Making Decisions from Data

Review

Terminology in quality control

– Capable versus in-control: distinct ideas, not mutually exclusive

– Trade-off: detection problem versus "false positives"

– Tracking averages to locate small deviations (p. 81)

Standard error

– SE (of something), SD(of something)
Typically reserve SE for SD of a "statistic" computed from data

– "Magic formula" – SE(avg of n items) = SD(1 item) / $\sqrt{n}$

– Less variation among averages than among individuals

– Estimating SE directly from collection of averages versus using n adjustment from SD of data

Control limits

	Process	**Sample**
– Process vs. sample features	μ, σ	Y-bar, s
– Limits for inspecting *one* shaft	$\mu \pm 3\sigma$	Y-bar ± 3 s
average of *five*	$\mu \pm 3\ \sigma/\sqrt{5}$	Y-bar ± 3 s/$\sqrt{5}$
average of *n*	$\mu \pm 3\ \sigma/\sqrt{n}$	Y-bar ± 3 s/$\sqrt{n}$

– Larger the number which are averaged, the less variation in the average from one to batch to the next.

– Consider how you would set the limits if you were tracking an average of all 100 (or all 400) and just had one sample to use to set the limits.

Using control charts

– Limits for tracking the SD of the process

– Importance of tracking both the mean and SD

– Review with *www.nyse.com* example for financial data.

Administrative Details

Reading

– Freedman et al. stories are great (polling, Ch 19 onward)

– Skim casebook and course pack prior to class

Assignment #1 this week

Key Application for Today

Finding the right balance of competing costs

– Pricing a textbook, using on-line re-pricing.

– Set the price too high: customers flee to competition

– Set the price too low: lost chance for higher profits

Definitions and Concepts

Elasiticity of demand (p 2-5 of course pack notes for this class)

– Price sensitive or price insensitive

– How do changes in price affect demand for a product?

Standard error

– As a measure of distance
Count the number of standard errors away from some contemplated reference value, and then employ the empirical rule.

– Associated use of normality
Use normality to describe the variation of an average

Discussion (lingering questions)

Why use averages in control charts?

– Normality is a better approximation for averages (CLT).

– Use of averages detects small changes faster.

– Avoid some of the problems with too many outside by chance alone. (i.e., reduce the problem of multiplicity)

Sampling variation

– How different might things be if I were to repeat it?

– How close would the average of my sample be to the average of another sample from the same population?

– Are results reproducible?

One sample versus many

– One way to approach the idea of a “sampling distribution”

– Conceptual approach, not the way to do things in practice

Why use n in the denominator of SE rather than n?

– SD is the square root of the variance.

– The divisor is n in a variance. We get a square root in the SE or SD(mean) when we convert back to the scale of the data.

Dynamic Pricing Experiment

Current conditions

– Sell a textbook for $50. Cost is $35, so profit $15/book sold.

– 21 out of 100 possible customers purchase the book.

Question

– Should the price be raised by 10% to $55?

Key unknown: elasticity of demand

– Elasticity controls how quantity sold reacts to changes in price.

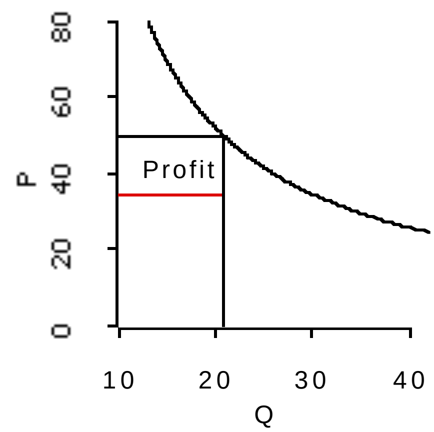

Maximize profits (area of rectangle)

Net Profit (price) = (quantity sold at price) (price – cost)

$$N(p) = Q_p\,(p - c)$$

Number sold depends on the price. If elastiticity is *e*, then quantity is often modeled as

$$Q_p = k\,P^e$$

Take derivatives of net profits with respect to the price, obtaining

$$dN(p)/dp = (e+1)\,k\,P^e - e\,k\,P^{e-1}\,c$$

To find the optimal, set the derivative to zero and solve for optimal price

$$p_{opt} = c\,e/(1+e)$$

Current price implies believe elastitity to be –10/3. Suppose it's more?

Some Alternatives

Current conditions

– Priced at $50 with 21/100 purchasing.

– Net profit expected per day for this one book is ($50-$35) 21 = **$315**.

Two hypotheses, two actions

– Suppose that we raise the price to $55. What will happen to sales?

Market Condition	New Demand	Profits	Action
Competitive H_c	14/100	14($55-35)=$280	Keep $50 price
Monopolist H_m	20/100	20($55-35)=$400	Raise price to $55

– How does one decide between the alternatives, especially when the same type of decision will be made for many products.

An Experiment

Gather some data

– Show book to sample of 100 customers at new $55 price.

– Find the proportion that are willing to purchase, the sample proportion.

Standard error of a proportion

– Proportions are averages!

– Formula: SE depends on the underlying proportion in the population

$$SE(\hat{p}) = \frac{\sigma}{\sqrt{n}} = \frac{\sqrt{p(1-p)}}{\sqrt{n}}$$

– Can we distinguish a monopolist market from a competitive market?
Draw the picture…

Finding a Decision Rule

Opportunity costs

– What might happen from the experiment?

	True State of Demand	
Decide	Competitive p = 0.14	Monopolist p=0.20
Competitive H_c		Don't raise prices, lose chance to improve profit, Loss of **$85**/day
Monopolist H_m	Raise prices, cause demand to fall. Loss of **$35**/day	

Prior knowledge

– Experience with other books indicates that
450 out of 500 are competitive.

– Aside: Bayes rule and odds ratios: how do data affect prior odds?

Run a further experiment

– Make further use of the database of prior experiences.

– We know that 450 out of 500 are competitive.

– What would happen if we tried to build a rule from samples of 100 customers that considered these other, similar books?

– Relevant?
We want a rule for pricing many products, not just one. We will be using such a procedure frequently, not just once.

Experiment I

Test a decision rule

– Consider the rule that classifies a book as "competitive" if 17 or fewer out of 100 potential customers purchases the book.

– What would happen if we applied this rule to samples of customers of the books we know to be competitive or monopolist?

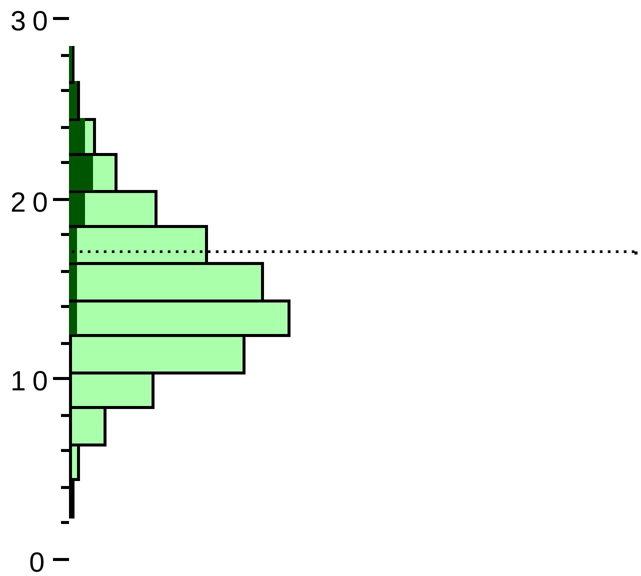

Errors

12 of the 50 monopolist are classified as competitive (dark below line)

Cost is 12($85) = $1020

72 of the 450 competitive are classified as monopolist (light above line)

Cost is 72($35) = $2520

Total opportunity cost of this rule is then **$3540**.

Best possible rule

– Can only shrink one type of error at the expense of raising the other type of error as the threshold moves.

– Optimal: Competitive if demand is 20/100 or smaller. Cost of $2565.

A better procedure

Use a larger sample to get a more accurate classification.

Experiment II

Larger sample

– Sample 400 customers of each of prior 500 books.

– Consider same rule based 17 or fewer.

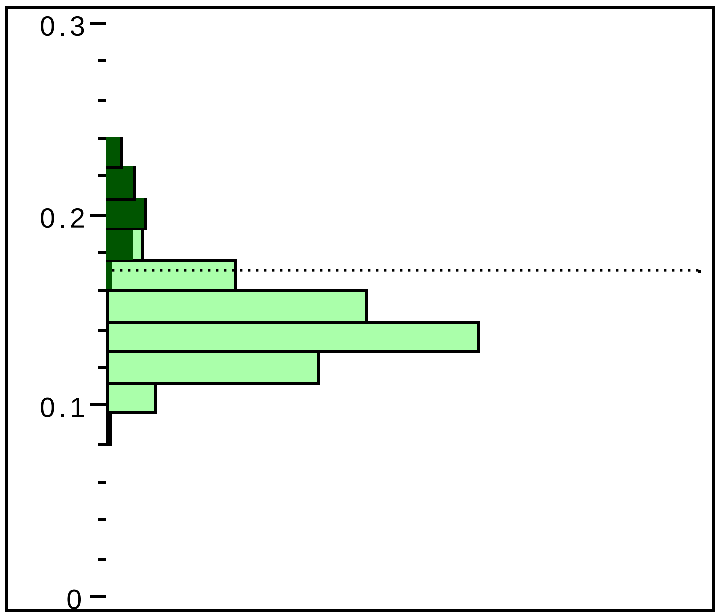

Benefit of larger sample

– Standard error is half of the size in the prior experiment, so less overlap of the two populations.

– Rule with threshold at 17.5 makes these errors:

3 monopolist classified as competitive	cost = 3($85) = $255
4 competitive classified as monopolist	cost = 4($35) = $140

so that the total cost is now $**395**

– Best cost using samples of 100 customers was $2565.

Role for models

– Use normal distribution to model variation of a proportion rather than sample from a historical data base.

Statistical Tests

Decision procedure

– Find a threshold such that if
sample proportion is above threshold –> monopolist H_m
sample proposition is below threshold –> competitive H_c

– Optimal threshold minimizes costs is *repeated* use of the decision rule.

Requirements for the procedure

– Simple world with two alternatives, p = 0.14 vs. p = 0.20.

– Prior experience that gives relative "belief" in these two states before we look at the data.

– Clear determination of costs of the errors.

– What does one do without these prerequisites?

Testing offers a "simplified" view of reality

– Choose a default decision, often the status quo.

– Only move from this state when data offer compelling evidence that this default choice is wrong.

– How much evidence is compelling, especially when one lacks a clear sense of the costs that are involved.

Null and alternative hypotheses, Type I and Type II errors

– In the book-selling example, the relative counts of the two books and the associated total costs for errors suggest choosing as a null hypothesis that the book is competitive.

– Still have both types of errors (calling competitive book monopolist and vice versa).

– These errors in testing are called type I and type II errors with the associated chances labeled α and β.