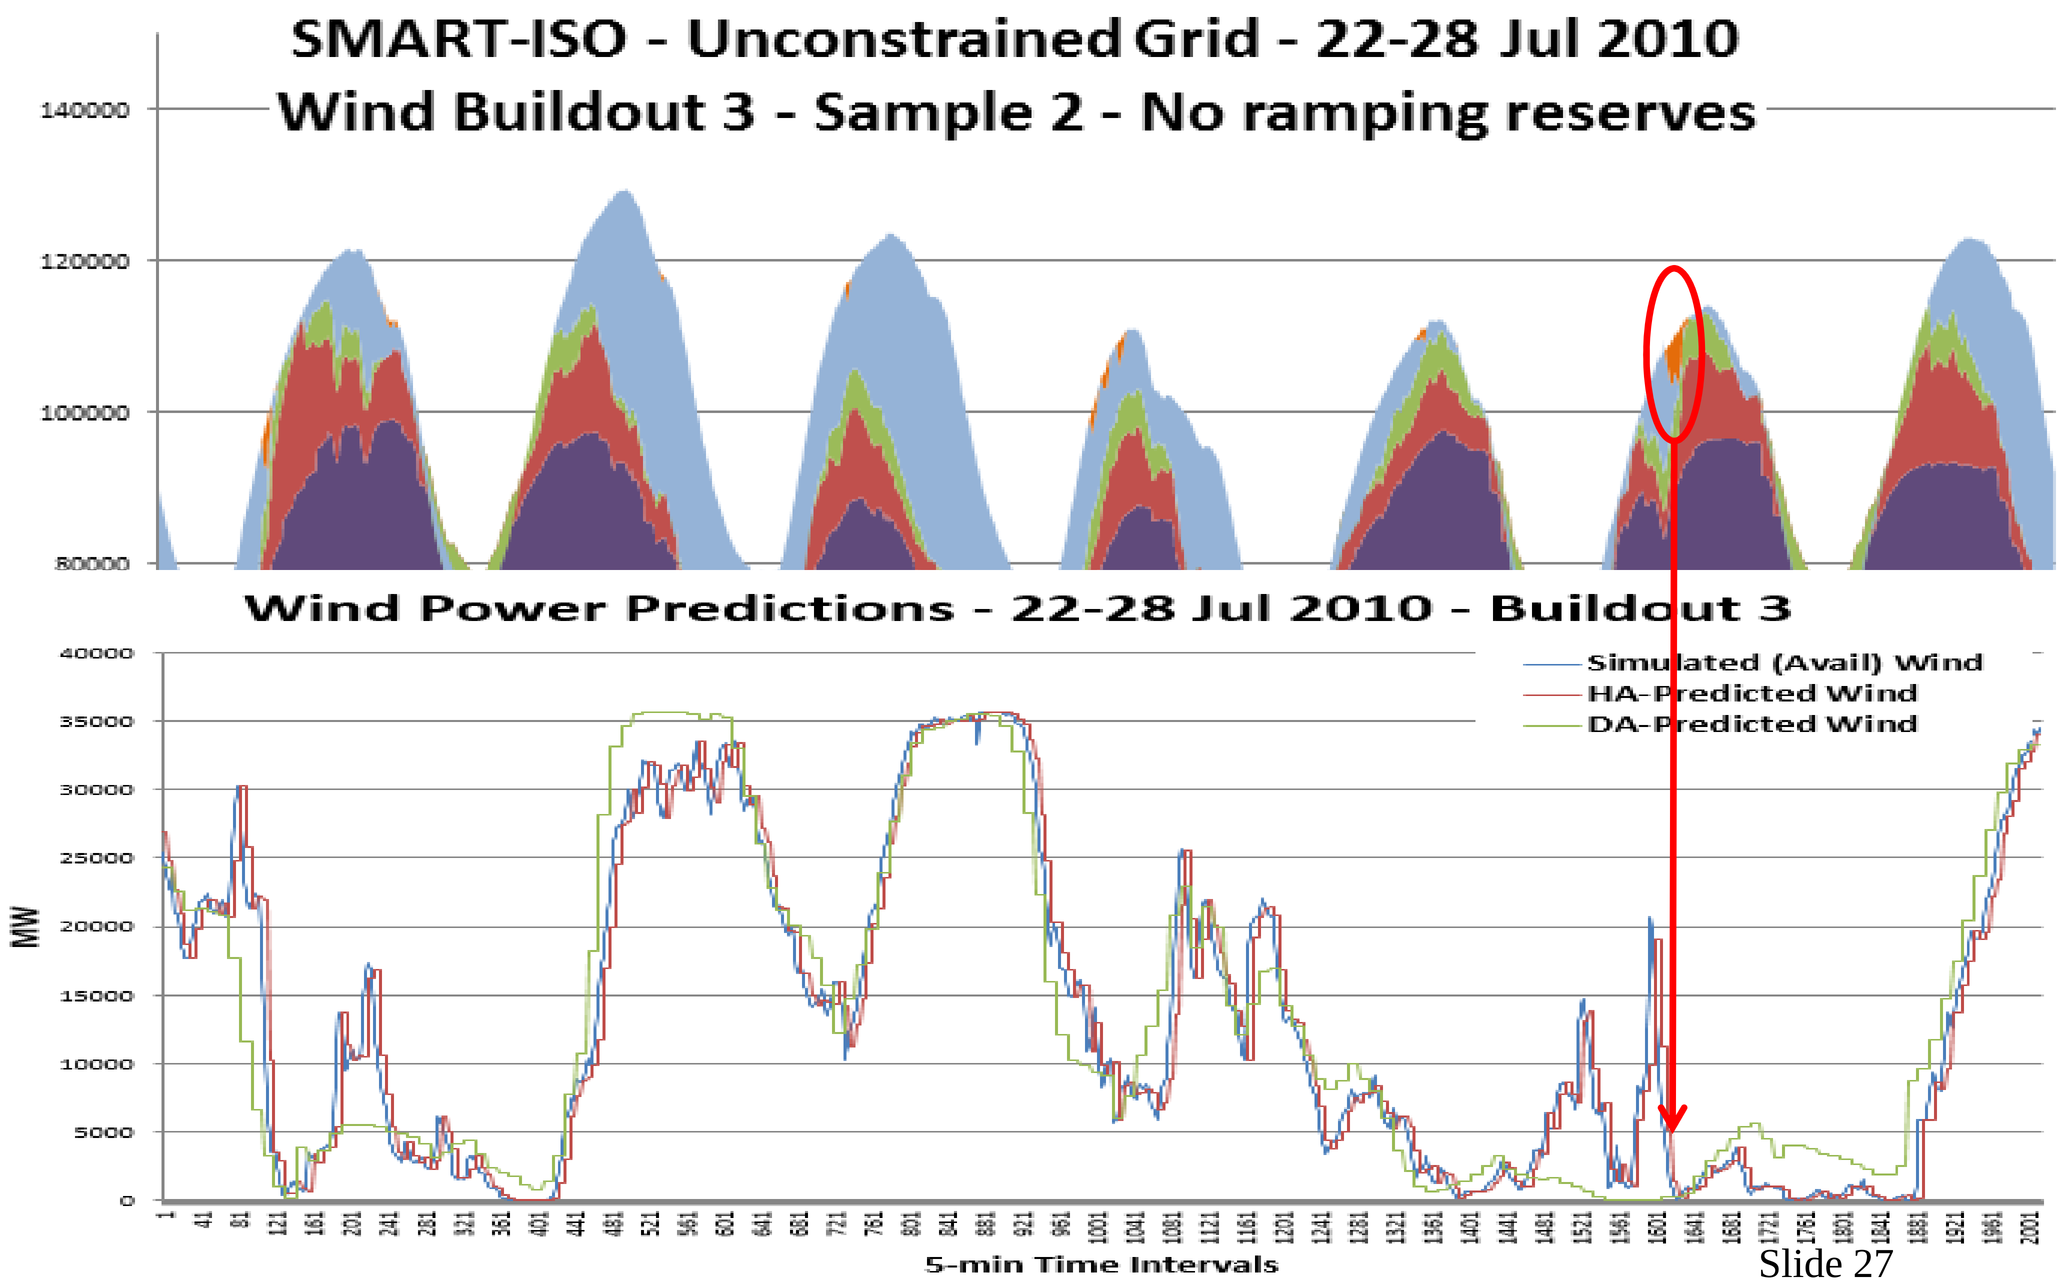

Stochastic lookahead policies

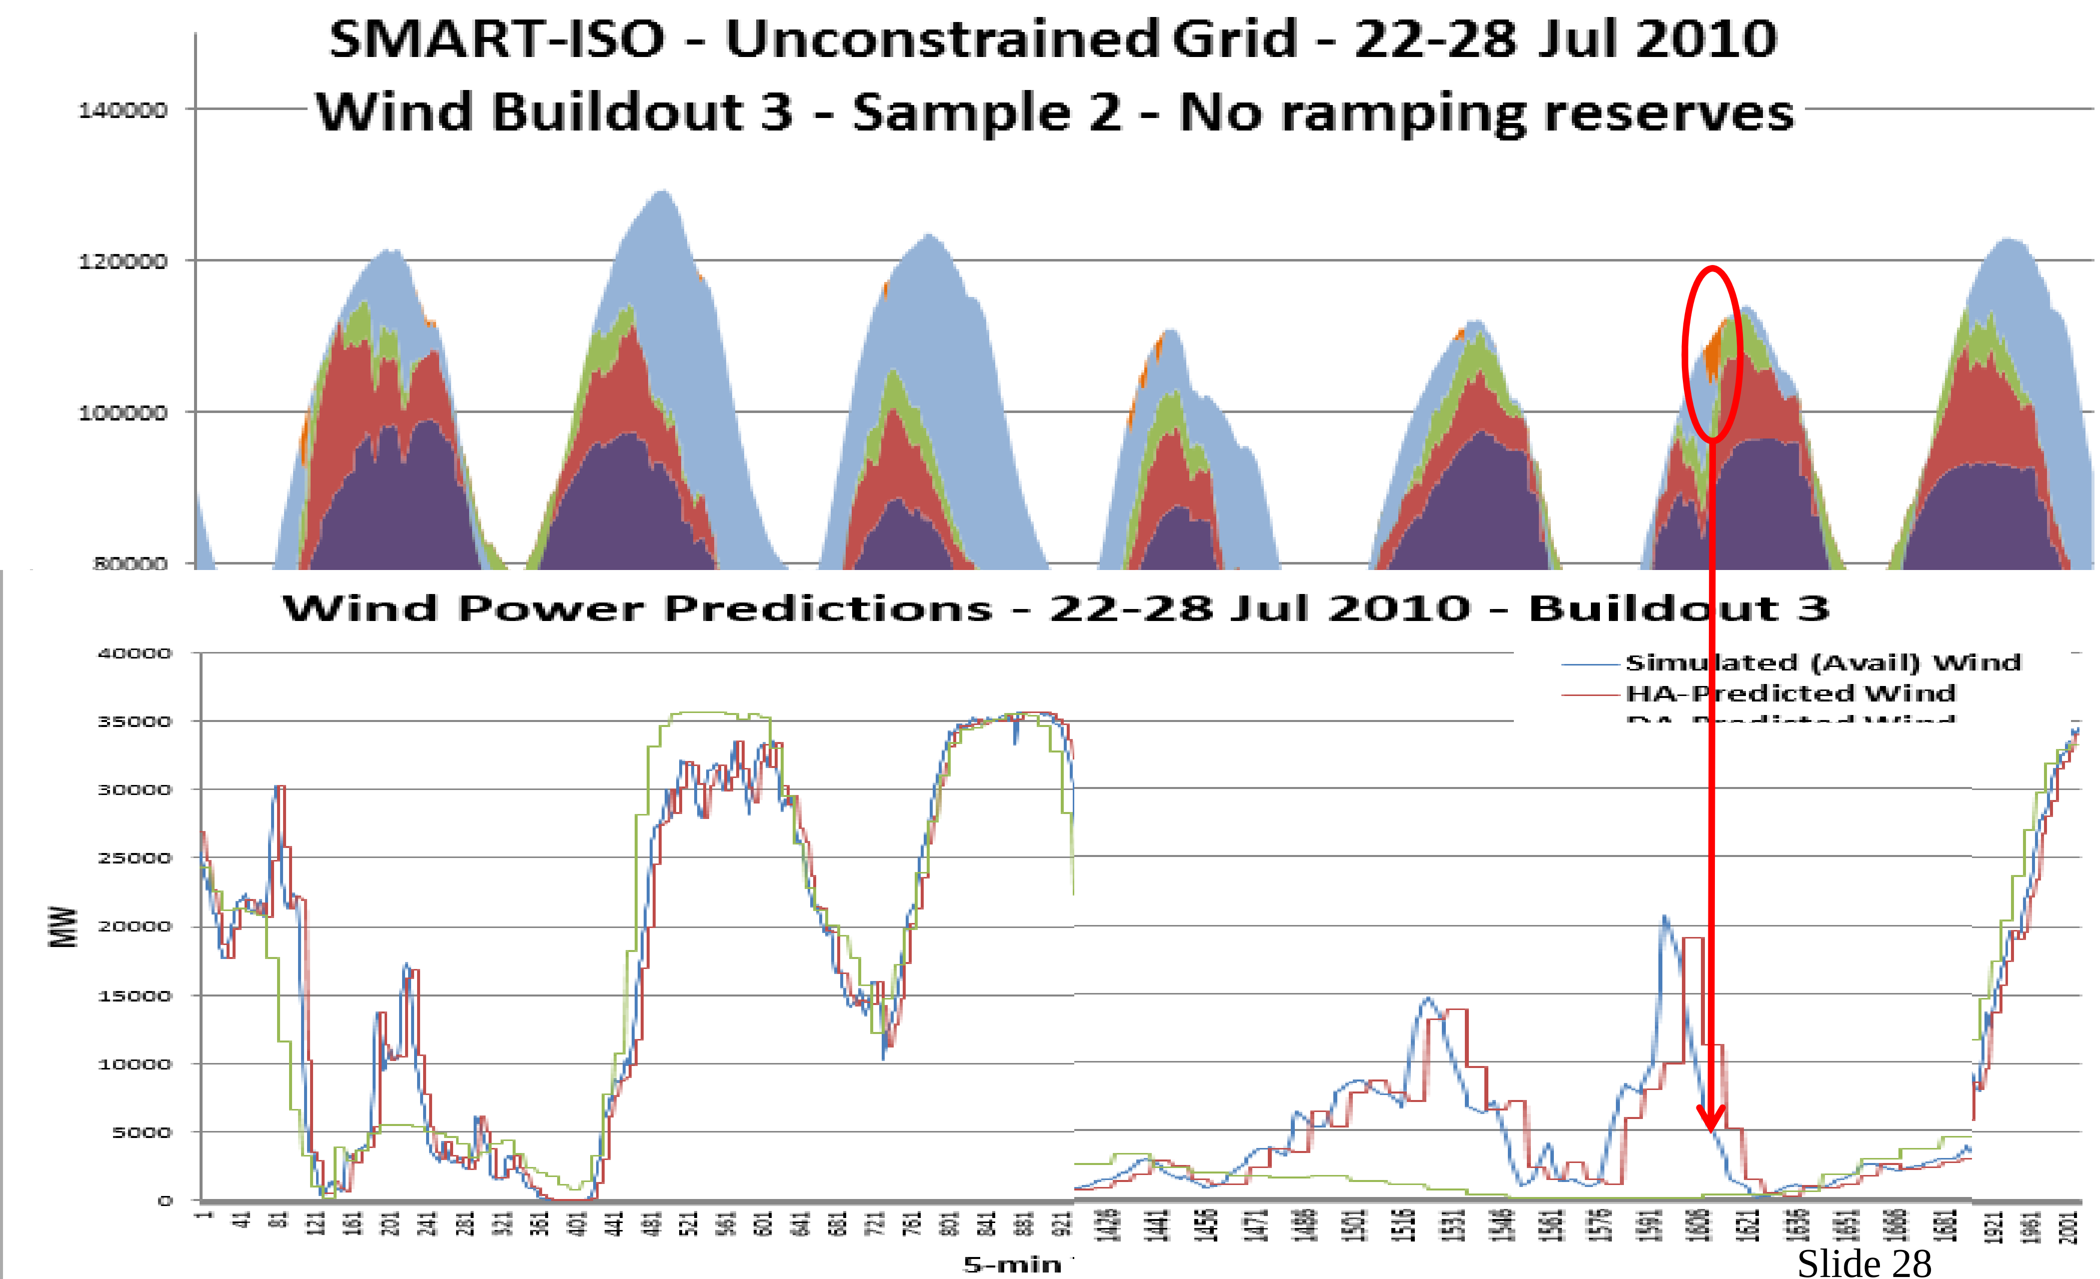

Stochastic lookahead policies

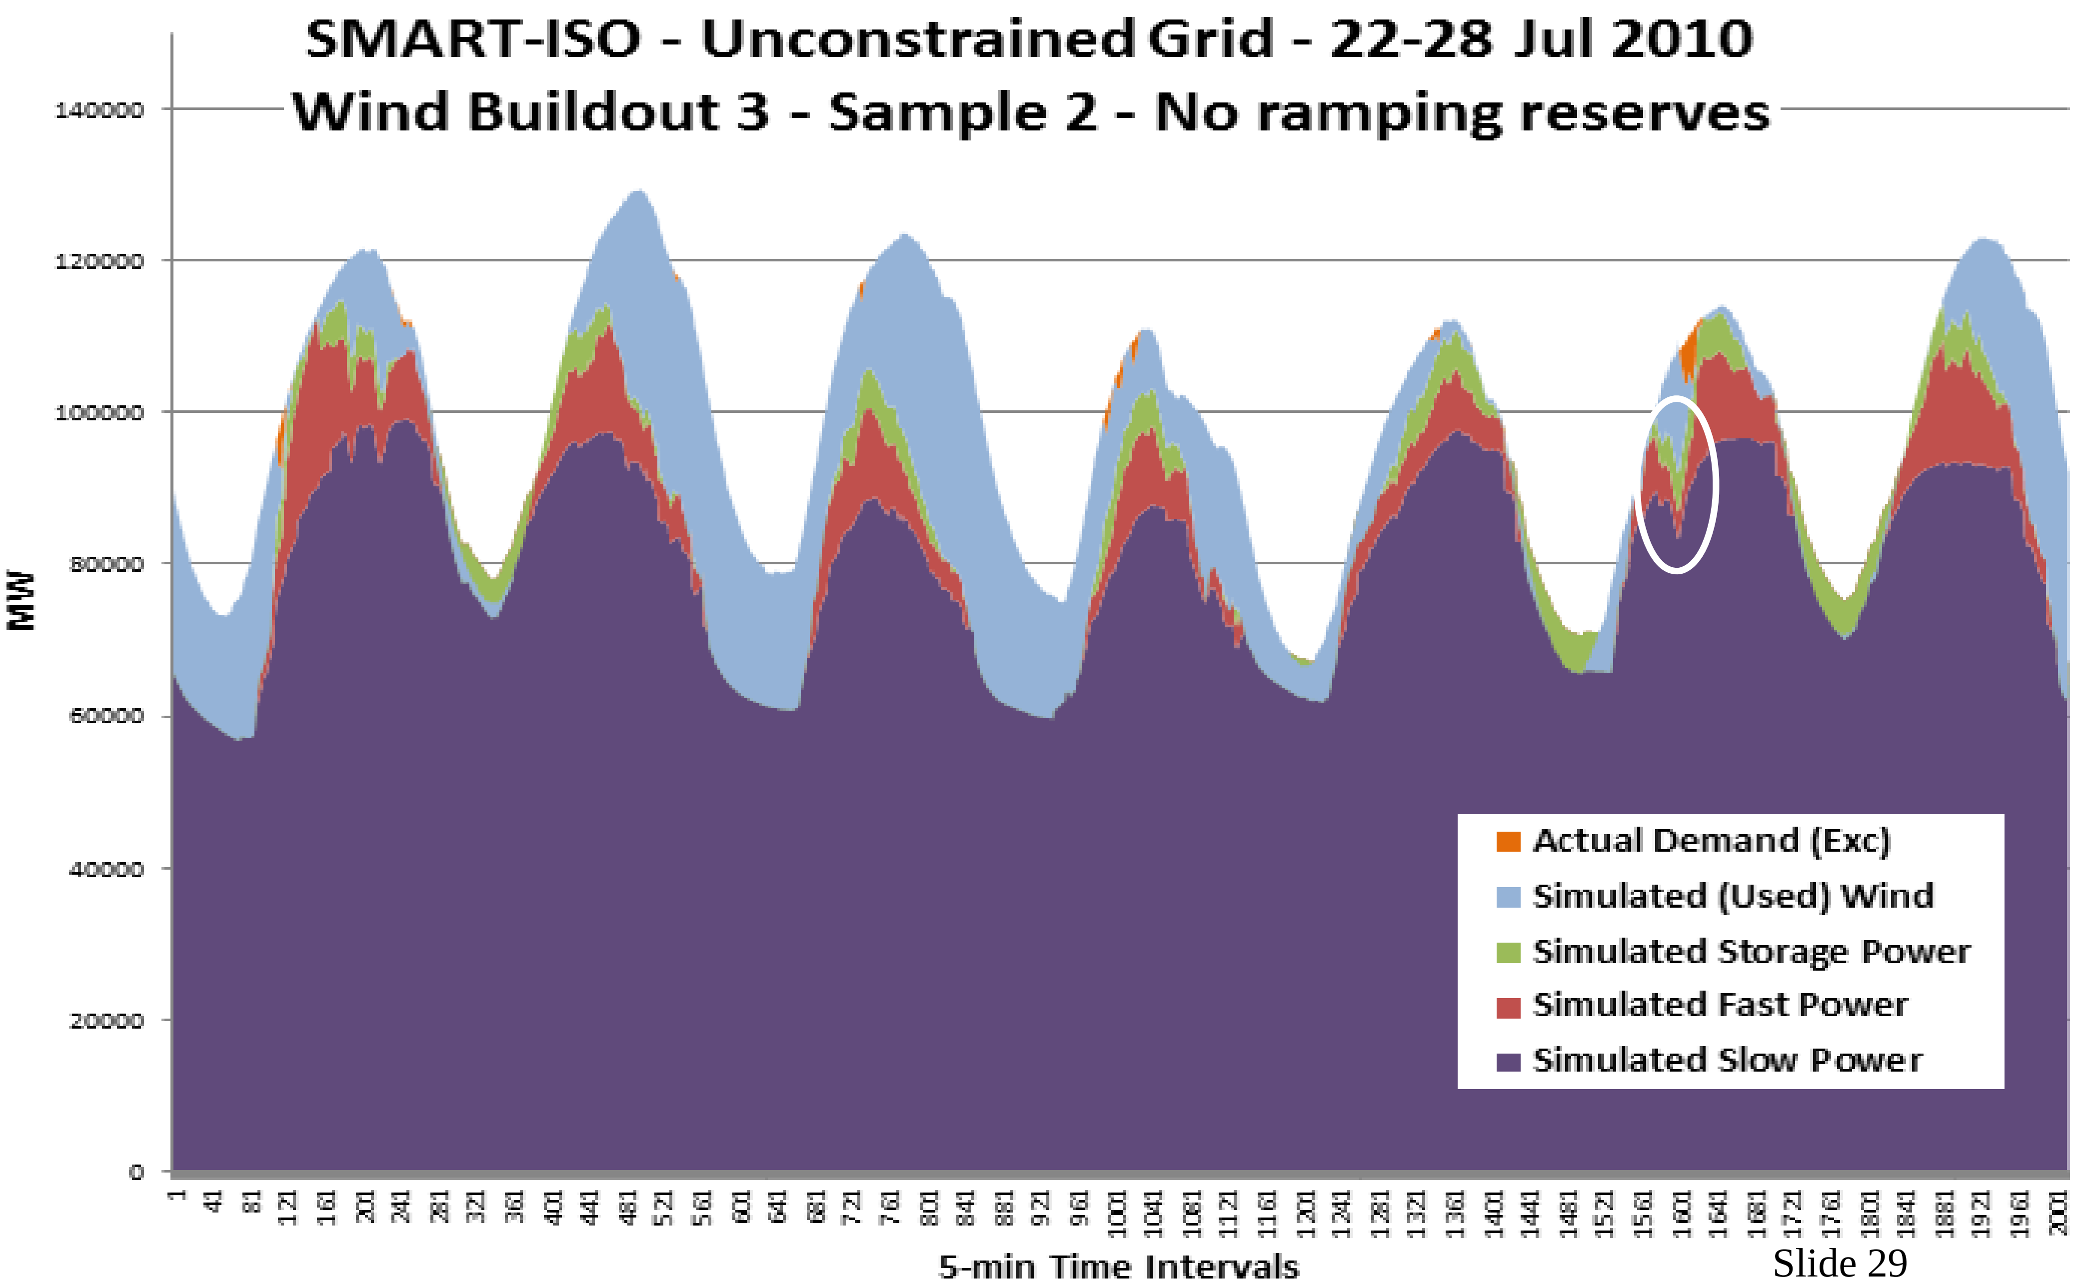

Stochastic lookahead policies

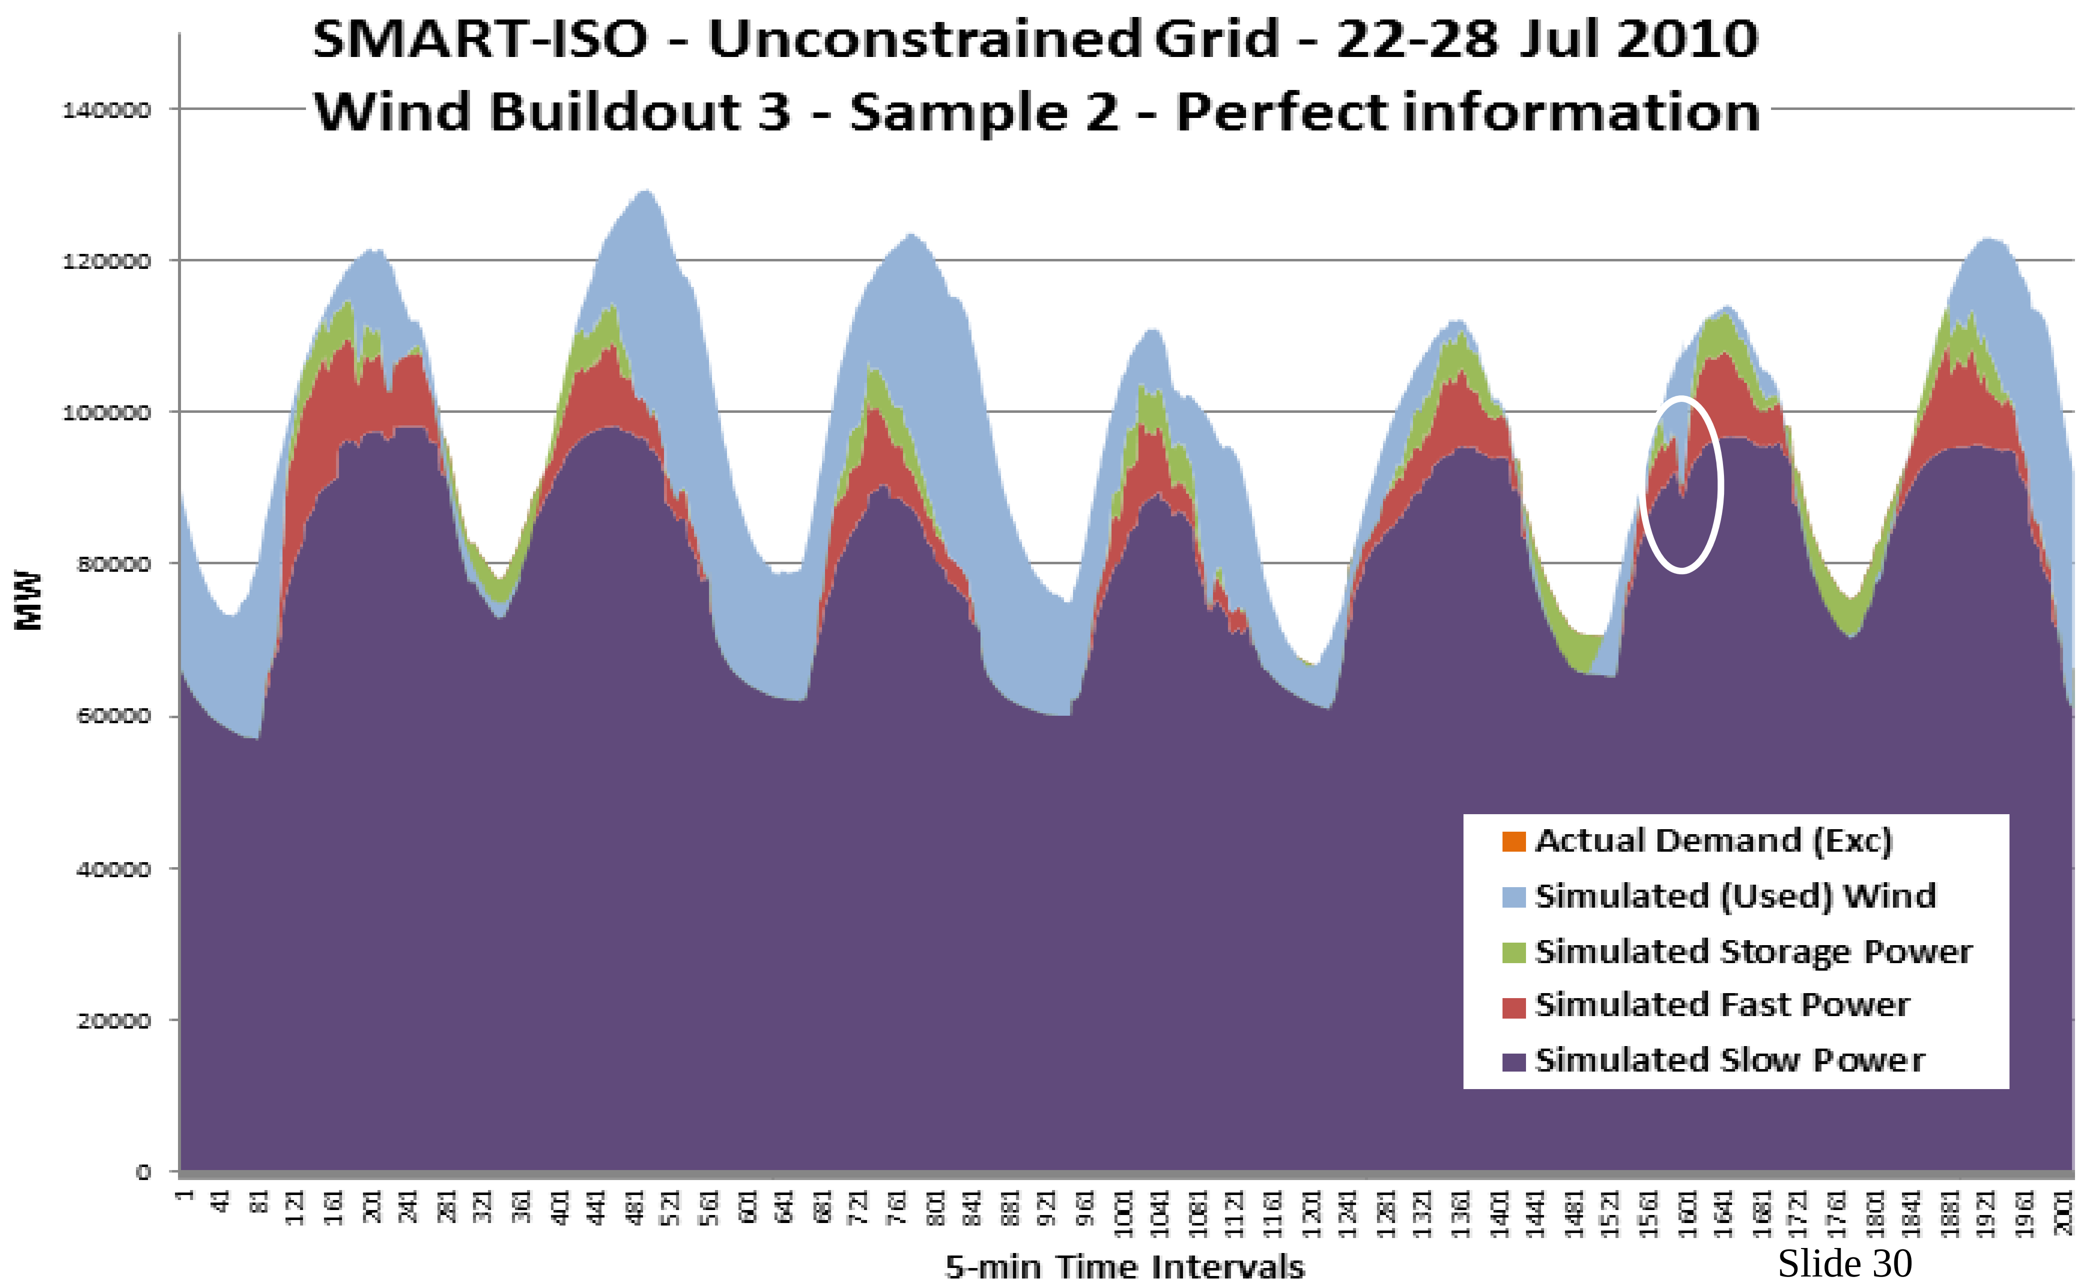

Stochastic lookahead policies

- The two-stage approximation

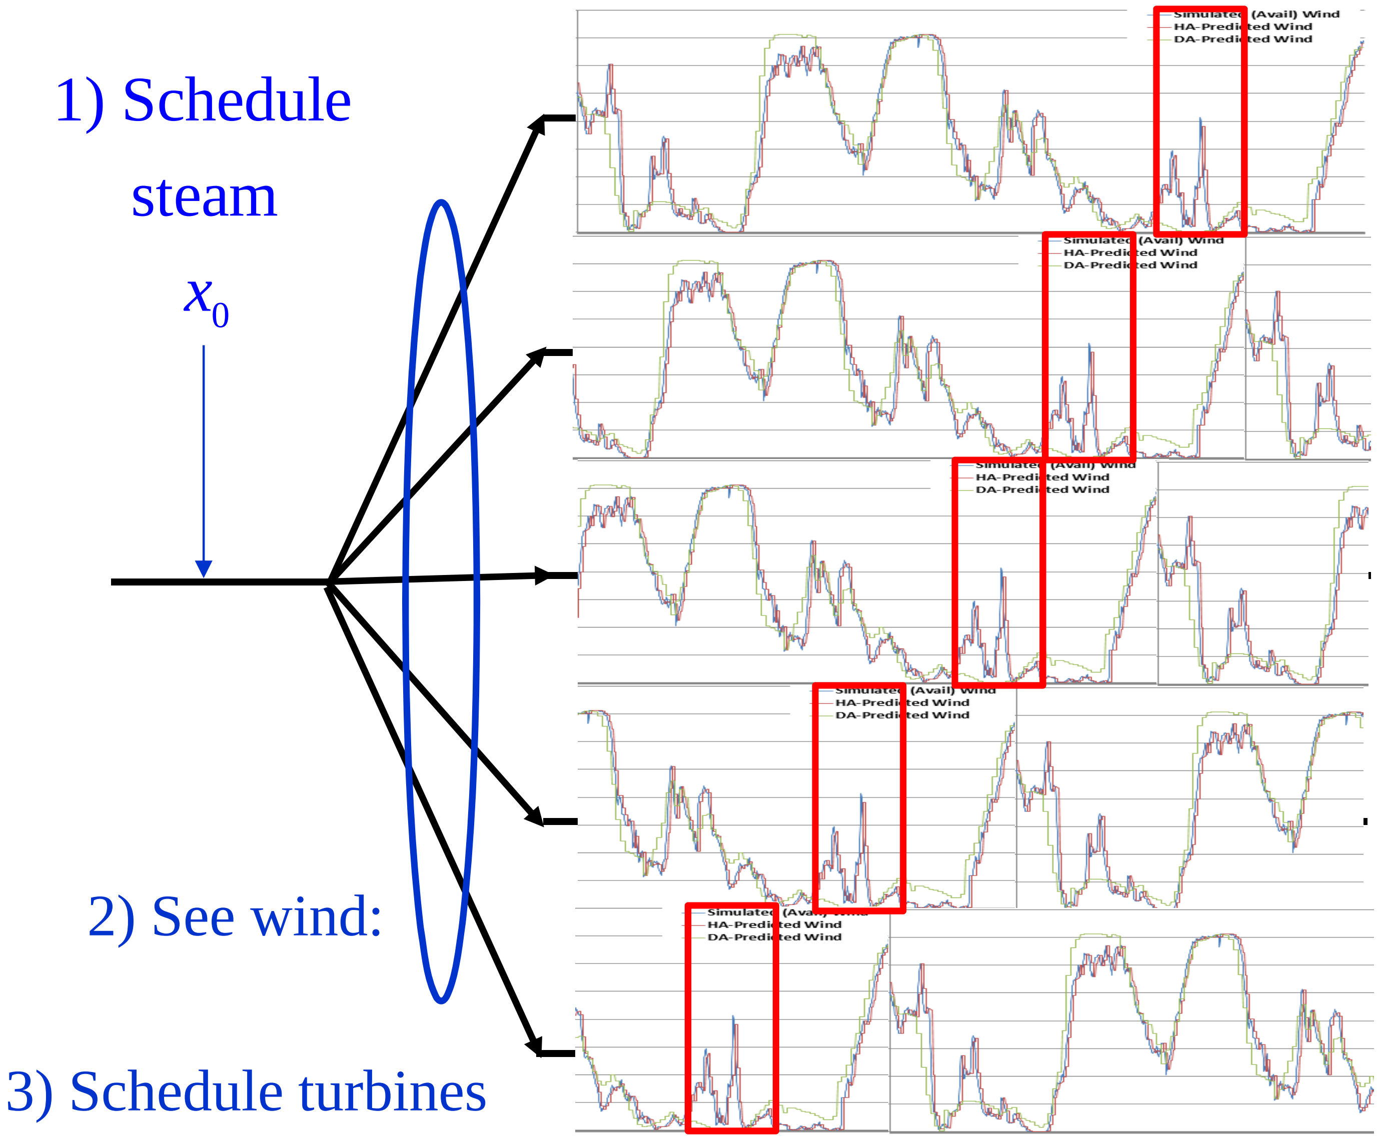

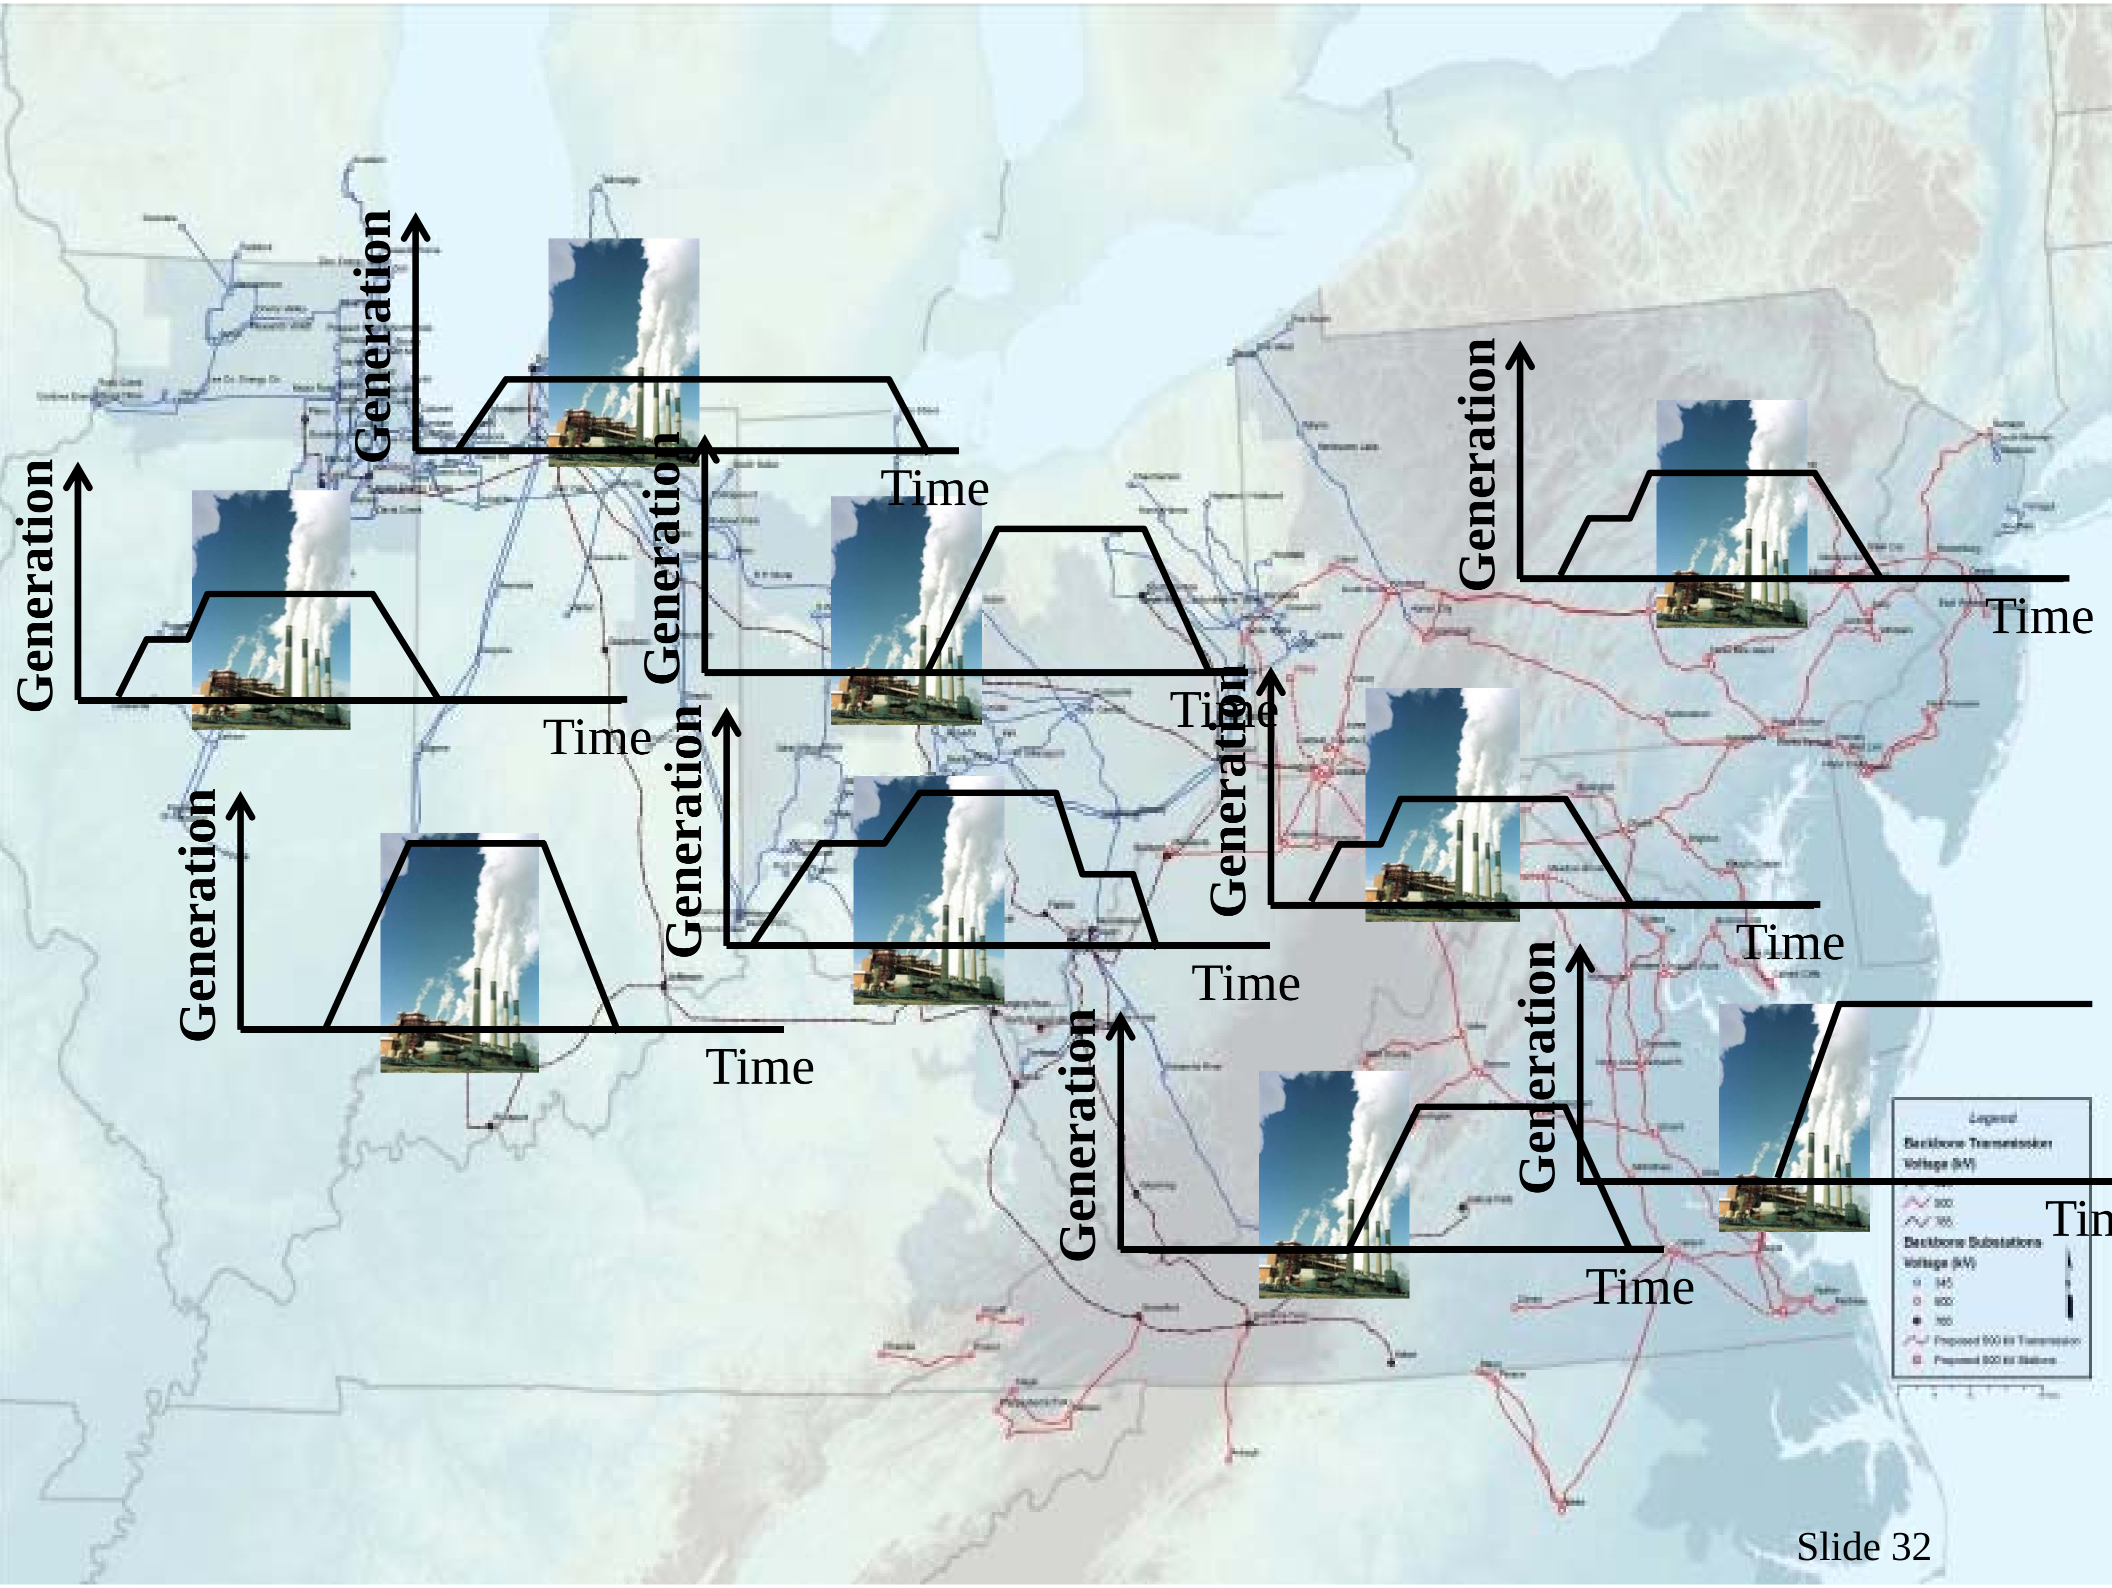
Generation
Time
Generation
Time
Generation
Time
Generation
Time
Generation
Time
Generation
Time
Generation
Time
Generation
Time
Generation
Time
Generation
Tim

Stochastic lookahead policies

- Observations
  - » As a general rule, we do not need a large number of scenarios to produce a robust policy.
  - » The problem here is that the day-ahead decision (steam generation) is a very high-dimensional vector. We need to schedule steam generators:
    - Across the grid (spatially)
    - Over the entire day (temporally)
  - » Scenarios have to create robust behavior across space, and across time.
  - » *It appears for this application, we may need a dramatically larger number of scenarios than is required in applications with lower-dimensional first stage decisions.*

Four classes of policies

1) Policy function approximations (PFAs)

» Lookup tables, rules, parametric functions

2) Cost function approximation (CFAs)

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

3) Policies based on value function approximations (VFAs)

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

4) Lookahead policies

» ***Deterministic lookahead:***

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

» ***Stochastic lookahead (e.g. stochastic trees)***

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

A hybrid lookahead-CFA policy

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

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

Steam generation

Gas turbines

  - » These decisions need to made with different horizons
    - Steam generation is made day-ahead
    - Gas turbines can be planned an hour ahead or less

A hybrid lookahead-CFA policy

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

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

Not defined!

Steam generation

Gas turbines

  - » These decisions need to made with different horizons
    - Steam generation is made day-ahead
    - Gas turbines can be planned an hour ahead or less

Designing a policy

- A robust lookahead-CFA policy
  - We imbed a policy for fast-response adjustments within a lookahead model for planning steam:

$$\min_{\substack{(x_{tt'})_{t'=1,\dots,24} \\ \pi}} \mathbb{E} \sum_{t'=t}^{t+48} C(x_{tt'}, Y^{\pi}(S_{tt'}))$$

Steam generation

- $x_{t,t'}$ is determined at time *t*, to be implemented at time *t'*
- $y_{t',t'}$ is determined at time *t'* by the policy $Y^{\pi}(S_{tt'})$

  - The challenge now is to adaptively estimate the ramping constraints $\theta = \left(\theta^{up}, \theta^{down}\right)$ and the policies $Y^{\pi}(S_{tt'})$.

Designing a policy

- A robust lookahead-CFA policy
  - We imbed a policy for fast-response adjustments within a lookahead model for planning steam:

$$\min_{\substack{(x_{tt'})_{t'=1,\dots,24} \\ \pi}} \mathbb{E} \sum_{t'=t}^{t+48} C\left(x_{tt'}, Y^{\pi}(S_{tt'})\right)$$

Real-time adjustment of hydro and thermal plants.
π is the type of policy.

- $x_{t,t'}$ is determined at time *t*, to be implemented at time *t'*
- $y_{t',t'}$ is determined at time *t'* by the policy $Y^{\pi}(S_{tt'})$

  - The challenge now is to adaptively estimate the ramping constraints $\theta = \left(\theta^{up}, \theta^{down}\right)$ and the policies $Y^{\pi}(S_{tt'})$.

Designing a policy

- A robust lookahead-CFA policy
  - » We imbed a policy for fast-response adjustments within a lookahead model for planning steam:

$$F_t(S_t|\theta) = \min_{\substack{(x_{tt'})_{t'=1,\ldots,24} \\ \pi}} \mathbb{E} \sum_{t'=t}^{t+48} C(x_{tt'}, Y^\pi(S_{tt'}))$$

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

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

    - $x_{t,t'}$ is determined at time *t*, to be implemented at time *t'*
    - $y_{t',t'}$ is determined at time *t'* by the policy $Y^\pi(S_{tt'})$

  - » The challenge now is to adaptively estimate the ramping constraints $\theta = \left(\theta^{up}, \theta^{down}\right)$ and the policies $Y^\pi(S_{tt'})$.

SMART-ISO "Regions"

Designing a policy

- We then have to tune the parameters of this policy in our *stochastic base model*.

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

Stochastic base model

$\pi = \theta = \left(\theta^{up}, \theta^{down}\right)$

Robust CFA policy

» The challenge now is to adaptively estimate the ramping constraints $\theta = \left(\theta^{up}, \theta^{down}\right)$.

SMART-ISO: Offshore wind study

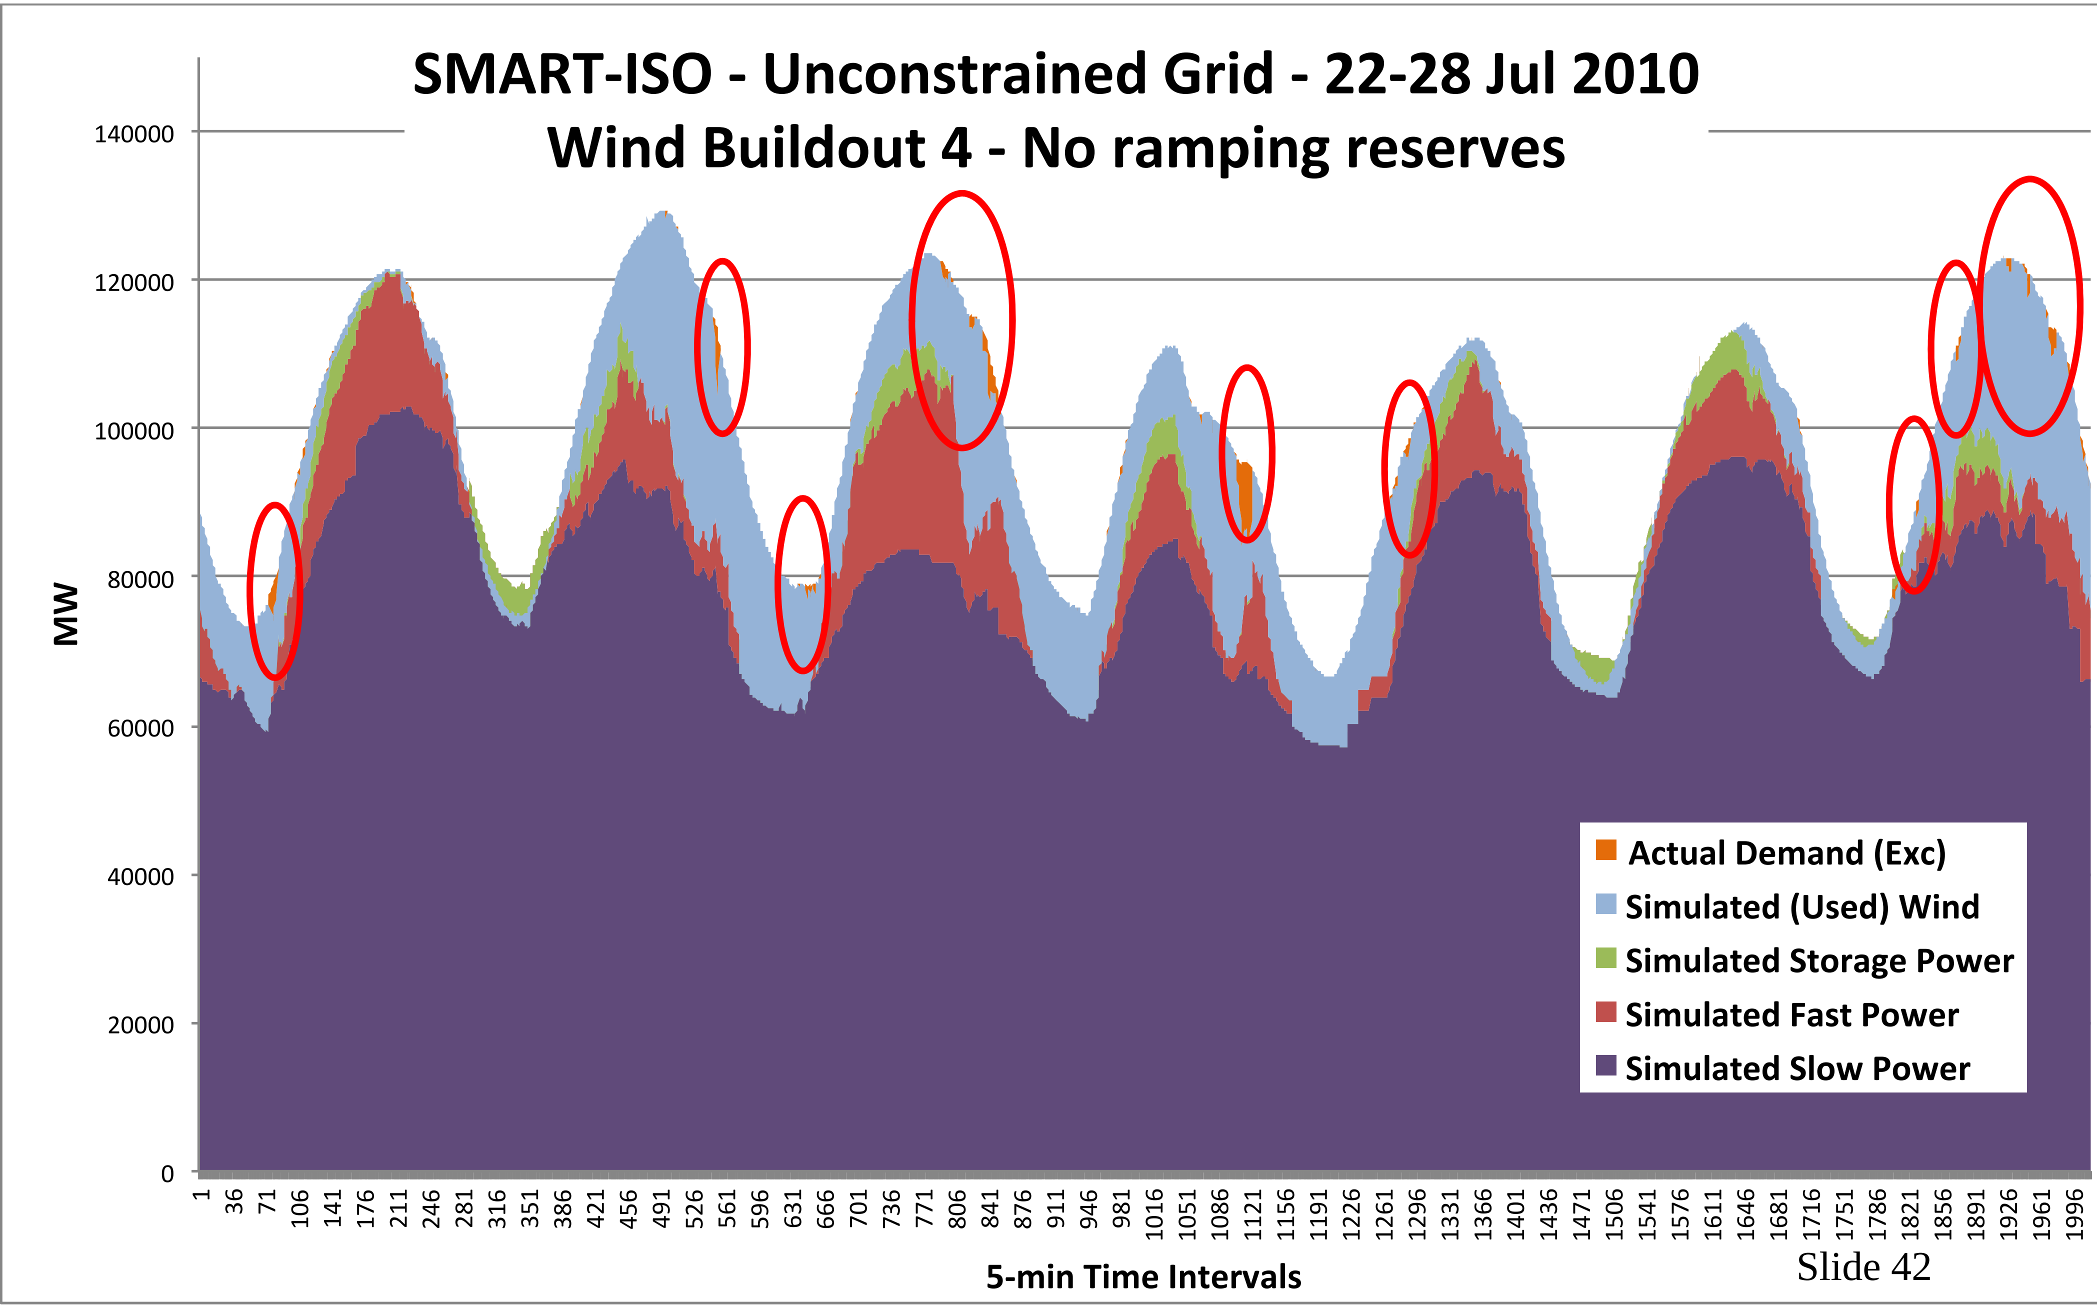

SMART-ISO: Offshore wind study

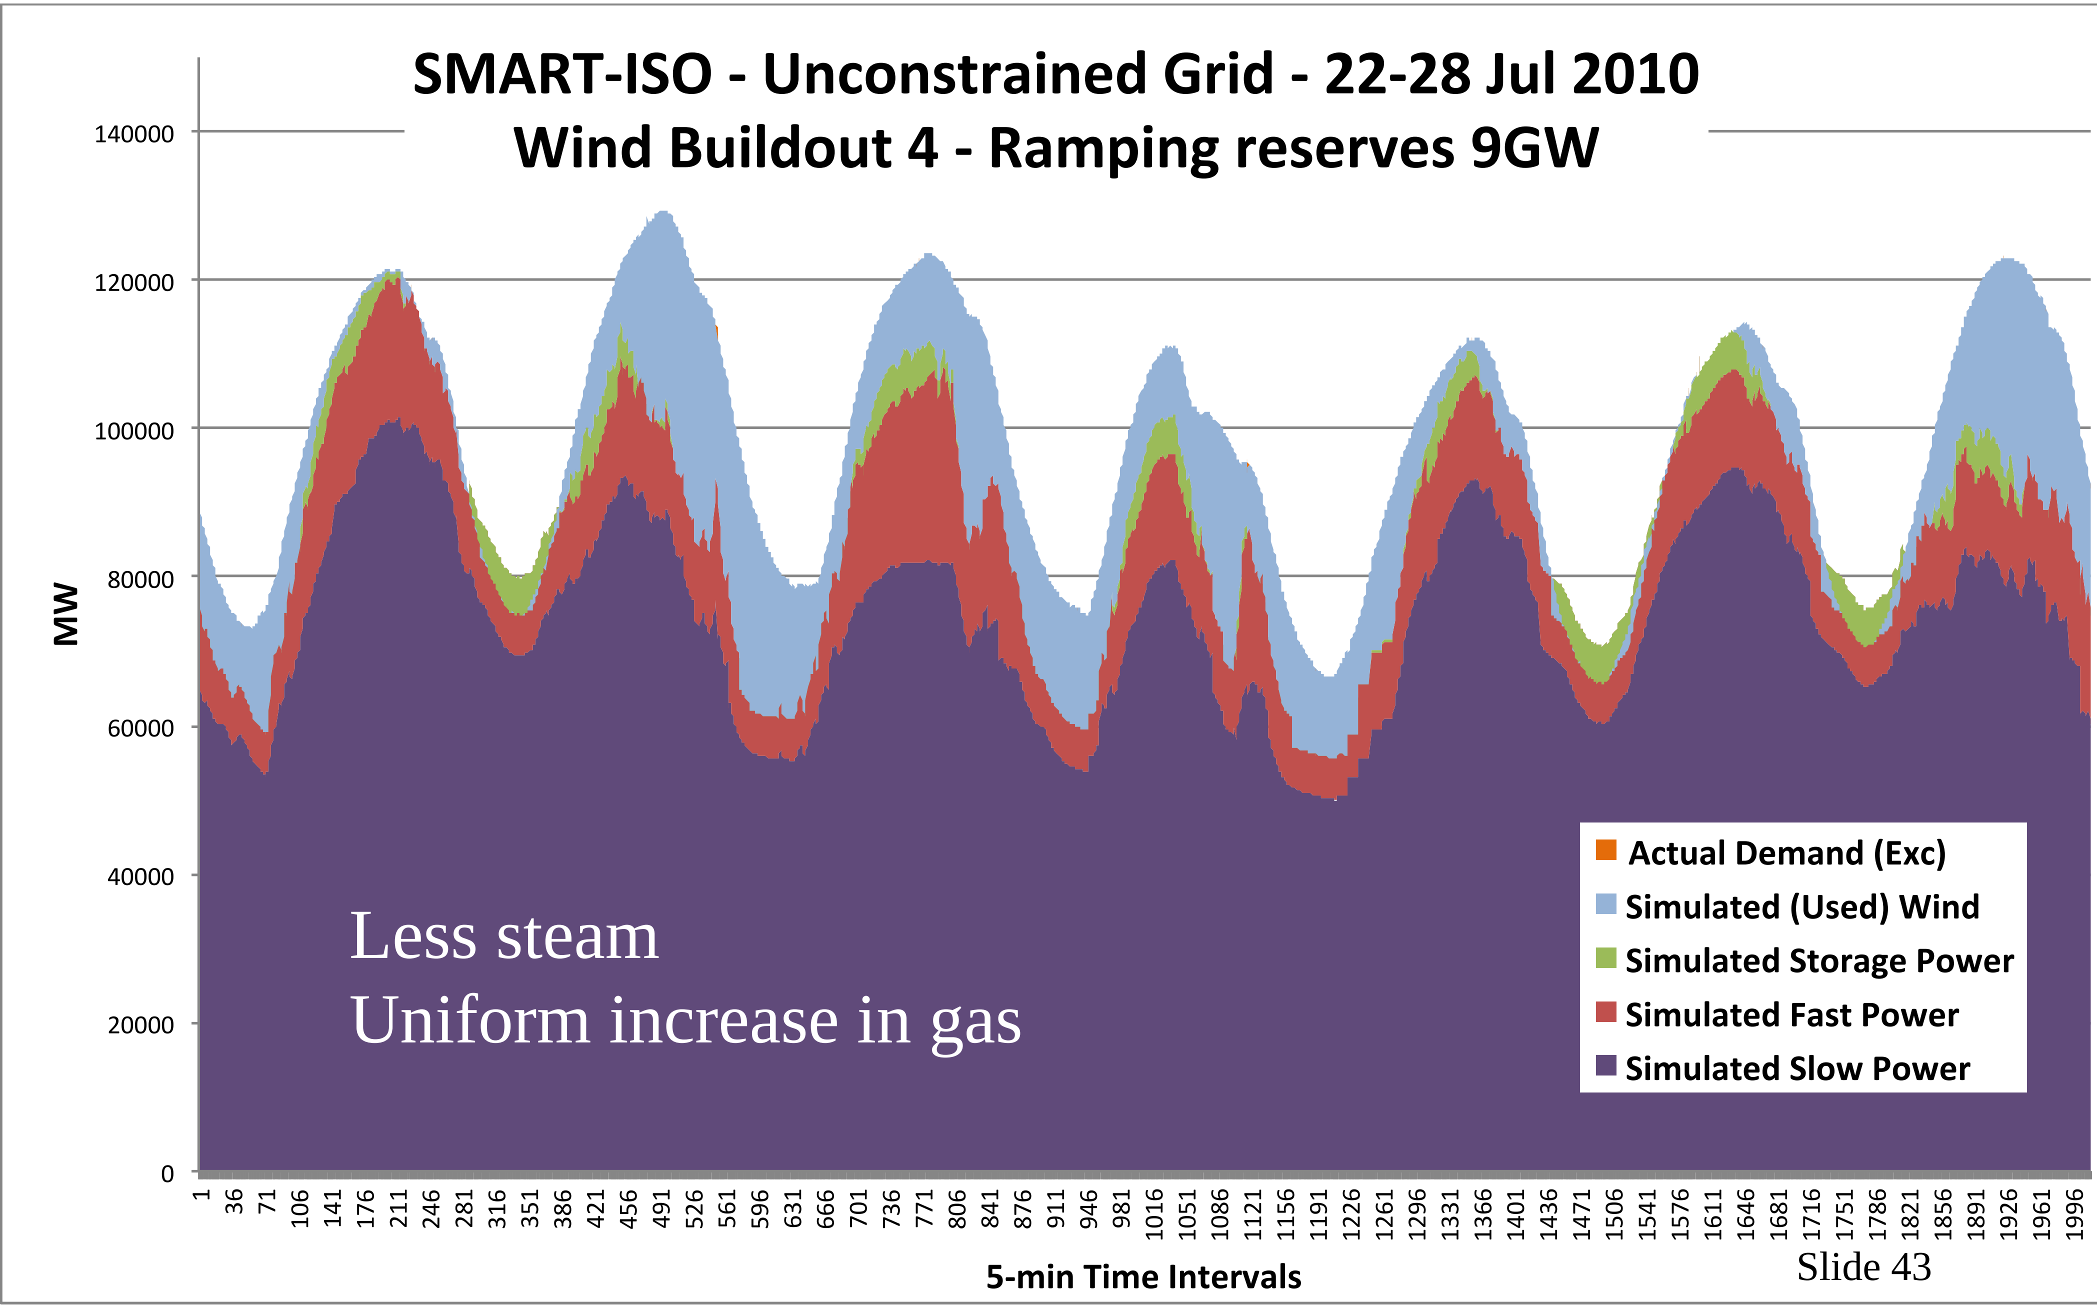

Lecture outline

- Perspectives on robust policies

Comparison of policies

- Robust cost function approximation
  - Requires solving modified deterministic lookahead model
  - Parameters have to be tuned, ideally with a highly realistic *base model* ("simulator")
  - User captures domain knowledge when specifying the structure of the CFA

- Stochastic lookahead
  - Requires solving approximate stochastic lookahead model
  - Tuning is generally not done (although in theory possible, but it is very expensive)
  - No need for user-specified parametric approximation (and no ability to incorporate domain knowledge)

Approximating distributions

- Stochastic lookahead model
  - » Uses approximation of the *information process* in the lookahead model

Parametric distribution (pdf)

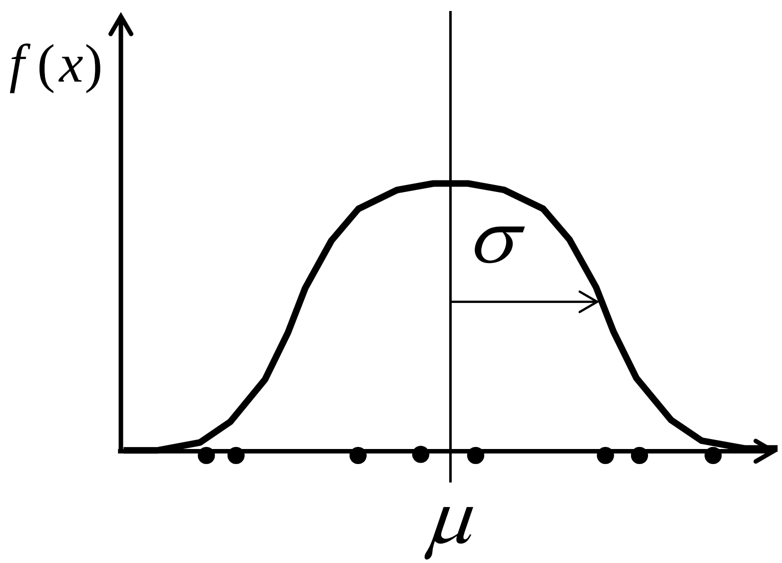

$$f(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\left(\frac{(x-\mu)^2}{\sigma^2}\right)}$$

Nonparametric distribution (cdf)

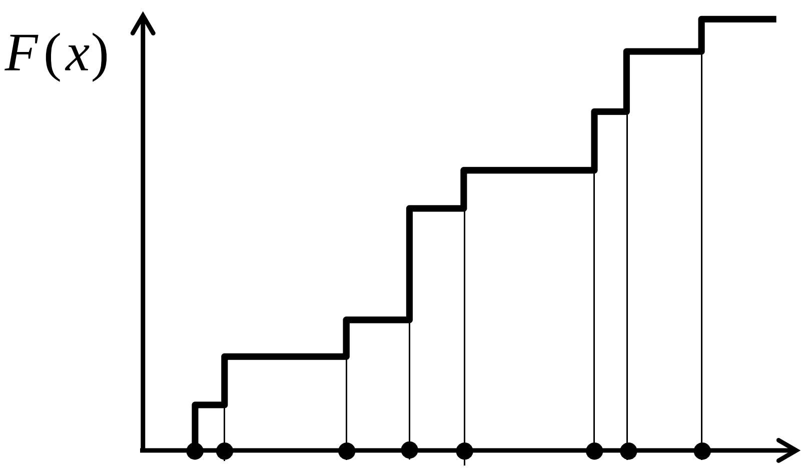

Approximating a function

- Parametric vs. nonparametric

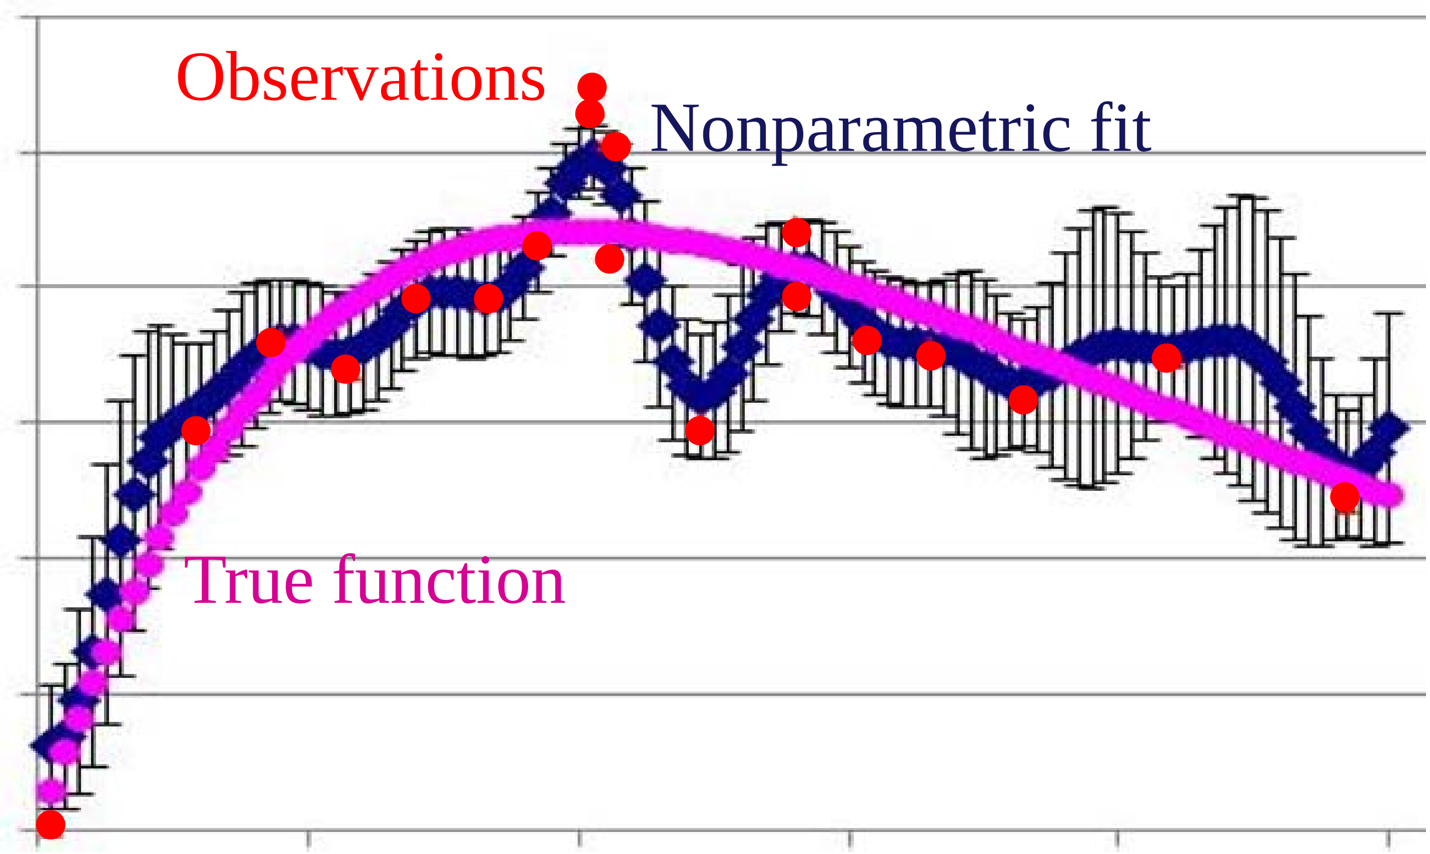

» We can use our understanding of the function to impose a shape.

Approximating a function

- Parametric vs. nonparametric

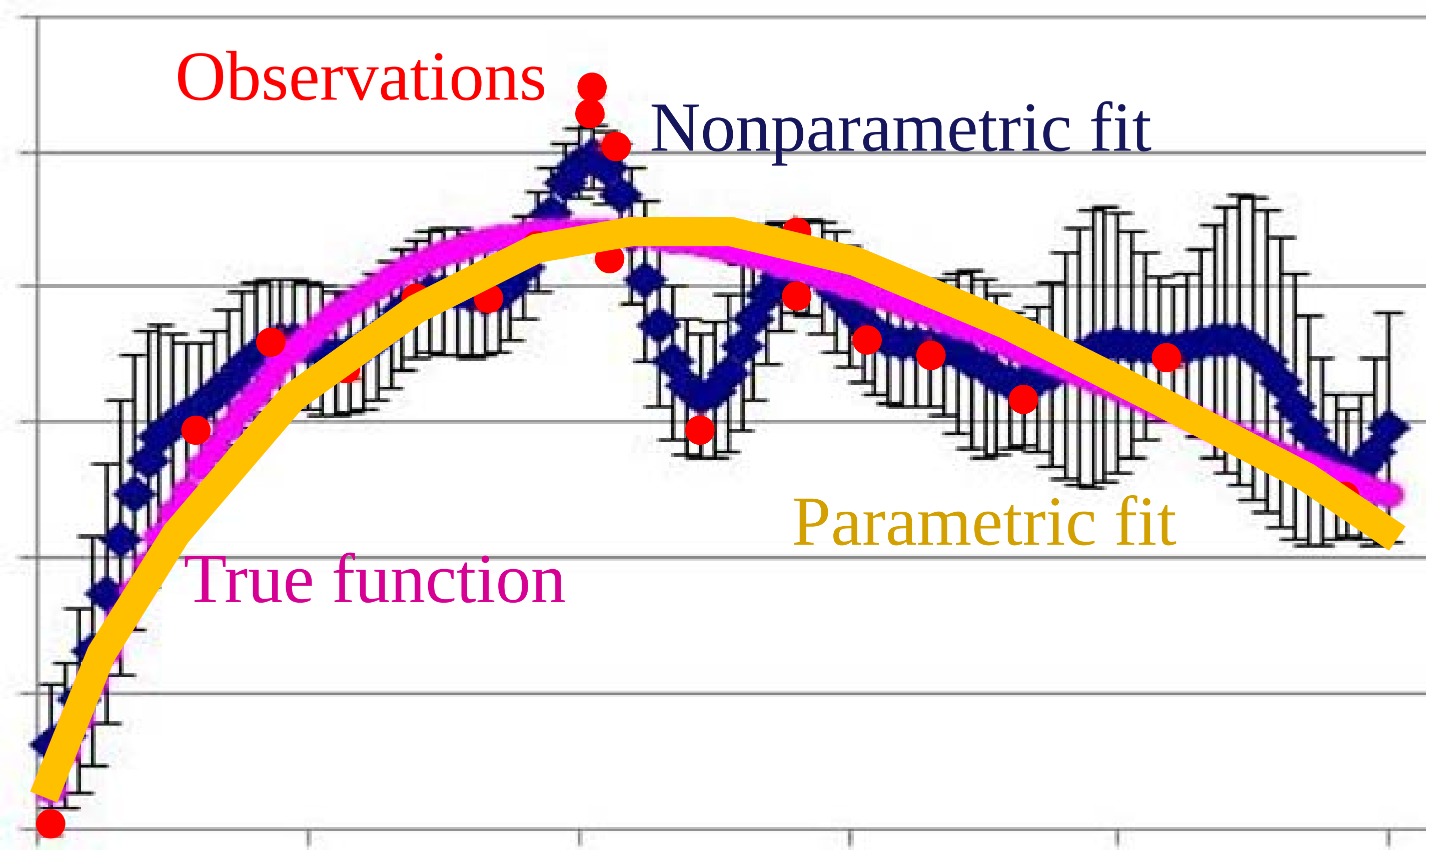

» We can use our understanding of the function to impose a shape.

Designing a policy

- A robust lookahead-CFA policy
  - We imbed a policy for fast-response adjustments within a lookahead model for planning steam:

$$F_t(S_t|\theta) = \min_{\substack{(x_{tt'})_{t'=1,\dots,24} \\ \pi}} \mathbb{E} \sum_{t'=t}^{t+48} C(x_{tt'}, Y^{\pi}(S_{tt'}))$$

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

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

  - It is easy to tune this policy when there are only two parameters $\theta = \left(\theta^{up}, \theta^{down}\right)$
  - But perhaps the ramping reserves should depend on other information?

Designing a policy

- A robust lookahead-CFA policy
  - We imbed a policy for fast-response adjustments within a lookahead model for planning steam:

$$F_t(S_t|\theta) = \min_{\substack{(x_{tt'})_{t'=1,\dots,24} \\ \pi}} \mathbb{E} \sum_{t'=t}^{t+48} C(x_{tt'}, Y^\pi(S_{tt'}))$$

$$x_{t,t'}^{\max} - x_{t,t'} \geq \theta^{up}(S_t) L_{tt'} \qquad \text{Up-ramping reserve}$$

$$x_{t,t'} - x_{t,t'}^{\max} \geq \theta^{down}(S_t) L_{tt'} \qquad \text{Down-ramping reserve}$$

  - The parameters might depend on a state variable that captures
    - Weather forecast (esp. change in weather)
    - Load forecast (indicates how close to capacity)
  - Now the ramping parameters are *functions*. ☹

Designing a policy

- How do we compare policies?
  - » Let's agree on a base model (a simulator) and compare different policies!

$$\bar{F}^{\pi} \approx \sum_{n=1}^{N} p(\omega^n) \sum_{t=0}^{T} \gamma^t C\Big(S_t(\omega^n), X^{\pi}(S_t(\omega^n))\Big)$$

    where ω represents a sample path capturing season, type of meteorology (stormy, calm), and stochastic variations around this base.

  - » Now compare
    - $\bar{F}^{Robust-CFA}(\theta^{Robust-CFA})$
    - $\bar{F}^{Stoch-LA}(\theta^{Stoch-LA})$

Designing a policy

- So we have a choice:
  - Stochastic lookahead model?
  - Robust cost function approximation tuned on a *stochastic base model* (the "simulator")?
- Stochastic lookahead model:
  - Lookahead model introduces several approximations:
    - Finite horizon, two-stages, limited scenarios, coarser temporal decomposition, unable to handle adaptive hedging
    - Computationally very hard to solve
    - Produces stochastic LMPs
  - Robust cost function approximation
    - Tuned using very realistic *base model* using parametric cost function approximation which captures domain knowledge.
    - Computationally comparable to current models
    - LMPs computed using existing practices.