



Predicting Demand in Computational Markets

Thomas Sandholm, Kevin Lai

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Motivation

Economic resource allocation offers many advantages

Important to predict

- application's resource demands

- other application's resource demands

- income

Consequences of poor prediction

- starvation, now or in the future

- over-saving → recession

- irrational behavior → mechanism disruption

Harder than Physical Markets

Fully virtualized markets are highly dynamic

- high transaction rate

- large scale

- low reaction time

Computation cannot be saved

- over-provisioned resources expire quickly

Problem



Given historical data about a resource's price, predict its future prices

Algorithms

Normal approximation

+ simple, compact model

Maximum entropy

+ simple, compact model, handles skewness

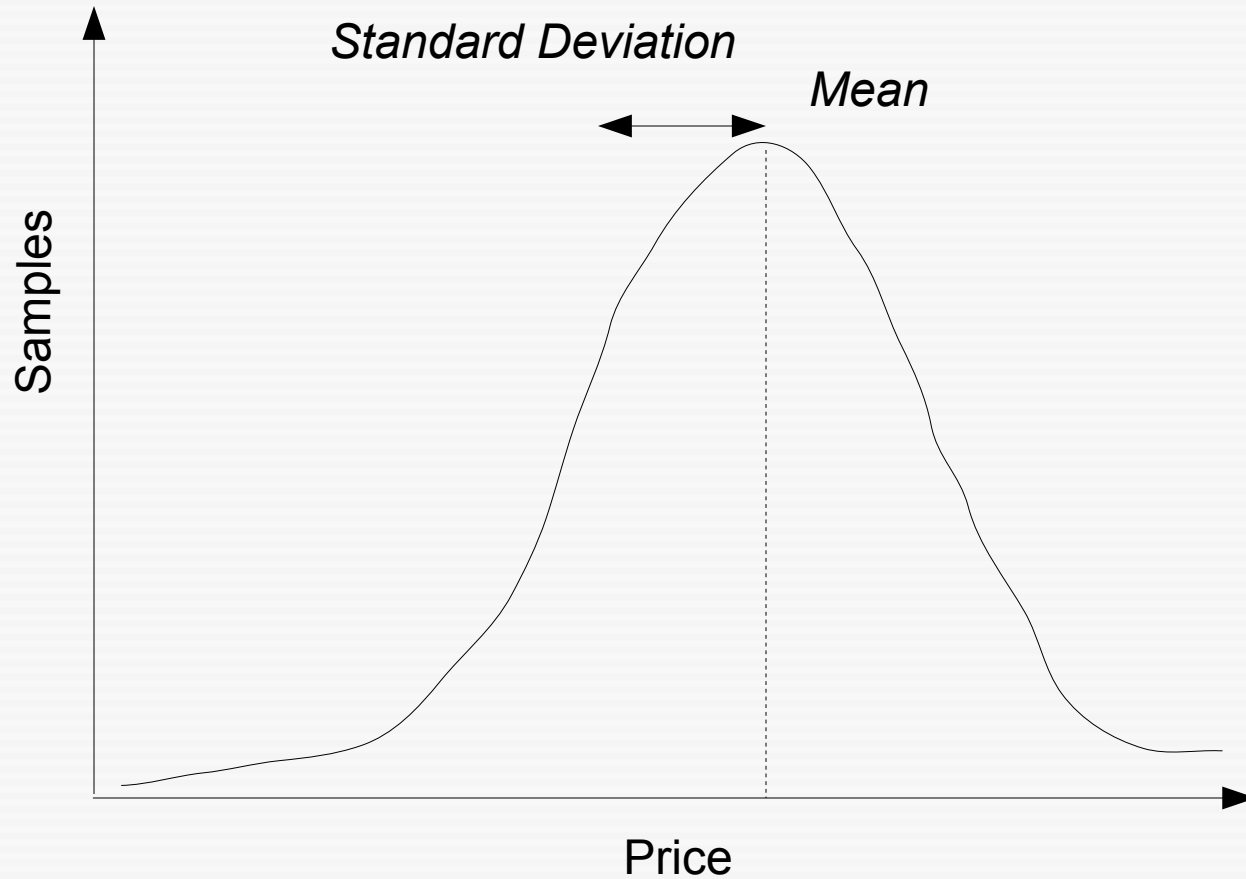
Histogram

+ multi-modal, handles asymmetry


Auto-regression

+ handles trends, repeating patterns

Normal Approximation



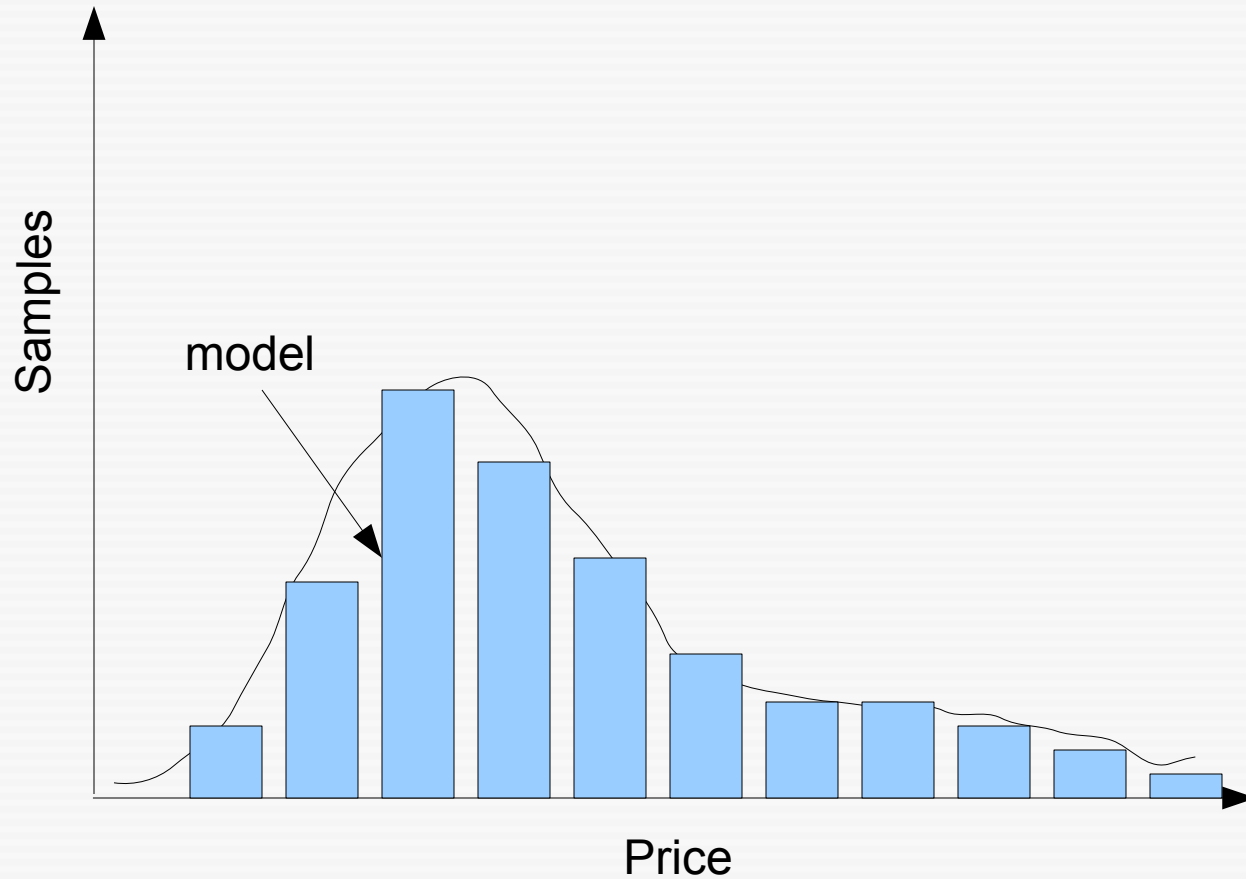
Maximum Entropy



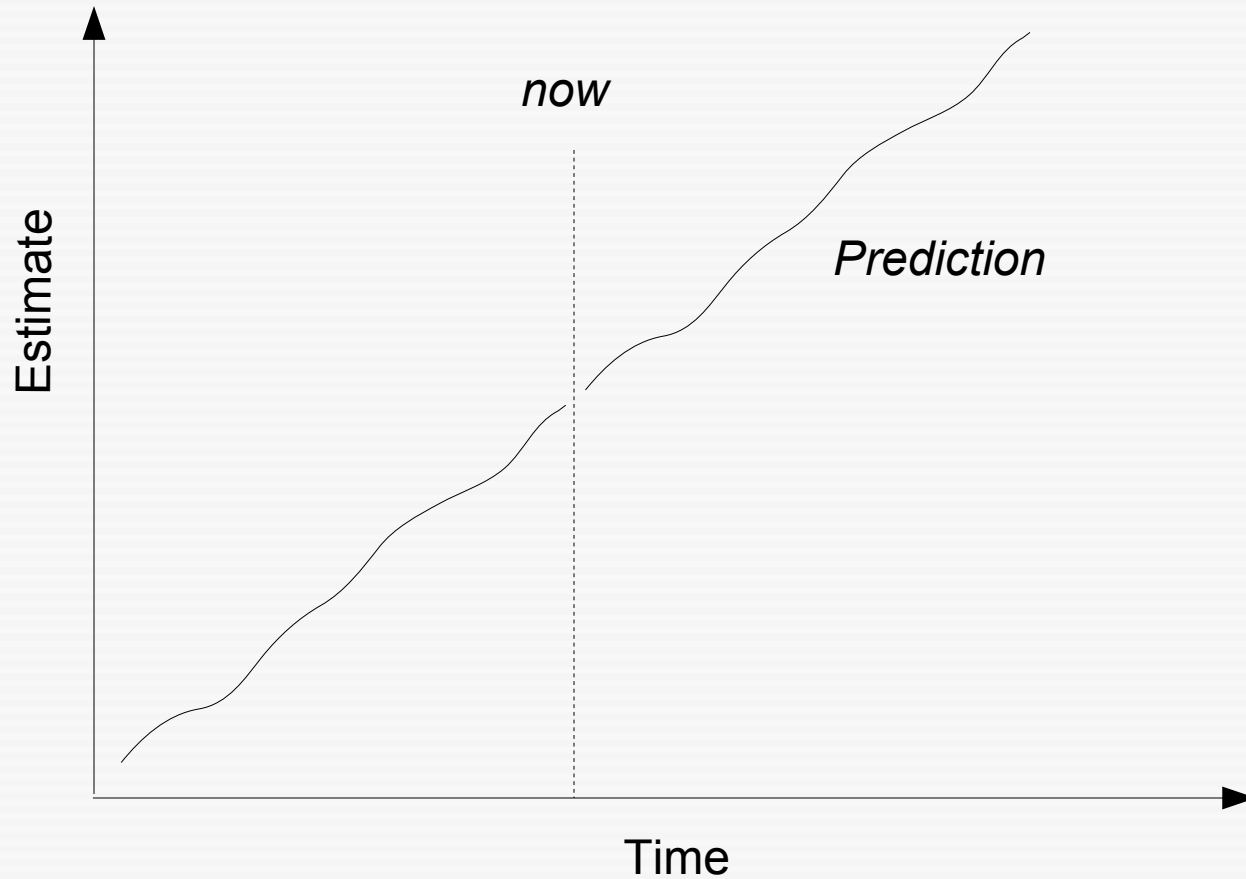
Histogram



Histogram



Auto-regression



Data Set



PlanetLab usage

- + highly dynamic (load changes second by second)
- + large scale (600 machines, ~600 users)
- +/- users do not specify price

Results

Normal approximation

inaccurate because of significant skewness

Maximum entropy

inaccurate in estimating distribution
accurate in estimating bounds

Histogram

inaccurate at low model sizes

Auto-regression

very high overhead, very inaccurate
trends are too short to be captured

Predicting Upper Bounds

$$P(\text{price} < \$1) = 90\%$$

Provider pricing

give tiered plans based on probability

User bidding

choose desired degree of predictability
bid as appropriate