Stochastic Models for Energy Resource Planning:
Sorting through the jungle of stochastic optimization

FERC Conference on Market Efficiency

June 28 2011

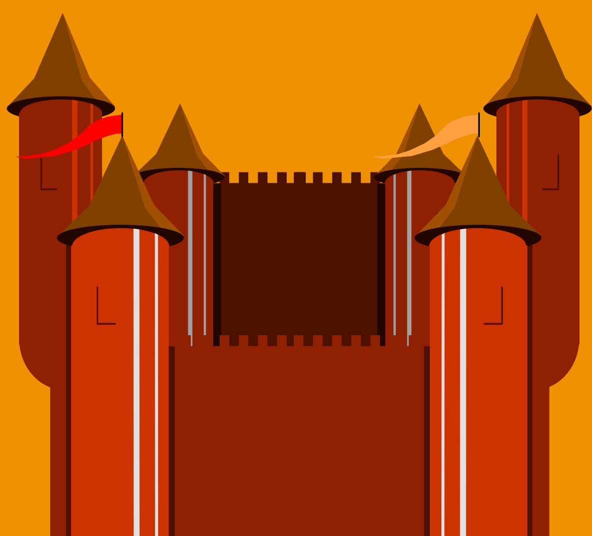

Warren B. Powell
PENSA Laboratory
Princeton University
http://www.castlelab.princeton.edu

The challenges of uncertainty

- The stochastic unit commitment problem
  - » Which power generating units should be committed when planning a day in advance?
  - » How should this problem be solved in the presence of significant levels of energy from renewables?
- Energy storage
  - » How do we manage energy storage in a real-time market to manage daily cycles?
  - » How do we plan regulation capacity in the day-ahead market?
- Balancing energy from renewables
  - » Variations from wind and solar can be regulated with a mixture of nuclear, natural gas (CC and CT), storage and demand response. But what is that mixture?
- Load curtailment and demand response
  - » When and where do we may requests to curtail loads on the system?
  - » How do we manage price signals?


Gray Buildings © 2008 Sanborn
Image © 2011 DigitalGlobe
©2010 Google
Imagery Dates: Oct 1, 2006 - Jun 18, 2010
40°45'20.50" N 73°59'08.54" W elev 567 ft
Eye alt 2007 ft

Intermittent energy sources

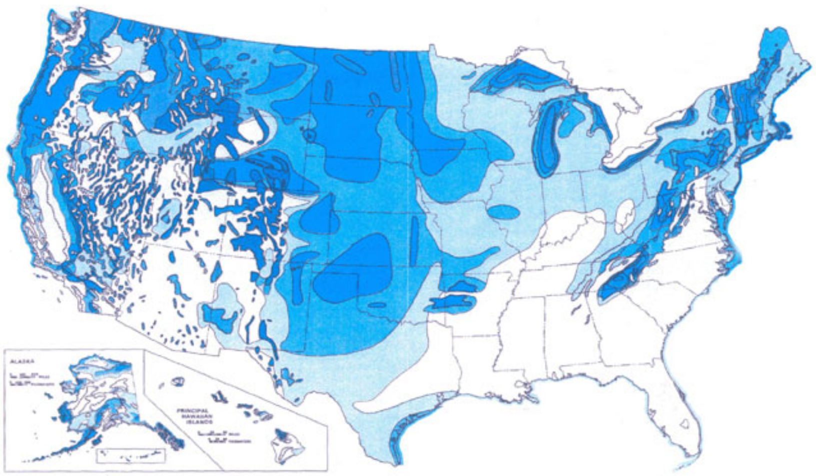

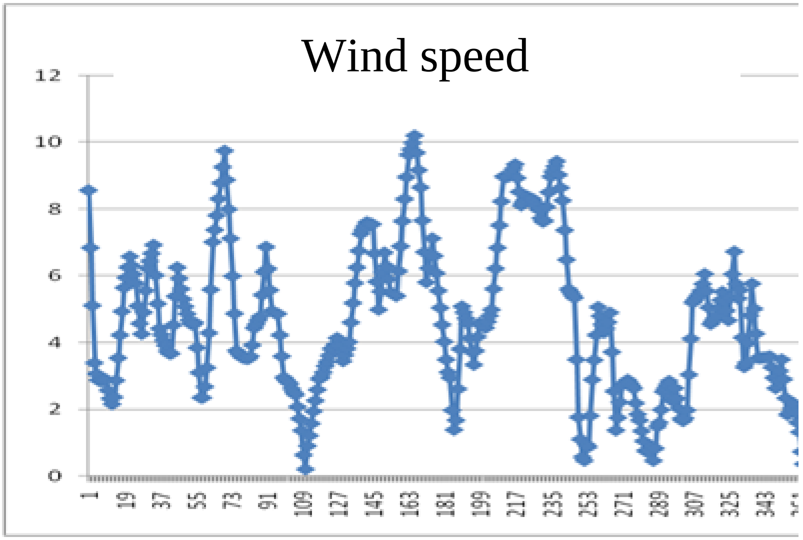

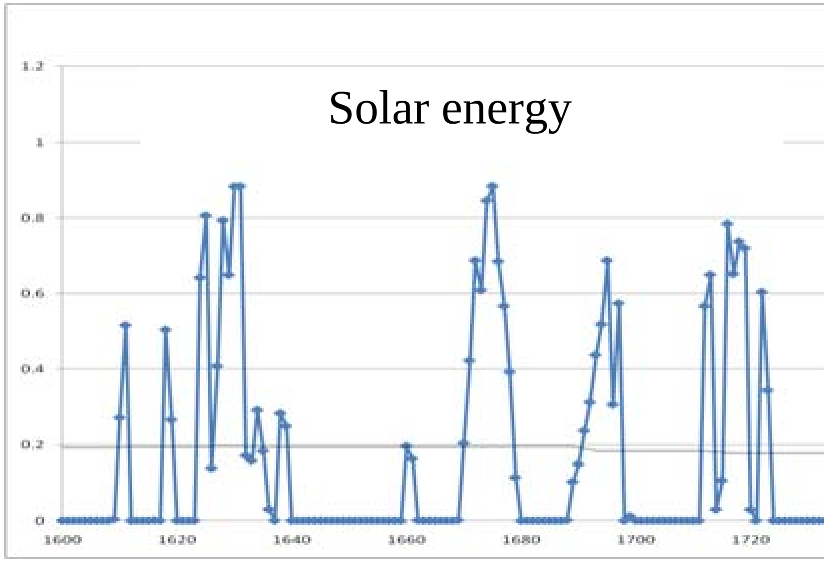

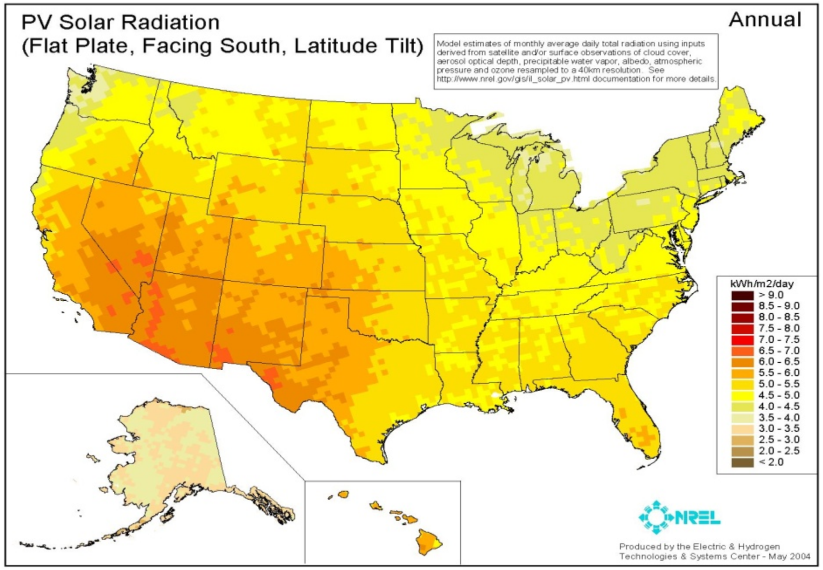

Energy storage

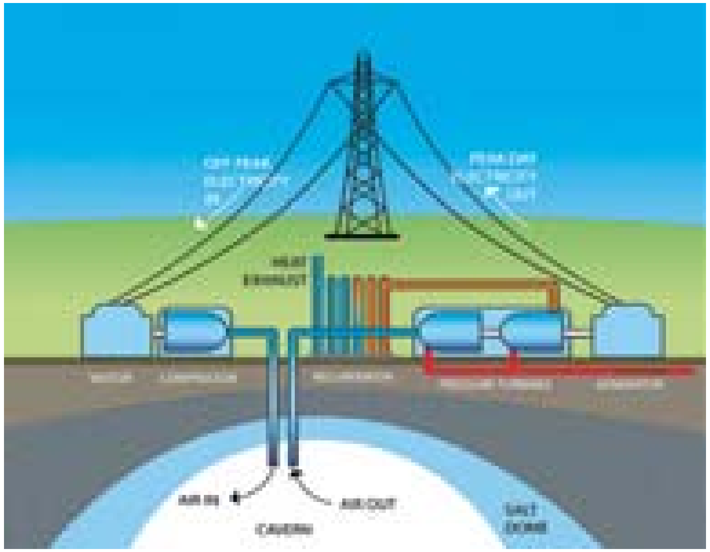

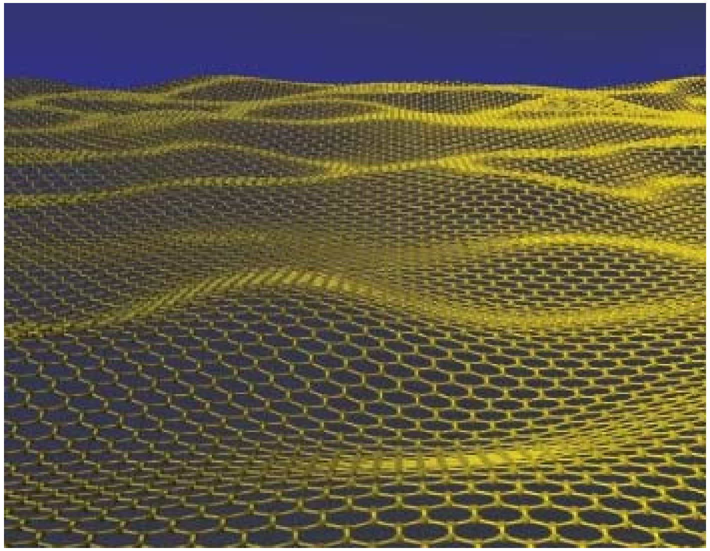

Energy storage portfolios

- Designing a dynamic storage control policy for portfolios of storage devices.

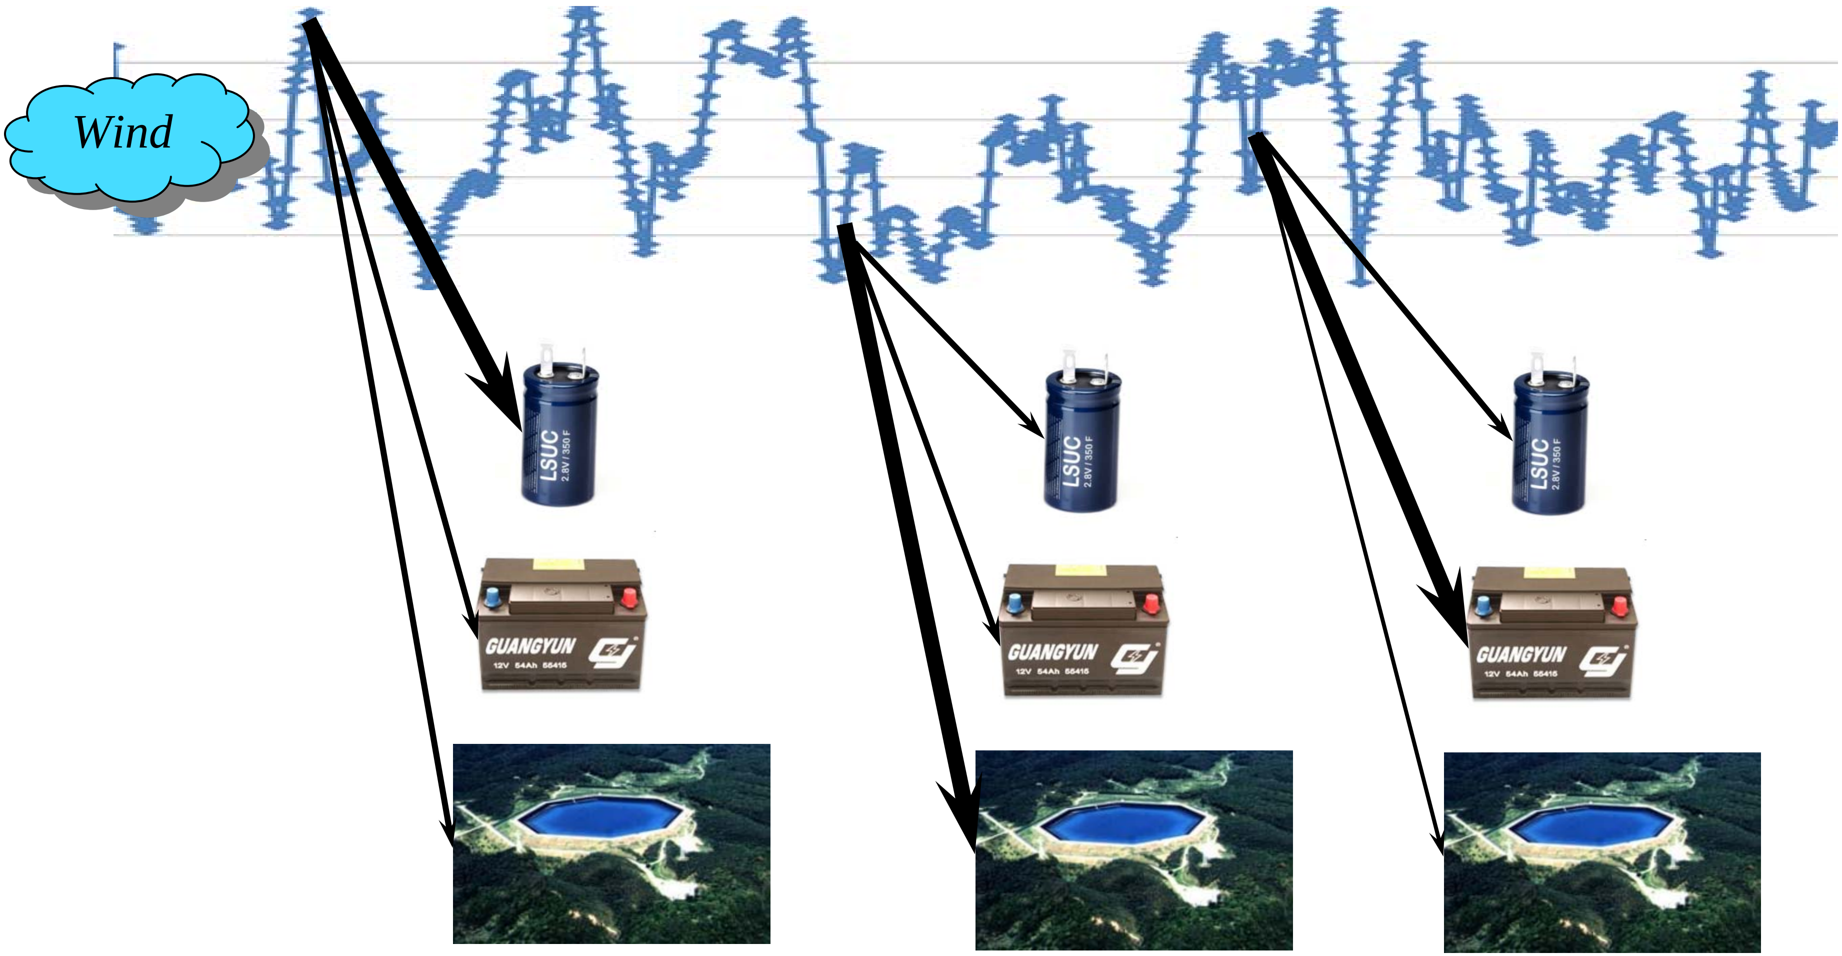

Electricity spot prices

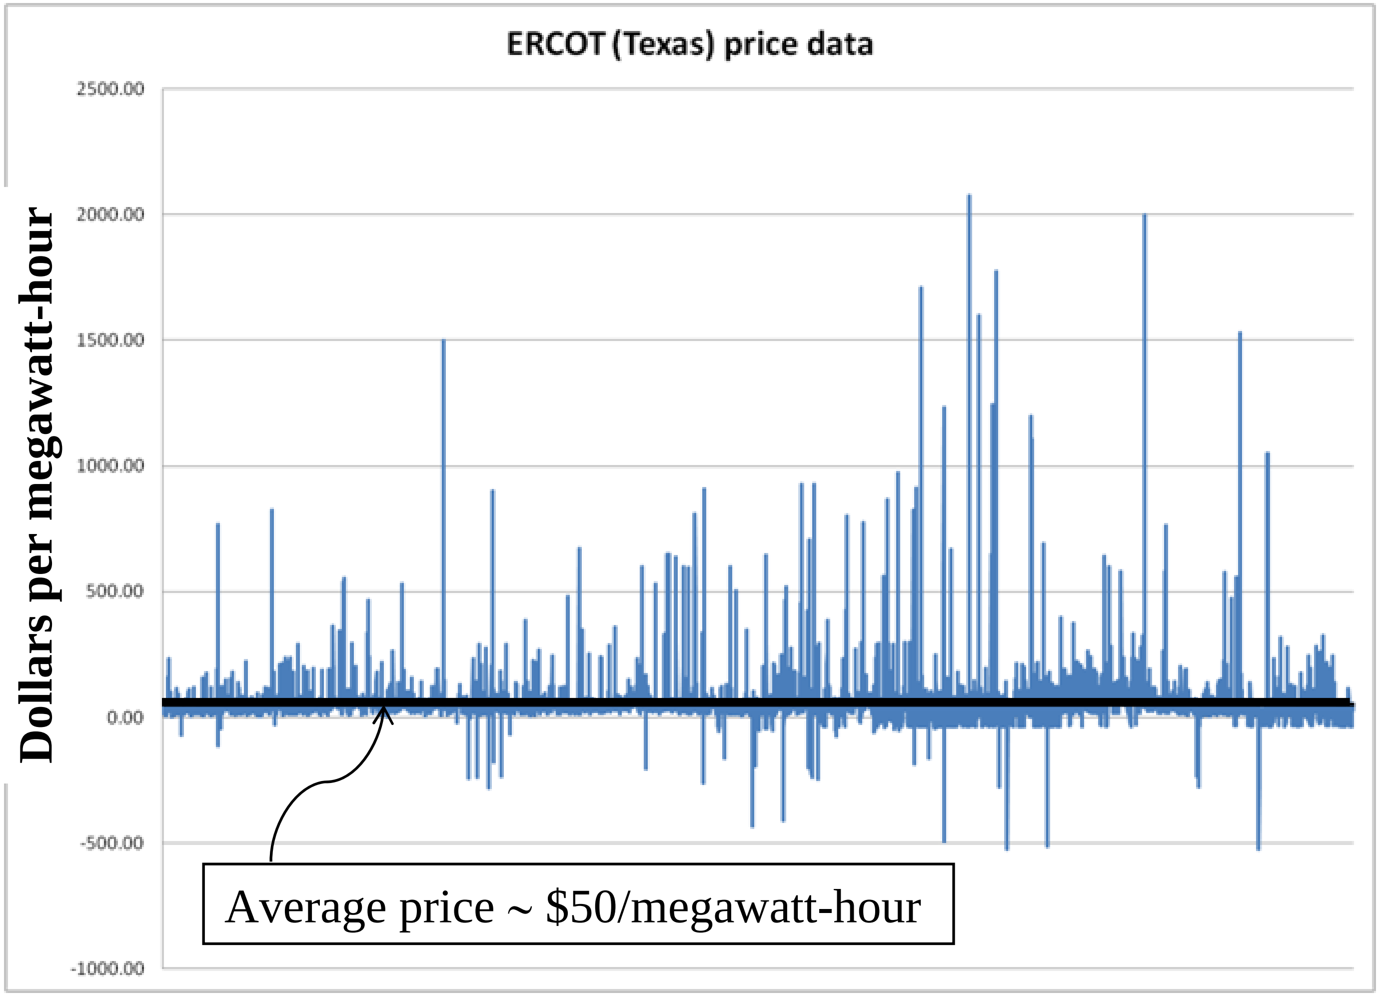

Energy resource modeling

- Need to plan long term energy investments…

Tax policy

Price of oil

Batteries

2010 2015 2020 2025 2030

Solar panels

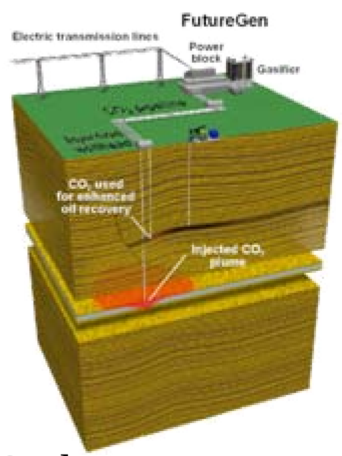

Carbon capture and sequestration

Climate change

Stochastic resource allocation

- The objective function

$$\max_{\pi} E^{\pi}\left\{\sum_{t}\gamma^{t} C\left(S_t, X^{\pi}(S_t)\right)\right\}$$

Expectation over all random outcomes

Contribution 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 have to find the best policy, which is a function that maps states to feasible actions, using only the information available when the decision is made.


Stochastic programming
Stochastic search
Model predictive control
Optimal control
Reinforcement learning
Q – learning
On-policy learning
Off-policy learning
Markov decision processes
Simulation optimization
Policy search

What is a policy?

- Policies come in four fundamental flavors:
  - » 1) Myopic policies
    - Take the action that maximizes contribution (or minimizes cost) for just the current time period:

$$X^M(S_t) = \arg\max_{x_t} C(S_t, x_t)$$

    - We can parameterize myopic policies with bonus and penalties to encourage good long-term behavior.

What is a policy?

- Policies come in four fundamental flavors:
  - » 2) Lookahead policies - Plan over the next T periods, but implement only the action it tells you to do now.
    - Deterministic forecast

$$X^M(S_t) = \arg\max_{x_t, x_{t+1}, \ldots, x_{t+T}} C(S_t, x_t) + \sum_{t'=t+1}^{T} C(S_{t'}, x_{t'})$$

    - Stochastic programming

$$X^M(S_t) = \arg\max_{x_t, \left(x_{t+1}, \ldots, x_{t+T}\right)(\omega)} C(S_t, x_t) + \sum_{\omega \in \Omega} p(\omega) \sum_{t'=t+1}^{T} C(S_{t'}(\omega), x_{t'}(\omega))$$

    - Rolling horizon procedures
    - Model predictive control
    - Rollout heuristics
    - Tree search

What is a policy?

- Policies come in four fundamental flavors:
  - » 3) Policy function approximations
    - Tabular
      - When in this state, take this action.
    - Regression models

$$X^M(S_t \mid \theta) = \theta_0 + \theta_1 S_t + \theta_2 \left(S_t\right)^2$$

    - Neural networks

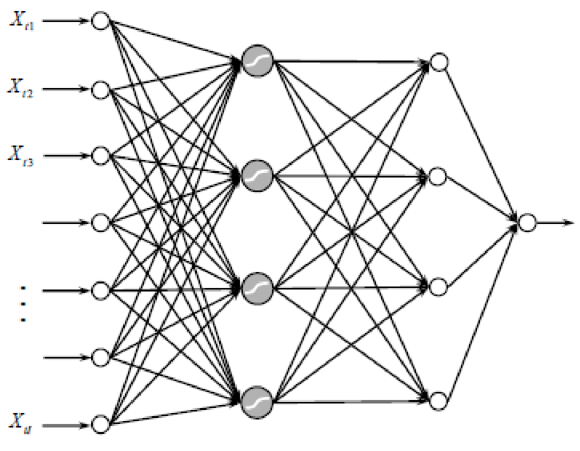

What is a policy?

- Policies come in four fundamental flavors:
  - 4) Policies based on value function approximations
    - Using the pre-decision state

$$X^M(S_t) = \arg\max_{x_t} \left(C(S_t, x_t) + \gamma E \bar{V}_{t+1}(S_{t+1}) \right)$$

    - Or the post-decision state:

$$X^M(S_t) = \arg\max_{x_t} \left(C(S_t, x_t) + \gamma \bar{V}_{t+1}\left(S_t^x(S_t, x_t) \right) \right)$$

    - The challenge is finding a good approximation of the value of being in a state.

Approximations

- There are three classes of approximation strategies (for policies and value functions):
  - » Lookup table
    - Given a discrete state, return a discrete action or value
  - » Parametric models
    - Linear models ("basis functions")
    - Nonlinear models
    - Neural networks
  - » Nonparametric models
    - Kernel regression
    - Dirichlet process-based models

Lecture outline

- Optimizing energy storage
- The stochastic unit commitment problem for PJM

Policy optimization

- Optimizing a policy for battery arbitrage

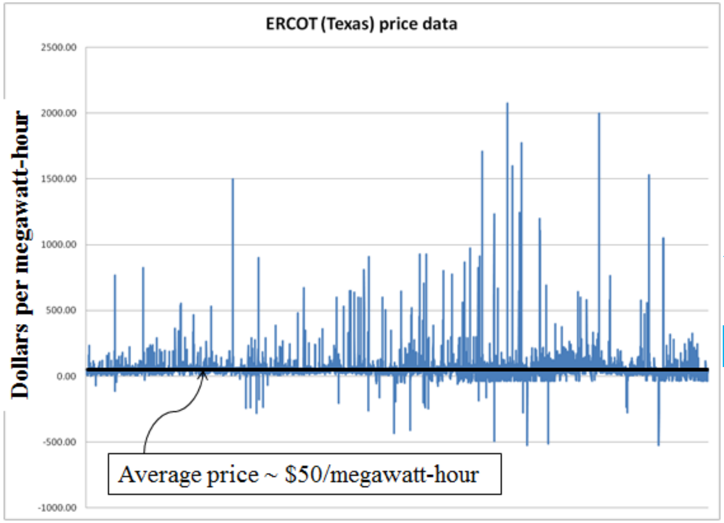

Optimizing storage

- Challenge: find a policy for charging and discharging the battery
  - » Strategy posed by the battery manufacturer: “Buy low, sell high”

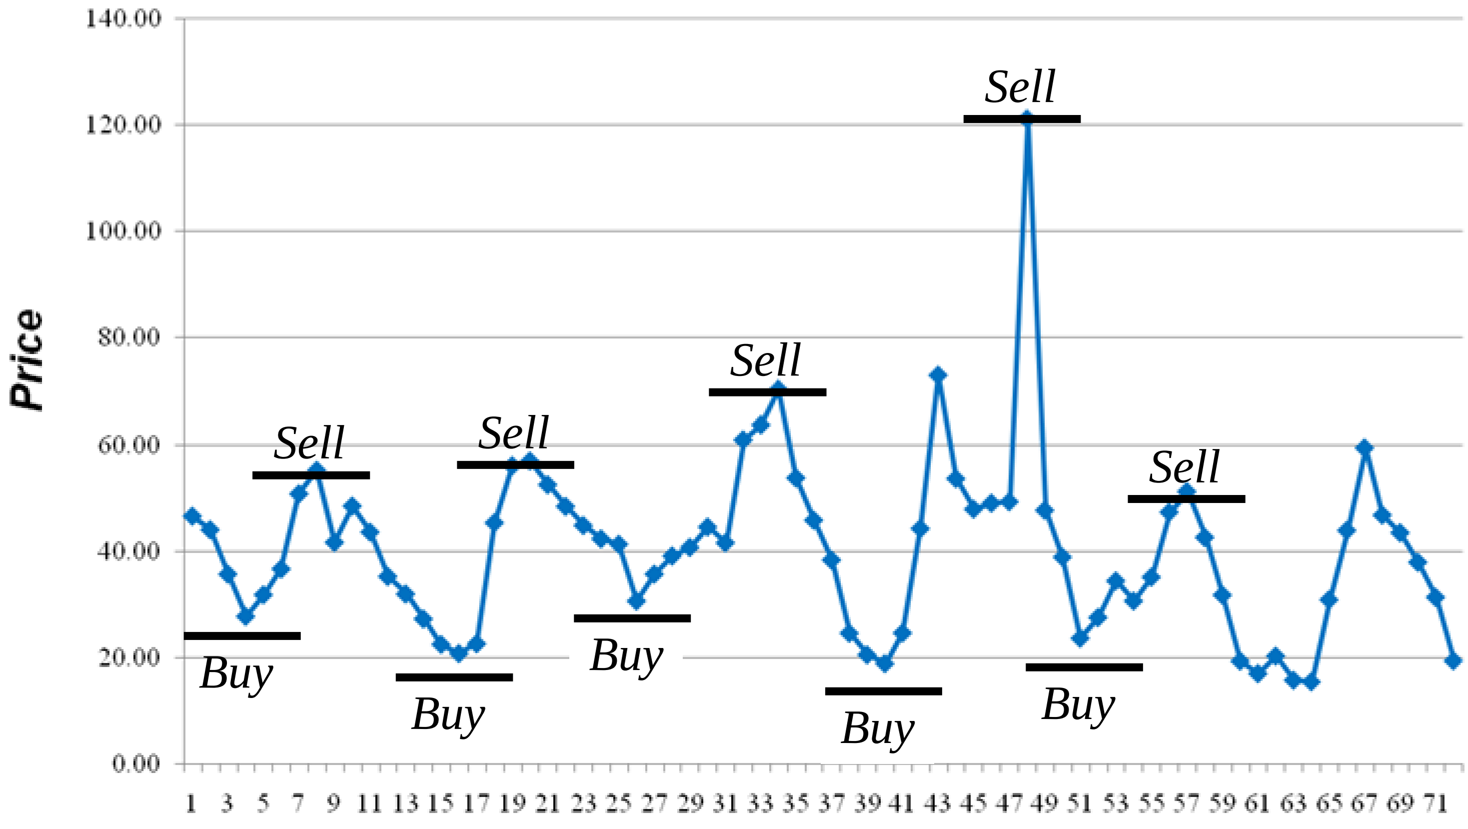

Decision making under uncertainty

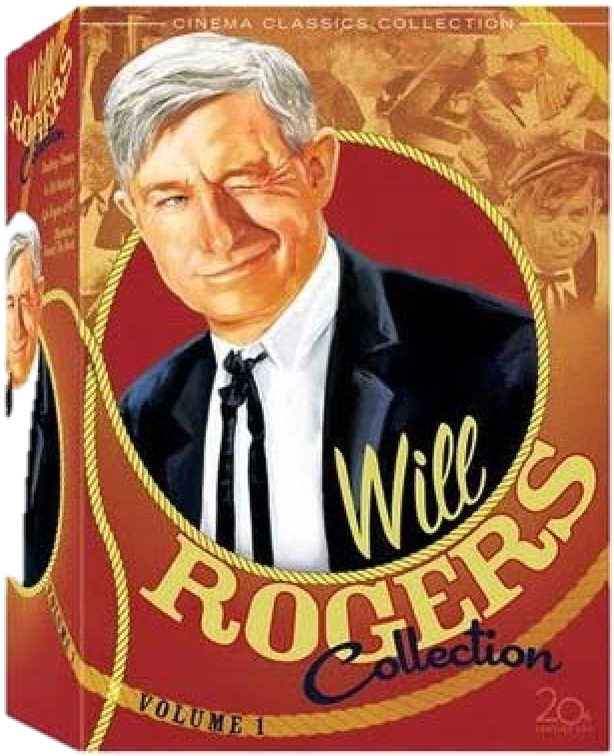

Don't gamble; take all your savings and buy some good stock and hold it till it goes up, then sell it. If it don't go up, don't buy it.

Will Rogers

It is not enough to model the *variability* of a process. You have to model the *uncertainty* – the flow of information.

Optimizing storage

- At NRG's request, we had to design a *simple*, *implementable* policy that did not cheat!

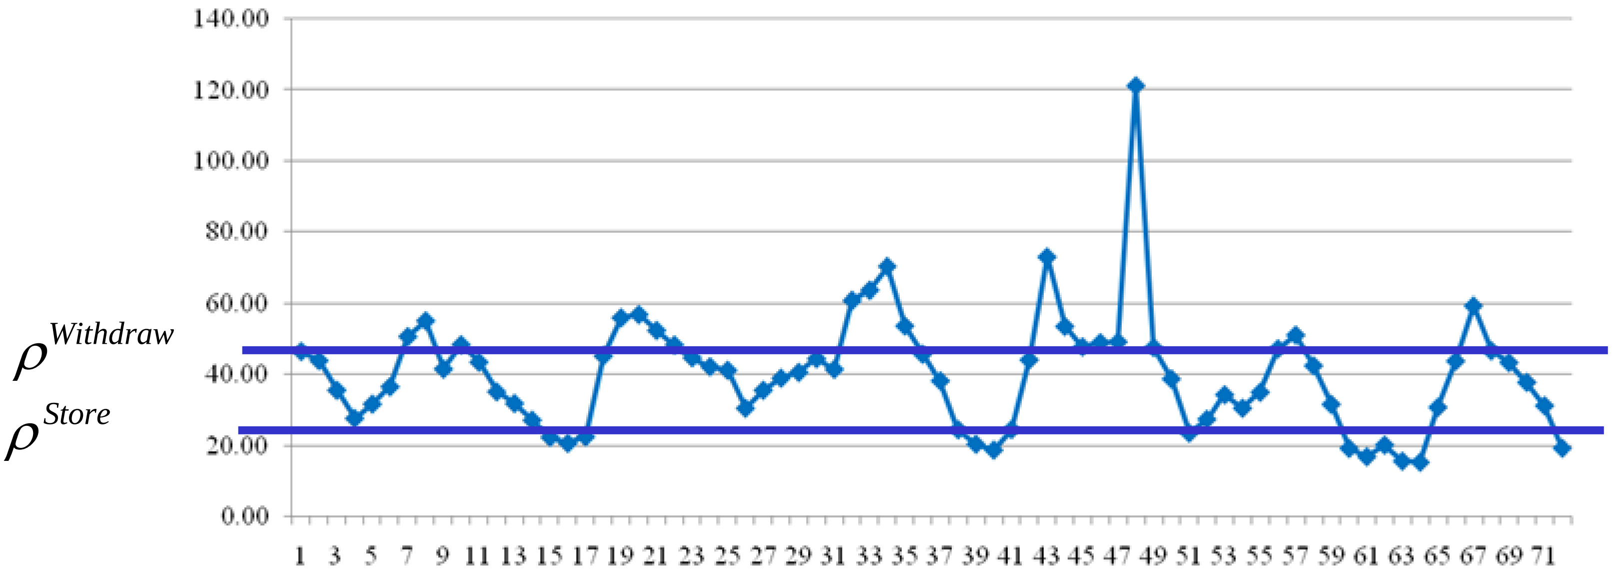

- We have developed a separate line of research in *optimal learning* to determine ρ^{Store} and $\rho^{Withdraw}$.

The exploration vs. exploitation problem

- The knowledge gradient:
  - » Assume you can make only one measurement, after which you have to make a final choice (the implementation decision).
  - » What choice would you make now to maximize the expected value of the implementation decision?

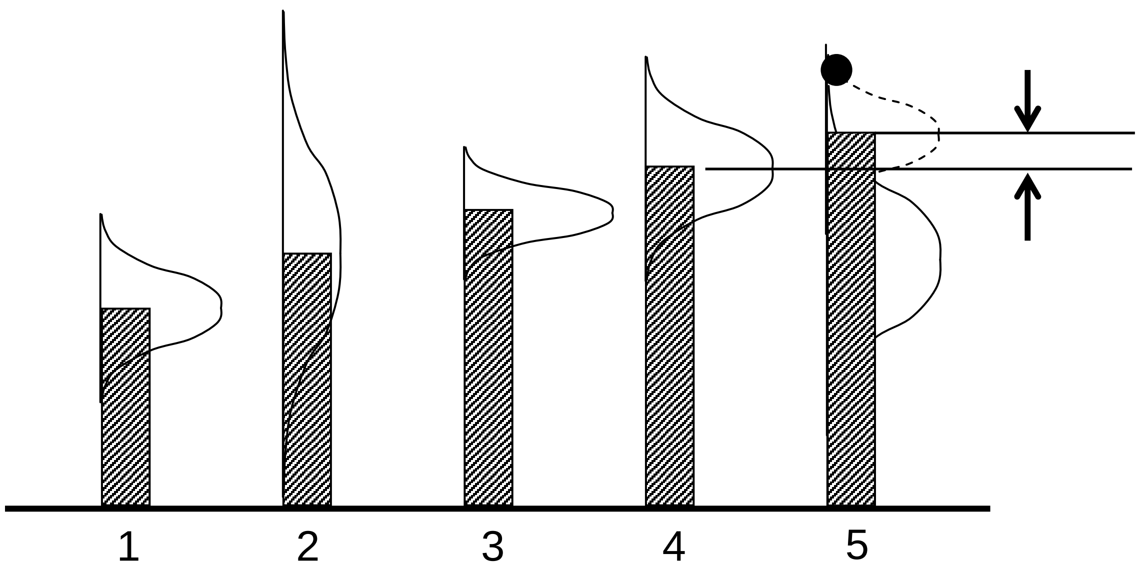

The exploration vs. exploitation problem

- The knowledge gradient
  - » The knowledge gradient is the expected value of a single measurement *x*, given by

$$\nu_x^{KG,n} = E^n \left\{ \max_y F(y, K^{n+1}(x)) \right\} - \max_y F(y, K^n)$$

Marginal value of measuring *x* (the knowledge gradient)

New optimization problem

Expectation over different measurement outcomes

Knowledge state

Implementation decision

Optimization problem given what we kno

Updated knowledge state given measurement *x*

  - » Knowledge gradient policy chooses the measurement with the highest marginal value.
  - » This can be viewed as a kind of coordinate ascent algorithm.

The knowledge gradient

- An important problem class involves *correlated beliefs* – measuring one alternative tells us something other alternatives.

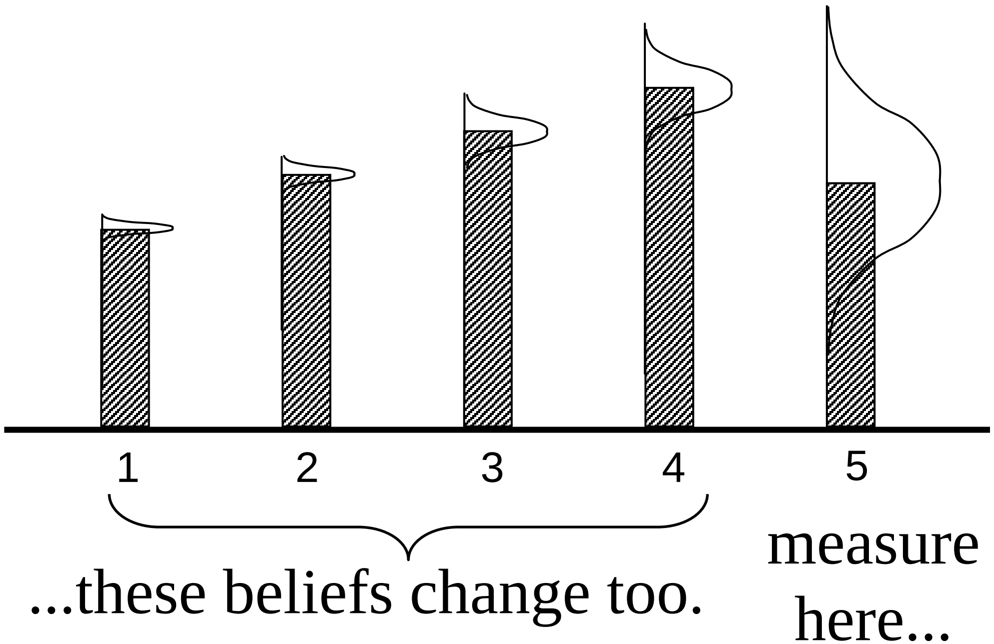

The knowledge gradient

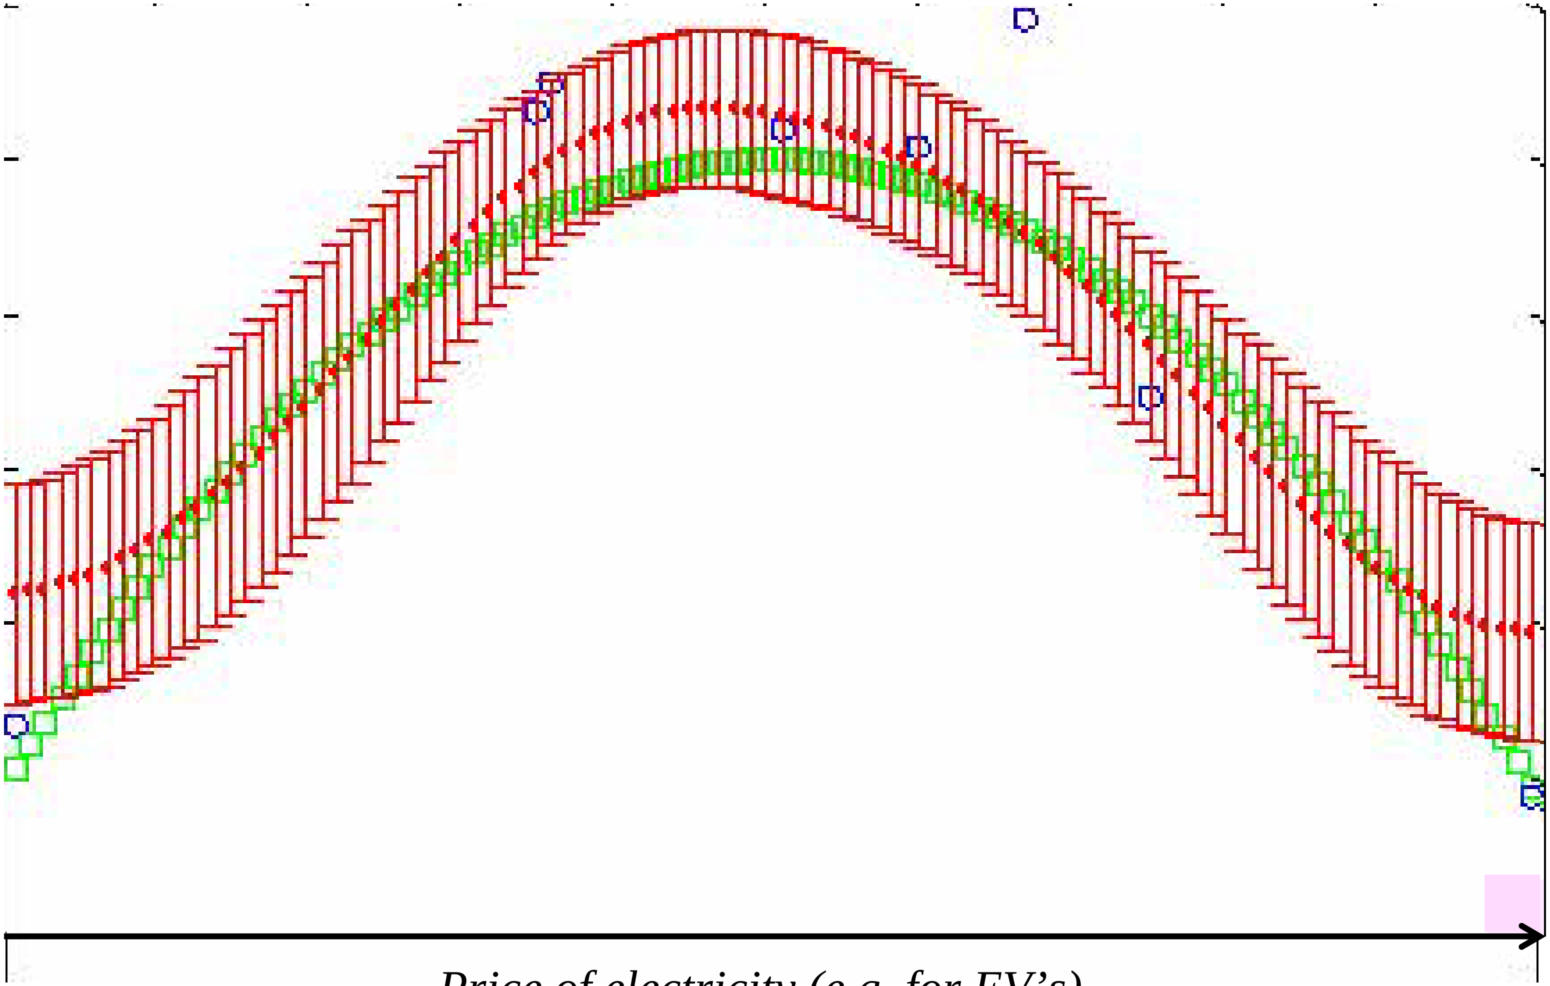

Price of electricity (e.g. for EV's)

Optimizing storage

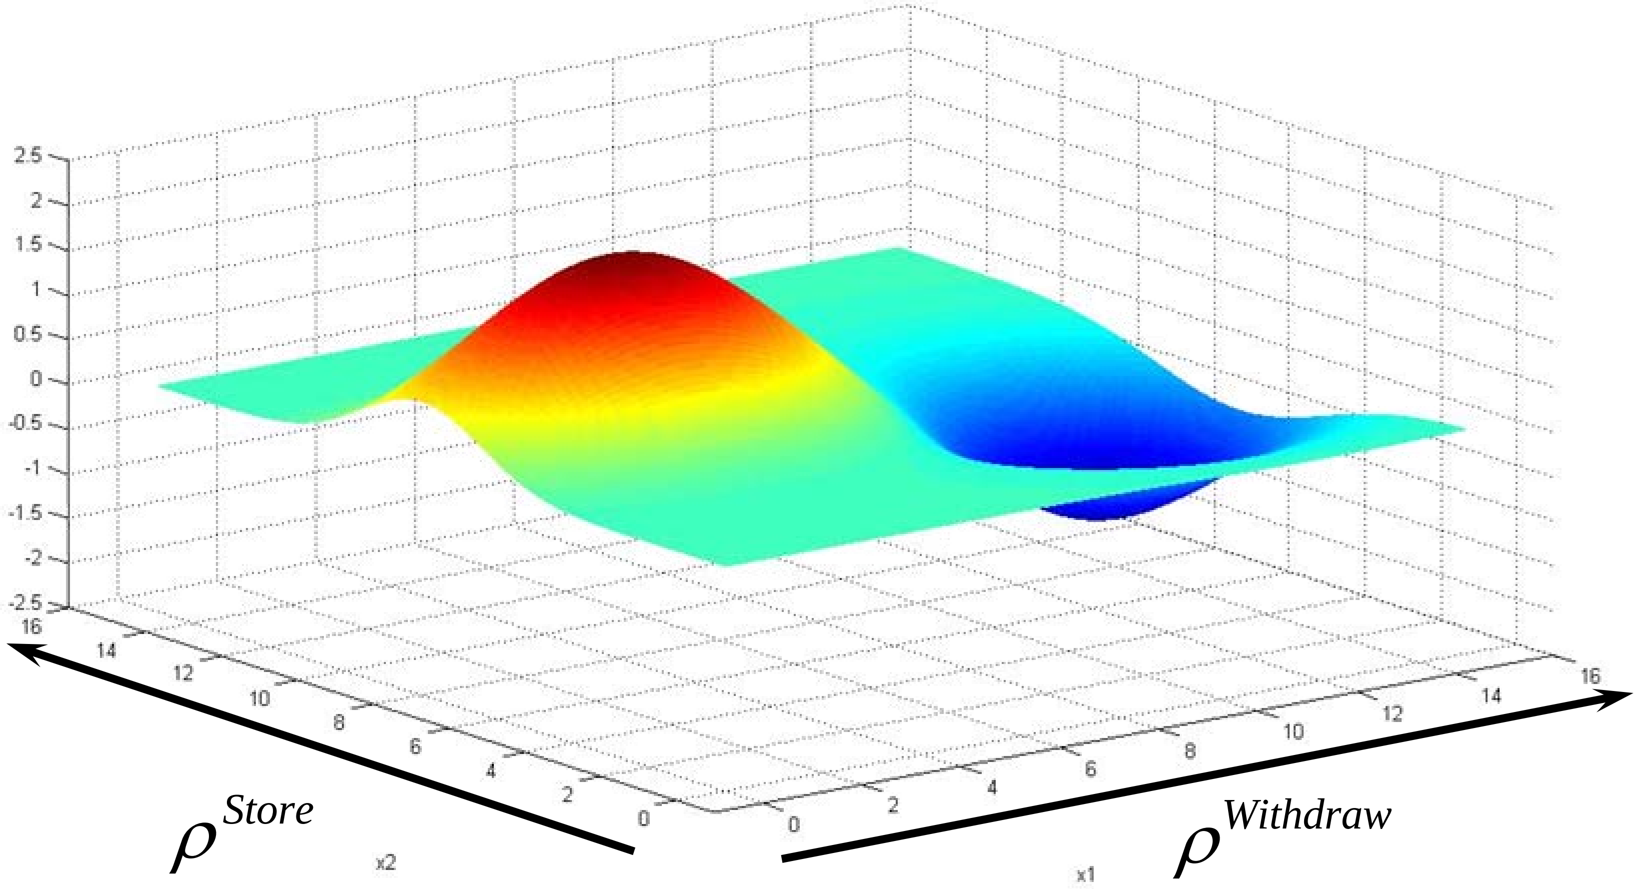

Optimizing storage

- Initially we think the value is the same everywhere:

Estimated value

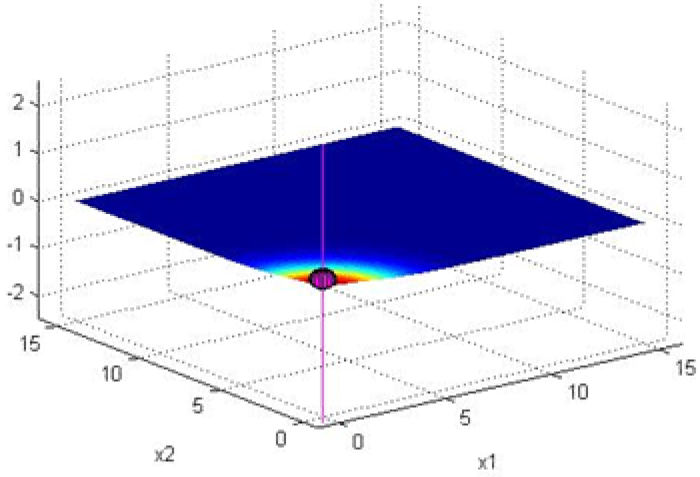

Knowledge gradient

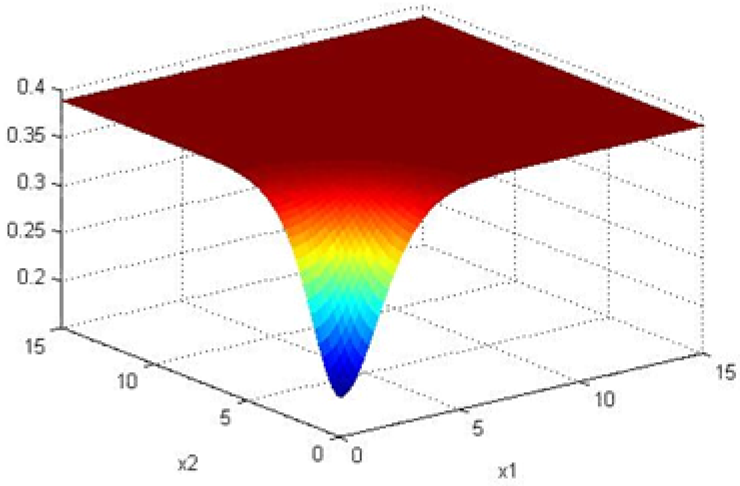

» We want to measure the value where the knowledge gradient is the highest. This is the measurement that teaches us the most.

Optimizing storage

- After four measurements:

Estimated value

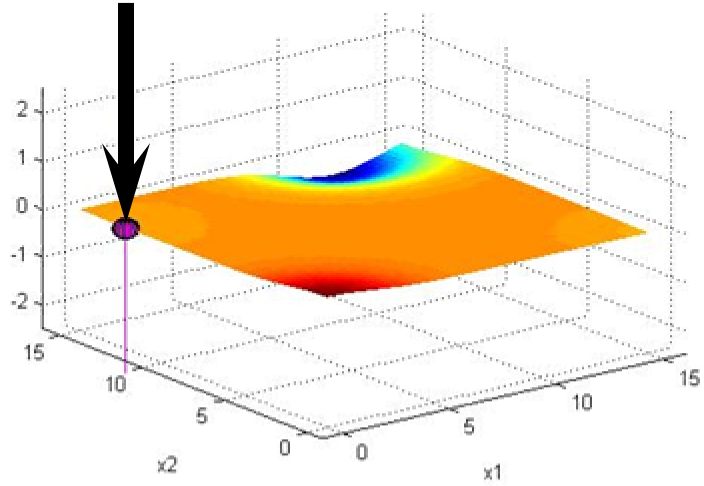

Knowledge gradient

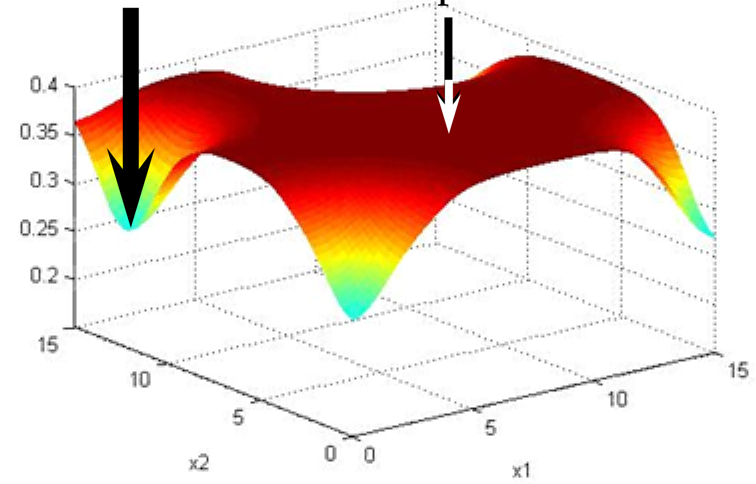

» Whenever we measure at a point, the value of another measurement at the same point goes down. The knowledge gradient guides us to measuring areas of high uncertainty.

Optimizing storage

- After five measurements:

Estimated value

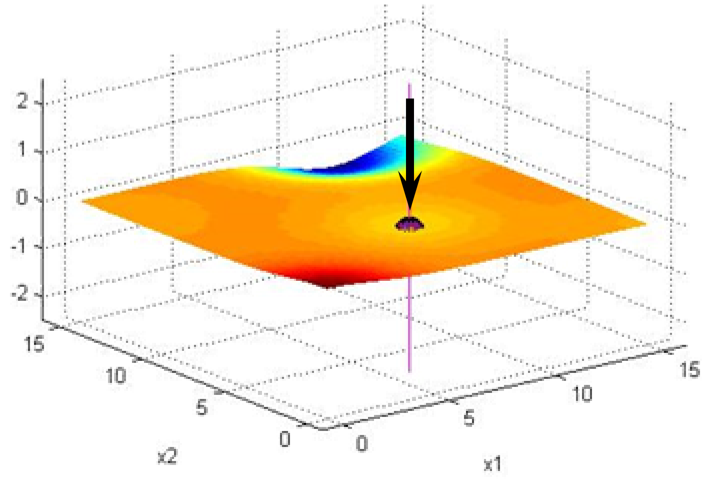

Knowledge gradient

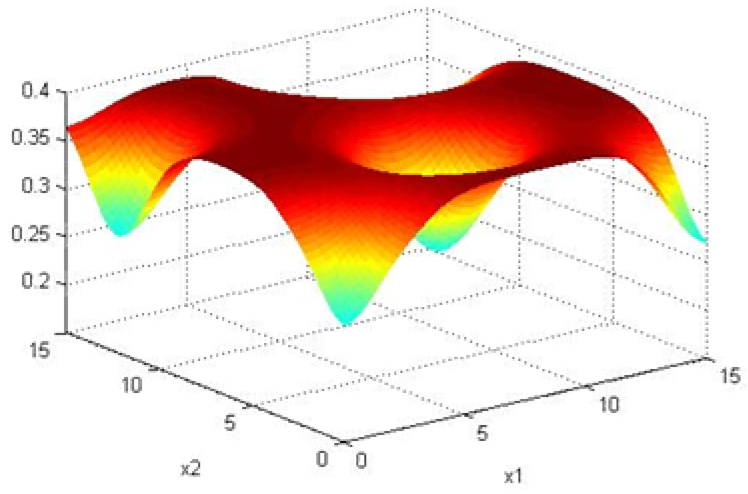

Optimizing storage

- After six samples

Estimated value

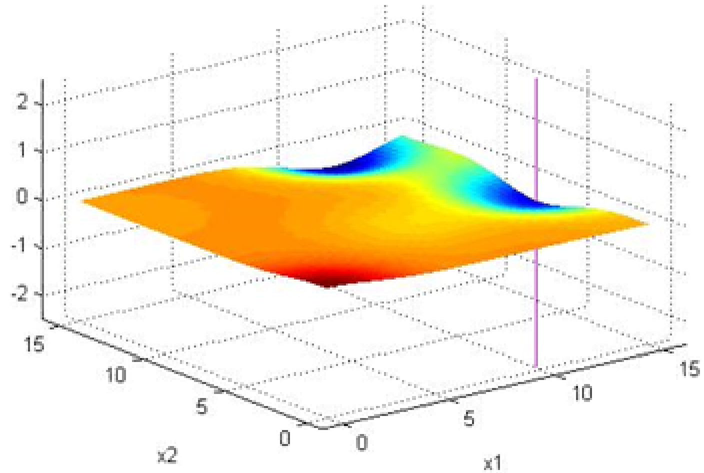

Knowledge gradient

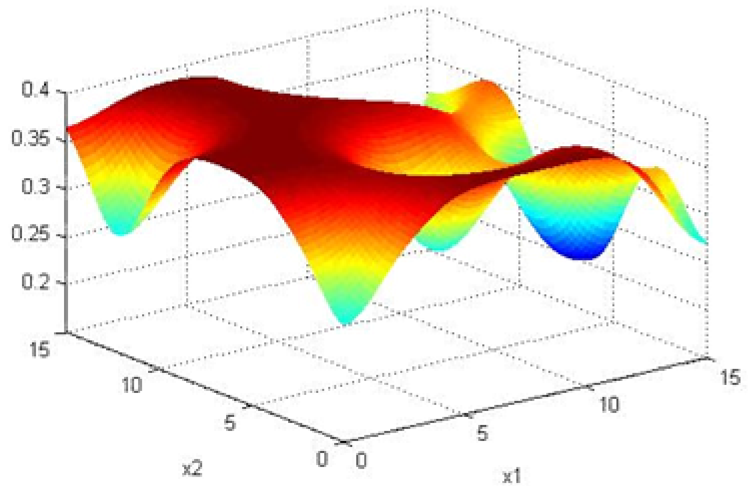

Optimizing storage

- After seven samples

Estimated value

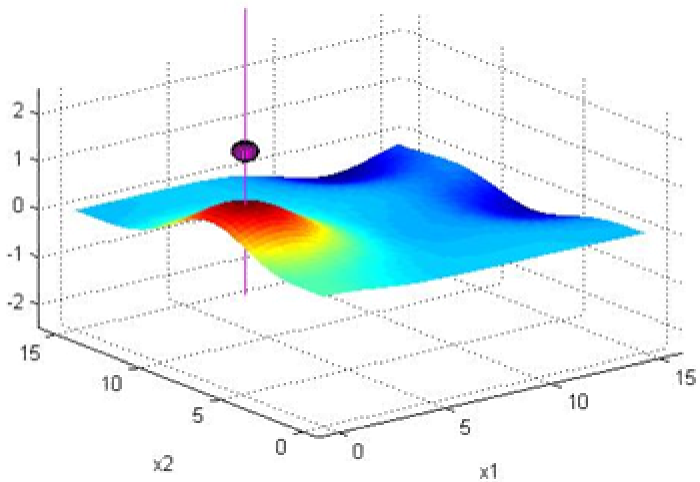

Knowledge gradient

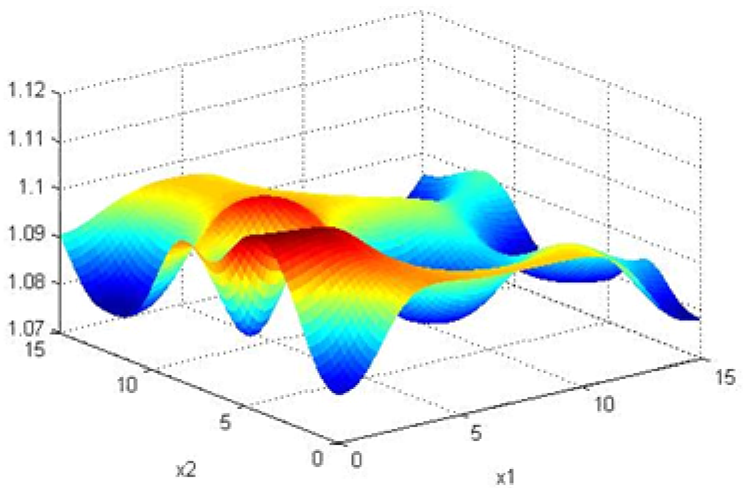

Optimizing storage

- After eight samples

Estimated value

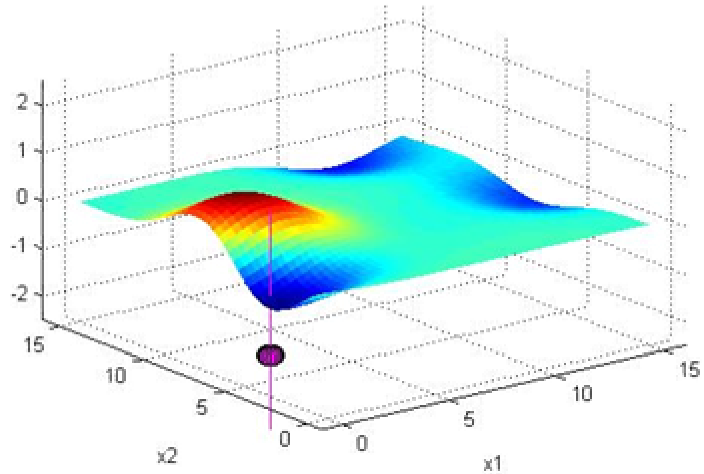

Knowledge gradient

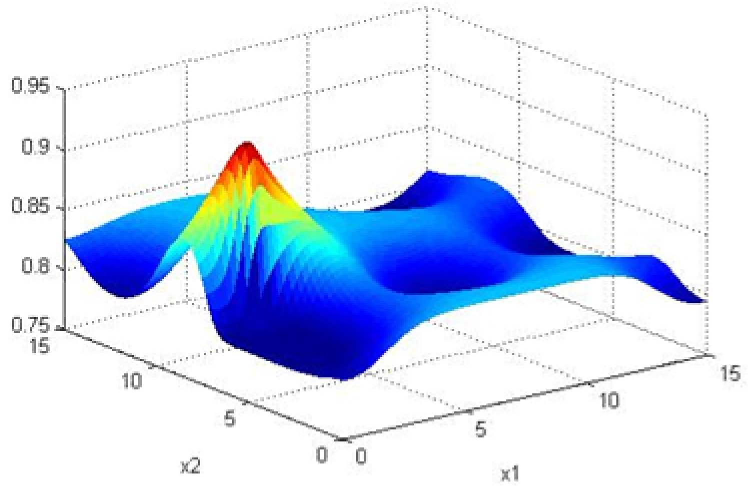

Optimizing storage

- After nine samples

Estimated value

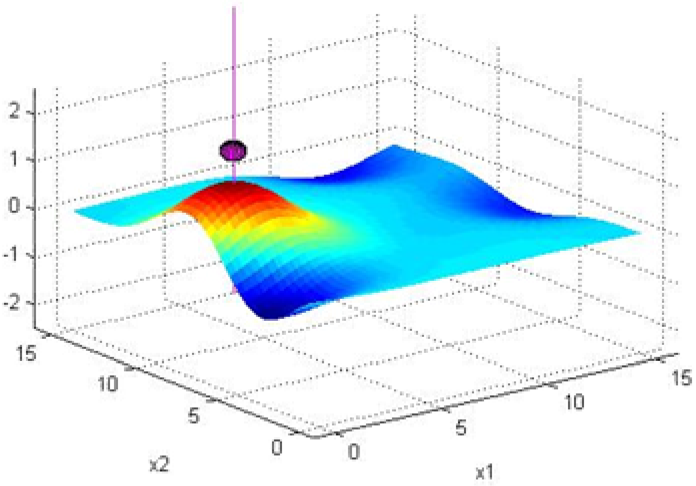

Knowledge gradient

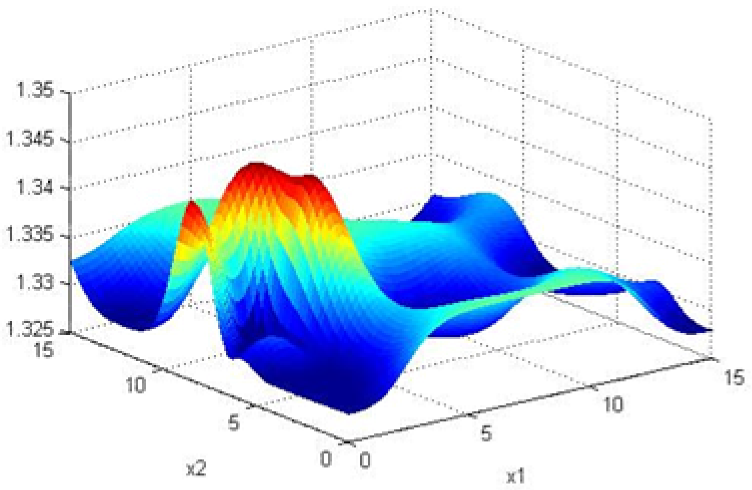

Optimizing storage

- After ten samples

Estimated value

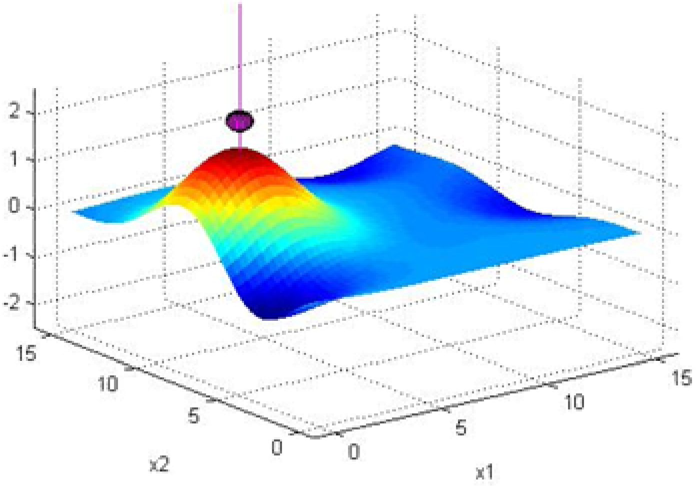

Knowledge gradient

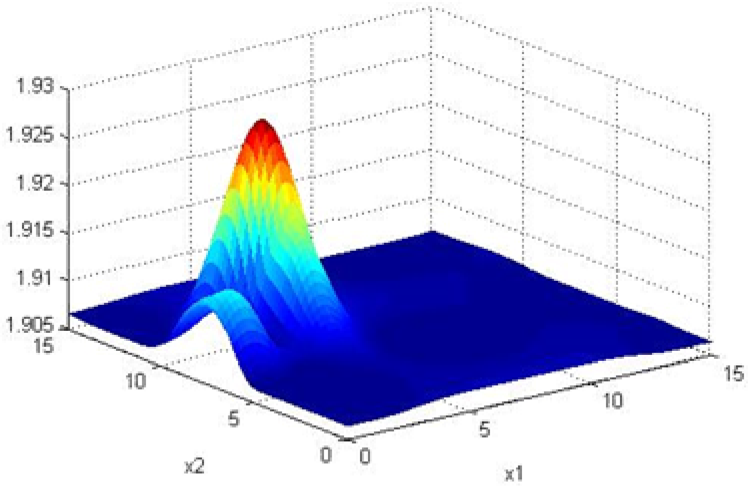

Optimizing storage

- After 10 measurements, our estimate of the surface:

Estimated value

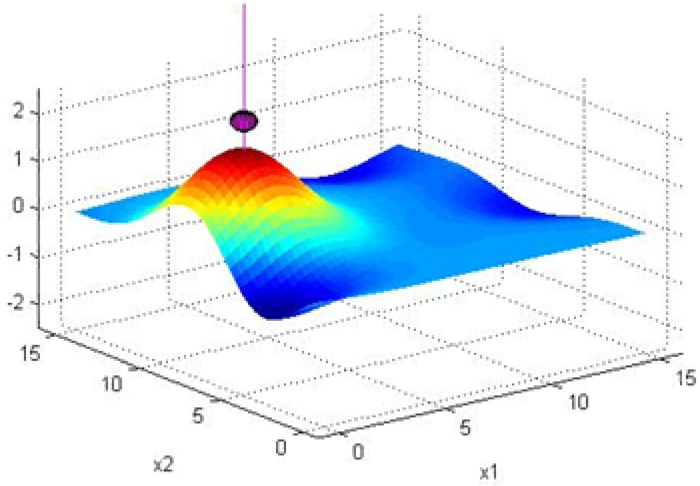

True concentration

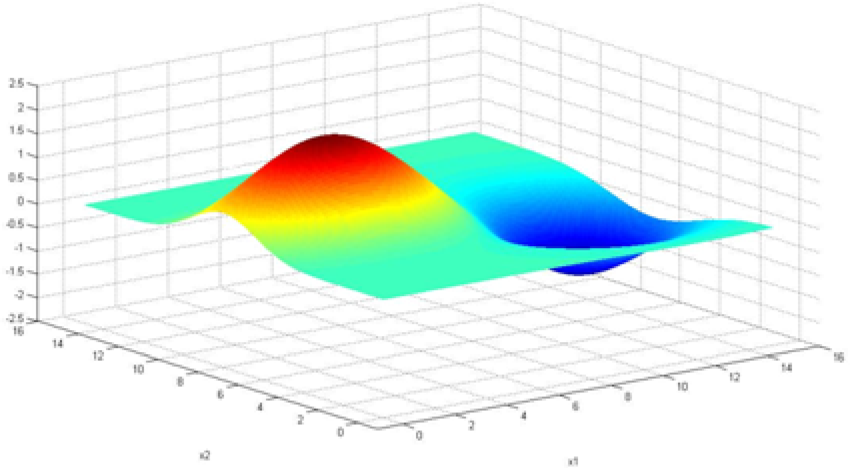

Optimizing storage

- After 10 measurements, our estimate of the surface:

Estimated value

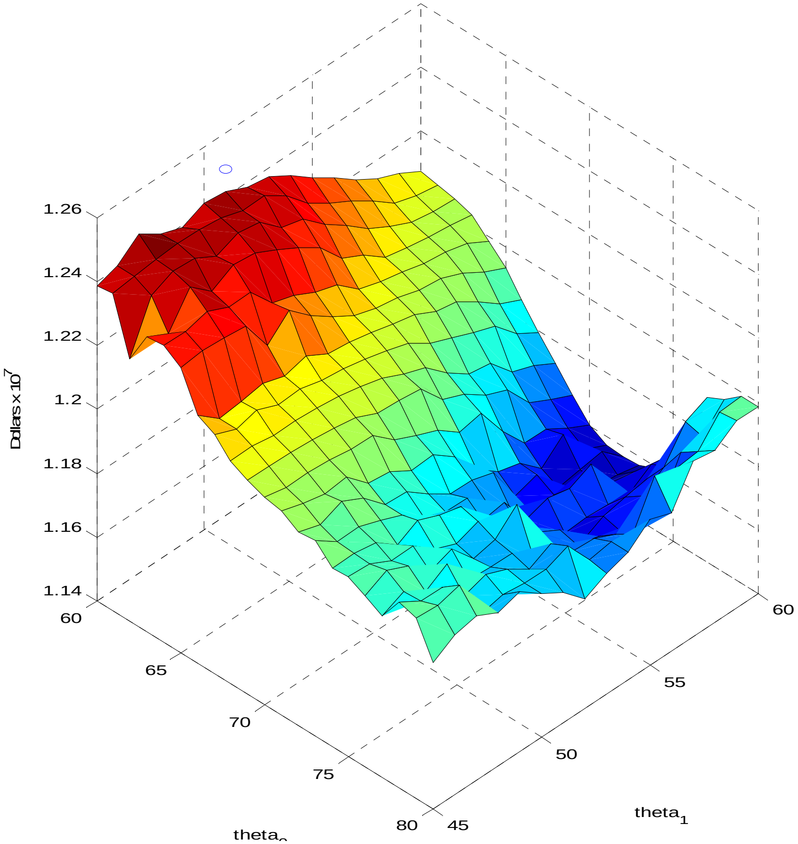

Knowledge gradient

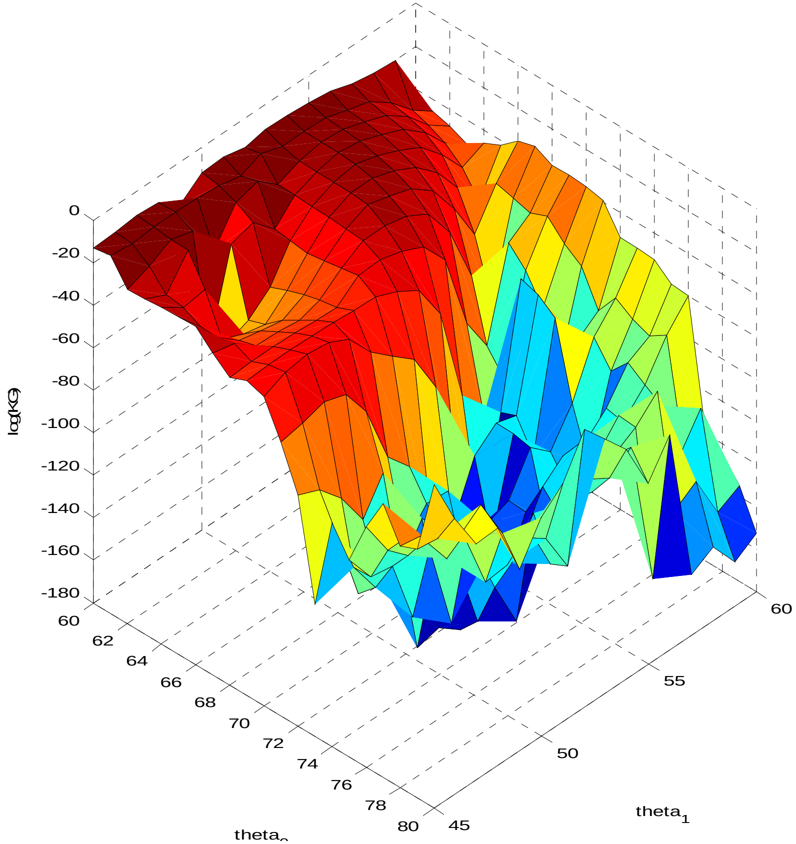

Optimizing storage

- The value of perfect information

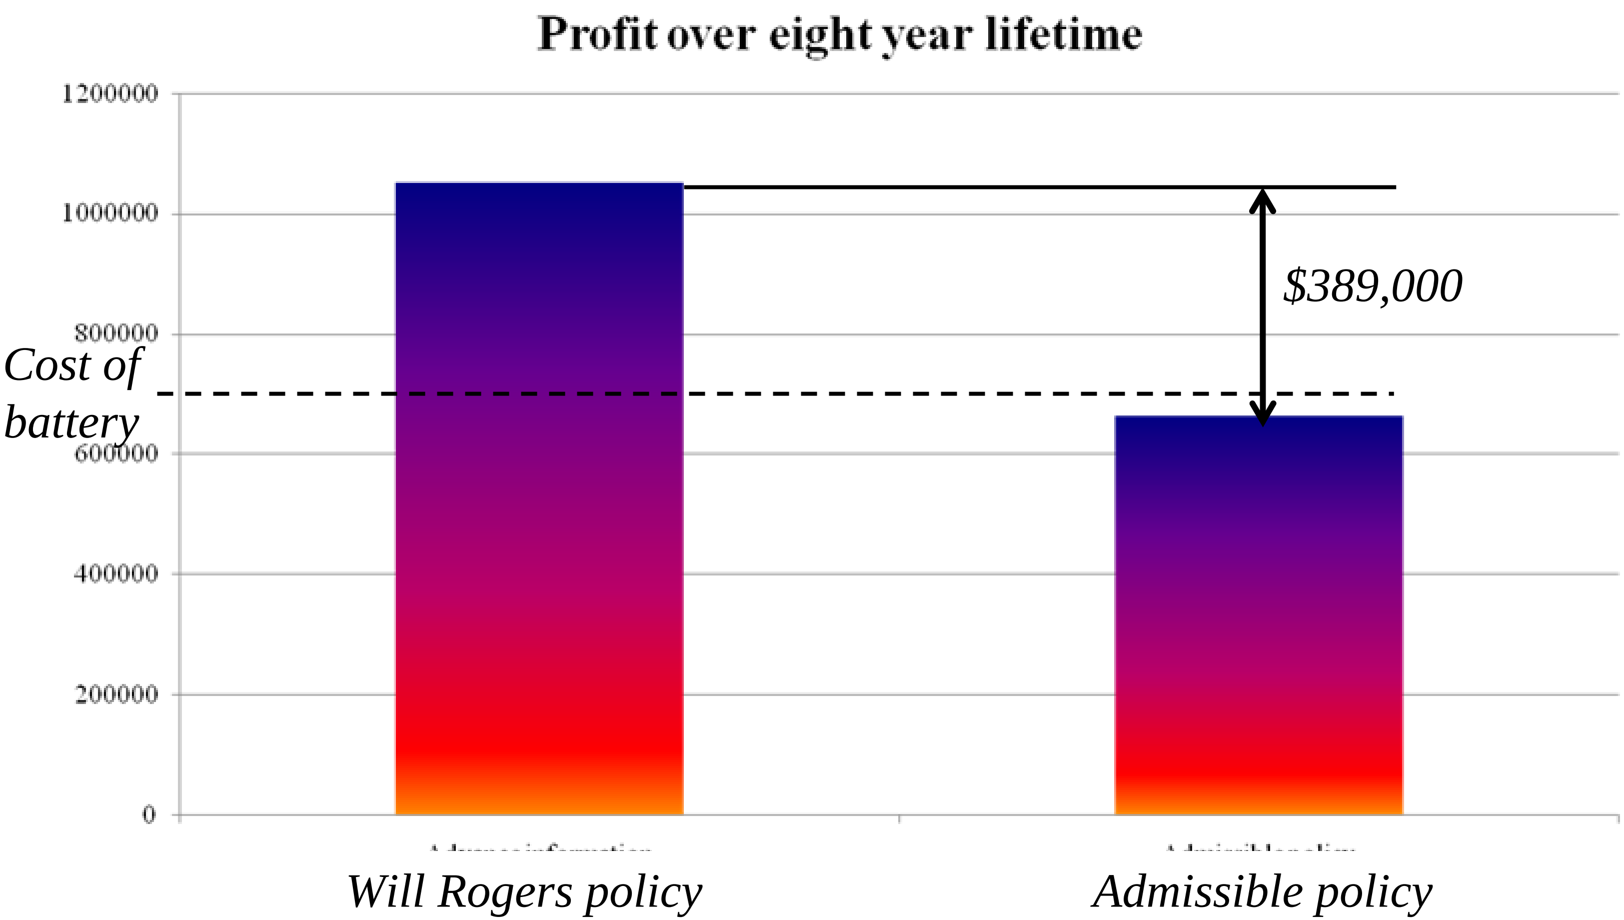

Lecture outline

- Optimizing energy storage
- The stochastic unit commitment problem for PJM


Legend
Backbone Transmission
Voltage (kV)
Backbone Substations
Voltage (kV)

The stochastic unit commitment problem

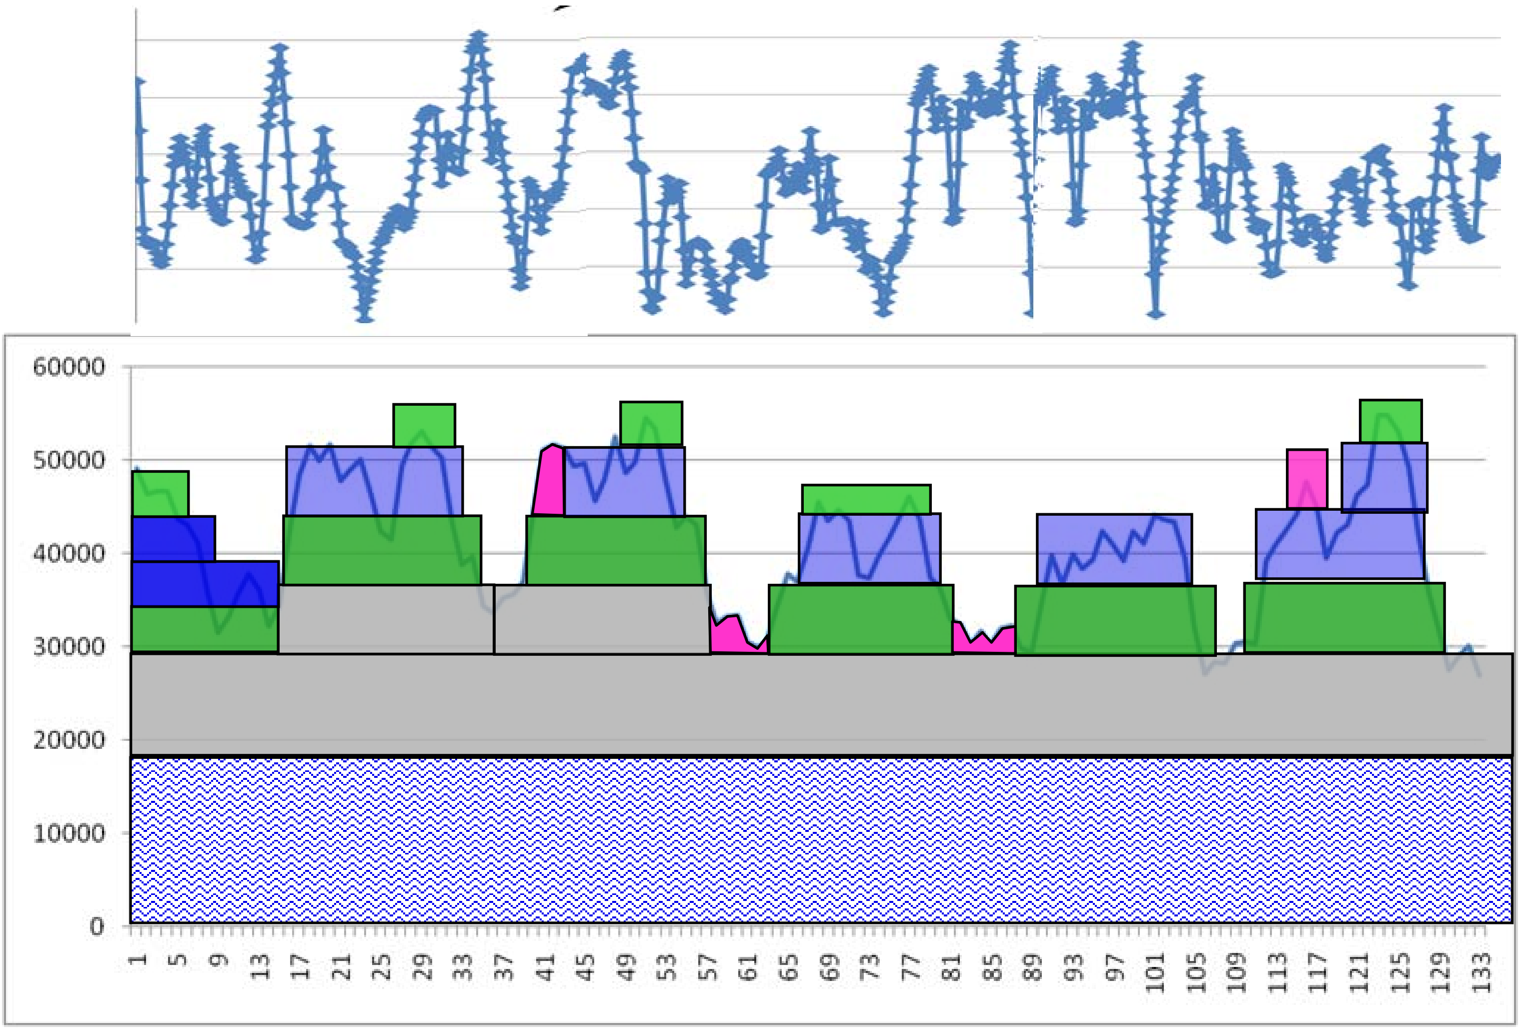

The stochastic unit commitment problem

- The day-ahead problem
  - » Determines which coal/nuclear/natural gas/…plants to turn on/off and when.
  - » Will use wind/solar when available and if needed.
  - » Requires point forecast of wind/solar/demand.

$$\min_{x_1,\ldots,x_{24}} \sum_{t=1}^{24} C(x_t)$$

….subject to numerous constraints, including integrality.

  - » The problem is solved using two adjustment parameters
    - p Fraction of generator capacity assumed (e.g. 93 percent)
    - q Quantile of wind forecast assumed for advance commitments.

The stochastic unit commitment problem

- The hour-ahead problem
  - Each hour, we can make modest adjustments. Plants that are “on” can be adjusted up and down.
  - Cannot turn coal plants on and off.
  - Actual may differ from forecast:
    - Wind may be higher or lower than forecast. If higher, may not be able to use it because of inability to scale back other sources.
    - Demand may exceed forecast. Wind/solar may fall below forecast. In this case, we find the least cost unit that can be scaled up quickly enough.

The stochastic unit commitment problem

- Modeling
  - » A deterministic model

$$\min_{x_1,\ldots,x_{24}} \sum_{t=1}^{24} C(x_t)$$

  - » Stochastic formulation – I

$$\min_{x_1^s,\ldots,x_{24}^s} \sum_{s=1}^{S} \sum_{t=1}^{24} p^s C(x_t^s) \qquad s = \text{"scenario"}$$

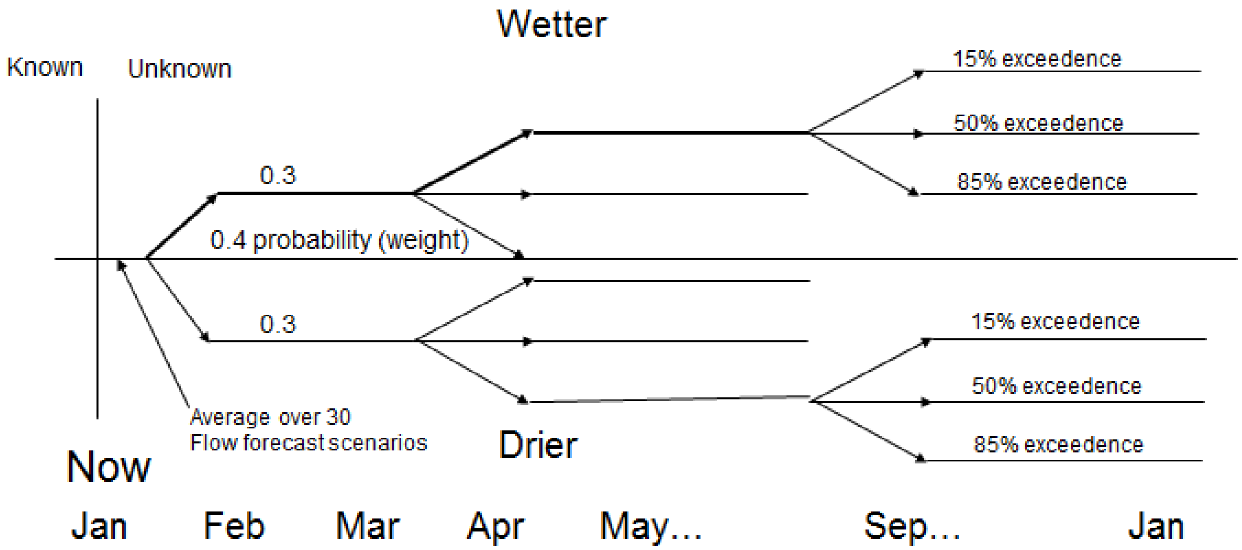

The stochastic unit commitment problem

- Modeling
  - » Stochastic formulation – II

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

$$y_{t',t'}$$

- $x_{t,t'}$ is determined at time t, to be implemented at time t'
- $y_{t',t'}$ is determined at time t', to be implemented at time $t'+1$

  - » Important to recognize information content
    - At time t, $x_{t,t'}$ is deterministic.
    - At time t, $y_{t',t'}$ is stochastic.

The stochastic unit commitment problem

- Matching supply to demand
  - » We have to find the best way to meet demand

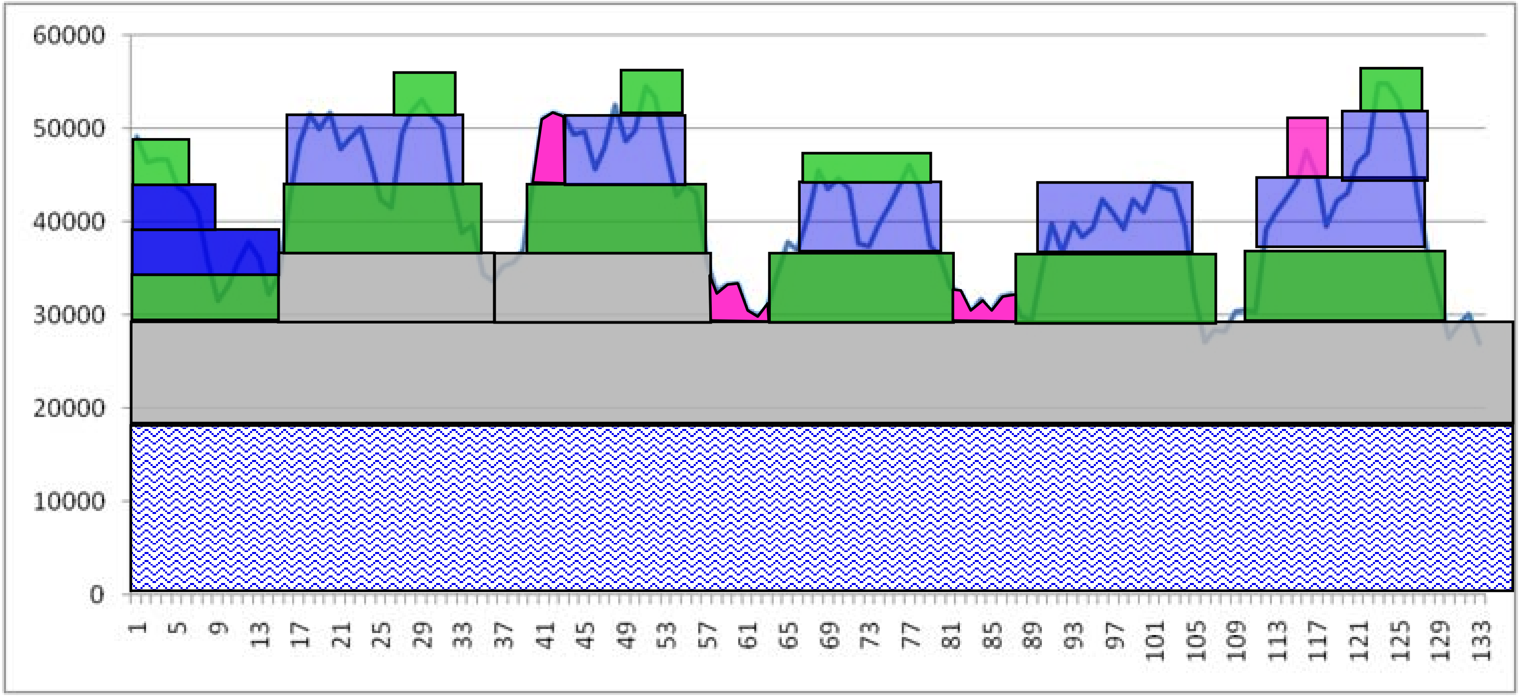

  - » …. Now we have to do it in the presence of significant levels of wind and solar energy.

The stochastic unit commitment problem

- The unit commitment problem
  - » Rolling forward with perfect forecast of actual wind, demand, …

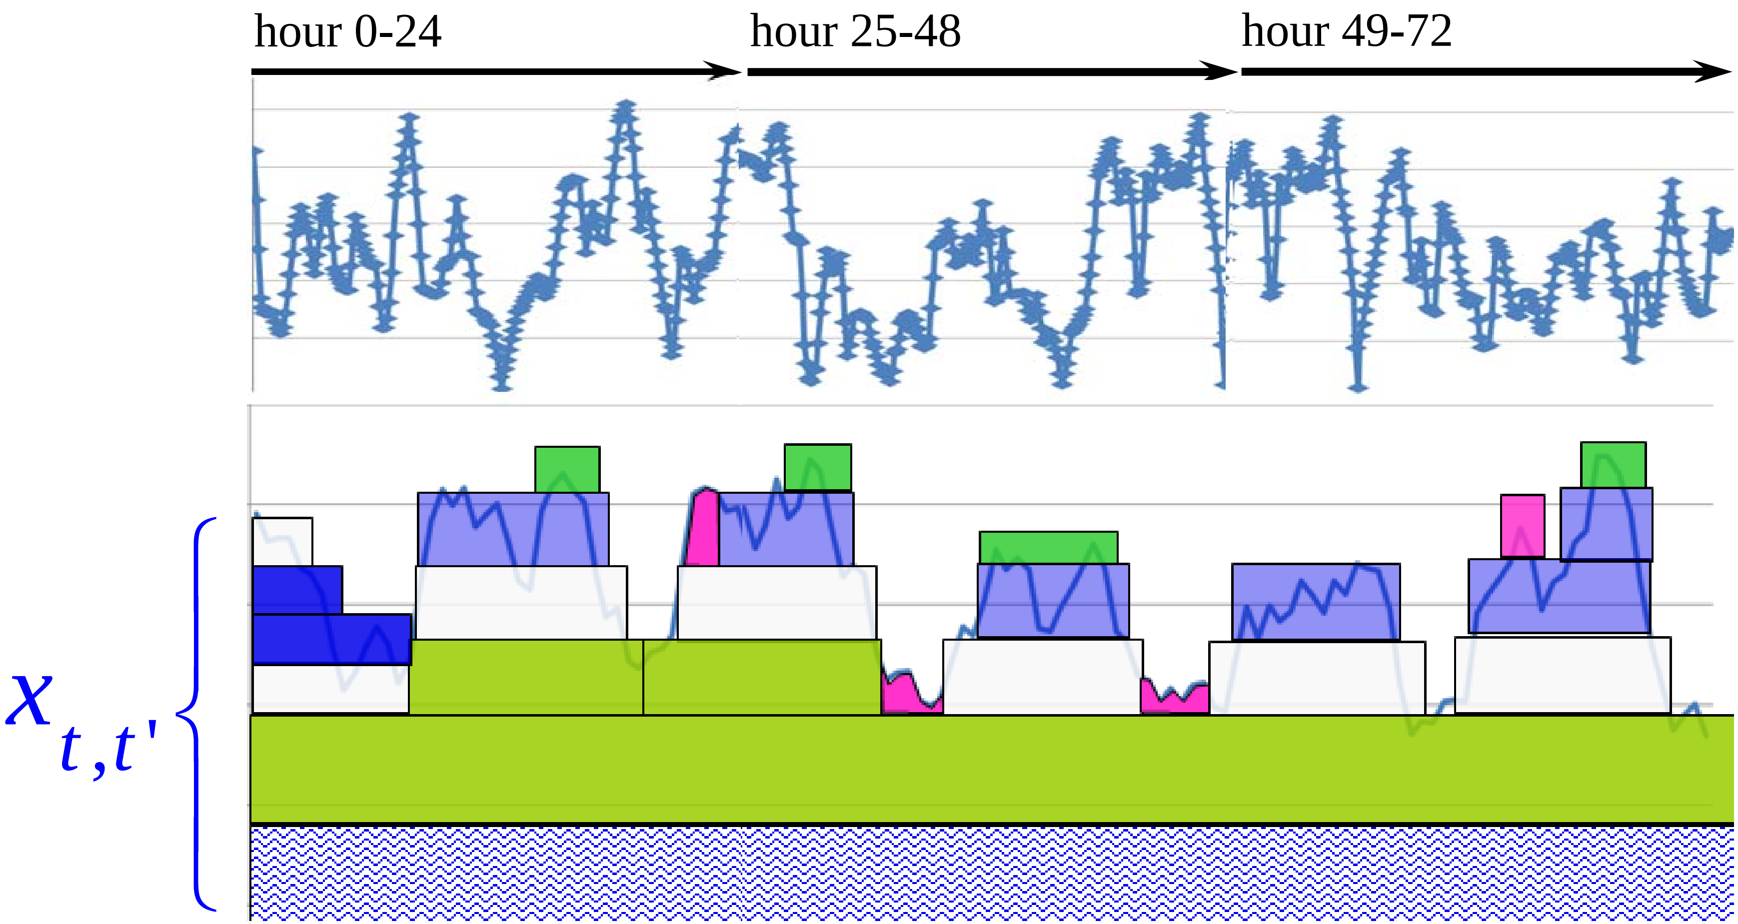

The stochastic unit commitment problem

- When planning, we have to use a *forecast* of energy from wind, then live with what *actually* happens.

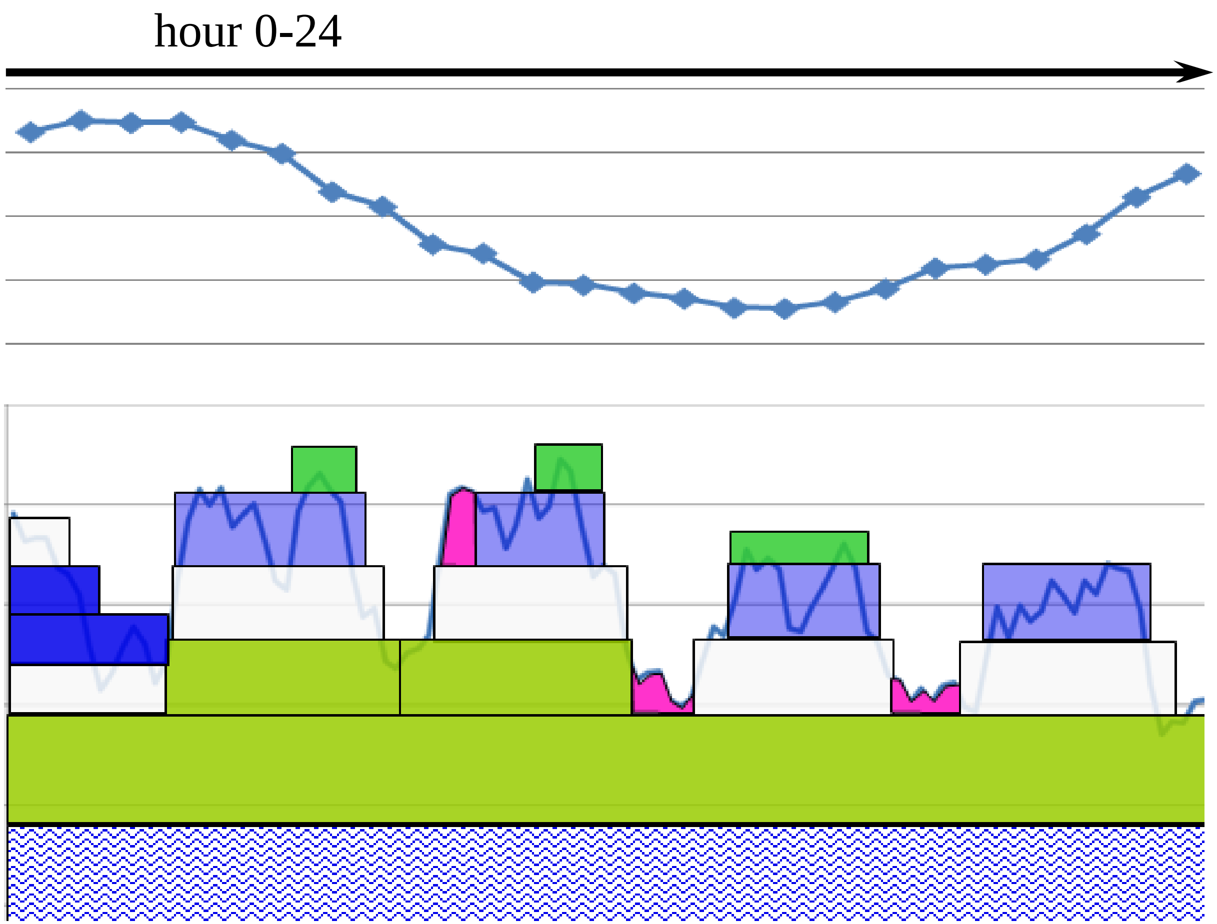

The stochastic unit commitment problem

- The unit commitment problem
  - » Stepping forward observing actual wind, making small adjustments

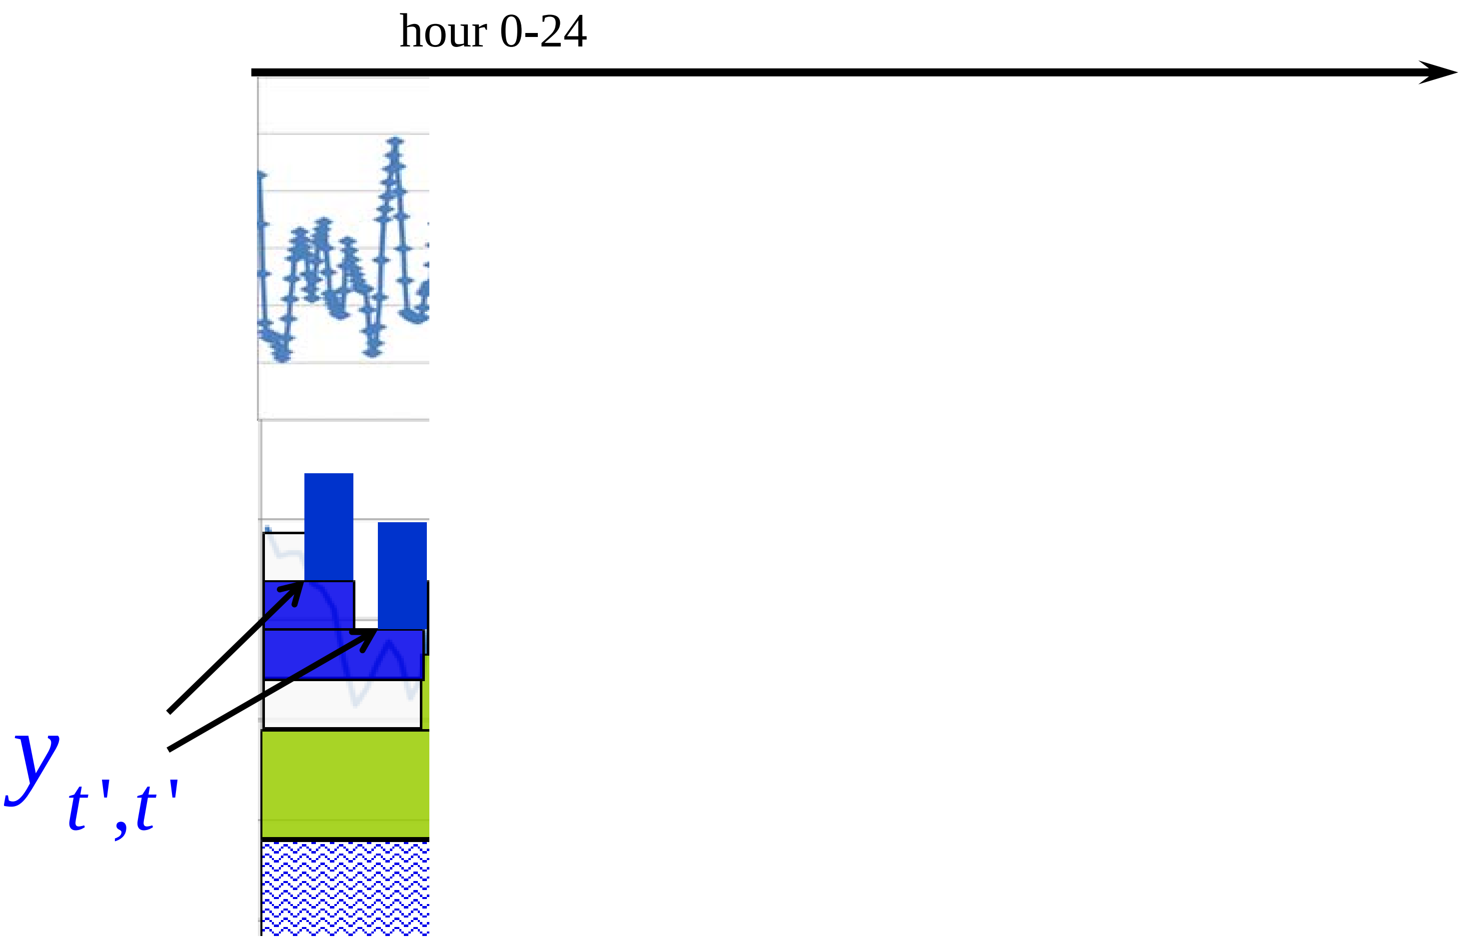

The stochastic unit commitment problem

- The unit commitment problem
  - » Stepping forward observing actual wind, making small adjustments

hour 0-24

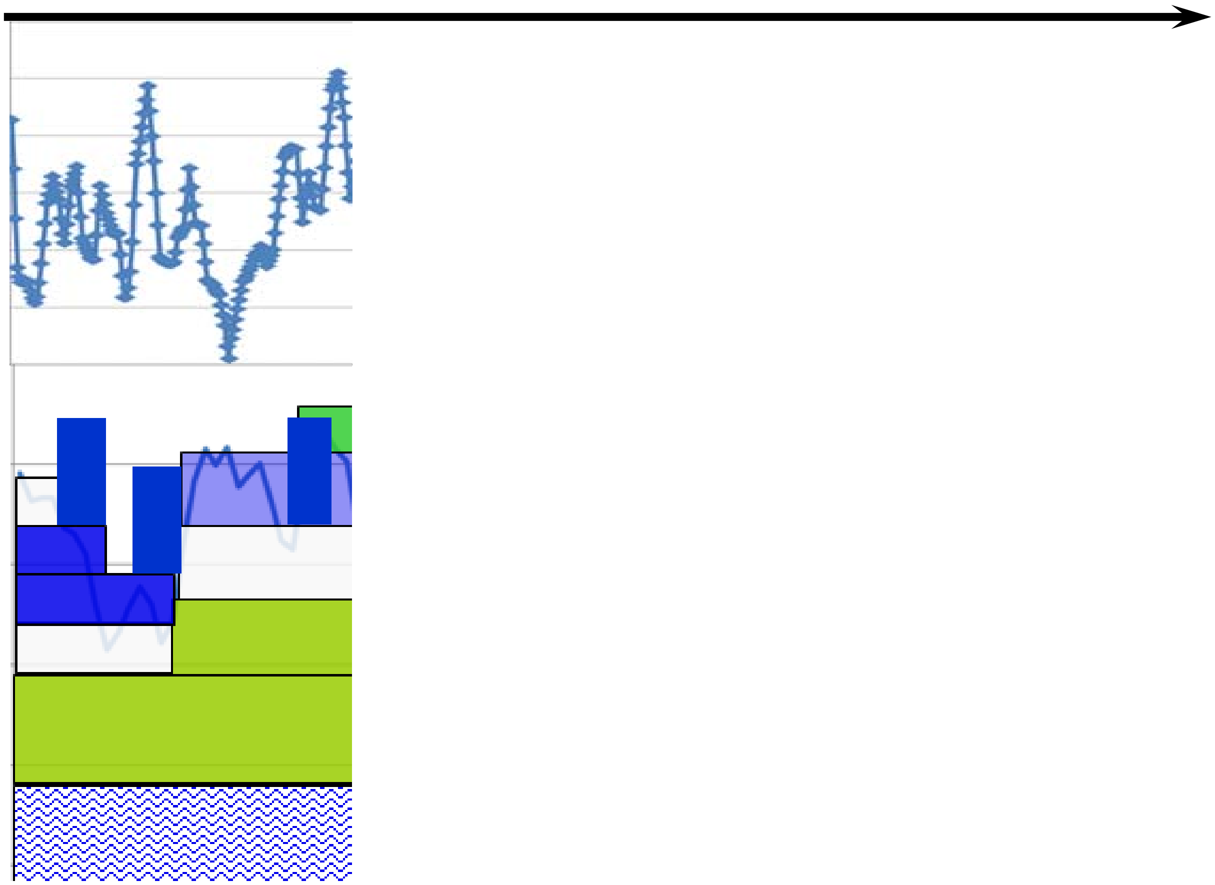

The stochastic unit commitment problem

- The unit commitment problem
  - » Stepping forward observing actual wind, making small adjustments

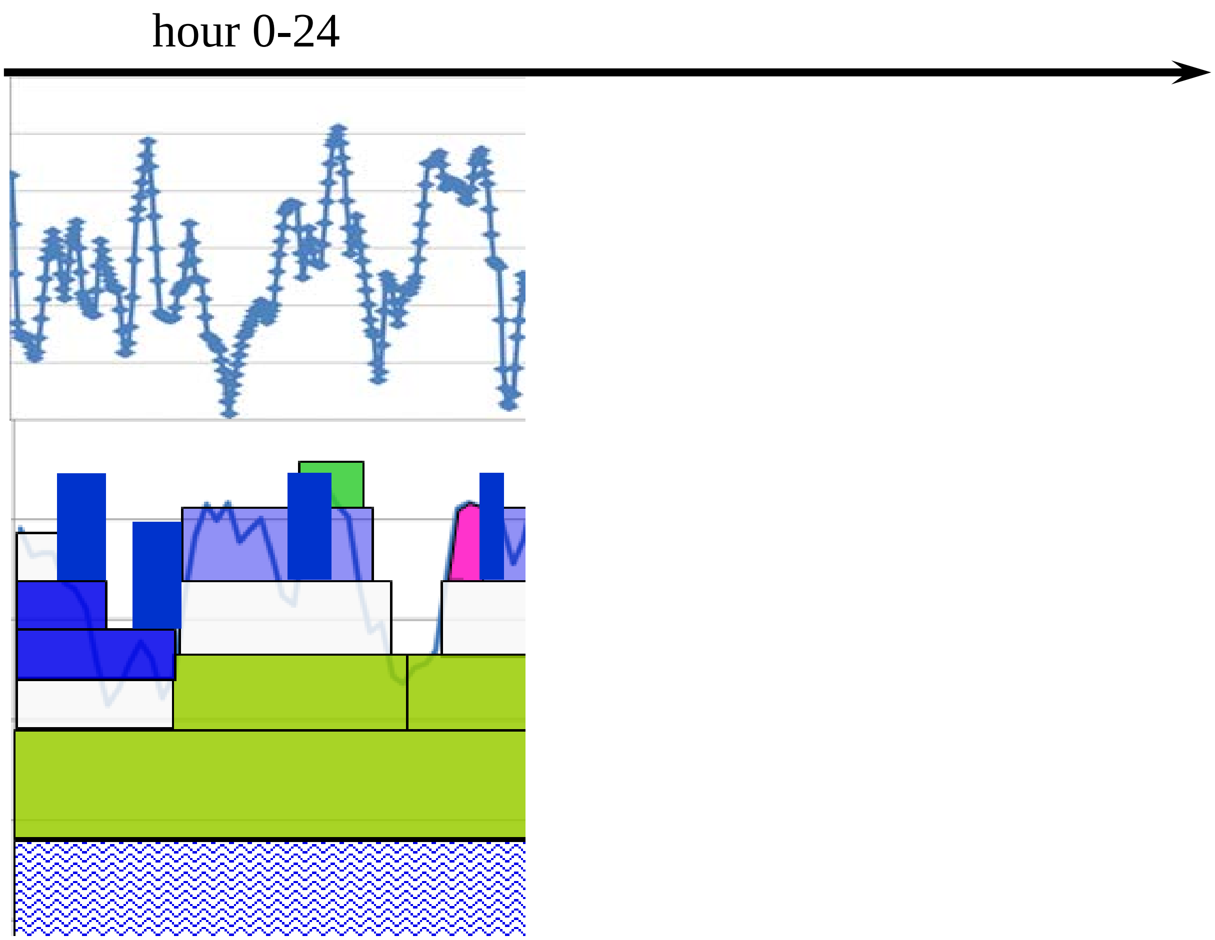

The stochastic unit commitment problem

- The unit commitment problem
  - » Stepping forward observing actual wind, making small adjustments

hour 0-24

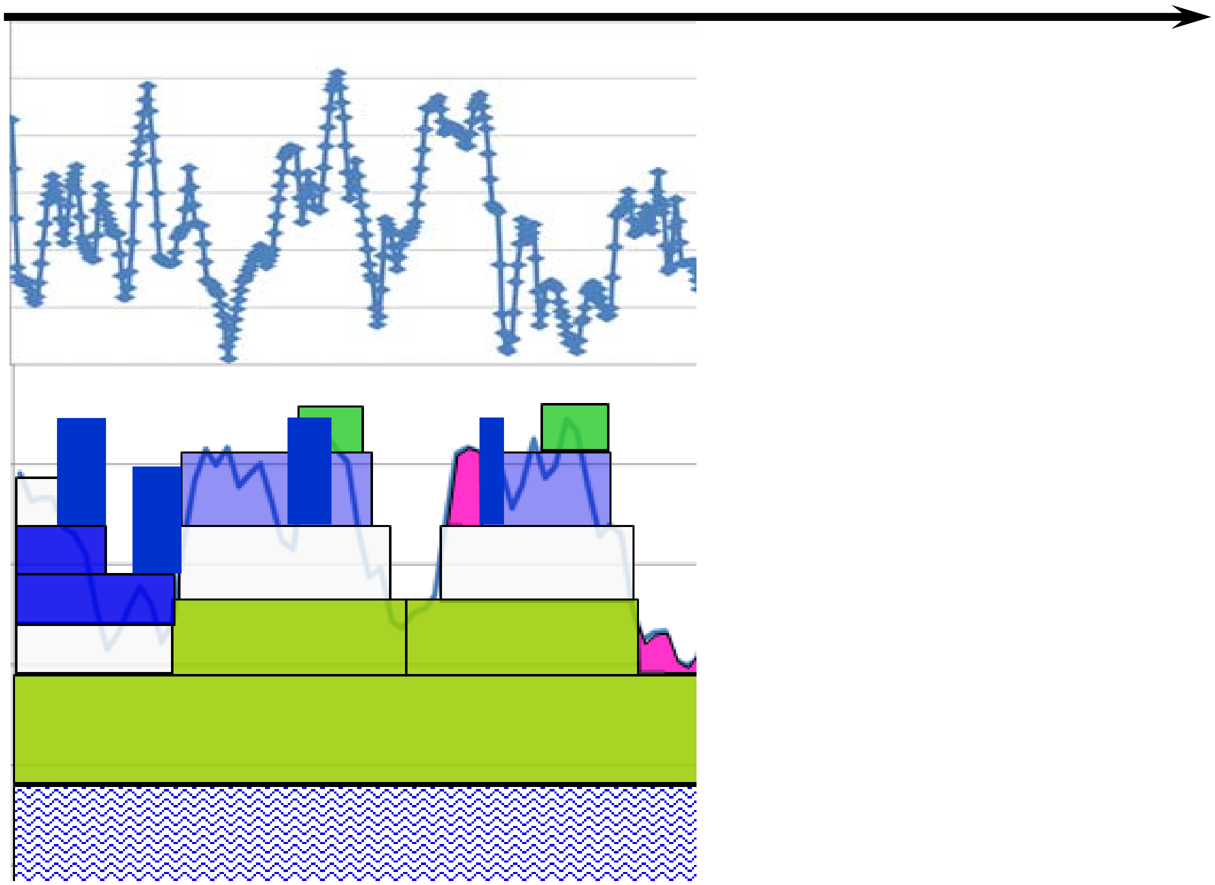

The stochastic unit commitment problem

- The unit commitment problem
  - » Stepping forward observing actual wind, making small adjustments

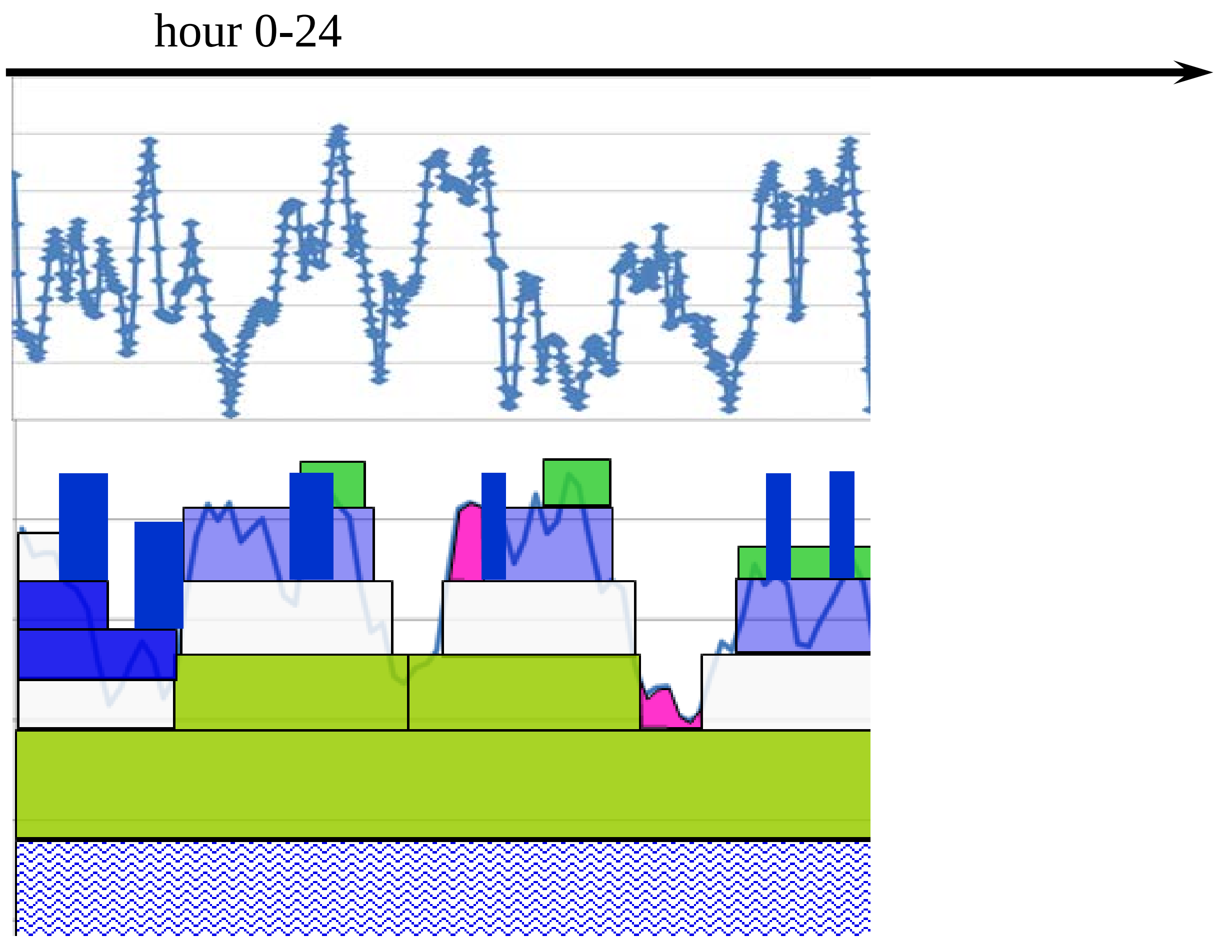

The stochastic unit commitment problem

- The unit commitment problem
  - » Stepping forward observing actual wind, making small adjustments

hour 0-24

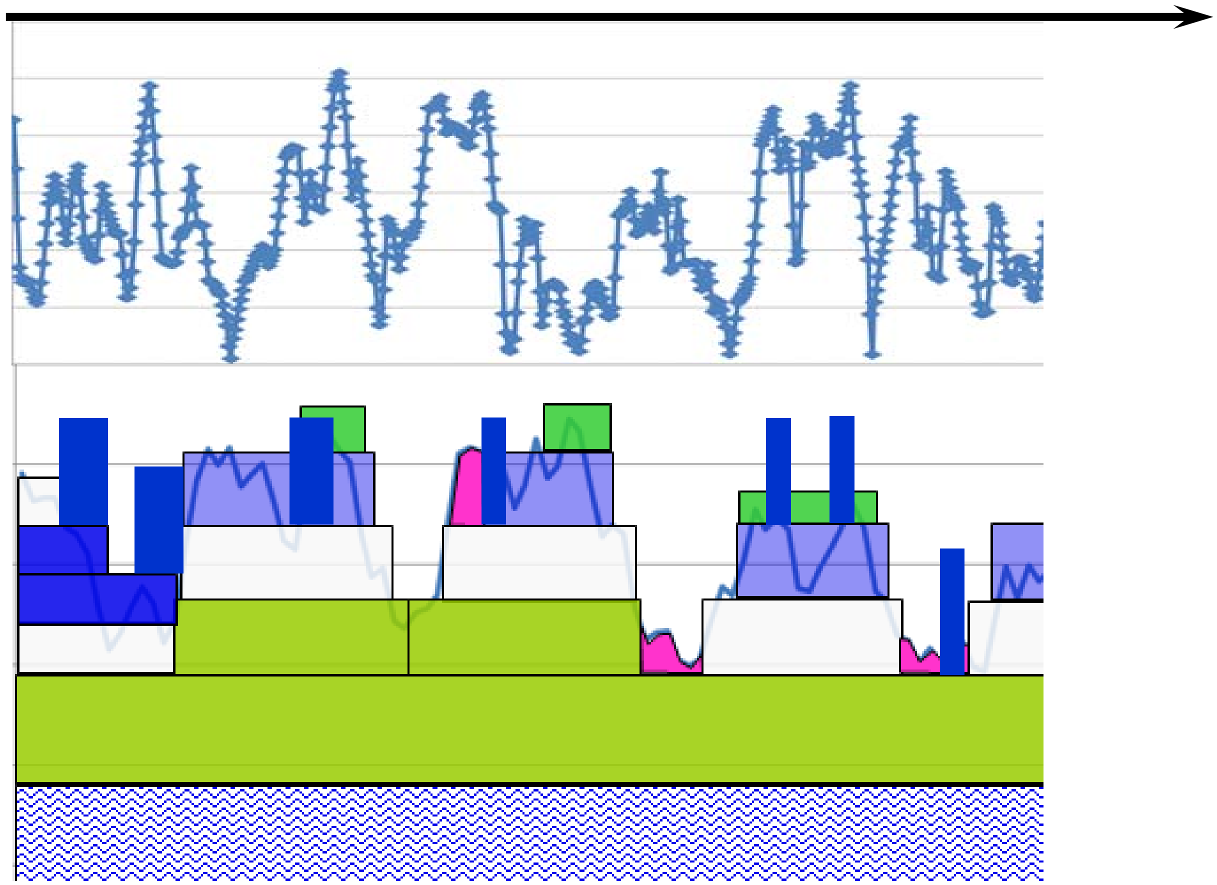

The stochastic unit commitment problem

- The value of optimizing *(p,q)*

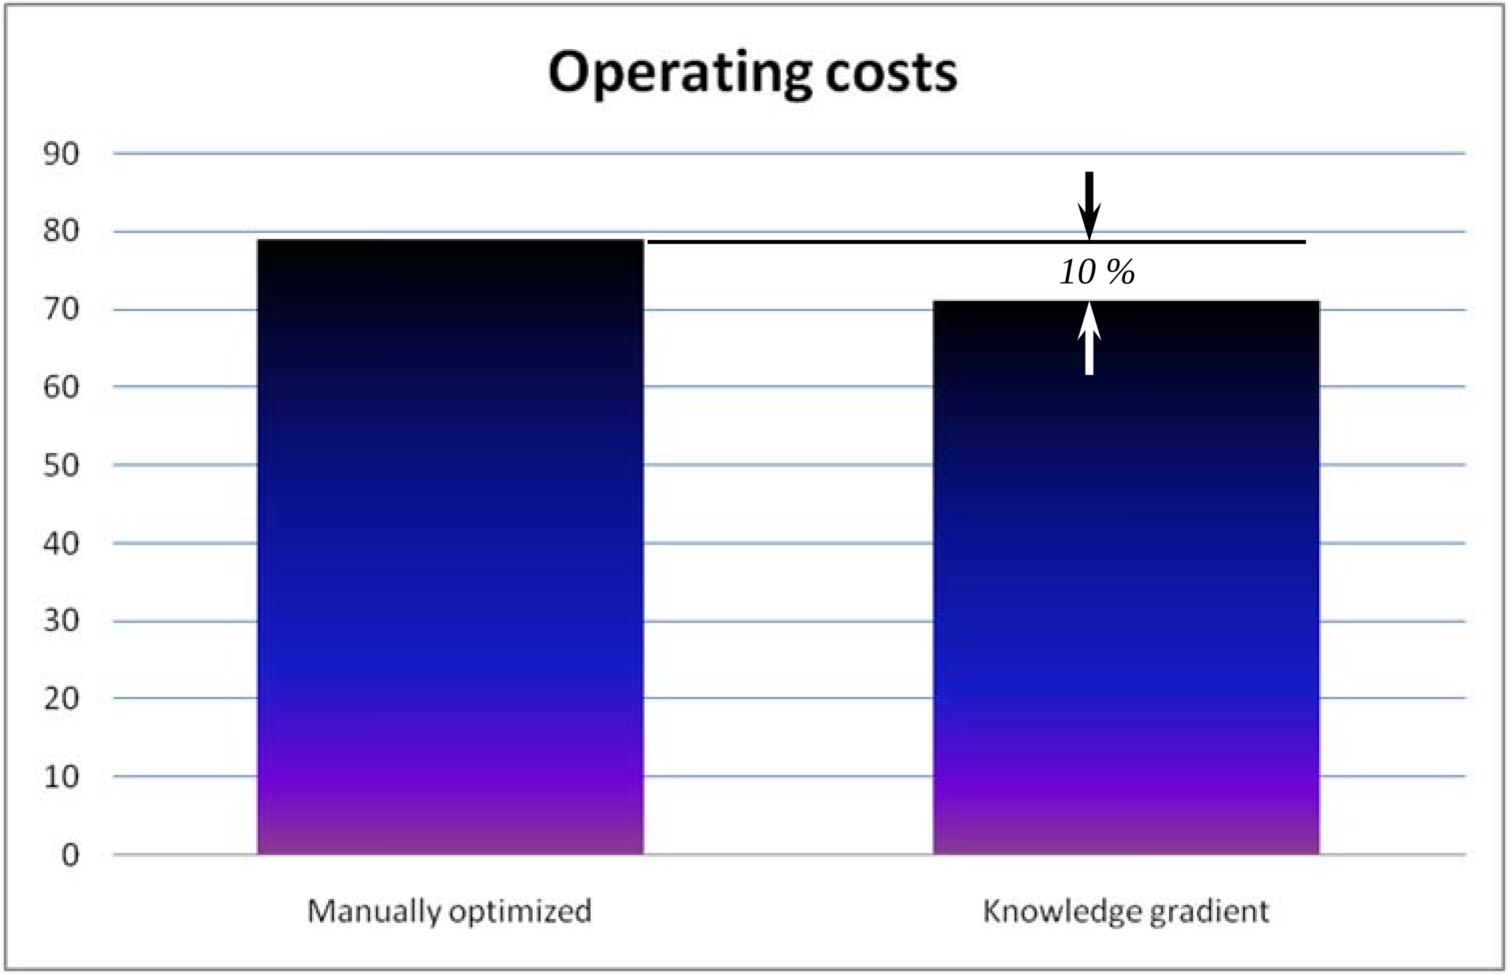

The stochastic unit commitment problem

- The effect of modeling uncertainty in wind

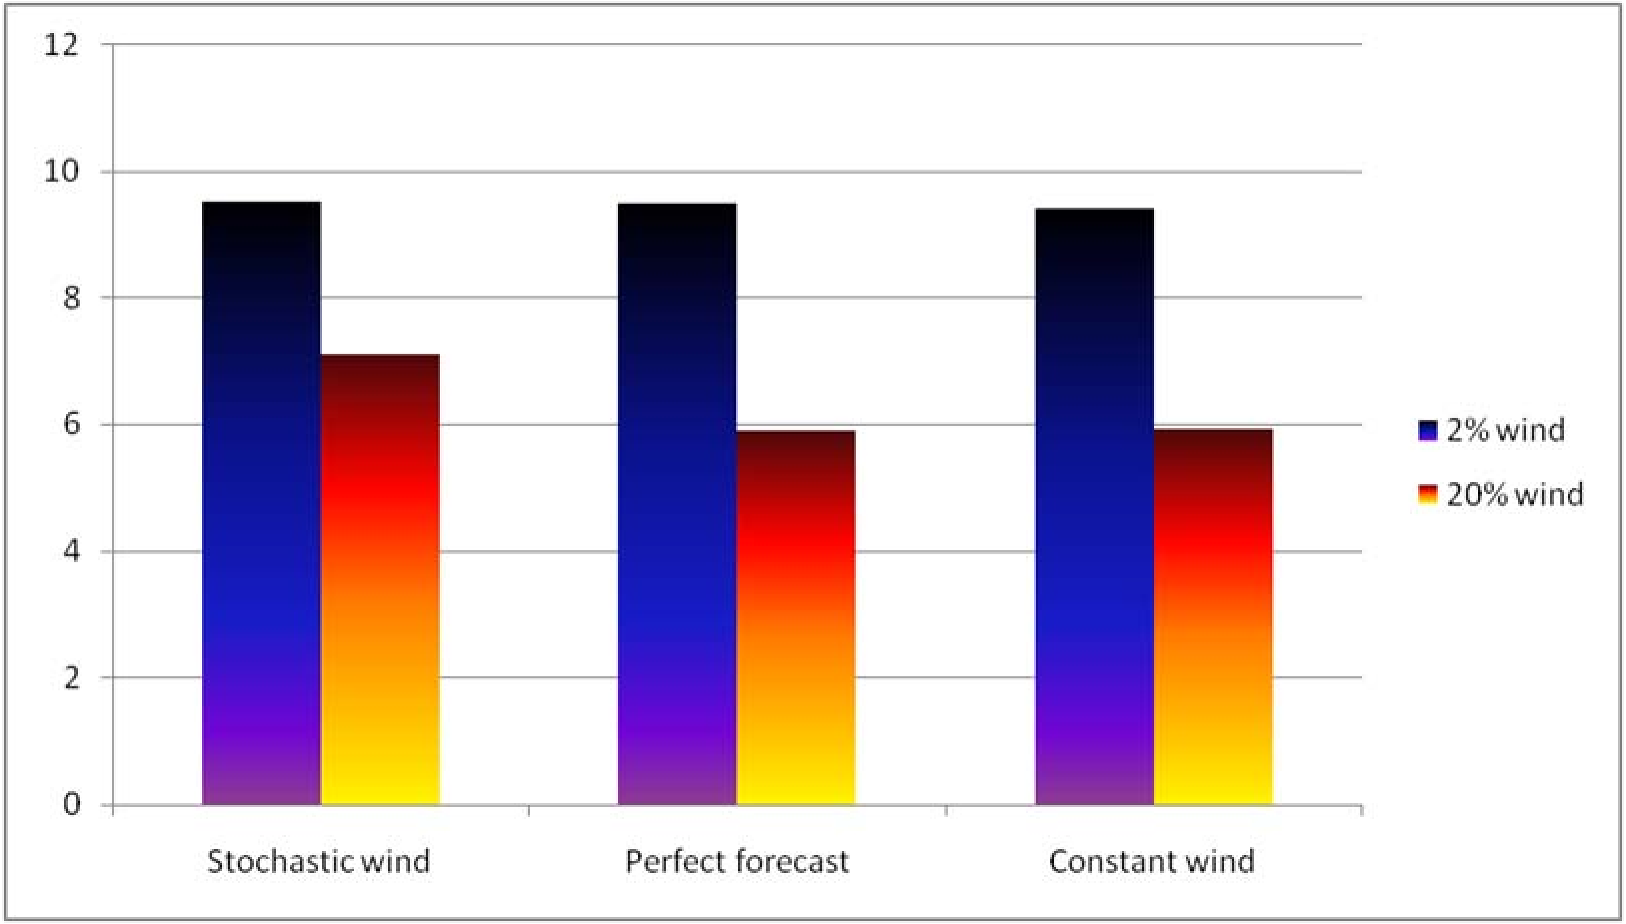

Conclusions

- The design of effective energy systems requires a careful understanding of how they are operated under realistic assumptions about what you know and when you know it.
- Accurate models of energy systems requires understanding not just the flow of ***physical systems***, but also the flow of ***information.***
- Making good decisions under uncertainty ("stochastic optimization") requires a balance of art (designing policies) and science (tuning policies).