Lecture 12: Testing via Slutsky, GARP, Aggregating Demand

1 Where we are

- Last time: properties of Hicksian demand,
  leading up to the Slutsky equation
- Today: using Slutsky to "test" our model;
  when is demand data rationalizable when it's finite data;
  aggregating demand across different individuals;
  recovering preferences from Marshallian demand;
  and a few other odds and ends
- First: any questions?

2 Using Slutsky to Test our Theory

- Recall the conclusion we reached earlier – the matrix of partial derivatives of Hicksian demand with respect to prices is symmetric and negative semi-definite
- That would be a nice conclusion to test, if we could observe someone's Hicksian demand
- While we don't expect to actually observe Hicksian demand,

  the Slutsky equation lets us infer its partial derivatives from something we can observe,

  Marshallian demand
- The Slutsky equation is $\frac{\partial x_i}{\partial p_j} = \frac{\partial h_i}{\partial p_j} - \frac{\partial x_i}{\partial w} x_j$; rewrite as

$$\frac{\partial x_i}{\partial p_j}(p, w) \quad + \quad \frac{\partial x_i}{\partial w}(p, w) x_j(p, w) \quad = \quad \frac{\partial h_i}{\partial p_j}(p, u)$$

- What's nice about this is:
  - The right-hand side is something we have very strong theoretical predictions about
  - The left-hand side is something we could potentially observe
- That is, if we observe a consumer's Marshallian demand everywhere, we can construct the matrix

$$D_p h(p, u) \quad = \quad \begin{bmatrix} \frac{\partial x_1}{\partial p_1} + \frac{\partial x_1}{\partial w} x_1 & \cdots & \frac{\partial x_1}{\partial p_k} + \frac{\partial x_1}{\partial w} x_k \\ \vdots & & \vdots \\ \frac{\partial x_k}{\partial p_1} + \frac{\partial x_k}{\partial w} x_1 & \cdots & \frac{\partial x_k}{\partial p_k} + \frac{\partial x_k}{\partial w} x_k \end{bmatrix}$$

- And we know that to be consistent with our theory, this matrix needs to be symmetric and negative semidefinite
- In fact, along with other easy-to-check conditions, that's also sufficient:
- **Theorem.** A differentiable demand function $x(p, w)$ is consistent with utility maximization given rational, locally non-satiated preferences if and only if:
  1. $x(p, w)$ is homogeneous of degree 0
  2. At every (p, w), $p \cdot x(p, w) = w$ (Walras' Law)
  3. At every (p, w), the matrix $[\frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} x_j]$ is symmetric and negative semidefinite
- To put it another way, our theory of rational choice predicts these properties –

  and nothing else!

- So this is nice – if we observe someone's entire demand function $x(p, w)$ at all prices and wealth levels, we can determine whether they're rational utility maximizers or not –

  by testing whether the Slutsky matrix coming from x is symmetric and negative semidefinite

- There are some empirical papers that do exactly that

- I don't know the empirical literature well,

  but I get the sense the early literature focused on symmetry of the Slutsky matrix,

  and tended to rely on estimating a parametric demand system and seeing if the parameters implied Slutsky-symmetry;

  some more recent work has included testing for negative semidefiniteness,

  and been more nonparametric

- I can at least point to you some relatively recent papers by authors I would trust:

  Arther Lewbel (1995), "Consistent Nonparametric Hypothesis Tests with an Application to Slutsky Symmetry," *Journal of Econometrics* 67

  B. Haag, S. Hoderlein and K. Pendakur (2009), "Testing and Imposing Slutsky Symmetry in Nonparametric Demand Systems," *Journal of Econometrics* 153 – focus on proper econometric testing

  Stefan Hoderlein (2011), "How Many Consumers Are Rational?" *Journal of Econometrics* 164 – focus on allowing for unobserved heterogeneity across consumers when using aggregate data

  Victor Aguiar and Roberto Serrano (2017), "Slutsky Matrix Norms: The Size, Classification, and Comparative Statics of Bounded Rationality," *Journal of Economic Theory* 172 – focus is to measure distance between observed Slutsky matrix and nearest matrix satisfying required conditions, and decompose violations into three pieces:

  "inattentiveness to changes in purchasing power" (a certain type of cycle existing in the graph formed as wealth and prices change in ways that should strictly increase indirect utility);

  "money illusion" (violations of homogeneity of degree 1, interpreted as an irrational focus on nominal rather than real prices);

  and actual violations of the Law of Compensated Demand

3 Recovering Preferences from $x(p, w)$

- So that's rationalizability:

  when are observations consistent with our model of rational choice?

  when they imply a Slutsky matrix which is symmetric and negative semidefinite

- (That's for the case of "infinite data" – when we observe the entire demand function;

  very shortly, we'll consider the question of finite observations)

- We can also think about whether we can recover preferences from observation of demand,

  or whether our model is identified...

- If we observe Marshallian demand $x(p, w)$ at every (p, w),

  and it satisfies these conditions to be rationalizable,

  can we recover the preferences that rationalize it?

- Turns out: Yes!

- But I'm going to skip this for now and come back at the end if there's time

4 Rationalizing Finite Data (GARP)

- So far, we've thought about when data is consistent with utility maximization if we know the entire Marshallian demand function –

  in a sense, if we had infinitely much data

- Next, what if we had only a finite number of observations?

- Think back to the first day of consumer theory,

  when we thought about preference and choice rules,

  and saw that a choice rule based on rational preferences had to satisfy the Weak Axiom of Revealed Preference:

  if $x, y \in A \cap B$, with $x \in C(A)$ and $y \in C(B)$, then $x \in C(B)$ and $y \in C(B)$

- The logic: if y was available when I chose x, we learn $x \succsim y$;

  and if x was available when I chose y, we learn $y \succsim x$

- We know that if we had observations for every possible choice problem,

  satisfying WARP is both necessary and sufficient for the data to be rationalizable

- But we also saw – the scissors-paper-rock example – that satisfying WARP is not enough if we don't observe choice from all choice sets

- So next, we need to figure out when finite data is rationalizable

- Marshallian demand is just choice from sets, where the sets happen to be budget sets

- If you see me demand bundle x at prices p and wealth w,

  you infer that I weakly prefer x to any other bundle I could have afforded

- And if we assume that preferences are locally non-satiated,

  then a bundle in the interior of my budget set can't be optimal,

  so you infer that I *strictly* prefer x to any bundle that was in the interior of my budget set

- For a simple example of what this rules out,

  suppose that at prices $(2, 1)$, I demand $(4, 2)$, but at prices $(1, 2)$, I demand $(2, 4)$

- In the first case, I chose to consume $(4, 2)$ when I could have afforded $(2, 4)$ and still had money left over – if preferences are LNS, this implies $(4, 2) \succ (2, 4)$

- In the second case, I did the opposite – demanded $(2, 4)$ when I could have afforded $(4, 2)$ and had money left over – so $(2, 4) \succ (4, 2)$

- Thus, these choices wouldn't be compatible with rational choice

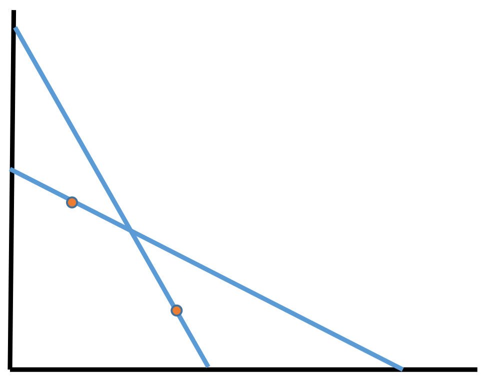

- Even if we observe Marshallian demand for all (p, w), though,

  we only observe choice on linear budget sets, not all possible choice sets,

- Plus, in real data, we're only going to see choice from a finite set of budget sets,

  not even all budget sets

- So WARP may not be a sufficient condition for rationalizability

- It turns out, there's a stronger condition which is sufficient –

  and it's basically what we get if we extend the logic of WARP via transitivity

- We'll assume throughout that preferences are LNS
- (If they weren't, our model would be vacuous –
  any data could be rationalized by flat preferences over everything)
- But we're **not** going to assume that demand is single-valued,[1]
  or that we observe the whole set $x(p, w)$ –
  for each (p, w) we observe, we assume we just see one element of $x(p, w)$

- some terminology first
- if a bundle x is chosen at prices p, and $p \cdot x' \leq p \cdot x$,
  then we say x is **directly revealed preferred** to x', or $x \succsim^D x'$
- that is, if we saw the consumer choose x, while x' was also in the consumer's budget set,
  we infer that if the consumer has rational preferences, $x \succsim x'$
  (or else they would have chosen x', or something else)
- or, $x \succsim^D x'$ if $x \in x(p, w)$ and $x' \in B(p, w)$

- to extend this relation via transitivity,
  we'll say x is **revealed preferred** to x', or $x \succsim^R x'$, if there's some string of bundles with

$$x \succsim^D x_1 \succsim^D x_2 \succsim^D \cdots \succsim^D x_n \succsim^D x'$$

- finally, we'll say x is **directly strictly revealed preferred** to x', or $x \succ^D x'$,
  if we saw the consumer choose x at prices p and $p \cdot x' < p \cdot x$,
  or if $x \in x(p, w)$ and $p \cdot x' < w$
- (since we assume preferences are LNS, if $x' \succsim x$ and $x' \in Int(B(p, w))$,
  there would have to be another bundle near x', also in the budget set,
  which was strictly preferred to x, so x could not have been chosen;
  so $x \in x(p, w)$ and $x' \in Int(B(p, w))$ requires $x \succ x'$)

[1] The corresponding condition that holds when Marshallian demand is assumed to be a singleton is called the Strong Axiom of Revealed Preference, or SARP.

- The **Generalized Axiom of Revealed Preference**, or GARP, says that for any x and x', if x is revealed preferred to x', then x' cannot be strictly directly revealed preferred to x.

Definition. *A series of observations* $\{(x^i, p^i)\}_{i=1}^N$ **satisfy the Generalized Axiom of Revealed Preference** *if for any two observations i and j, if* $x^i \succsim^R x^j$ *then* $x^j \not\succ^D x^i$.

- This is equivalent to not having a "loop" of direct revealed preference relations

$$x_1 \succsim^D x_2 \succsim^D \cdots \succsim^D x_n \succsim^D x_1$$

where at least one of the links is strict revealed preference,

or a string of inequalities

$$p^1 \cdot x^1 \geq p^1 \cdot x^2, \qquad p^2 \cdot x^2 \geq p^2 \cdot x^3, \qquad \cdots, \qquad p^{n-1} \cdot x^{n-1} \geq p^{n-1} \cdot x^n, \qquad p^n \cdot x^n \geq p^n \cdot x^1$$

with at least one strict inequality

- It should be obvious that if preferences are rational, any data generated must satisfy GARP
- The question is whether this is "if and only if"

- One other definition, which may be obvious: for any set of observations $(p^t, w^t, x^t)_{t=1,2,\ldots,T}$, we'll say the utility function u rationalizes the data if

$$x^t \quad \in \quad \arg\max u(x) \quad \text{subject to} \quad p^t \cdot x \leq w^t$$

for every t – that is, if the observations are consistent with rational choice given u

- (Also note that we don't need to observe w,

since we assume LNS and therefore infer $w^t = p^t \cdot x^t$.)

The punchline:[2]

Theorem (Afriat's Theorem). *Let (p^i, x^i), $i = 1, 2, \ldots, n$ be a finite number of observations of price vectors and consumption bundles. The following conditions are equivalent:*

1. *There exists a locally nonsatiated utility function that rationalizes the data*
2. *The data satisfy GARP*
3. *There exist numbers $\{U^i, \lambda^i\}_{i=1}^n$, $\lambda^i > 0$, such that for every $\{i, j\}$,*

$$U^i \quad \leq \quad U^j + \lambda^j p^j (x^i - x^j)$$

4. *there exists a LNS, continuous, concave, monotonic utility function that rationalizes the data*

- So, for any finite set of observations,

  if they satisfy GARP, they're compatible with rational choice;

  and if they fail GARP, they're not

[2] For the proof, see Varian (1982), "The Nonparametric Approach to Demand Analysis," *Econometrica* 50(4), or the papers by Afriat that he cites.

- Why do we need GARP, not just WARP?
- That is, isn't the pairwise relation enough, like in producer theory with WAPM?

  Why do we need to extend it with transitivity?
- Example:[3]

prices	demand
$(1, 1, 2)$	$(1, 0, 0)$
$(2, 1, 1)$	$(0, 1, 0)$
$(1, 2, 1.5)$	$(0, 0, 1)$

- The first observation implies $(1, 0, 0) \succsim (0, 1, 0)$;

  but we never see $(0, 1, 0)$ chosen when $(1, 0, 0)$ is affordable, since at the second observation, $(1, 0, 0)$ is outside of the budget set,

  so we never get a direct contradiction
- The second observation implies $(0, 1, 0) \succsim (0, 0, 1)$;

  but again, we never see $(0, 0, 1)$ chosen when $(0, 1, 0)$ is affordable
- Finally, the third observation implies $(0, 0, 1) \succ (1, 0, 0)$,

  but we never see $(1, 0, 0)$ chosen when $(0, 0, 1)$ is affordable
- So we need transitivity to make the argument –

  we know the data can't be rationalized,

  and they fail GARP, but they don't fail WARP

- (With the producer problem, the firm was always choosing from the same production set as prices changed,

  so a pairwise comparison of choices was enough.

  With the consumer problem, a consumer has the same preferences across observations,

  but both prices and wealth vary,

  so in some sense the consumer is facing different problems at each price vector)

[3] Borrowed from Ichiro Obara's lecture slides from UCLA

- Afriat's theorem gives us a funny second result: if data can be rationalized at all, it can be rationalized by a utility function that's monotone and concave
- Now, if u is concave, then it's quasi-concave, and if u is quasi-concave, then the preferences it represents are convex
- So this is saying, any data that's consistent with rational choice, is consistent with monotone, convex preferences

- So... does this mean rational preferences have to be monotone and convex?
- No, of course not
- You can have whatever preferences you like – they may be monotone, they may be convex, they may not be
- What this is saying is that with observed choice data, we'll never be able to detect whether your preferences have these properties or not
- Whatever choices you make, we could find monotone and convex preferences that would lead to the same choices
- Let's see why

- First, suppose you had non-monotone preferences
- That is, suppose there were two bundles x and x', with $x' \gg x$ but $x \succsim x'$
- How would we observe this? We would need to see you choose x when x' was also available
- But that means finding a consumer problem (p, w) such that $x \in x(p, w)$ while $p \cdot x' \leq w$
- But since $x \ll x'$, if $p \cdot x' \leq w$, $p \cdot x < w$
- And we know that with LNS preferences, you'll never choose x when it doesn't exhaust your budget
- So the exact observation we would need to learn that your preferences are not monotone, is impossible to get if your preferences are LNS
- (To put it another way, the non-monotonicity in your preferences would be in a region where you weren't choosing anyway, so we'd have no way to observe it from choices)

- Similarly, suppose you have non-convex preferences

- Convex preferences say that if $x \succsim y$ and $x' \succsim y$, then $x^t = tx + (1-t)x' \succsim y$ as well

- So if preferences are *not* convex, there must be some x, x', y, and t such that

  $x \succsim y$, $x' \succsim y$, and $y \succ x^t$

- Or by transitivity, $x \succ x^t$ and $x' \succ x^t$

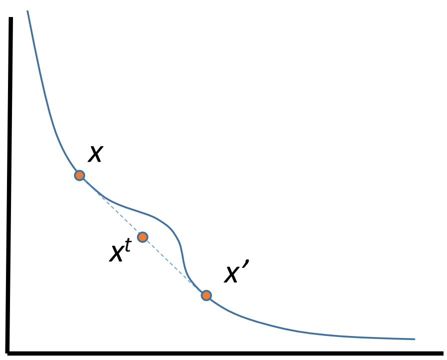

- Now, let's think about what we'd need to observe to *figure out* that both $x \succ x^t$ and $x' \succ x^t$

- First, to learn that $x \succ x^t$, we'll need a data point where x is chosen, even though x^t is in the interior of the budget set

- And second, to learn that $x' \succ x^t$, we'll need a data point where x' is chosen, even though x^t is in the interior of the budget set

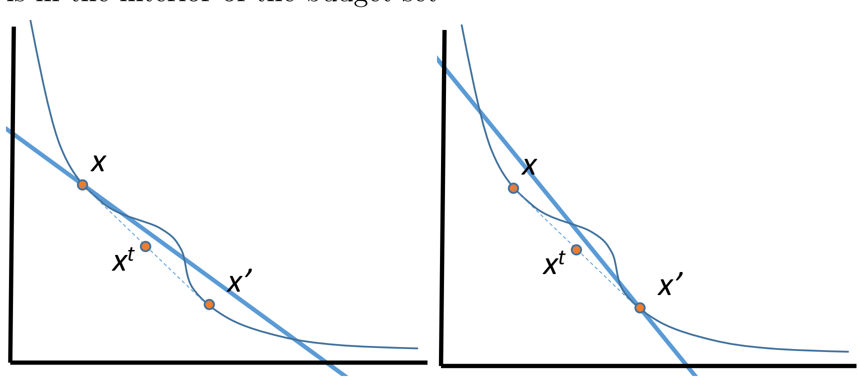

- But in the first case, if x is on the budget line and x^t is in the interior, then x' is in the interior too

  So we can only observe x being chosen if $x \succ x'$

- And likewise, we can only see x' chosen in the second case if $x' \succ x$

- So with LNS preferences, there's simply no way to generate the data we would need to learn that your preferences aren't convex

- They might be, they might not be, we just won't know

5 Testing GARP

- Christopher Chambers and Federico Echenique (2016), *Revealed Preference Theory*, Cambridge University Press (Econometric Society Monographs)

  ch. 5 is all about "practical issues" of testing GARP,

  and an overview of the empirical literature

- One challenge is that GARP is a binary condition –

  however many observations you have, your data collectively either satisfy it or don't;

  and in a rich dataset, you always expect some noise,

  so you would expect it to be violated

- There's a need, therefore, to find a way to quantify whether violations are significant –

  either statistically, or economically

- Afriat proposed an index, now called Afriat's Efficiency Index,

  to measure "how close" a consumer was to satisfying GARP,

  basically saying how much we need to "deflate" an individual's choices by to satisfy GARP,

  with 1 being data that satisfies GARP already, and lower numbers meaning larger violations

- Others have modified this index

- Chambers and Echenqiue propose a different idea, the "money pump index,"

  which measures the severity of a GARP violation by how much money someone could make by arbitraging your "irrational" choices

- (If your choices violate GARP, it's because you sometimes "overpay" for a bundle you could have bought cheaper a different time;

  they measure how much money someone could make basically buying this bundle cheap and selling it to you expensive,

  swapping around your choices)

- Under some assumptions, they turn this measure of "how irrational" you are into a statistical test

- Another issue is power
- People often use cross-sectional data,
  looking at observations of different households at the same time
- In those cases, wealth (or expenditure) varies a lot more than prices
- This means budget sets are usually nested,
  and it's almost impossible for GARP to be violated,
  so using GARP as a test has very little power
- I'll mention one solution in a minute

- A few interesting empirical papers they mention:
  - Dowrick and Quiggin (1994), "International Comparisons of Living Standards and Tastes: A Revealed-Preference Analysis," *American Economic Review* 84.1
    cross-country comparisons (60) using 1980 national consumption data –
    consistent with GARP, seven consumption categories
  - Landsburg (1981), "Taste Change in the United Kingdom, 1900-1955," *Journal of Political Economy* 89.1
    time series data, weakness of test given rising income
  - Blundell, Browning and Crawford (2003), "Nonparametric Engel Curves and Revealed Preference," *Econometrica* 71.1
    cross-sectional data, introduced Engel curve adjustment to solving rising-income problem
- And a few experimental ones:
  - Andreoni and Miller (2002), "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism," *Econometrica* 70.2
    (testing whether overly generous behavior in a "dictator game" can be explained by "rational" altruistic preferences – basically, yes)
  - Harbaugh, Kraus and Berry (2001), "GARP for Kids: On the Development of Rational Choice Behavior," *American Economic Review* 91.5
    (everyone's pretty close to rational; 11 year olds are more rational than 7 year olds, and about as rational as college-age adults)
  - Choi, Kariv, Müller and Silverman (2014), "Who Is (More) Rational?" *American Economic Review* 104.6 – linking "rationality" to socioeconomics
    (men, well-educated, and rich choose "more rationally";
    higher Afriat efficiency index is a predictor of higher household wealth)

6 Aggregating Demand

- Other than the experimental papers, most of the papers above use *aggregate* data –
  they observe total consumption for a group of individuals (maybe even an entire country),
  but test for rationality as if it were the demand of a single individual
- So... is this a valid exercise?
- Should we expect demand from a group of rational individuals to satisfy GARP,
  as it would if it were the choices of a single rational individual?

- In producer theory, we saw that *production aggregates* –
  if each firm is maximizing profits based on its own production set,
  you showed on the first homework of the semester,
  it will look just like profit maximizing behavior of a single firm.
- Does the same thing hold for consumers?

- The answer is, generally no, but yes in a few special cases.
- For general utility functions, individual consumers can be perfectly rational,
  but generate data that is not rationalizable at the aggregate level
- However, demand does aggregate in two special cases,
  which you'll be exploring on the next homework

- (The problem is summed up well in the intro to the Blundell/Browning/Crawford paper:
  "...If we do reject revealed preference conditions on aggregate data,
  we have no way of assessing whether this is due to a failure at the micro level
  or to the inappropriate aggregation across households that do satisfy the integrability conditions but who have different nonhomothetic preferences.")

- The reason demand doesn't aggregate has to do with varying wealth
- In producer theory, each firm had a fixed production set,
  and prices were varying across observations
- In consumer theory, each individual has fixed preferences,
  but prices *and individual budget levels* can vary across observations
- But because Marshallian demand has wealth effects,
  aggregate demand depends on the distribution of wealth, not just total wealth,
  which wouldn't be picked up in aggregate data

- Easiest to understand if we think about two consumers,
  one who loves steak, one who loves sushi
- Suppose we have two observations
- At the first observation, steak guy had a lot of money, sushi guy didn't,
  so aggregate demand was for a lot of steak and a little sushi
- At the second observation, steak guy had a little money, sushi guy had a lot,
  so aggregate demand was for a lot of sushi and a little steak
- But that's basically independent of prices
- If the price of steak was a little higher on the first day,
  and the price of sushi was a little higher on the second day,
  aggregate demand would appear irrational –
  the two consumers together demanded lots of steak when steak was expensive,
  and lots of sushi when sushi was expensive,
  violating GARP – even though both consumers were perfectly rational[4]

[4] We can generate a similar counterexample even with consumers who have identical preferences. Consider two individuals with utility functions $u(x_1, x_2) = \frac{x_1}{10} + (\frac{x_2}{10})^2$. You can verify that optimal behavior is to spend one's whole budget on good 2 if $w > \frac{10p_2^2}{p_1}$ and one's whole budget on good 1 otherwise. Consider two observations, one at prices $p = (10, 9)$ when the consumers have budgets $w_1 = 80$ and $w_2 = 70$; and a second at prices $(9, 10)$ when the consumers have wealth $w_1 = 120$ and $w_2 = 30$. Aggregate demand would be $(15, 0)$ at the first observation and $(3\frac{1}{3}, 12)$ at the second, which we can confirm would fail GARP.

- Again, this contrasts with producer theory –

  on the very first homework, you all proved that if each individual firm is maximizing profits,

  industry data will be rationalizable as well

- The difference is that the consumer problem has *wealth effects* –

  consumers behave differently at different wealth levels –

  which means that the distribution of wealth among consumers, not just the aggregate wealth level, matters

- At our two observations, total wealth was the same;

  but in this case, consumers rationally demand good two only when they're sufficiently wealthy,

  so shifting the distribution of wealth shifted demand,

  in a way that was inconsistent with GARP

6.1 so when *does* demand aggregate?

- So, when *does* demand aggregate?

  when can we aggregate individual demand choices,

  and get something consistent with a single person's choices according to rational preferences?

- Since the "problem" we just saw had to do with wealth effects,

  it's not surprising that if we eliminate wealth effects, we can get demand to aggregate

- There are two ways we can do this:

  1. give each consumer quasilinear utility (not necessarily the same quasilinear utility)
  2. give each consumer identical, homothetic preferences
     (preferences are homothetic if they "scale up and down" –
     if $x \succsim y$ if and only if $\lambda x \succsim \lambda y$ for any $\lambda > 0$)

- There's a more general condition, but these are the two main cases that satisfy it.

Proposition. Suppose there are n consumers, and consumer i has indirect utility function

$$v_i(p, w_i) \quad = \quad a_i(p) + b(p) w_i$$

Then aggregate demand

$$X \quad = \quad \sum_{i=1}^{n} x^i(p, w_i)$$

is the same as the demand of a single consumer with indirect utility function

$$V(p, W) \quad = \quad \left(\sum_{i=1}^{n} a_i(p) \right) + b(p) W$$

when $W = \sum_i w_i$.

- $v_i = a_i(p) + b(p) w_i$ is called the "Gorman Form"
- note that a_i can differ across i, but $b(\cdot)$ cannot
- So basically, if your consumers have indirect utility functions which are linear in wealth and all have the same slope,

  then you can add up all their demand, treat them like a single person, and everything works

- But otherwise, you generally can't

- The two special cases I mentioned give indirect utility functions that take this Gorman form, and they're both preferences you'll work with on Homework 6
- *Homothetic* preferences are preferences where $x \succsim y$ if and only if $\lambda x \succsim \lambda y$ for every $\lambda > 0$
- You're showing that homothetic preferences can be represented by a utility function which is homogeneous of degree 1, i.e., $u(\lambda x) = \lambda u(x)$
- If utility is homogeneous of degree 1, indirect utility turns out to be linear in wealth,

$$v(p, w) \quad = \quad \tilde{v}(p) w$$

- So if all your consumers had identical preferences, and those preferences were homothetic, then indirect utility takes the Gorman form and demand aggregates

- And you'll also show that this holds for *quasilinear* utility,

  since quasilinear preferences lead to an indirect utility function

$$v(p, w) \quad = \quad a(p) + w$$

  meaning consumers don't even have to have the *same* quasilinear preferences;

  so as long as each consumer's preferences are quasilinear, demand aggregates

- But those are the main cases where this works;

  otherwise, individual wealth effects create enough of a problem that "market demand" doesn't have to behave like the choices of a single, rational representative consumer

7 Recovering Preferences from $x(p, w)$

- We said earlier that if Marshallian demand is homogeneous of degree 0,

  satisfies Walras' Law, and gives a Slutsky matrix that's symmetric and negative semidefinite,

  then it's consistent with rational choice

- So if we observe Marshallian demand $x(p, w)$ at every (p, w),

  and it satisfies these conditions,

  can we recover preferences that rationalize it? Yes!

- Recall first that Marshallian demand is determined by preferences, not utility –

  the solution to the consumer problem is the same for any u representing the same preferences

- So if we start with demand, the best we can hope to do is recover ordinal preferences,

  not a unique cardinal utility function

- Which leaves us with a degree of freedom – we can normalize cardinal utility however we want, since we're going to recover the same preferences however we do it

- So let's fix a reference price vector p^0,

  and normalize utility such that $v(p^0, w) = w$, or $e(p^0, u) = u$

- (That is, since there's no natural scale for utility,

  we can pick a reference price vector, and just define the cardinal level of utility as the amount of money it takes to afford it at prices p_0)

- Now fix a value of u; for each good i, we know that

$$\frac{\partial e}{\partial p_i}(p, u) \quad = \quad h_i(p, u) \quad = \quad x_i(p, e(p, u))$$

- Since $x_i(\cdot, \cdot)$ is a known function (it's observed),

  this is a differential equation where the unknown variable is e

- And since we can do this for each good i,

  we have a system of k differential equations in e,

  plus the initial condition $e(p^0, u) = u$

- It turns out that symmetry of the Slutsky matrix is exactly the condition that this system has a solution (also known as "integrability" conditions);

  and negative semi-definiteness guarantees that the solution e is concave (which it has to be)

- Since we can do this separately for each value of u,

  we've now got a concave **expenditure function** e, defined for all p and u,

  that's consistent with our observed Marshallian demand

- Once we know the expenditure function $e(p, u)$, we can recover preferences from that

- MWG do it one way, very much analogous to how we constructed the outer bound Y^O from the profit function in producer theory[5]

- An alternative I prefer[6] is, once we have the expenditure function, to recover the indirect utility function as the solution to

$$e(p, v(p, w)) \quad = \quad w$$

- And then once we have v, define the utility function by

$$u(x) \quad = \quad \min_{p \gg 0} v(p, 1) \qquad \text{subject to} \qquad p \cdot x = 1$$

- The logic is that if x is optimal at prices p and wealth $w = 1$,

  then $p \cdot x = 1$ (by Walras' Law) and $u(x) = v(p, 1)$;

  and for any other price vector p' such that $p' \cdot x = 1$,

  x is feasible, so $v(p', 1) \geq u(x)$

- So for any x that's optimal for some price level, $u(x)$ must be $\min_{p: p \cdot x = 1} v(p, 1)$

- (For x that are never chosen, this may be optimistic relative to the truth –

  there could be x that are never chosen that have $u(x) < v(p, 1)$ for every p with $p \cdot x = 1$ –

  this just gives a utility function that's consistent with the observed choices.)

[5]Assign utility by $u(x) \geq u$ if and only if $p \cdot x \geq e(p, u)$ for every $p \gg 0$. The proof this works is on MWG p 77
[6]from Varian (1992), *Microeconomic Analysis* (3rd ed.), pp 129-130

- Since we're inferring preferences from choices,

  there's no guarantee we match the original preferences on bundles that never get chosen;

  but we are guaranteed that we've found preferences that generate the Marshallian demand function we started with.

- (Since we're doing all this under the assumption that Marshallian demand is single-valued and differentiable,

  we're more or less thinking about the case where preferences are strictly convex.

  Like with production, if preferences are strictly convex,

  then every strictly-positive bundle will sometimes be optimal,

  and so we'll fully recover the "right" preferences this way.

  If preferences have some non-convexities, we can never learn about preferences over the bundles that are never optimal,

  all we can do is find a set of preferences that would rationalize the observed Marshallian demand.)

8 Welfare effects of price changes

Skipped in lecture – please read on your own!

(Just some definitions you're supposed to know)

- We'll wrap up consumer theory thinking about measuring the welfare effect of price changes

- So, some big policy thing happens, and it results in some price changes from p^0 to p^1
- How do we measure the results?
- In some sense, what people care about is

$$v(p^1, w) - v(p^0, w)$$

  the change in utility they can achieve
- But that can't give us a sensible answer, because it depends on cardinal utility – different utility functions representing the same preferences give different answers, and we know we have no way to infer peoples' cardinal utility from choices

- Instead, we can focus on the change in expenditure functions
- That is, if we can calculate

$$e(p^1, \cdot) - e(p^0, \cdot)$$

  we can say, how much more or less money does it take now, to achieve a particular level of utility?
- This is based purely on preferences – not the utility representation – so it's a better-posed question; and since the answer is in dollars, it allows us to add up across people

- Of course, there's still the question of what target utility level to use, and there are two obvious candidates – the original one, and the new one

8.1 Compensating Variation

- Let's consider a "good" change – a shift to lower prices, meaning it's now cheaper to afford the same utility level as before, or we can now achieve a higher level than we could prior
- **Compensating variation** measures how much more cheaply we can afford the old level of utility now: if we let $u^0 = v(p^0, w)$ and $u^1 = v(p^1, w)$, then

$$CV = e(p^0, u^0) - e(p^1, u^0)$$

- CV also answers the question, now that prices have changed,

  how much money can we take away from you, and leave you as well off as you were before

- Since $e(p^0, u^0) = w$, we can rewrite the previous as

$$w - CV = e(p^1, u^0)$$

  so it follows that

$$v(p^1, w - CV) = v(p^1, e(p^1, u^0)) = u^0$$

  so after a good price change, this is how much money you could lose and still be as well-off as you were before the change[7]

8.2 Equivalent Variation

- **Equivalent Variation** measures how much more cheaply you can now afford the *new* level of utility, relative to what it would have cost you at the old price level
- That is,

$$EV = e(p^0, u^1) - e(p^1, u^1)$$

- Or, how much money could we have given you *instead of* the price change,

  to make you equally happy?

- That is, since $e(p^1, u^1) = w$,

$$w + EV = e(p^0, u^1)$$

  and therefore

$$v(p^0, w + EV) = v(p^0, e(p^0, u^1)) = u^1$$

[7]Since $e(p^0, u^0) = w = e(p^1, u^1)$, CV can also be written as $CV = e(p^1, u^1) - e(p^1, u^0)$, because we're again measuring how much money could you give up right now (at the new prices) and still afford your old lifestyle.

8.3 CV versus EV

- We can think of CV as your maximum willingness to pay to "buy" a price change –

  you could receive the price change, give up CV, and still be as well off as before

- EV might be more appropriate to measure how much better off you are when a change has already happened –

  "this change had the same utility benefit as if you had given me \$50" –

  while CV might be easier for considering lots of alternative policies, since it uses the common baseline u^0

- To see the difference more clearly, think about a dramatic example – how much has quality of life improved since the Middle Ages

- Life expectancy was short – partly because a lot of things we take for granted today weren't available, like antibiotics and clean water

- If you were rich enough, you ate well, but there wasn't a lot on TV

- We can think of prices of many as having come down a lot – perhaps even from infinity –

  if we fix wealth at, say, an upper-middle-class level

- Compensating Variation asks, if you're upper-middle-class today, how much money could you give up, and still afford the quality of life of someone upper-middle-class in the Middle Ages

- Probably an awful lot – you could be pretty poor today, and still have a two-generations-old iPhone, penicillin when you need it, and an apartment that stays warm in the winter,

  but there should be some amount of money we could take away from you, to the point where you're destitute and starving,

  and make you no better off than you would have been in the Middle Ages

- Equivalent Variation asks, if we went back in time to the Middle Ages, how much money would we have had to give you, to let you afford your current 2019 level of utility

- Quite possibly infinite – there's no amount of wealth in 1400 that would have bought you an iPhone, or a Netflix account, or a life expectancy over 60

  (The average life expectancy of English and Scottish kings from 1000 to 1600 was $\sim$ 50 years)

- As we'll see, for normal goods, Equivalent Variation tends to be larger than Compensating Variation

8.4 Changes in just one price

- Suppose that only the price of good i changes, with $p_i^1 < p_i^0$
- We can calculate compensating variation as

$$CV = e(p^0, u^0) - e(p^1, u^0) = \int_{p_i^1}^{p_i^0} \frac{\partial e}{\partial p_i}(p_i, p_{-i}, u^0) dp_i = \int_{p_i^1}^{p_i^0} h_i(p, u^0) dp_i$$

- And we can do the same with equivalent variation:

$$EV = e(p^0, u^1) - e(p^1, u^1) = \int_{p_i^1}^{p_i^0} \frac{\partial e}{\partial p_i}(p_i, p_{-i}, u^1) dp_i = \int_{p_i^1}^{p_i^0} h_i(p, u^1) dp_i$$

- Now, since price went down, we know $u^1 > u^0$
- We can show that if good i is a normal good (x_i increasing in w), then h_i is increasing in u, and so $EV > CV$
- And if good i is an inferior good, then h_i is decreasing in u, so $EV < CV$

8.5 One more measure – consumer surplus

- So, for changes in a single price, CV and EV can be calculated by integrating under the Hicksian demand curve for that good, one at the old utility level and one at the new utility level
- Of course, just for the hell of it, we could integrate under Marshallian demand instead;

  we call that Consumer Surplus
- Consider the following, which shows the demand just for good i as a function of p_i

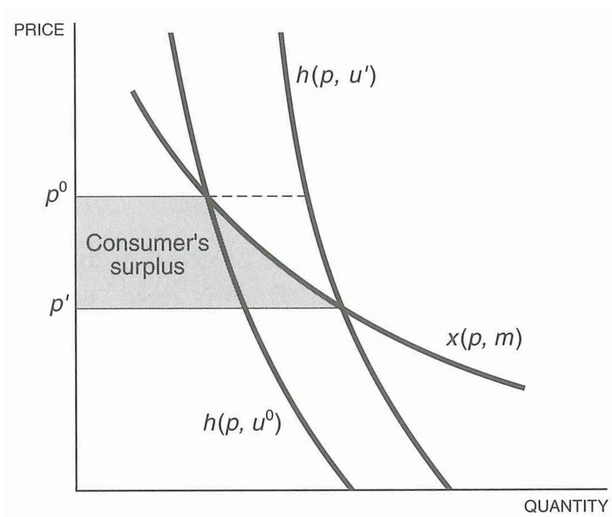

(This diagram is for a normal good; for an inferior good, things would be reversed)

- This, I think, initially appealed because we can observe Marshallian demand but not Hicksian demand, so we have a better shot at a direct calculation
- If i is either a normal or an inferior good, x_i will be between the two Hicksian curves, so CS will be between the other two, and maybe a reasonable estimate for either
- And if there are no wealth effects, then $h(p, u)$ doesn't depend on u, and EV, CV, and CS would all be identical
- But consumer surplus is an interesting measure in its own right

- Consumer Surplus is a standard measure of how much extra value consumers get from a single product – basically, the difference between the price they pay and the price they would have been willing to pay
- If we think of the demand curve for one good as the aggregate demand, each point on the demand curve represents a single consumer with that exact willingness to pay
- The area between a horizontal price line and the demand curve, then – the area above the demand – is the consumer surplus
- And when a price goes down, the change in CS is clear visually

- Two points we already mentioned:
  1. If good i is a normal good, then

     $$EV \quad \geq \quad \Delta CS \quad \geq \quad CV$$

     so we can use EV and CV to put bounds on CS, or we can use CS to approximate either of them
  2. If preferences are quasilinear, then there are no wealth effects (Hicksian demand for good i is the same at u^1 as it was at u^0), and so

     $$EV = CV = \Delta CS$$