Demand Modeling

UCSD MGT 100 Week 3

Kenneth C. Wilbur and Dan Yavorsky

Demand Curves

- Theory
- Challenges
- What firms do

The elasticities, e and η, call for more extended explanation. In a general way these terms have reference to the responsiveness of buyers or sellers to changes in price. As an economic concept the picture is that of a "market" made up of an indefinite number of competing buyers and an indefinite number of competing sellers, the latter holding in their possession an indefinite quantity of a certain article. Under the concept of *demand* it is believed that if at a given instant of time the sellers had thrown on the market a definite portion of their stock, that portion would all have been taken at a certain definite price. If, however, *at the same instant,* they had offered more, the price would have been less, and, if less, the price would have been more. That is, for an offering of any portion of the stock there is, *at that instant,* a definite price at which that portion will be absorbed. Likewise, with reference to *supply* it is supposed that at a given price a definite quantity will be forthcoming from sellers. If, at that instant, the price had been higher more would have been forthcoming; if lower, less. To give mathematical definiteness to the concept, the coefficient of elasticity may be defined as the ratio of the percentage change in quantity to the percentage change in price, and may be represented by the expression:

$$\begin{matrix} e \text{ [elasticity of supply]}^{3} \text{ or} \\ \eta \text{ [elasticity of demand]}^{3} \end{matrix} = \frac{\dfrac{\Delta x}{x}}{\dfrac{\Delta y}{y}}$$

Since, with reference to supply, an increase in price is accompanied by an increase in quantity, $\frac{\Delta x}{x}$ and $\frac{\Delta y}{y}$

[3] In this expression Δx (read delta x) means the increase or decrease in quantity; Δy, the increase or decrease in price.

will have the same sign and hence e will be positive. With reference to demand, however, since at an increase in price a smaller quantity will be taken, $\frac{\Delta x}{x}$ and $\frac{\Delta y}{y}$ will have opposite signs and hence η will be negative.

From what has been said it is obvious that supply conditions at any instant of time may be represented by

FIGURE 1. TYPICAL SUPPLY AND DEMAND CURVES.

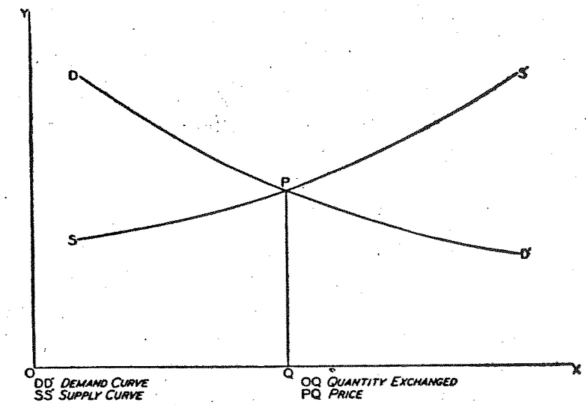

an ascending, and demand conditions by a descending, curve. The point where the curves intersect determines the price and quantity exchanged at that instant. This is shown in Figure 1.

The economic concept of elasticity supposes different experiments with prices in the same market at a single instant of time. Obviously such experiments cannot be made. Actual observations must be made at different times and during the period between observations conditions both of supply and demand may change. Indeed,

Inverse Demand Curve

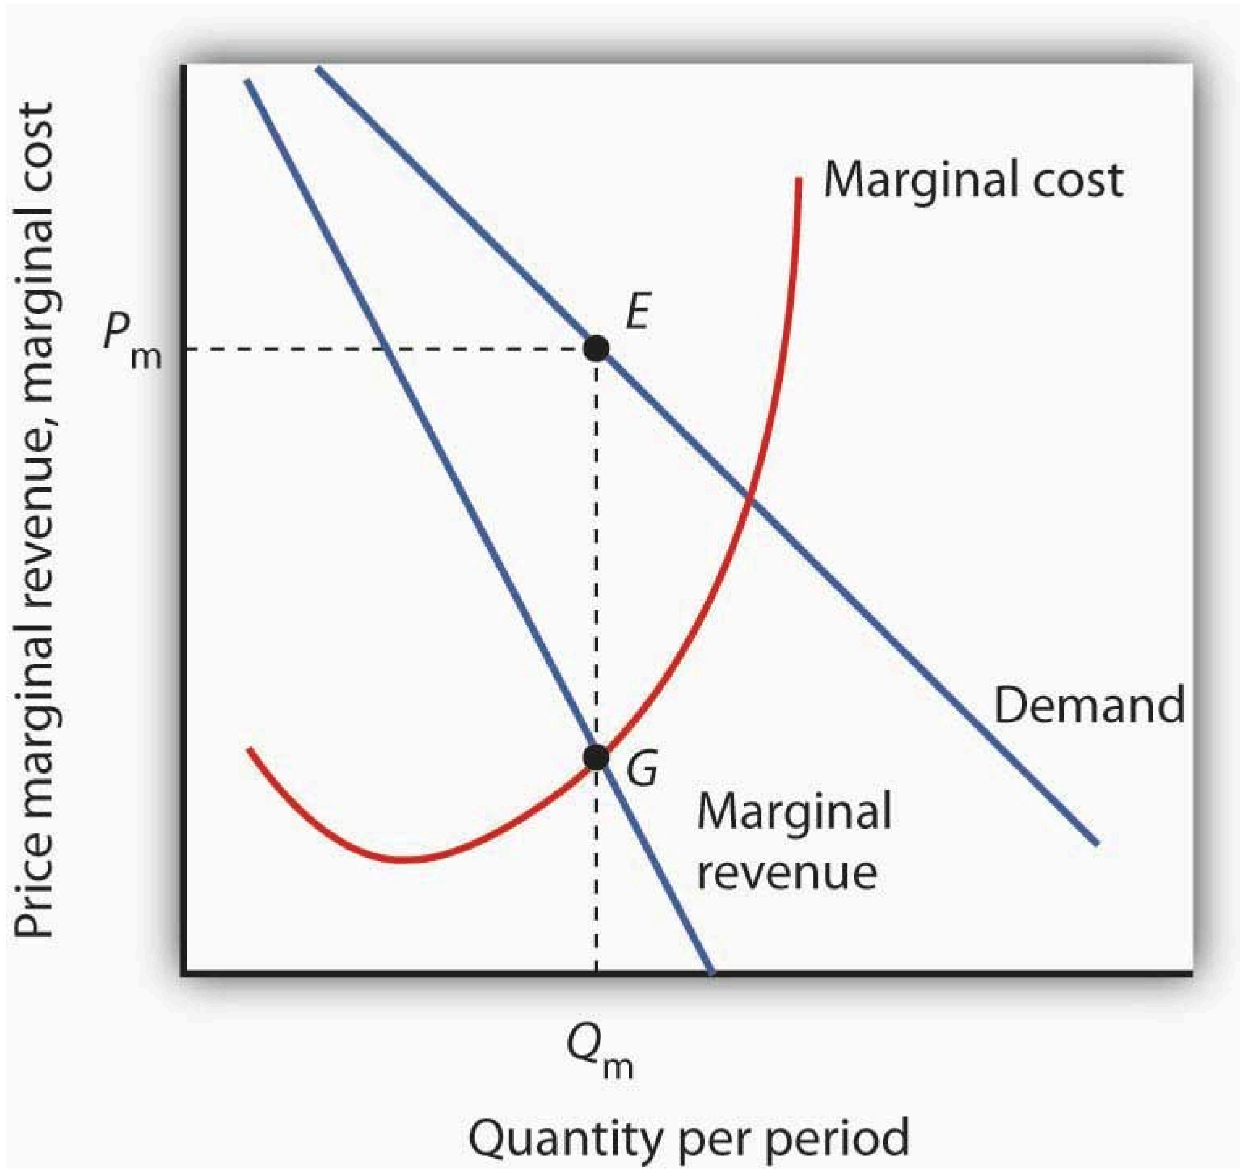

Demand Curves: Theory

- Useful theoretical concept that summarizes market response to price

```
   - Why is it useful?
```

- Often taught with perfect competition

```
   - Typically assumes stable, competitive, frictionless markets w free
entry, full information, no differentiation
   - Model predicts zero LR economic profits
   - Any evidence?
```

- Also, taught with monopoly i.e. market power

```
   - What is market power? How would we measure it?
```

Demand Curves: Challenges

- Idealized demand curves are hard to estimate (why?)
- Where do product attributes come in?

```
  - What if we don't observe all attributes? "Price endogeneity"
```

- Many things can predict demand

```
  - preferences, information, advertising, quality, match value, quality,
complements, substitutes, competitor prices, entry, taxes and other
policies, retail distribution, nature of equilibrium, stockpiling, consumer
income, ...
```

- What do we need to estimate Demand?

```
  - Observable, exogenous variation in costs or price
  - Otherwise, "price endogeneity" will bias demand estimates
```

How firms learn demand

- Market research
  - Conjoint analysis, customer interviews, simulated purchase environments
- Expert judgments, e.g. salesforce input
- Cost-driven price adjustments
  - Often one-sided
- Demand modeling with archival data
- Price experiments
  - Market tests, digital experiments, bandits, digital coupons
- Best practice: triangulation

Price Experiments

- Ideally, the best way to learn demand, because you create exogenous price variation; but…
  - Competitors & consumers can observe price variation
  - May change purchase timing, stockpiling or reference prices
  - Competitors, distribution partners or suppliers may react
- Hence, experimenting to learn demand may change future demand

Demand Modeling: Pros

- Relatively inexpensive for large organizations
- Fully compatible with price experiments
  - `Either/or framing would be a false dichotomy: "Yes-and"`
- Confidential, fast
- Depends on real consumer choices, i.e. revealed preferences
  - `As opposed to stated preferences`
- Enables demand predictions at counterfactual prices
- Enables predictions of competitor pricing response
- Prediction accuracy: evaluable after price changes

Demand Modeling: Cons

- Requires data, exogenous price variation, time, effort, training, commitment, trust, organizational buy-in
- Always subject to untestable modeling assumptions
- Requires the near future to resemble the recent past
    - To be fair, all predictive analytic techniques require these 3

Do Demand Models Work?

Multinomial Logit

- history, math, properties, discussion

Multinomial Logit (MNL)

- Time tested, famous, popular demand model
- Ported to econ from stats & psych by McFadden (1974, 1978, 1981; “Logistic regression’’)

Table 1. Prediction Success Table, Journey-to-Work

(Pre-BART Model and Post-BART Choices)

Cell Counts	Predicted Choices				
Actual Choices	**Auto Alone**	**Carpool**	**Bus**	**BART**	**Total**
Auto Alone	255.1	79.1	28.5	15.2	378
Carpool	74.7	37.7	15.7	8.9	137
Bus	12.8	16.5	42.9	4.7	77
BART	9.8	11.1	6.9	11.2	39
Total	352.4	144.5	94.0	40.0	631
Predicted Share	55.8%	22.9%	14.9%	6.3%	
(Std. Error)	*(11.4%)*	*(10.7%)*	*(3.7%)*	*(2.5%)*	
Actual Share	59.9%	21.7%	12.2%	6.2%	

Multinomial Logit (MNL)

- Let i index consumers, $j = 1, \ldots, J$ products, and t index choice occasions
- Assume each i gets indirect utility u_{ijt} from product j in market t:

$$u_{ijt} = x_{jt}\beta - \alpha p_{jt} + \epsilon_{ijt}$$

- Then, assuming each i picks the j that maximizes u_{ijt}, has unit demand, and that ϵ_{ijt} ~*i.i.d.*$EV_1(0,1)$, market share is

$$Prob.\{u_{ijt} > u_{ikt} \forall k \neq j\} \equiv s_{jt} = \frac{e^{x_{jt}\beta - \alpha p_{jt}}}{\sum_{k=1}^{J} e^{x_{kt}\beta - \alpha p_{kt}}}$$

With N_t consumers, $q_{jt}(\vec{x};\vec{p}) = N_t s_{jt}$. See Train 2009 sec 3.10 for proof

- Estimating α and β enables us to predict *every product's quantity response* to a change in *any product's attributes* x_{jt} or price p_{jt}

MNL: Common factors

- Suppose γ_t indicates how popular the category is in market t, so utility is $u_{ijt} = \gamma_t + x_{jt}\beta - \alpha p_{jt} + \epsilon_{ijt}$. Then market share becomes

$$s_{jt} = \frac{e^{\gamma_t + x_{jt}\beta - \alpha p_{jt}}}{\sum_{k=1}^{J} e^{\gamma_t + x_{kt}\beta - \alpha p_{kt}}}$$

$$s_{jt} = \frac{e^{\gamma_t} e^{x_{jt}\beta - \alpha p_{jt}}}{e^{\gamma_t} \sum_{k=1}^{J} e^{x_{kt}\beta - \alpha p_{kt}}}$$

$$s_{jt} = \frac{e^{x_{jt}\beta - \alpha p_{jt}}}{\sum_{k=1}^{J} e^{x_{kt}\beta - \alpha p_{kt}}}$$

- Similar exercise applies to individual-specific intercepts γ_i, or individual-time interactions γ_{it}. Only differences across products affect predicted market shares

MNL: Share differences

- We set product $j = 1$ utility to $u_{i1t} = \epsilon_{i1t}$ to normalize utility, i.e.

$$s_{1t} = \frac{1}{\sum_{k=1}^{J} e^{x_{kt}\beta - \alpha p_{kt}}}$$

- Hence for every $j \neq 1$, consider an affine transformation of market shares:

$$ln(s_{jt}) - ln(s_{1t}) = x_{jt}\beta - \alpha p_{jt}$$

- Looks like a regression equation… we sometimes add an error ξ_{jt} to represent unobserved product attributes:

$$ln(s_{jt}) - ln(s_{1t}) = x_{jt}\beta - \alpha p_{jt} + \xi_{jt}$$

- We use the $(J-1)T$ *differences* in market shares to estimate demand parameters

MNL Estimation

Define $y_{ijt} \equiv 1\{i \text{ chose } j \text{ in } t\}$. I.e., $y_{ijt} = 1$ iff i choose j at t; otherwise $y_{ijt} = 0$.

1. Maximum likelihood: Assume the probability of observation $\{i, j, t\}$ is $s_{jt}^{y_{ijt}}$; then lik. is $L(\beta) = \prod_{\forall i,j,t} s_{jt}^{y_{ijt}}$; & choose parameters to maximize log lik:

$$\sum_{\forall i,j,t} y_{ijt} \ln s_{jt}$$

2. Method of Moments: Choose α and β to solve some vector of "moments," i.e. interactions between exogenous observables z and error terms, e.g.

$$\sum_{\forall j,t} z_{jt}\left(\frac{\sum_{\forall i} y_{ijt}}{N_t} - s_{jt}\right) = 0$$

3. Linearize the model, choose parameters to minimize the sum of square errors

MNL Goodness-of-fit Statistics

- Discrete choice models predict *choice probabilities* rather than *choices*, because utility is always unobserved; hence nonstandard fit statistics
- Predicted outcomes are inherently stochastic, so limited predictive ability

1. Likelihood Ratio Test: (not R-sq)

$$\rho = 1 - \frac{lnL(\hat{\beta})}{lnL(0)}$$

- As $L(\hat{\beta}) \rightarrow 1, lnL(\hat{\beta}) \rightarrow 0, \rho \rightarrow 1$
- As $lnL(\hat{\beta}) \rightarrow lnL(0), \rho \rightarrow 0$

```
-   Heuristic: 0.2-0.4 is pretty good
```

2. Hit Rate: % of individuals for whom most-probable choice was actually chosen
3. R-sq using prediction errors at the jt level

MNL Pros

1. Microfounded, i.e. behavioral predictions are consistent with a clearly specified theory of consumer choice

```
    - Theory is utility maximization
    - Economists widely believe that microfounded models are more
generalizable than purely statistical models*
```

2. Extensible to accommodate preference heterogeneity

```
    - We'll cover 3 types of extensions in heterogeneous demand modeling
```

3. Likelihood function is globally concave in the parameters, ensuring fast and reliable estimation

```
    - Remember our local vs. global optimum discussion?
```

MNL Cons

1. Assuming $\epsilon_{ijt} \sim i.i.d. EV_1(0,1)$ is convenient but unrealistic

```
 - More likely, more similar products would experience more similar demand
shocks
 - Alternatives exist but can be computationally expensive
```

2. Analyst selects the choice set $j = 1, \ldots, J$, market size N_t, attributes x_{jt}, and price structure p_{jt}.

```
 - What's a j? What's a t? What's in x? How do we measure p? Who's in N?
 - "Tuning factors" or "Analyst degrees of freedom"
```

3. Market share derivatives depend on market shares alone (IIA; see Train Sec. 3.6)

4. Price Endogeneity

```
 - Affects all demand models, not just MNL
```

IIA: Deeper dive

- Famous example from McFadden (1974)
- Suppose you estimate demand for transportation with three options: {Blue Bus, Red Bus, Car}, each with 33% market share
- Now suppose you painted the red buses blue and want to predict market shares in the new choice set: {Blue Bus, Car}
- MNL will predict Blue Bus and Car shares of 50%, not 67% and 33% (why?)

IIA: Remedies

IIA is testable & usually rejected by data

Common remedies:

1. Extend the model to impose structure on choice set, e.g. Nested Logit or Ordered Logit
2. Change $\epsilon_{ijt} \sim i.i.d. EV_1(0,1)$, e.g. Multivariate Probit with correlated errors
3. Change model structure so IIA property does not obtain, e.g. heterogeneous logit

Price endogeneity: 35 INSANE explanations

- #4 will SHOCK you. Like, comment, subscribe
- Covered on the exam…ask lots of questions

If it can be shown that during a period of time covered by two or more observations either curve remains fixed while the other moves to right or left, price-output data will reveal points on the curve that remains fixed. This should be obvious from the analysis given above but may be illustrated by a diagram. (Figure 3.)

FIGURE 3. PRICE-OUTPUT DATA REVEAL—
(A) SUPPLY CURVE
(B) DEMAND CURVE

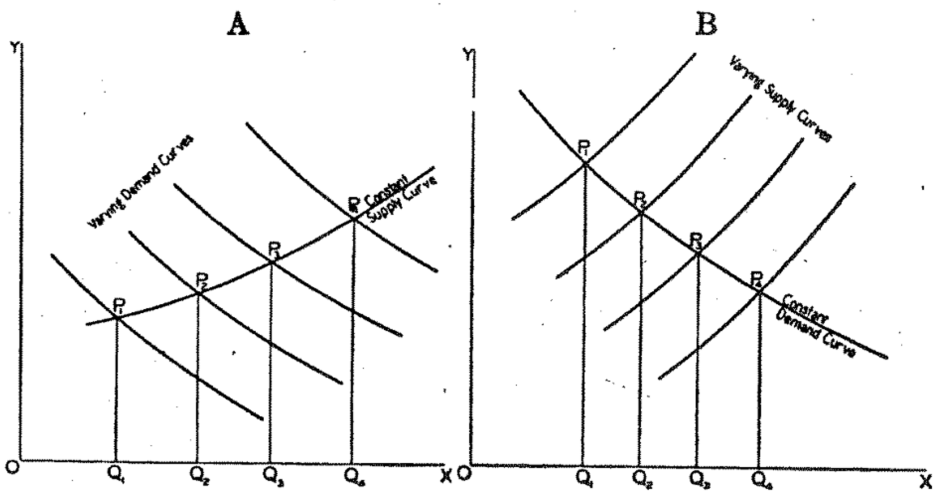

If both supply and demand conditions change, price-output data yield no direct information as to either curve. (Figure 4.)

FIGURE 4. PRICE-OUTPUT DATA FAIL TO REVEAL EITHER SUPPLY OR DEMAND CURVE.

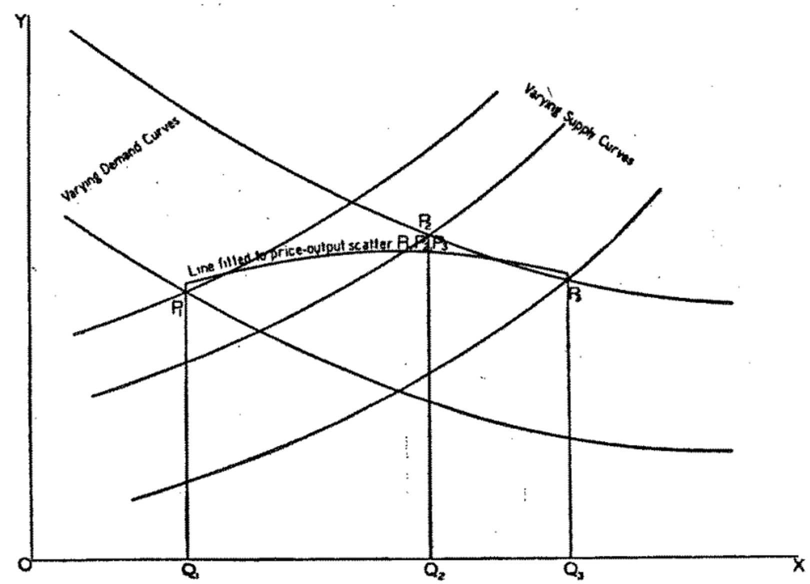

Unfortunately for our problem, the case represented by Figure 4 is the more common, and even if either curve does remain fixed during the period covered by the observations there is no certain way of knowing this fact in advance.[5]

2. Fundamental issue

- Demand model is a causal price-quantity relationship
  Yet observed prices may correlate with other demand and supply determinants
  Exogenous price variation req’d to distinguish correlation from causation (“identification”)

```
- Price endogeneity is a "data problem" not a "model problem"
- Can be hard to verify empirically--needed data is missing--but widely
believed important
- Implies wrong demand slope, biased demand predictions
```

- Sign of bias depends on unobserved correlation

```
- If corr(price,unobs)<0 --> estimated demand is "too flat" or too elastic
- If corr(price,unobs)>0 --> "too steep" or too inelastic
- Affects all demand models, not just MNL
```

3. E-Commerce Example

- Imagine an unobserved demand shock, such as a viral Instagram post, increases Amazon product awareness and sales
- Sales spike, inventory drops, automated pricing system increases price to monetize remaining inventory
- What do data show? Corr(sales, price) > 0 !!
  - Common enough to be unsurprising when this happens

4. An Analogy

- Foot size data significantly predict reading comprehension among children!

```
   - In fact, age causes both foot size
   - And age causes reading comprehension
   - Without age, you cannot accurately estimate the causal relationship
between foot size and reading comprehension. You can only measure the
correlation
```

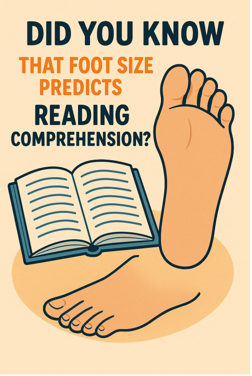

5. Sports example

- You have data on Lakers ticket prices and sales before and after the Luka trade
- Prices and sales are both higher after the trade
- Does this mean that higher price caused higher sales?

6. Example: Digital systems

- Uber surge pricing:
  Positive Demand shocks increase price
  Negative Supply shocks increase price
- System adjusts the price without knowing the causes

  `Many digital inventory-based pricing systems are similar`


AT&T
12:41 AM
49%
SURGE PRICING
Demand is off the charts! Rates have increased to get more Ubers on the road.
7.0x
THE NORMAL RATE
$175 MINIMUM FARE
+
$8.75 / MIN & $35 / MILE

7. Shrinkflation

Shrink the package, maintain the price
Also, "Skimpflation" : reduce ingredient intensity

8. Multiple determinants

- Price causes quantity demanded; but so do other things

describe the problem of multiple determinants in 100 words or less

Multiple Determinants of a Problem

Reasoned for 1m 30s

In many fields, phenomena are shaped by numerous interconnected factors, making it difficult to isolate and measure the contribution of each cause. This is the problem of multiple determinants: the potential overlap, interactions, and confounding among factors can complicate robust explanations and predictions. Researchers must use sophisticated methods such as multiple regression, structural models, or experimental designs to tease apart these influences. Failing to address multiple determinants can lead to incomplete or erroneous conclusions about why something happens. Scrupulous study design, rigorous data collection, and sound analysis are essential to mitigate these challenges.

Source:

Freedman, D. A. (2009). *Statistical Models: Theory and Practice*. Cambridge University Press.

9. Oft-unobserved price correlates

1. Changes in consumer preferences, income, market size
2. Retail distribution, prominence, stocking
3. Digital marketing, including search ads, display ads, affiliates, influencers, coupons
4. Competitor prices, preference shocks, retail, dig mktg

Any may correlate with equilibrium prices, leading to endogeneity biases if left uncontrolled

10. General model interpretation

11a. Simulation

11b. Simulated data generating process

```
# 1. Simulation parameters and correlated cost shocks
set.seed(14)                        # for reproducibility
n_periods    <- 100                 # number of periods
market_size  <- 100                 # total market size (e.g. number of customers)
rho          <- 0.9                 # influence of costshock1 on costshock2

# Demand model parameters
alpha       <- 0.2                  # price sensitivity (common across products)
intercept1  <- 9                    # baseline utility for product 1
intercept2  <- 9                    # baseline utility for product 2
# (Outside option utility is normalized to 0)

# Simulate cost shocks for the two firms (correlated)
shock1 <- rnorm(n_periods)
shock2 <- rho * shock1 + (1 - rho) * rnorm(n_periods)
```

```
> # 3. OLS regressions for Firm 1's demand
> model_naive <- lm(Q1 ~ price1, data = data)
> summary(model_naive)

Call:
lm(formula = Q1 ~ price1, data = data)

Residuals:
     Min       1Q   Median       3Q      Max 
-1.31867 -0.34908 -0.00637  0.32919  1.19378 

Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept) 51.78942    0.17660  293.26   <2e-16 ***
price1      -0.59737    0.05565  -10.73   <2e-16 ***
```

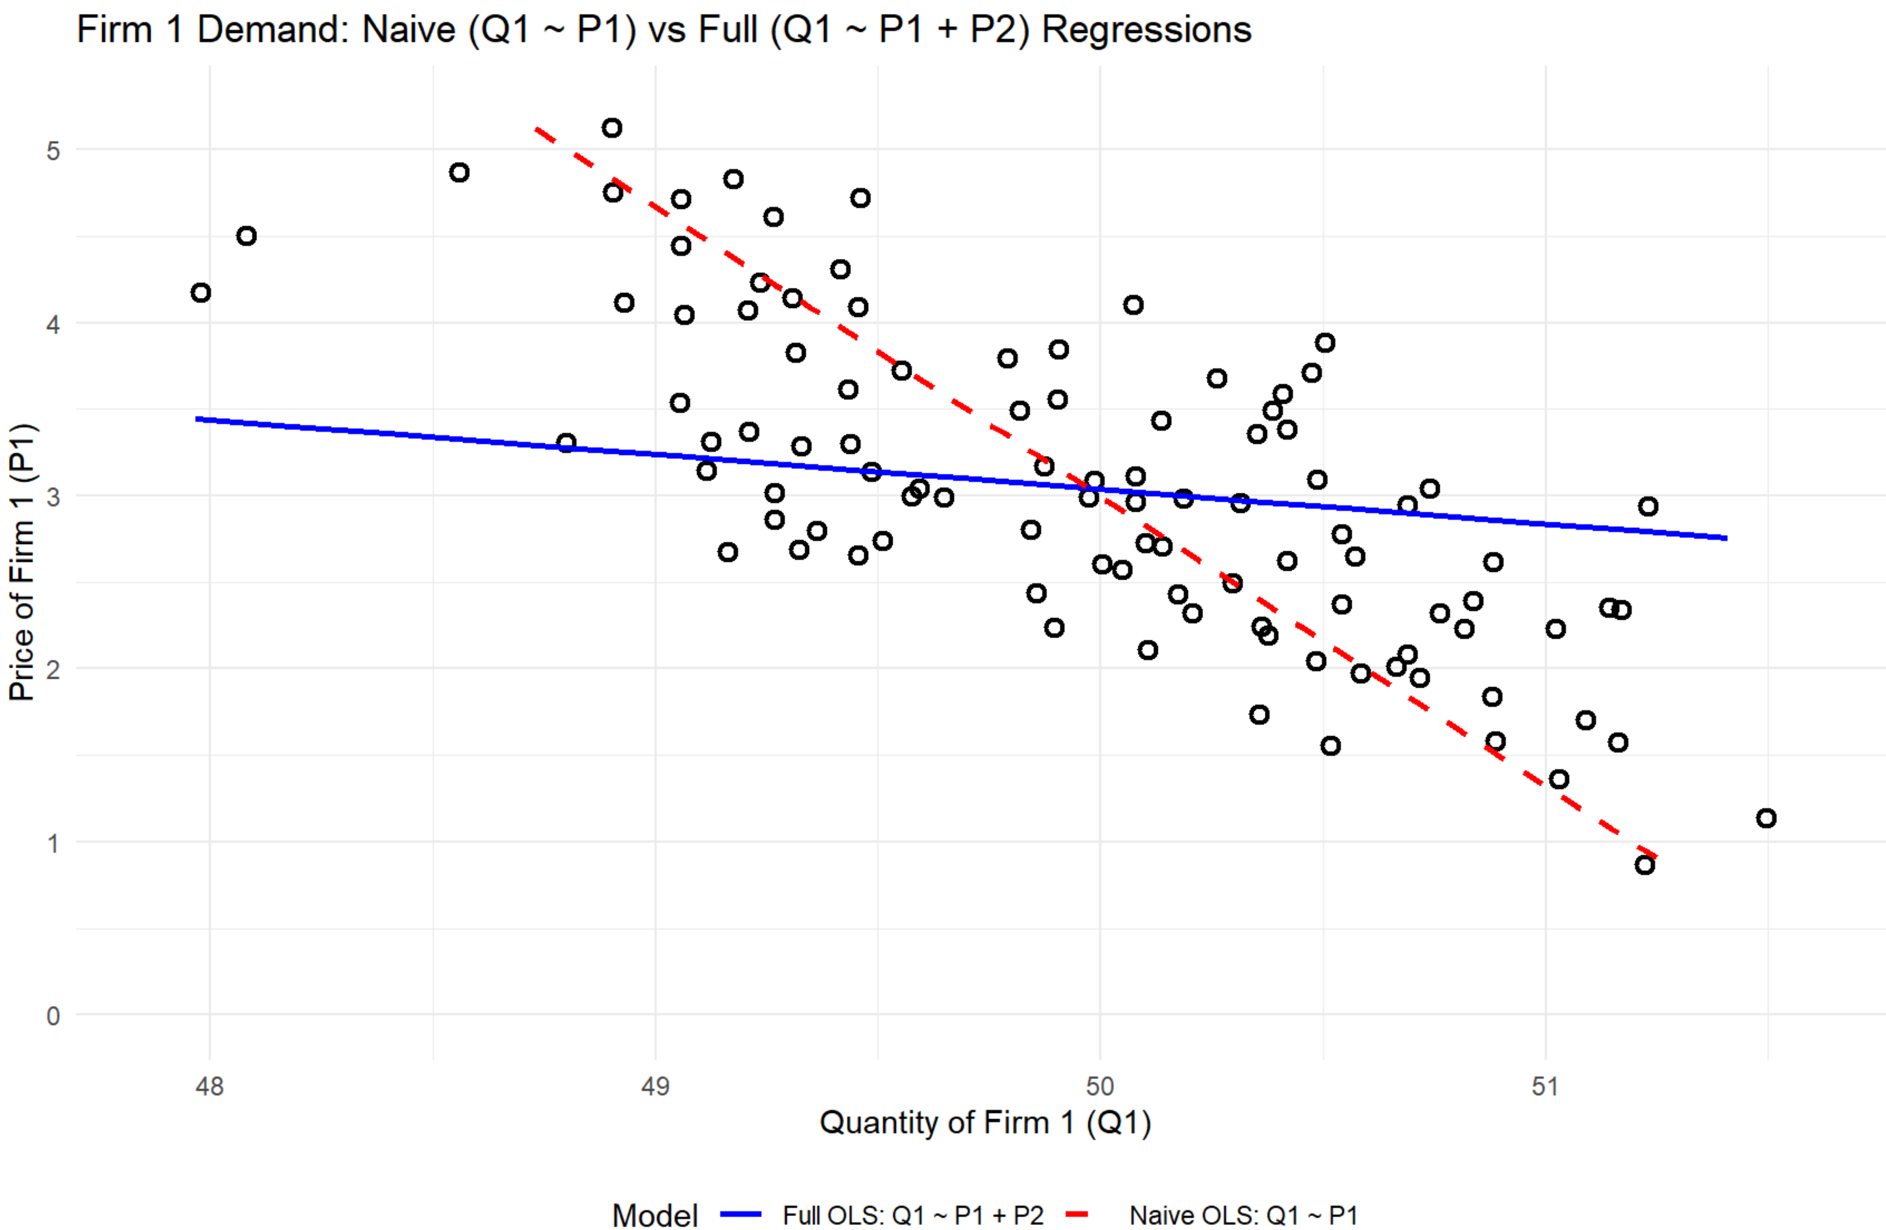
Firm 1 Demand: Naive (Q1 ~ P1) vs Full (Q1 ~ P1 + P2) Regressions
Price of Firm 1 (P1)
0
1
2
3
4
5
48
49
50
51
Quantity of Firm 1 (Q1)
Model
Full OLS: Q1 ~ P1 + P2
Naive OLS: Q1 ~ P1

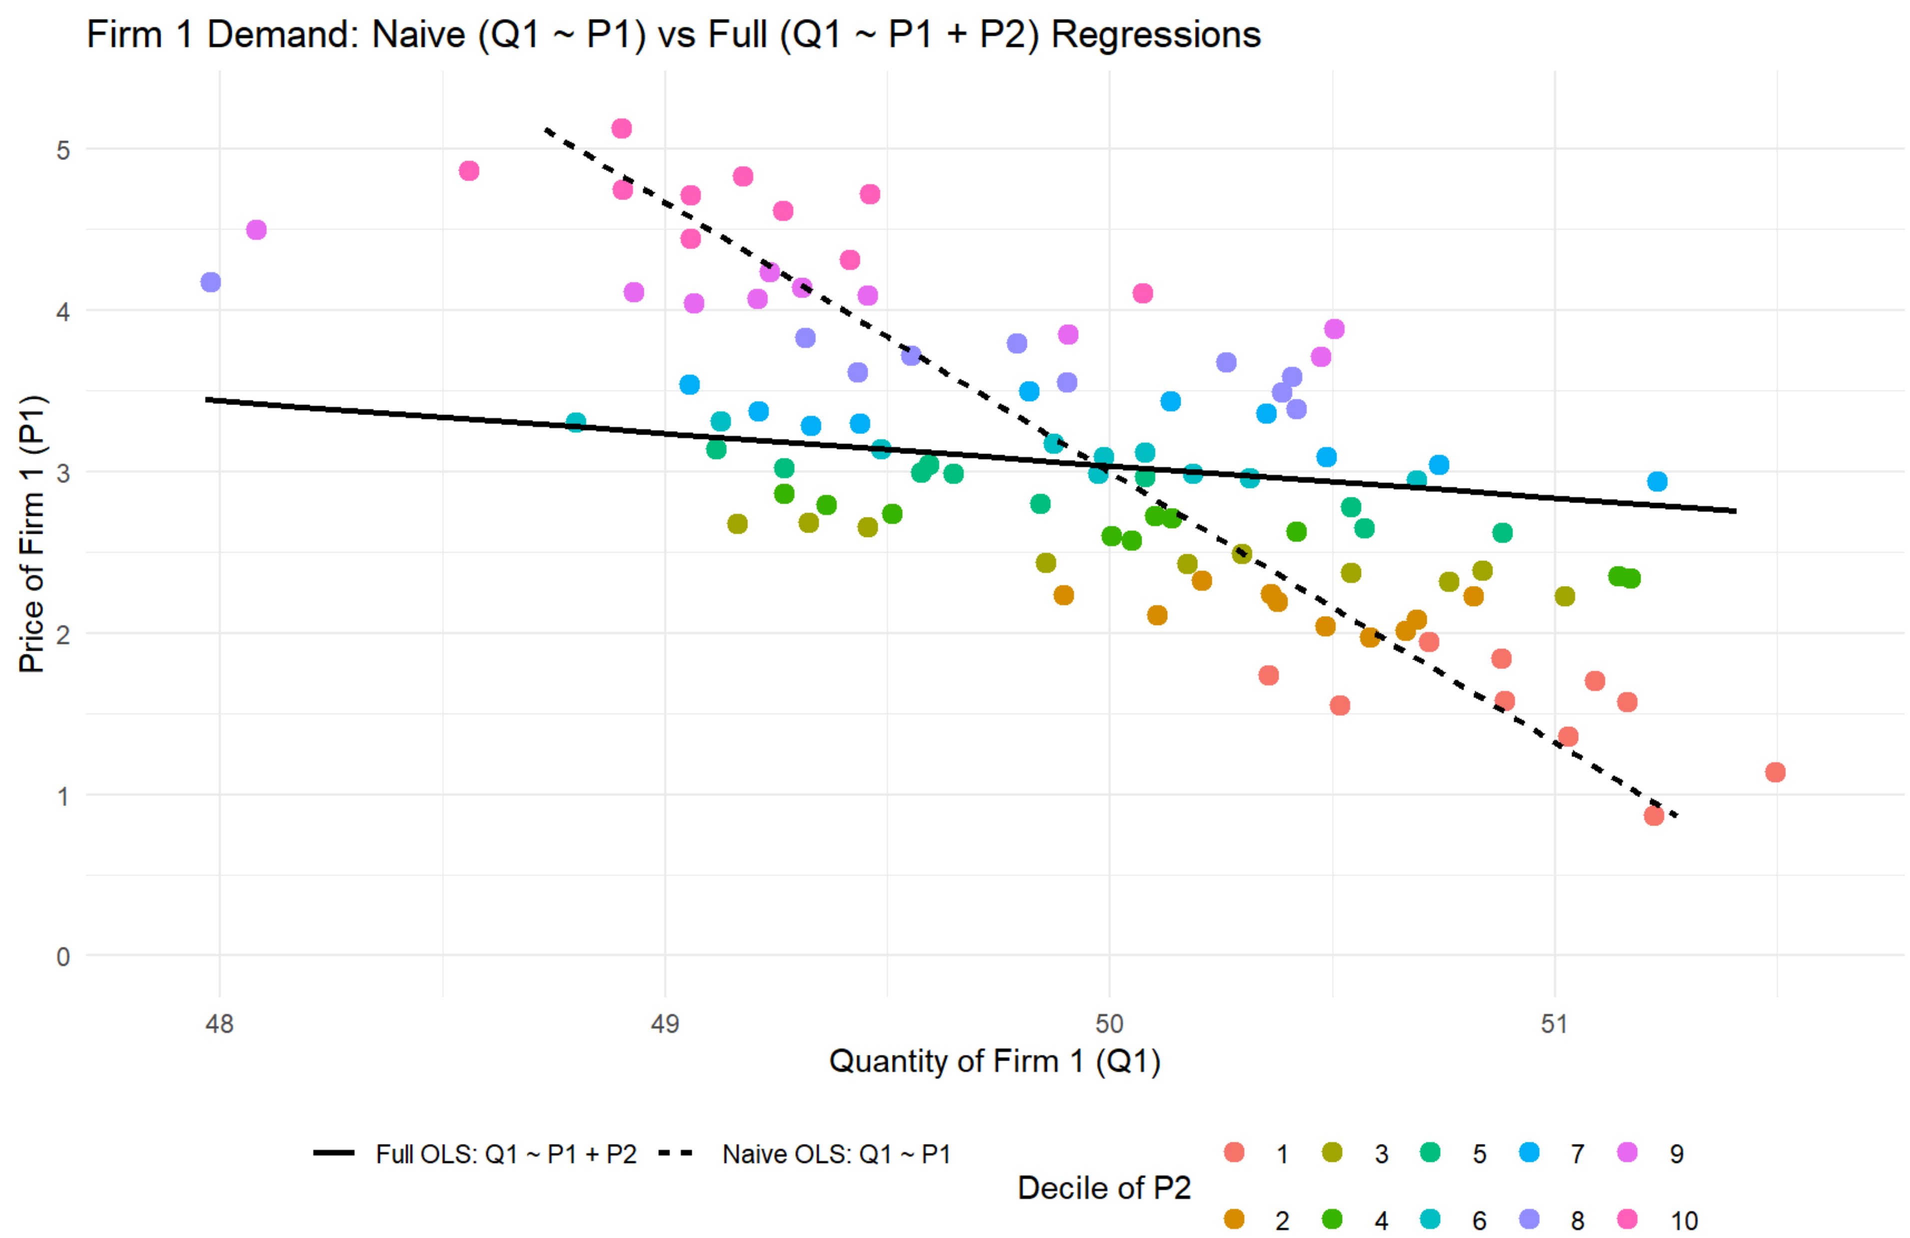

Firm 1 Demand: Naive (Q1 ~ P1) vs Full (Q1 ~ P1 + P2) Regressions
Price of Firm 1 (P1)
0
1
2
3
4
5
48
49
50
51
Quantity of Firm 1 (Q1)
Full OLS: Q1 ~ P1 + P2
Naive OLS: Q1 ~ P1
Decile of P2
1
2
3
4
5
6
7
8
9
10

Common solutions

- Experiments

```
Randomizing price eliminates confounding with unobservables
Gold standard, but has drawbacks mentioned earlier
```

- Quasi-experiments using archival data

```
1. Instrumental variables
2. Regression discontinuities
3. Natural experiments
4. Difference-in-differences
5. Synthetic controls
6. Double/debiased machine learning
These approaches are beyond the scope of this class
```

- Model the price-setting process

```
But without exogenous price variation, this is difficult to evaluate
```

- Best practice: Triangulate

Exogenous smartphone price variation

- Smartphone discounts are randomly assigned,
  thus we have exogenous price variation to identify β
  - `This class focuses on how we can use demand models`
  - `Endogeneity remedies: Future metrics classes or graduate study`
  - `Treat the topic as a demand modeling risk to be understood`

- Describe price endogeneity in your own words

```
Generate a novel example related to price and demand
Explain how to resolve it
```

Class script

- Wrangle data
- Estimate MNL
- Interpret parameters and SEs
- Assess model fit

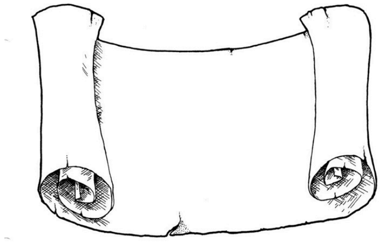

Wrapping up

Recap

- Demand modeling enables data-driven sales predictions for counterfactual prices and attributes, facilitating profit maximization
- MNL is popular bc it is powerful, microfounded and tractable
- MNL has limitations (eg, IIA), but is extesnible by modeling heterogeneity
- Price endogeneity is a data problem that biases demand parameter estimates when price is correlated with unobserved supply or demand shifters. Usually resolved with exogenous price variation

Going further

- Supply, Demand, and the Instrumental Variable: Lessons for Data Scientists from the Economist's Toolbox
- Causal inference in economics and marketing (Varian 2016)
- Discrete Choice Analysis with R
- Microeconometric models of consumer demand by Dube (2018)
- Empirical Models of Demand and Supply in Differentiated Products Industries by Gandhi & Nevo