

On the Value of Dynamic Pricing

7th Annual INFORMS Revenue Management and Pricing Section Conference
June 28 & 29, 2007

Miguel Sousa Lobo

Duke University, Fuqua School of Business, m1obo@duke.edu

Is there value in dynamic pricing?

Empirical observation:

Widely used in some industries, non-existent in others.

Is it only due to institutional inertia, business constraints (operational, advertising, legal)?

Literature:

Different models, wide difference in reported value (from numerical simulations).

What does 'dynamic pricing' mean? Change selling price...

... for reasons other than uncertainty about the demand model:

- Deterministic: Seasonal pricing, product lifecycle management, skimming (segmentation over time), penetration pricing
- Pricing for capacity: To deal with variability in demand

... to deal with uncertainty about the demand model:

- Model learning: Bayesian updating
- Price exploration: Sacrifice short-term revenue for informativeness

... by mechanisms other than posted price (auctions, etc.)

Dynamic pricing here:

Changing a posted price over time
to deal with uncertainty in demand.

This uncertainty can be either

period-specific uncertainty (change price to manage capacity), or
model uncertainty (change price due to learning or in order to learn).

Outline

- Model, problem typology
 - Capacitated vs. uncapacitated
 - Magnitude uncertainty vs. elasticity uncertainty
- Learning, informativeness, value of information, bounds
- Policy recommendations/predictions for each of the four cases

Uncapacitated

Period $k = 1, 2, \dots, T$ (here $T = 2$)

Price p_k

Demand X_k

$$\text{maximize } \mathbf{E} \sum_{k=1}^T p_k X_k \quad (\text{over adapted } p_k)$$

Capacitated

Inventory on-hand Y_k (initial inventory Y_1)

$$\text{maximize } \mathbf{E} \sum_{k=1}^T p_k (X_k \wedge Y_k) \quad (\text{over adapted } p_k)$$

$$\text{where } Y_{k+1} = Y_k - (X_k \wedge Y_k)$$

Model uncertainty and period-specific uncertainty

Model uncertainty, $\theta : \Omega \rightarrow \Theta$ ($\Theta = \{1, 2\}$ for two realizations)

Period-specific uncertainty, ξ_k

$X_k : \Omega \rightarrow \mathbf{N}_0$, measurable $\sigma(p_k, \theta, \xi_k)$

X_k auto-correlated through model uncertainty

Conditional on model realization, demand is independent over time

Period-Specific Uncertainty

Here: $X_k | p_k, \theta \sim \text{Poisson}(\lambda_\theta(p_k))$

$$\lambda_\theta(p_k) = \mathbf{E}_{\xi_k}(X_k | p_k, \theta)$$

$$\lambda(p_k) = \mathbf{E}_\theta \lambda_\theta(p_k)$$

Model Uncertainty

Linear (or locally linearized) response to price

$$\lambda_{\theta}(p_k) = b_{\theta} - a_{\theta}p_k.$$

Two-parameter model, need to consider joint distribution of a_{θ} and b_{θ} .

Are all directions for the uncertainty in the vector (a, b) equivalent?

Simplest model that preserves key features:

Two time periods.

Two equally-probable model realizations, each linear,

$$\lambda_1 = b_1 - a_1p \text{ and } \lambda_2 = b_2 - a_2p.$$

How does shape of joint distribution impact:

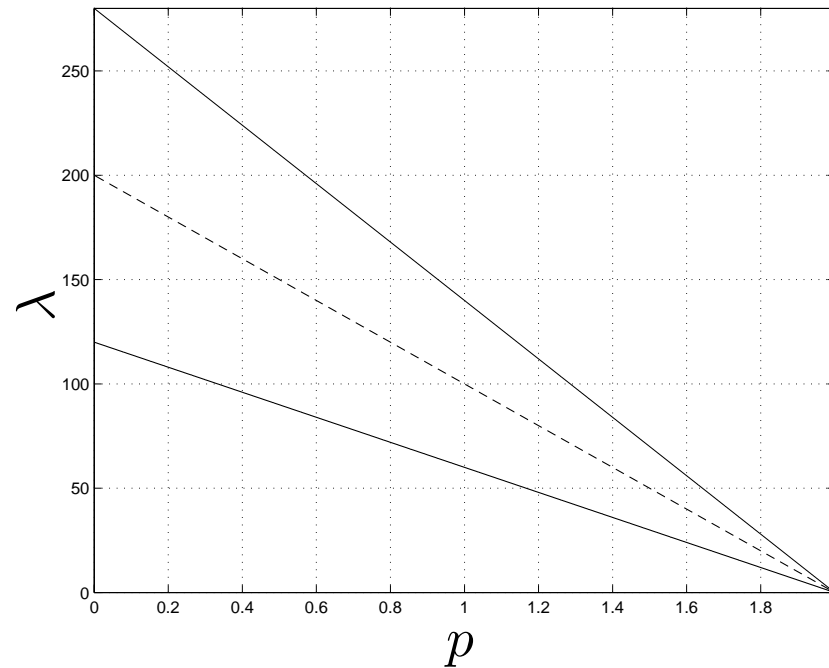
The cost and ability to learn about θ (cost in immediate revenue)?

The value of learning about θ (reward in future revenue)?

And how does this interact with capacitated vs. uncapacitated?

Magnitude, or multiplicative uncertainty $\lambda = (b - ap)(1 \pm \varepsilon)$

Consumer response to price is known, uncertainty about market size.

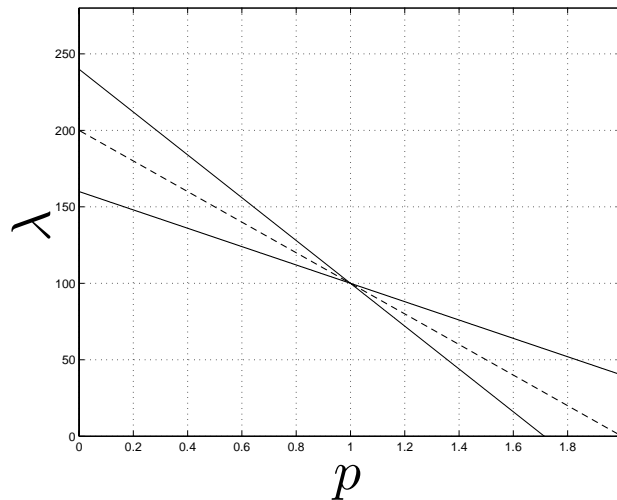


Number of potential customers is uncertain.

Distribution of willingness-to-pay is known (uniform for linear).

Elasticity uncertainty $\lambda = \lambda_0 - a(1 \pm \varepsilon)(p - p_0)$

Mean demand is well known at some historical price (assumed to be the myopic-optimal price), response to price changes is uncertain.



Common structure:

Uncertainty about both market size and consumer response
 $\Rightarrow \sim$ two-dimensional uncertainty

Uncertainty at a given price removed by repeated pricing at that level
 $\Rightarrow \sim$ one-dimensional uncertainty

Learning

θ is the common uncertainty accross periods

Prior model probability: $\mathbf{P}(\theta = 1) = \mathbf{P}(\theta = 2) = \frac{1}{2}$

Posterior model probability: $\psi(p, x) = \mathbf{P}(\theta = 1 | p_1 = p, X_1 = x)$

Bayesian updating:

$$\begin{aligned}\psi(p, x) &= \frac{\mathbf{P}(x|p, \theta = 1)\mathbf{P}(\theta = 1)}{\mathbf{P}(x|p, \theta = 1)\mathbf{P}(\theta = 1) + \mathbf{P}(x|p, \theta = 2)\mathbf{P}(\theta = 2)} \\ &= \frac{\frac{\lambda_1^x}{x!}e^{-\lambda_1}}{\frac{\lambda_1^x}{x!}e^{-\lambda_1} + \frac{\lambda_2^x}{x!}e^{-\lambda_2}}.\end{aligned}$$

Expectation of the posterior model probability

$$\mathbf{E}_X \psi(p, X) = \frac{1}{2}, \text{ for any } p$$

Variance of the posterior model probability

Prior measure of the expected informativeness of the observation.

Mean demand given price p under each model realization

$$\lambda_1(p) = \mathbf{E}_X(X|p, \theta = 1) \quad \lambda_2(p) = \mathbf{E}_X(X|p, \theta = 2)$$

Write $\lambda_1 = \lambda - \frac{1}{2}\delta\sqrt{\lambda}$ and $\lambda_2 = \lambda + \frac{1}{2}\delta\sqrt{\lambda}$

(with Poisson, δ is the ratio between the separation of the means and the average standard deviation: $\delta = \Delta\lambda/\sigma$)

$$\mathbf{Var}_X \psi(p, X) = \frac{1}{8} \sum_{x=0}^{\infty} \frac{\lambda^x}{x!} e^{-\lambda} \frac{\left(\left(1 - \frac{\delta}{2\sqrt{\lambda}}\right)^x e^{\frac{1}{2}\delta\sqrt{\lambda}} - \left(1 + \frac{\delta}{2\sqrt{\lambda}}\right)^x e^{-\frac{1}{2}\delta\sqrt{\lambda}} \right)^2}{\left(1 - \frac{\delta}{2\sqrt{\lambda}}\right)^x e^{\frac{1}{2}\delta\sqrt{\lambda}} + \left(1 + \frac{\delta}{2\sqrt{\lambda}}\right)^x e^{-\frac{1}{2}\delta\sqrt{\lambda}}}$$

with λ and δ functions of p

Bounds on $\text{Var}_X \psi(X)$

$$\frac{1}{4} \left(1 - e^{-\frac{1}{4}\delta^2}\right) \leq \text{Var}_X \psi(X) \leq \frac{1}{2} \left(\frac{1}{1 + e^{-\delta^2}} - \frac{1}{2}\right).$$

The upper bound is from lower limit on Poisson rate, highest variance for a given δ .

The lower bound is given by variance with Gaussian (upper limit on Poisson rate), then lower bound on Gaussian variance.

Can also do polynomial approximations and bounds:

$$\text{Var}_X \psi(X) = \frac{1}{8}\delta^2 - \frac{1}{24}\delta^4 + O(\delta^6)$$

Myopic Price

Uncapacitated: $p_k = \frac{\hat{b}_{k-1}}{2\hat{a}_{k-1}}$

Capacitated: $p_k = \frac{\hat{b}_{k-1} - Y_k / (T - k + 1)}{\hat{a}_{k-1}}$

Expected Cost of Deviation from Myopic Price in First Period

$$\Delta R_1 = -\hat{a} (\Delta p_1)^2, \quad \text{where} \quad p_1 = \frac{\hat{b}}{2\hat{a}} + \Delta p_1$$

Expected Value of Information in Second Period

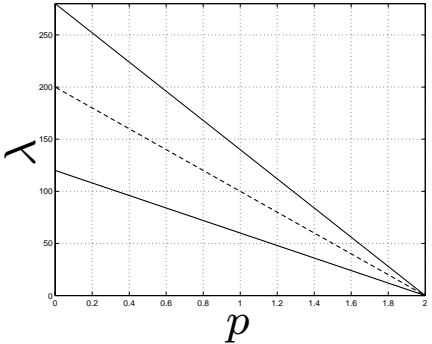
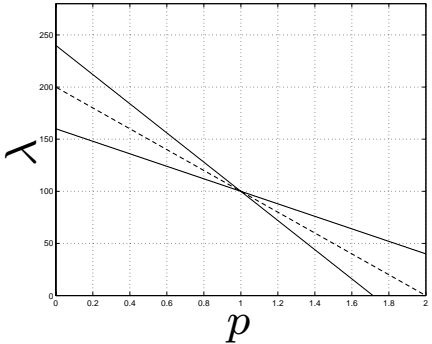
For each case: compute expected increase in second-period revenue as a function of first-period price informativeness.

Example: for uncapacitated w/ elasticity uncertainty, the expected second-period gain as a function of first-period price informativeness is

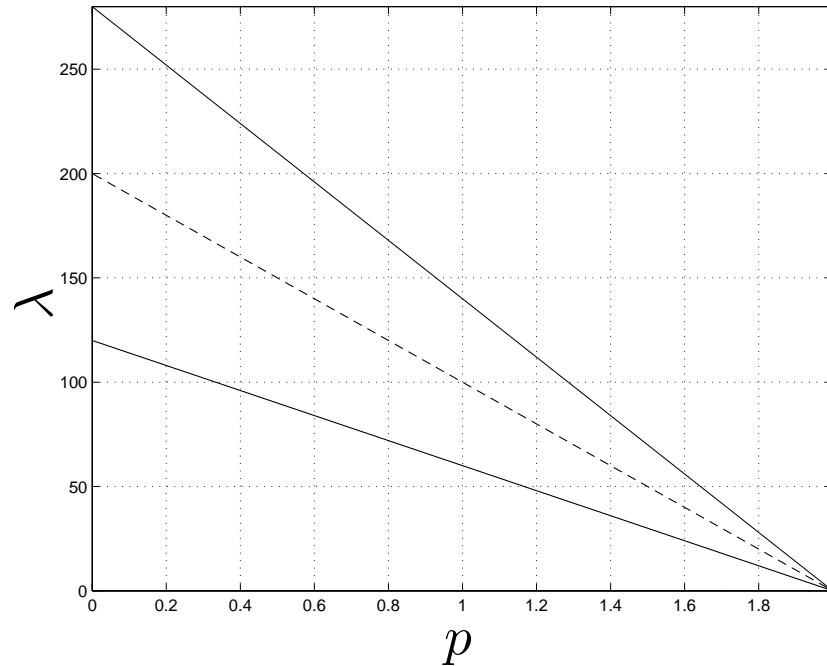
$$\frac{\Delta R_i}{\mathbf{E}R_0} = \left(\frac{\Delta a}{2}\right)^2 \mathbf{Var} \psi + \text{h.o.t.}$$

more precisely:
$$\frac{1}{1 - \left(\frac{\Delta a}{2}\right)^2} \leq \frac{\Delta R_i / \mathbf{E}R_0}{\left(\frac{\Delta a}{2}\right)^2 \mathbf{Var} \psi} \leq \frac{1}{\left(1 - \left(\frac{\Delta a}{2}\right)^2\right)^2}$$

Four cases

	<p>Magnitude uncertainty</p> 	<p>Elasticity uncertainty</p> 
<p>Uncapacitated $\max. \mathbf{E} \sum_{k=1}^T p_k X_k$</p>	<p>Case 1</p>	<p>Case 2</p>
<p>Capacitated $\max. \mathbf{E} \sum_{k=1}^T p_k (X_k \wedge Y_k)$</p>	<p>Case 3</p>	<p>Case 4</p>

Case 1: Uncapacitated w/ Magnitude Uncertainty



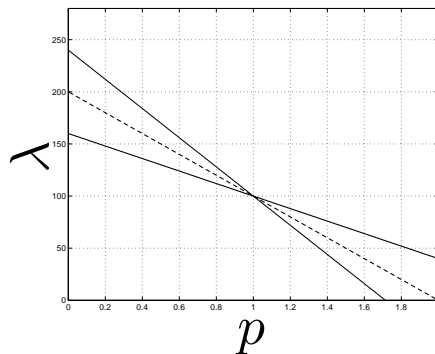
$$p_{\theta=1}^* = p_{\theta=2}^* = \frac{1}{2} b_1/a_1 = \frac{1}{2} b_2/a_2$$

Optimal price independent of $\theta \Rightarrow$ Information has no value

Uncapacitated \Rightarrow No need to adjust price to manage capacity

Static price is optimal

Case 2: Uncapacitated w/ Elasticity Uncertainty



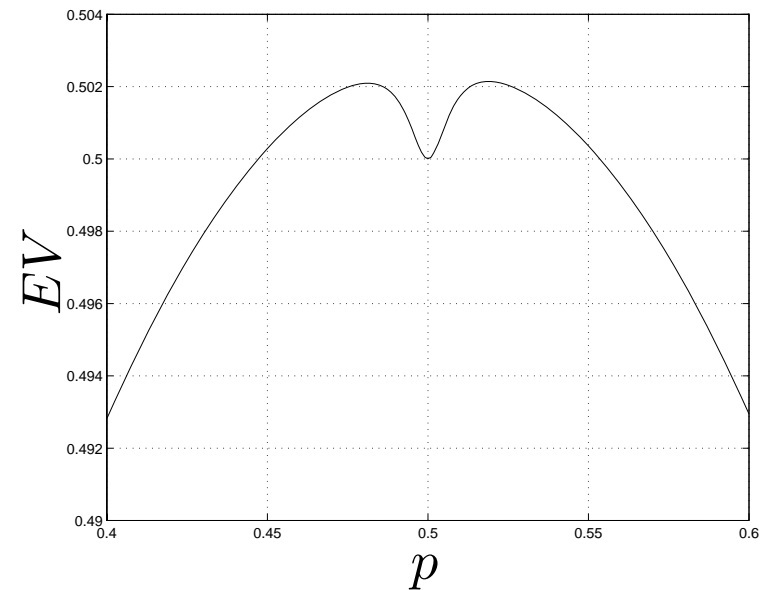
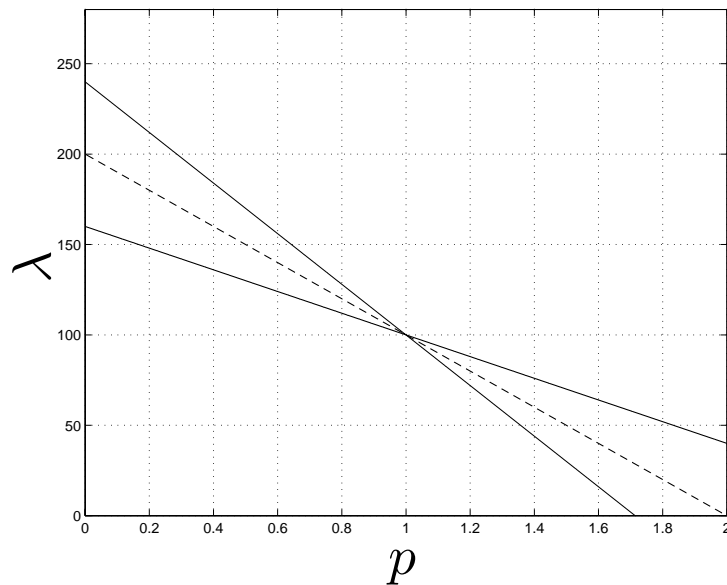
Assume historical pricing is consistent with beliefs
 \Rightarrow Confounding price equals myopic-optimal price

$$\lambda = q_0 - \frac{q_0}{p_0} (1 \pm \varepsilon)(p - p_0)$$

the price elasticity of demand is $e = \frac{dq/q}{dp/p} = -\frac{q_0}{p_0} \frac{p}{q}$

so that at p_0 we have $e = -1$.

Exploration/exploitation tradeoff



No learning at confounding price.

Deviation from myopic price allows for learning,
which allows for more accurate pricing in second period,
which leads to higher expected revenue.

Too large a deviation: better learning, but too high cost in first period,
overall loss in expected revenue.

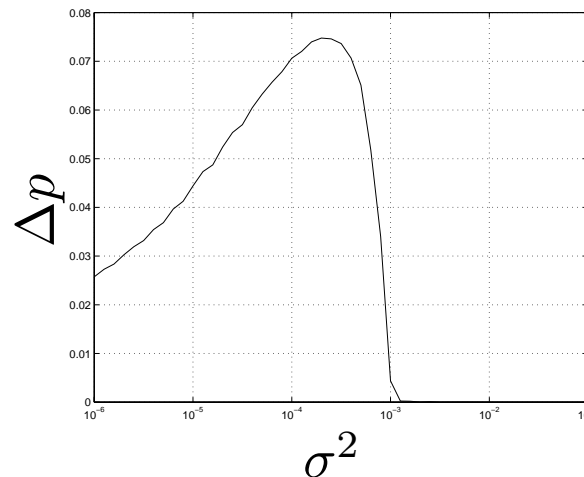
By comparing lower-order terms in Δp_1 (and using concavity of higher-order terms), we find it is optimal to deviate from static price when model uncertainty is high and period-specific variance is low:

$$(\Delta e)^2 > \sqrt{2} p^* \left(\frac{\sigma}{\lambda} \right)^3$$

with $\Delta e = \Delta a / (2a)$ the uncertainty in the price elasticity of demand.

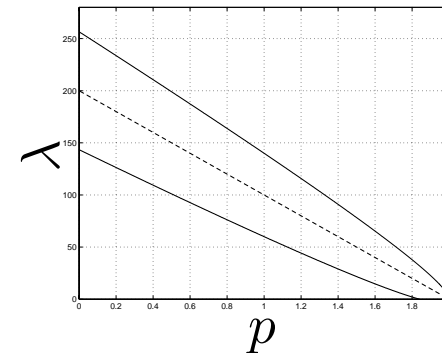
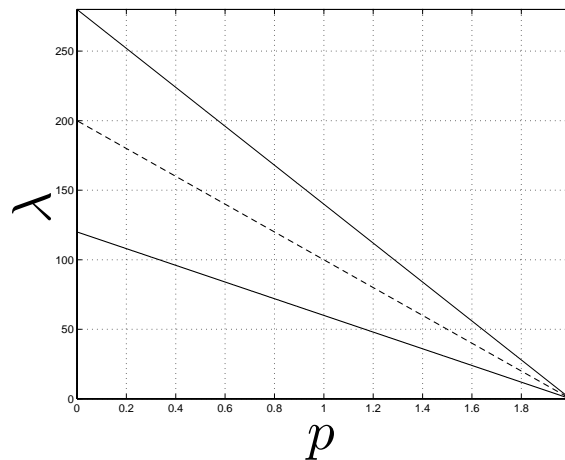
Approximations and bounds for:

- Optimal price exploration (magnitude of deviation from myopic)



- Expected gain from dynamic pricing, etc.

Case 3: Capacitated w/ Magnitude Uncertainty



Changing the price results in small changes in δ

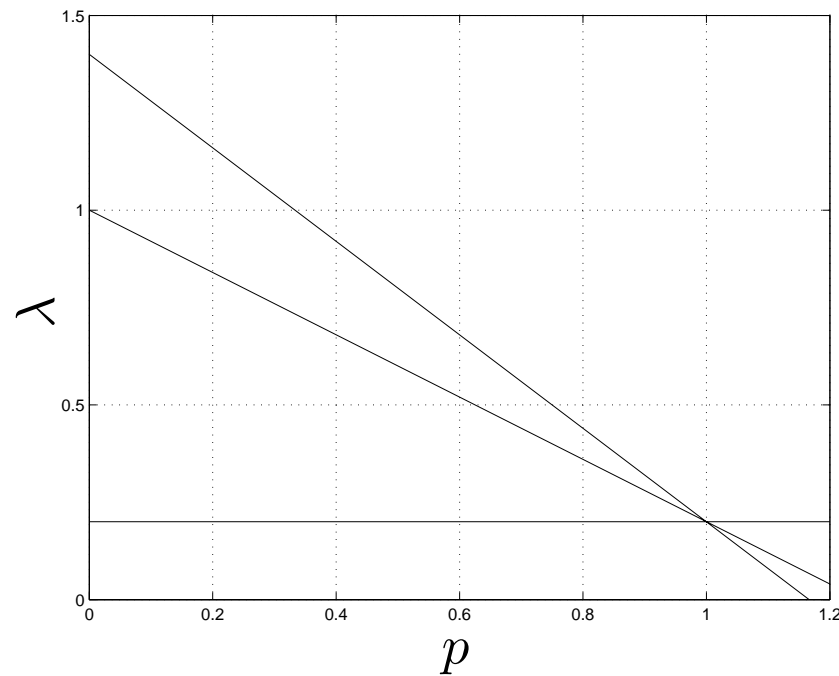
⇒ Deviations from myopic price provide minimal gain in learning.

Deviations have significant cost ⇒ Myopic price is nearly optimal.

In 2nd period: update model estimate, price for remaining capacity.

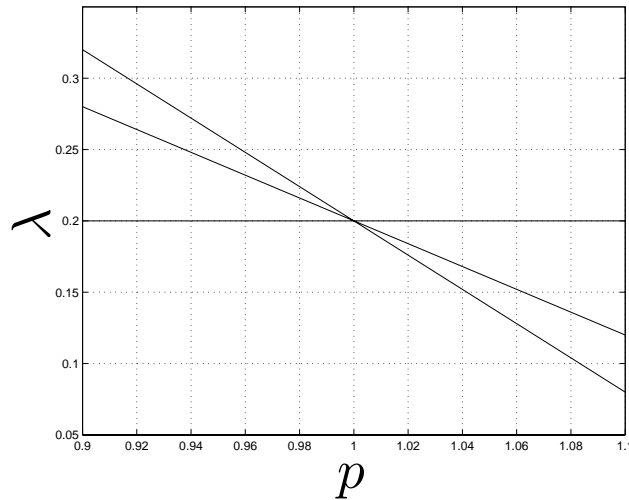
(Exact with modified model where $\delta(p) = \frac{\Delta\lambda}{\sqrt{\lambda}}$ is constant.)

Case 4: Capacitated w/ Elasticity Uncertainty



Assume demand well-known at myopic price, $p = (Y/T - b)/a$.

Confounding quantity is initial capacity divided by number of time periods.



If period-specific variance is low, little value in learning

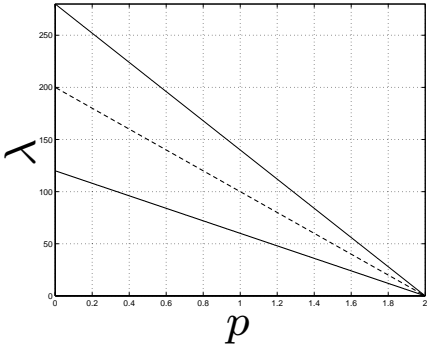
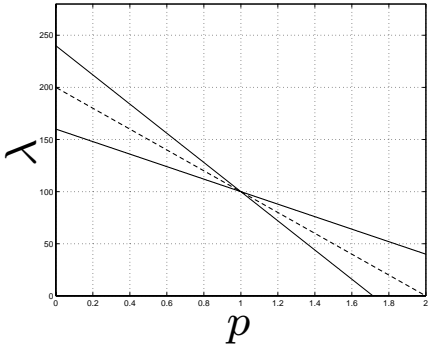
- Second period capacity will be near confounding quantity
- Model information has low value

If period-specific variance is high, learning is too costly

- λ is small, δ is small unless price deviation is large
- Learning requires large deviations, high cost

\Rightarrow Never optimal to deviate

Policy recommendations/predictions

	<p>Magnitude uncertainty</p> 	<p>Elasticity uncertainty</p> 
<p>Uncapacitated $\max. \mathbf{E} \sum_{k=1}^T p_k X_k$</p>	<p>Static pricing</p>	<p>Price exploration (active learning)</p>
<p>Capacitated $\max. \mathbf{E} \sum_{k=1}^T p_k (X_k \wedge Y_k)$</p>	<p>Pricing for capacity (passive learning)</p>	<p>Pricing for capacity (no learning)</p>

Other possible extensions

Optimal timing of the price move

1st period of length τ , 2nd period $T - \tau$ (τ chosen at $t = 0$)

Compute approximations, tight bounds for optimal τ

Example, uncapacitated with elasticity uncertainty:

as exploration becomes more valuable $\tau \rightarrow 0$,

as exploration becomes less valuable $\tau \rightarrow T/2$,

hence no impact on conditions for exploration to be of value

Multiple periods

Generally hard problem, maybe extend bounds for some cases?

Continuous time?

Information loss, model drift

Value of reserving capacity (right not to sell), or τ as a stopping time