

Lecture Note 3 - Theory of Choice and Individual Demand

David Autor, MIT Economics and NBER

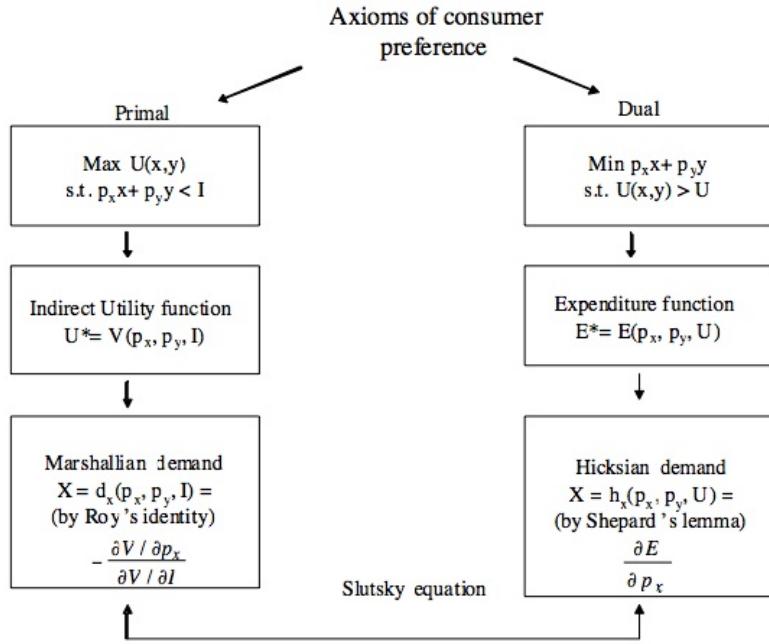
14.03/14.003 Microeconomic Theory and Public Policy, Fall 2025
(Updated on September 15, 2025)

1 Theory of consumer choice

Agenda for next several lectures

1. Utility maximization
2. The Carte Blanche principle
3. Indirect utility function
4. Application: The demand for subsidized health insurance (Finkelstein, Hendren, Shepard, 2019)
5. Expenditure function
6. Relationship between Expenditure function and Indirect utility function
7. Demand functions
8. Application: The employment effects of the Earned Income Tax Credit (EITC) (Eissa and Leibman, 1996)
9. Income and substitution effects
10. Normal and inferior goods
11. Compensated and uncompensated demand (Hicksian, Marshallian)
12. Application: Giffen goods and subsistence consumption (Jensen and Miller, 2008)

Theory roadmap for the next couple lectures:



2 Utility maximization subject to budget constraint

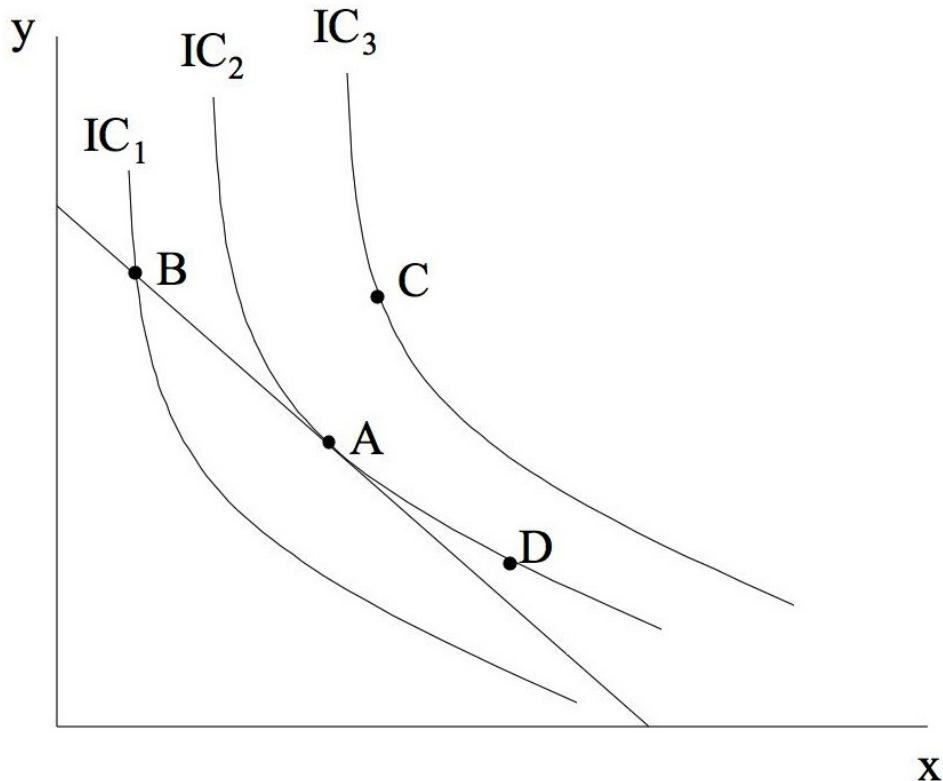
Ingredients

- Utility function (preferences)
- Budget constraint
- Price vector

Consumer's problem

- Maximize utility subject to budget constraint.
- Characteristics of solution:
 - Budget exhaustion (non-satiation)
 - For most solutions: psychic trade-off = monetary payoff
 - Psychic trade-off is MRS
 - Monetary trade-off is the price ratio
- From a visual point of view utility maximization corresponds to point A in the diagram below

- The slope of the budget set is equal to $-\frac{p_x}{p_y}$
- The slope of each indifference curves is given by the MRS

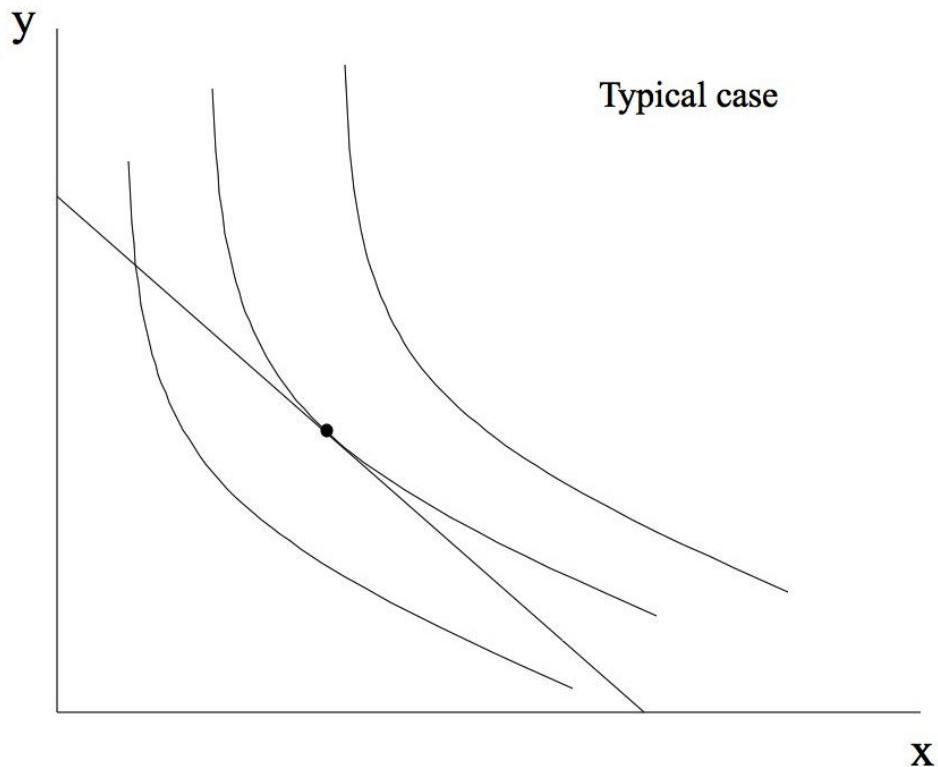


- We can see that $A \succ B$, $A \succ D$, $C \succ A$. Why should one choose A ?

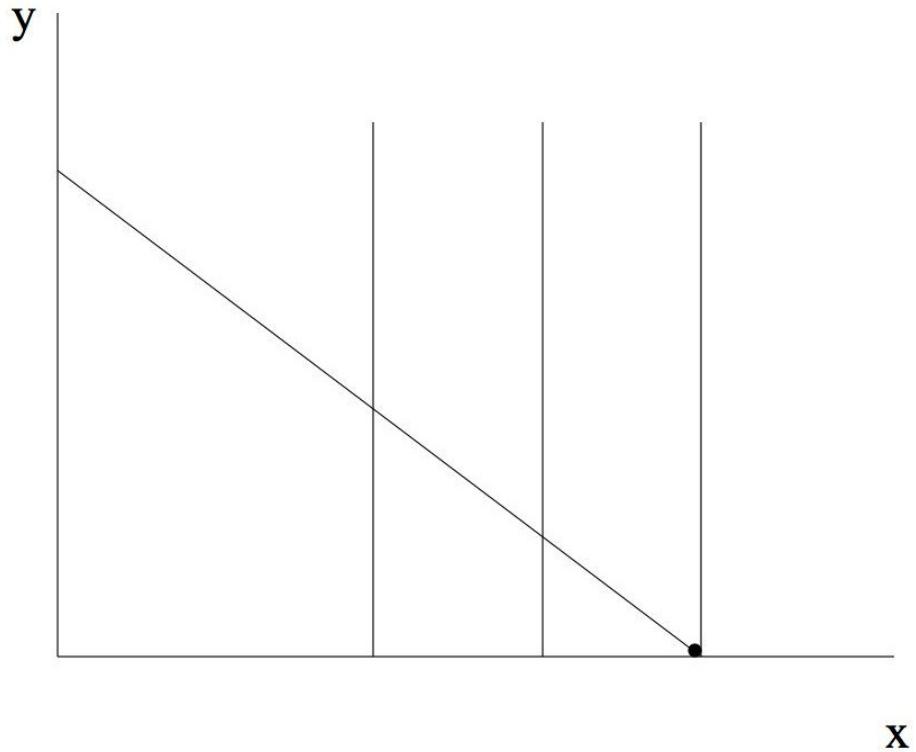
2.1 Interior and corner solutions

There are two types of solution to this problem, interior solutions and corner solutions

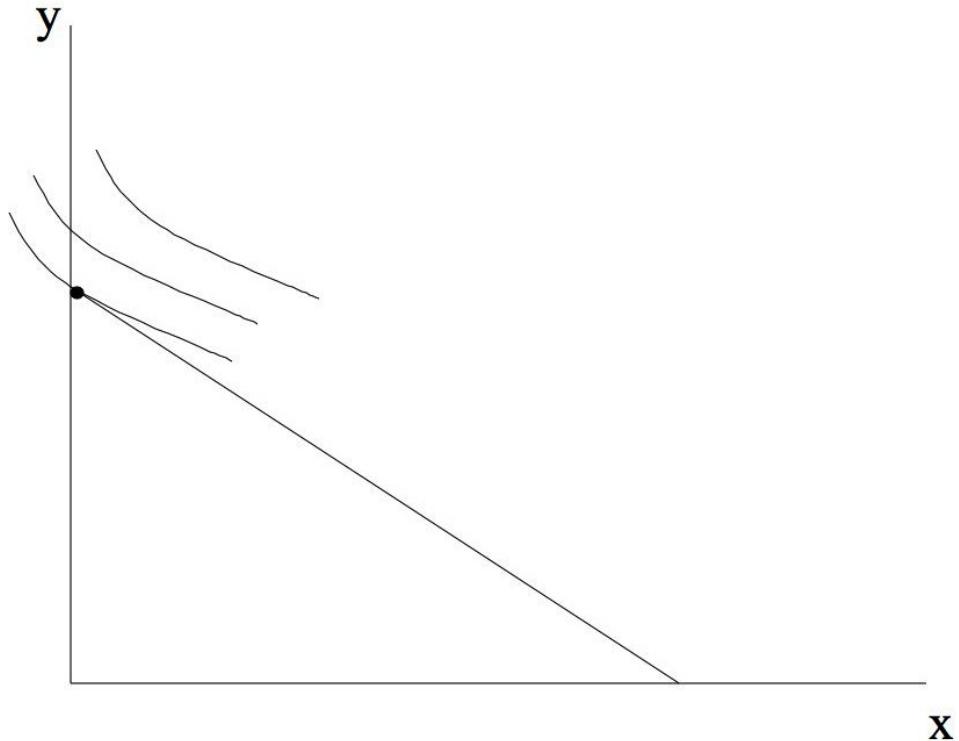
- The figure below depicts an interior solution



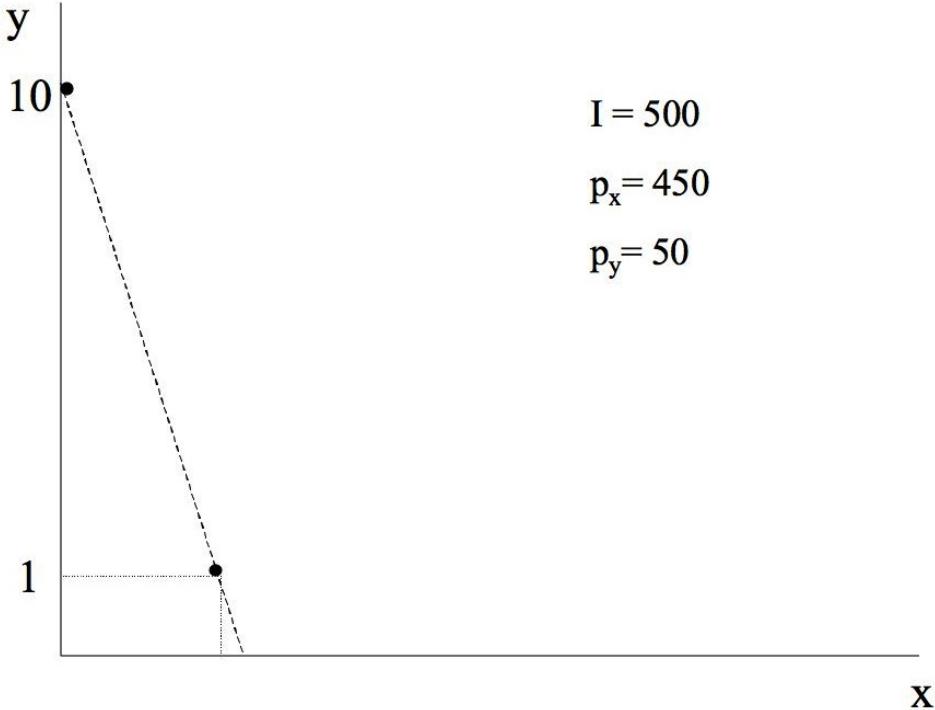
- The next figure depicts a corner solution. In this specific example the shape of the indifference curves means that the consumer is indifferent to the consumption of good y . Utility increases only with consumption of x . Thus, the consumer purchases x exclusively.



- In the following figure, the consumer's preference for y is sufficiently strong relative to x that the the psychic trade-off is always lower than the monetary trade-off. (This must be the z case for many products that we don't buy.)



- What this means is that the corners (more precisely, the axes), serve as constraints. The consumer would prefer to choose a bundle with negative quantities of x and positive quantities of y . That's not feasible. But to solve the problem using the Lagrangian method, we'd need to specifically impose these non-negativity constraints to prevent a non-sensical solution.
- Another type of “corner” solution can result from indivisibilities the bundle (often called integer constraints).



- Given the budget and set of prices, only two bundles are feasible—unless the consumer could purchase non-integer quantities of good x . We normally abstract from indivisibility.
- Going back to the general case, how do we know a solution exists for consumer, i. e. how do we know the consumer can choose? The axiom of completeness guarantees this. Every bundle is on some indifference curve and can therefore be ranked: $A^I B, A \succ B, B \succ A$.

2.2 Mathematical solution to the Consumer's Problem

- Mathematics

$$\begin{aligned}
 & \max_{x,y} U(x, y) \\
 \text{s.t. } & p_x x + p_y y \leq I \\
 & L = U(x, y) + \lambda(I - p_x x - p_y y) \\
 1. \quad & \frac{\partial L}{\partial x} = U_x - \lambda p_x = 0 \\
 2. \quad & \frac{\partial L}{\partial y} = U_y - \lambda p_y = 0 \\
 3. \quad & \frac{\partial L}{\partial \lambda} = I - p_x x - p_y y = 0
 \end{aligned}$$

- Rearranging (1) and (2):

$$\frac{U_x}{U_y} = \frac{p_x}{p_y}$$

This means that the psychic trade-off is equal to the monetary trade-off between the two goods.

- Equation (3) states that budget is exhausted (non-satiation).
- Also notice that:

$$\begin{aligned}\frac{U_x}{p_x} &= \lambda \\ \frac{U_y}{p_y} &= \lambda\end{aligned}$$

- What is the meaning of λ ?

2.3 Interpretation of λ , the Lagrange multiplier

- At the solution of the Consumer's problem (more specifically, an interior solution), the following conditions will hold:

$$\frac{\partial U / \partial x_1}{p_1} = \frac{\partial U / \partial x_2}{p_2} = \dots = \frac{\partial U / \partial x_n}{p_n} = \lambda$$

This expression says that at the utility-maximizing point, the next dollar spent on each good yields the same marginal utility.

- So what is $\frac{dU^*}{dI}$? Return to Lagrangian:

$$\begin{aligned}L &= U(x, y) + \lambda(I - p_x x - p_y y) \\ \frac{\partial L}{\partial x} &= U_x - \lambda p_x = 0 \\ \frac{\partial L}{\partial y} &= U_y - \lambda p_y = 0 \\ \frac{\partial L}{\partial \lambda} &= I - p_x x - p_y y = 0 \\ \frac{dL}{dI} &= \left(U_x \frac{\partial x^*}{\partial I} - \lambda p_x \frac{\partial x^*}{\partial I} \right) + \left(U_y \frac{\partial y^*}{\partial I} - \lambda p_y \frac{\partial y^*}{\partial I} \right) + \lambda\end{aligned}$$

By substituting $\lambda = \frac{U_x}{p_x} \Big|_{x=x^*}$ and $\lambda = \frac{U_y}{p_y} \Big|_{y=y^*}$, we see that both expressions in parenthesis are zero.

- We conclude that:

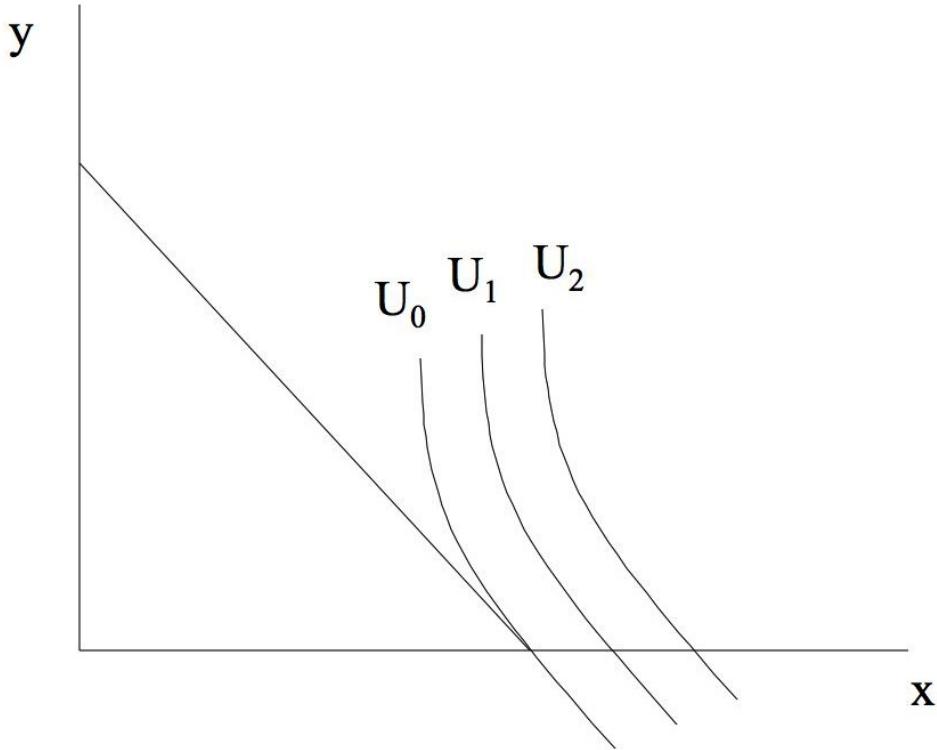
$$\frac{dL}{dI} = \frac{\partial L}{\partial I} = \lambda$$

λ equals the “shadow price” of the budget constraint, i.e. it expresses the quantity of utils that could be obtained with the next dollar of consumption. Note that this expression only holds when $x = x^*$ and $y = y^*$. If x and y were not at their optimal values, then the total derivative of L with respect to I would also include additional cross-partial terms. These cross-partial terms are zero at $x = x^*$ and $y = y^*$.

- What does the “shadow price” mean? It’s essentially the “utility value” of relaxing the budget constraint by one unit (e.g., one dollar).
- Note that this shadow price is not uniquely defined since it corresponds to the marginal utility of income in “utils”—an ordinal value. Thus, the shadow price is defined only up to a monotonic transformation.
- We could also have determined that $dL/dI = \lambda$ without calculations by applying the envelope theorem. Note that the envelope theorem for constrained problems says that $\frac{dU^*}{dI} = \frac{\partial L}{\partial I} = \lambda$. At the utility maximizing solution to this problem, x^* and y^* are already optimized and so an infinitesimal change in I does not alter these choices. Hence, the effect of I on U depends only on its direct effect on the budget constraint and does not depend on its indirect effect (due to re-optimization) on the choices of x and y . This “envelope” result is only true in a small neighborhood around the solution to the original problem.

2.4 Corner solutions

- When at a corner solution, consumer buys zero of some good and spends the entire budget on other goods.
- What problem does this create for the Lagrangian?



- The problem above is that a point of tangency doesn't exist for positive values of y . Hence we also need to impose "non-negativity constraints": $x \geq 0, y \geq 0$. This will not be important for problems in this class, but it's easy to add these constraints to the maximization problem.

2.5 An Example Problem

- Consider the following example problem:

$$U(x, y) = \frac{1}{4} \ln x + \frac{3}{4} \ln y$$

- Notice that this utility function satisfies all axioms:

1. Completeness, transitivity, continuity [these are pretty obvious]
2. Non-satiation: $U_x = \frac{1}{4x} > 0$ for all $x > 0$. $U_y = \frac{3}{4y} > 0$ for all $y > 0$. In other words, utility rises continually with greater consumption of either good, though the rate at which it rises declines (diminishing marginal utility of consumption).
3. Diminishing marginal rate of substitution:
 - Along an indifference curve of this utility function: $\bar{U} = \frac{1}{4} \ln x_0 + \frac{3}{4} \ln y_0$.

- Totally differentiate: $0 = \frac{1}{4x_0}dx + \frac{3}{4y_0}dy$.
- Which provides the marginal rate of substitution $-\frac{dy}{dx}|_{\bar{U}} = \frac{U_x}{U_y} = \frac{4y_0}{12x_0}$.
- The marginal rate of substitution of x for y is increasing in the amount of y consumed and decreasing in the amount of x consumed; holding utility constant, the more y the consumer has, the more y he would give up for one additional unit of x .
- Example values: $p_x = 1$, $p_y = 2$, $I = 12$. Write the Lagrangian for this utility function given prices and income:

$$\begin{aligned}
 & \max_{x,y} U(x,y) \\
 \text{s.t. } & p_x x + p_y y \leq I \\
 L &= \frac{1}{4} \ln x + \frac{3}{4} \ln y + \lambda(12 - x - 2y) \\
 1. \quad \frac{\partial L}{\partial x} &= \frac{1}{4x} - \lambda = 0 \\
 2. \quad \frac{\partial L}{\partial y} &= \frac{3}{4y} - 2\lambda = 0 \\
 3. \quad \frac{\partial L}{\partial \lambda} &= 12 - x - 2y = 0
 \end{aligned}$$

- Rearranging (1) and (2), we have

$$\begin{aligned}
 \frac{U_x}{U_y} &= \frac{p_x}{p_y} \\
 \frac{1/4x}{3/4y} &= \frac{1}{2}
 \end{aligned}$$

- The interpretation of this expression is that the MRS (psychic trade-off) is equal to the market trade-off (price-ratio).
- What's $\frac{dL}{dI}$? As before, this is equal to λ , which from (1) and (2) is equal to:

$$\lambda = \frac{1}{4x^*} = \frac{3}{8y^*}.$$

The next dollar of income could buy one additional x , which has marginal utility $\frac{1}{4x^*}$ or it could buy $\frac{1}{2}$ additional y , which provides marginal utility $\frac{3}{4y^*}$ (so, the marginal utility increment is $\frac{1}{2} \cdot \frac{3}{4y^*}$).

- It's important that $dL/dI = \lambda$ is defined in terms of the optimally chosen x^*, y^* . Unless we are at these optima, the envelope theorem does not apply. In that case, dL/dI would also depend on the cross-partial terms: $(U_x \frac{\partial x}{\partial I} - \lambda p_x \frac{\partial x}{\partial I}) + (U_y \frac{\partial y}{\partial I} - \lambda p_y \frac{\partial y}{\partial I})$.
- Incidentally, you should be able to solve for the prices and budget given, $x^* = 3, y^* = 4.5$.

- Having solved that, you can verify that $\frac{1}{4x^*} = \frac{3}{8y^*} = \lambda$. That is, at prices $p_x = 1$ and $p_y = 2$ and consumption choices $x^* = 3, y^* = 4.5$, the marginal utility of a dollar spent on either good x or good y is identical.

2.6 Lagrangian with Non-negativity Constraints [Optional]

$$\begin{aligned}
 & \max U(x, y) \\
 \text{s.t. } & p_x x + p_y y \leq I \\
 & y \geq 0 \\
 L & = U(x, y) + \lambda(I - p_x x - p_y y) + \mu(y - 0) \\
 \frac{\partial L}{\partial x} & = U_x - \lambda p_x = 0 \\
 \frac{\partial L}{\partial y} & = U_y - \lambda p_y + \mu = 0 \\
 \mu y & = 0
 \end{aligned}$$

- Final equation above implies that $\mu = 0, y = 0$, or both. (This is called a “complementary slackness” condition; either the constraint is slack, implying $\mu = 0$, or the constraint is binding, implying that $y = 0$, and so in either case, we have that the product $\mu y = 0$.)
- We then have three cases.

- $y = 0, \mu \neq 0$ (since $\mu \geq 0$ then it must be the case that $\mu > 0$)

$$\begin{aligned}
 U_y - \lambda p_y + \mu & = 0 \rightarrow U_y - \lambda p_y < 0 \\
 \frac{U_y}{p_y} & < \lambda \\
 \frac{U_x}{p_x} & = \lambda
 \end{aligned}$$

Combining the last two expressions:

$$\frac{U_x}{U_y} > \frac{p_x}{p_y}$$

This consumer would like to consume even more x and less y , but she cannot.

- $y \neq 0, \mu = 0$

$$\begin{aligned}
 U_y - \lambda p_y + \mu & = 0 \rightarrow U_y - \lambda p_y = 0 \\
 \frac{U_y}{p_y} & = \frac{U_x}{p_x} = \lambda
 \end{aligned}$$

Standard FOC, here the non-negativity constraint is not binding.

3. $y = 0, \mu = 0$

Same FOC as before:

$$\frac{p_x}{p_y} = \frac{U_x}{U_y}$$

Here the non-negativity constraint is satisfied with equality so it doesn't distort consumption.

3 Indirect Utility Function

- For any:
 - Budget constraint
 - Utility function
 - Set of prices

We obtain a set of optimally chosen quantities:

$$\begin{aligned} x_1^* &= x_1(p_1, p_2, \dots, p_n, I) \\ &\dots \\ x_n^* &= x_n(p_1, p_2, \dots, p_n, I) \end{aligned}$$

So when we say

$$\max U(x_1, \dots, x_n) \text{ s.t. } PX \leq I$$

we get as a result:

$$U(x_1^*(p_1, \dots, p_n, I), \dots, x_n^*(p_1, \dots, p_n, I)) \equiv V(p_1, \dots, p_n, I).$$

We call $V(\cdot)$ the “Indirect Utility Function.” This is the value of maximized utility under given prices and income.

- So remember the distinction:
 - Direct utility: utility from consumption of (x_1, \dots, x_n)
 - Indirect utility: utility obtained when facing (p_1, \dots, p_n, I)

- Example

$$\begin{aligned}
\max U(x, y) &= x^{0.5}y^{0.5} \\
\text{s.t. } p_x x + p_y y &\leq I \\
L &= x^{0.5}y^{0.5} + \lambda(I - p_x x - p_y y) \\
\frac{\partial L}{\partial x} &= .5x^{-0.5}y^{0.5} - \lambda p_x = 0 \\
\frac{\partial L}{\partial y} &= .5x^{0.5}y^{-0.5} - \lambda p_y = 0 \\
\frac{\partial L}{\partial \lambda} &= I - p_x x - p_y y = 0
\end{aligned}$$

- We obtain the following:

$$\lambda = \frac{.5x^{-0.5}y^{0.5}}{p_x} = \frac{.5x^{0.5}y^{-0.5}}{p_y},$$

which simplifies to:

$$x = \frac{p_y y}{p_x}.$$

- Substituting into the budget constraint gives us

$$\begin{aligned}
I - p_x \frac{p_y y}{p_x} - p_y y &= 0 \\
p_y y &= \frac{1}{2}I, \quad p_x x = \frac{1}{2}I
\end{aligned}$$

$$x^* = \frac{I}{2p_x}, \quad y^* = \frac{I}{2p_y}$$

Half of the budget goes to each good.

- Thus, for a consumer with $U(x, y) = x^{0.5}y^{0.5}$, budget I , and facing prices p_x and p_y will choose x^* and y^* and obtain utility:

$$U(x^*, y^*) = \left(\frac{I}{2p_x}\right)^{0.5} \left(\frac{I}{2p_y}\right)^{0.5}.$$

Thus, the indirect utility for this consumer is

$$V(p_x, p_y, I) = U(x^*(p_x, p_y, I), y^*(p_x, p_y, I)) = \left(\frac{I}{2p_x}\right)^{0.5} \left(\frac{I}{2p_y}\right)^{0.5}$$

- Why bother calculating the indirect utility function? It saves us time. Instead of recalculating the utility level for every set of prices and budget constraints, we can plug in prices and income to get consumer utility. This comes in handy when working with individual demand functions. Demand functions give the quantity of goods purchased by a given consumer as a function of prices and income (or utility).

4 The Carte Blanche Principle

- One immediate implication of consumer theory is that consumers make optimal choices for themselves given prices, constraints, and income. [Generally, the only constraint is that they can't spend more than their income, but we'll see examples where there are additional constraints.]
- This observation gives rise to the Carte Blanche principle: consumers are always weakly better off receiving a cash transfer than an in-kind transfer of identical monetary value. [Weakly better off in that they may be indifferent between the two.]
- With cash, consumers have Carte Blanche to purchase whatever bundle of goods or services they can afford – including the good or service that alternatively could have been transferred to them in-kind.
- Prominent examples of in-kind transfers given to U.S. citizens include Supplemental Nutrition Assistance Program (SNAP—formerly known as Food Stamps), housing vouchers, health insurance (Medicaid), subsidized educational loans, child care services, job training, etc. [An exhaustive list would be long indeed.]
- Economic theory suggests that, relative to the equivalent cash transfer, these in-kind transfers serve as *constraints* on consumer choice. If consumers are rational, constraints on choice cannot be beneficial.
- For example, consider a consumer who has income $I = 100$ and faces the choice of two goods, food and housing, at prices p_f, p_h , each priced at 1 per unit. The consumer's problem is

$$\begin{aligned} & \max_{f,h} U(f, h) \\ \text{s.t. } & f + h \leq 100 \end{aligned}$$

- The government decides to provide a housing subsidy of 50. This means that the consumer can now purchase up to 150 units of housing but no more than 100 units of food. The consumer's problem is:

$$\begin{aligned} & \max_{f,h} U(f, h) \\ \text{s.t. } & f + h \leq 150 \\ & h \geq 50. \end{aligned}$$

- Alternatively, if the government had provided 50 dollars in cash instead, the problem would be:

$$\begin{aligned} & \max_{f,h} U(f, h) \\ \text{s.t. } & f + h \leq 150. \end{aligned}$$

- The government's transfer therefore has two components:
 1. An expansion of the budget set from I to $I' = I + 50$.
 2. The imposition of the constraint that $h \geq 50$.
- The canonical economist's question is: why do both (1) and (2) when you can just do (1) and potentially improve consumer welfare at no additional cost to the government? (Of course, I don't expect you to accept this argument as gospel truth. But it's a good default position—better, perhaps, than the alternative default that it's better for the government to dictate choices to consumers than to allow them to make them for themselves.)

4.1 An application: The demand for subsidized health insurance

Relative to most other industrialized countries, the U.S. has a unique set of institutional arrangements for providing healthcare and health insurance.

- The U.S. spends a far larger share of Gross Domestic Product on healthcare than any other country. The OECD estimates that the U.S. spent almost 17% of GDP on healthcare in 2019. By way of comparison, Germany spent 11.8% of GDP on healthcare in 2019, the highest among all OECD countries. The OECD-wide average was 9.0%. The U.S. spends essentially twice the OECD-average, and about half again as much as the next spendiest¹ country!
- Health insurance for working-age adults is primarily provided through employers as a 'fringe benefit' rather than through either a public insurance system or a direct-to-household system. As a result, when U.S. workers lose their jobs, they are also at substantial risk of losing their health insurance—which seems like a suboptimal insurance system. (At age 65, U.S. adults become eligible for Medicare, which is a generous, largely publicly-paid insurance program.)
- Due to its high cost and (in many cases) the work-contingent provision, a substantial fraction of U.S. adults lacks health insurance. In 2019, this fraction was estimated at 10.9% among non-elderly Americans ages 0-64. In 2013, this fraction was 16.8%, but there was a dramatic fall after 2013. (Do you know why?)

The U.S. federal government and many state governments subsidize health insurance premiums to boost health insurance enrollment among low-income individuals and households. Subsidized enrollees pay less than their full insurance premium—which should be roughly equal to their *expected* healthcare cost—and the government pays the difference. (I emphasize *expected* healthcare costs because realized healthcare costs may be much higher *or* much lower than expected—that's what the insurance policy is all about. We'll talk much more about insurance later in the semester.)

¹Note: *spendiest* is a word in the Scrabble dictionary.

Is subsidizing health insurance premiums for low income households a good policy? This is a complex question, and there are many things you might want to consider before you reach a conclusion. I won't enumerate what those might be, but we'll discuss some of them in one class. One thing you'd certainly want to know is how much consumers value these health insurance subsidies. Stated differently, how much would they be willing to pay (WTP) for the insurance policy were it not subsidized. Willingness to pay is a measure of the psychic value that consumers place on having the insurance policy relative to alternative uses of the money. For example, let's say that insurance policy costs \$400 per month and a consumer would be willing to pay \$399 for it. She won't buy the policy at its sticker price. But given a small subsidy (\$1.01) she would do so, though she should be close to indifferent.

The 2019 *AER* paper by Finkelstein, Hendren, Shepard provides evidence from the Massachusetts Commonwealth Care program on what low-income consumers are WTP for health insurance.² Commonwealth Care is a subsidized insurance market for low income families in Massachusetts. Figure 2 of the paper, reproduced below, provides key details on how the subsidy policy works. The x -axis in this figure, *FPL*, is household income as a percentage of the Federal Poverty Line. A household at 100% FPL is exactly at the poverty line, and a household at 200% FPL is at twice that line. The FPL is quite low. In 2022, the U.S. annual FPL for a family of four was \$27,250. It is difficult to believe that a family of four could cover rent in Cambridge at that annual income level, let alone food, electricity, heat, clothing, transportation, education, and health insurance.

²The paper actually provides evidence on a *number* of deeply important questions about the demand for insurance. We will not focus on most of those questions at present, though we may return to them later in the semester.

Panel A. Premiums for cheapest plan, 2009–2013 Panel B. Prices, subsidies, and premiums in 2011

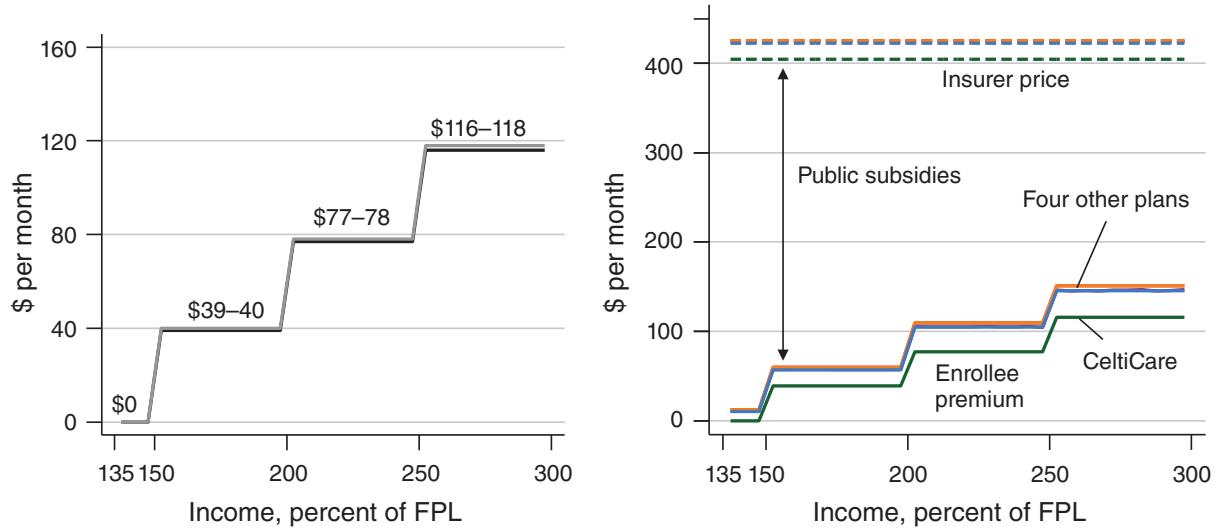


FIGURE 2. INSURER PRICES AND ENROLLEE PREMIUMS IN COMM CARE MARKET

Notes: Panel A plots enrollee premiums for the cheapest plan by income as a percent of FPL, noting the thresholds (150 percent, 200 percent, and 250 percent of FPL) where the amount increases discretely. The black lines show the values that applied in 2009–2012; the gray lines show the (slightly higher) values for 2013. Panel B shows insurer prices (dotted lines) and enrollee premiums (solid lines) for the five plans in 2011. In this year, four insurers set prices within \$3 of a \$426/month price cap, while CeltiCare set a lower price (\$405) and therefore had lower enrollee premiums.

In 2009 – 2013, the cost of an insurance policy under Commonwealth Care was within rounding error of \$400. The out of pocket cost for consumers eligible for Commonwealth care ranged from \$0 for consumers with incomes below 150% FPL to \$116 for consumers with incomes exceeding 250% FPL. The discontinuity points in this subsidy schedule turn out to be very useful for economic analysis. Here's why. It's very likely that WTP for healthcare rises with income. So, consumers with 200% FPL have higher WTP for insurance on average than consumers with 100% FPL. As we move along the x-axis, WTP is rising while subsidies are falling. This makes it hard to know how much consumers would be willing to pay for health insurance at a given income level. We cannot simply compare consumers at 150% FPL paying the \$40 premium to those at 250% FPL paying the \$116 premium to determine the effect of the subsidy on demand. Almost surely, consumers with FPL 250% would be more likely to buy the policy than consumers with FPL 150%, with or without the subsidy.

That's where the discontinuities in (see technical discussion below). We don't expect WTP for health insurance to jump discreetly when FPL crosses some arbitrary threshold—but prices do. Hence, we can use these discontinuities to compare consumers with almost the same incomes who face different prices for health insurance. For example, by comparing consumers who are just above versus just below 200% of FPL, we can determine what fraction are willing to pay, at least \$40 but less than \$77. Using these comparisons at multiple discontinuities, we can construct an estimate of

WTP for the population of eligible consumers (i.e., their demand curve).

Figure 5 provides some key evidence: the number of consumers purchasing insurance falls *sharply* at each discontinuity. The fraction purchasing falls by 37% as income crosses the 200% FPL level. This implies that 37% of consumers value the policy at at least \$40 per month but less than \$77 per month. Question: what assumption am I implicitly invoking about the ‘exchangeability’ of consumers on the left and right-hand sides of the 200% FPL threshold when I make this inference? Can you express this assumption using potential outcomes notation?

Panel A. Average monthly enrollment by income

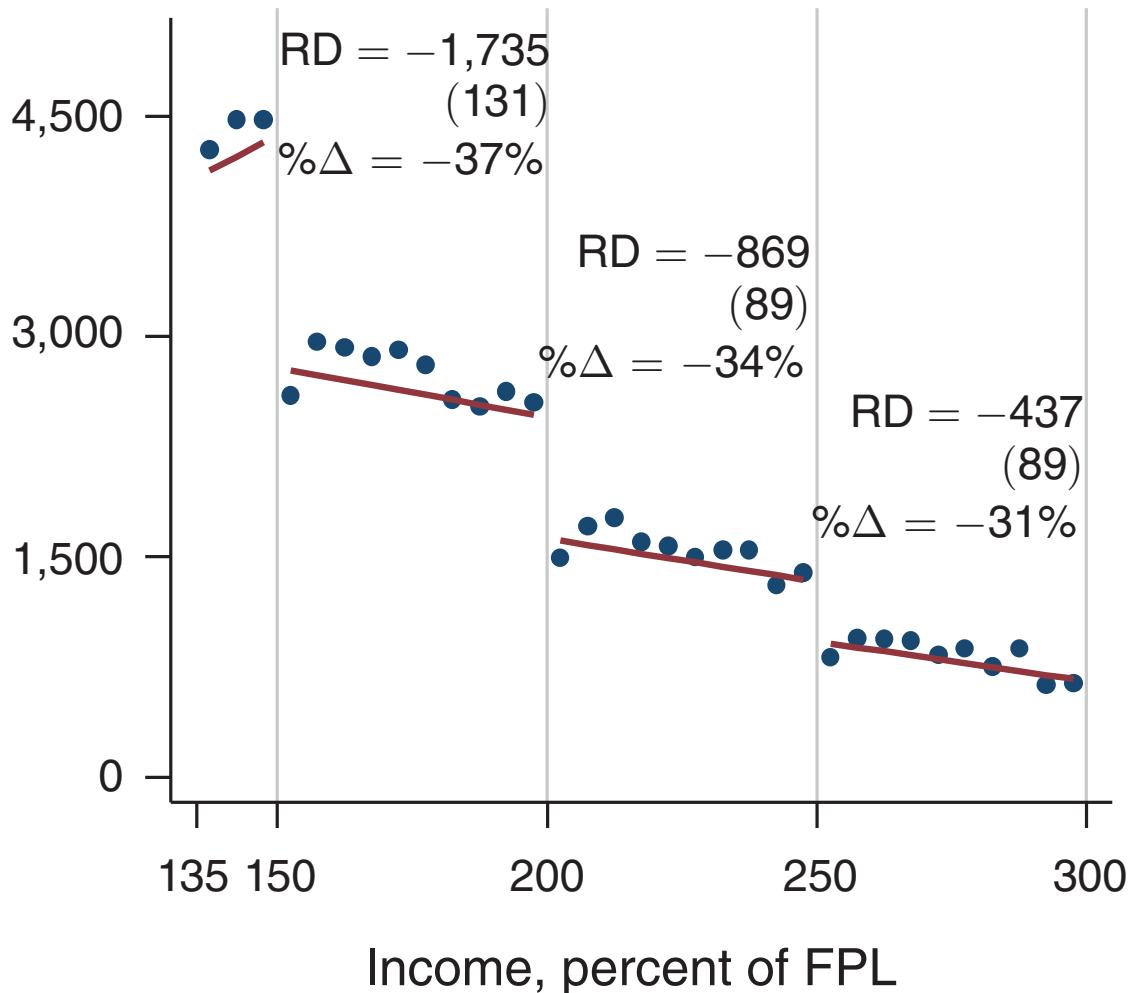


FIGURE 5. COMM CARE ENROLLMENT AND AVERAGE INSURER COSTS, 2009–2013

Notes: The figure shows discontinuities in enrollment and average insurer costs at the income thresholds (150 percent, 200 percent, and 250 percent of FPL) at which enrollee premiums increase (see Figure 2). Panel A shows average enrollment in CommCare (total member-months, divided by number of months) by income over the 2009–2013 period our data span. Panel B shows average insurer medical costs per month across all CommCare plans over the same period. In each figure, the dots represent raw values for a 5 percent of FPL bin, and the lines are predicted lines from our linear RD specification in equation (1). RD estimates and robust standard errors (in parentheses) are labeled just to the right of each discontinuity; percent changes relative to the value just below the discontinuity are labeled as %Δ =.

You may wonder: why does the number of insurance purchasers slope downward between the discontinuity thresholds (e.g., between 151% and 199% of FPL) even though incomes are rising in this range. All else equal, shouldn't higher incomes raise demand for health insurance? The answer is yes. Figure 4 of FHS clarifies what's going on: the number of Massachusetts residents who are

eligible for Commonwealth Care is falling as incomes rise—probably because many of those with higher incomes are currently employed and hence ineligible for the subsidies. If you squint at this figure, you can see that the take-up rate of Commonwealth Care (the share purchasing among those eligible) is generally *rising* with FPL, as we would expect, but then jumps down discreetly at each discontinuity in the subsidy schedule.

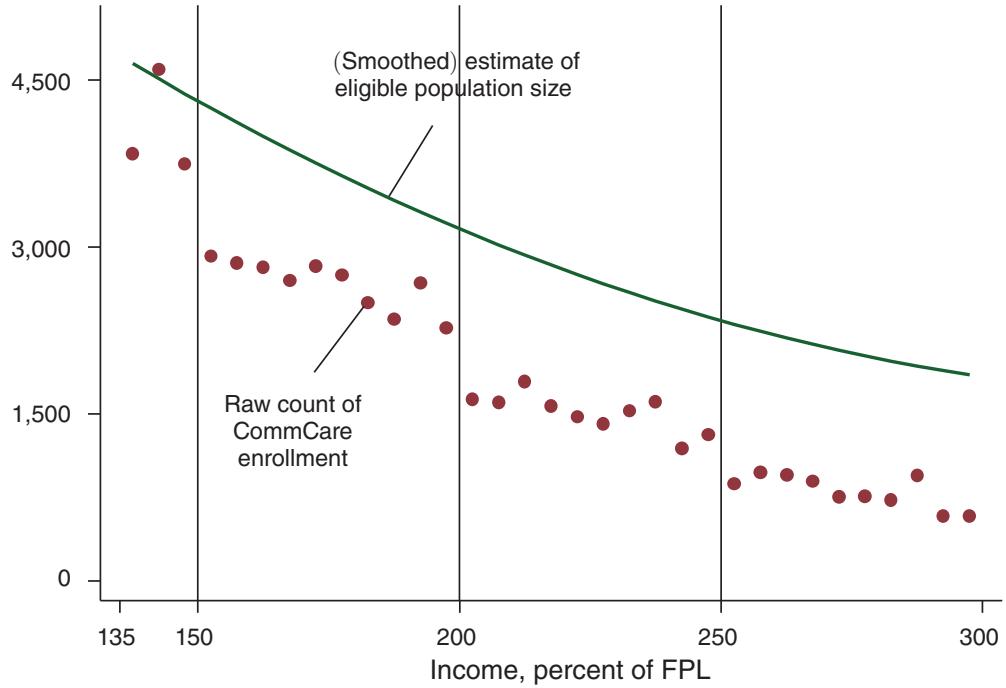


FIGURE 4. ELIGIBLE AND ENROLLED POPULATION, 2011

Notes: Figure shows our (smoothed) estimate of the CommCare-eligible population in 2011 (based on ACS data), and raw enrollment counts in CommCare in 2011 by bins of 5 percent of the FPL.

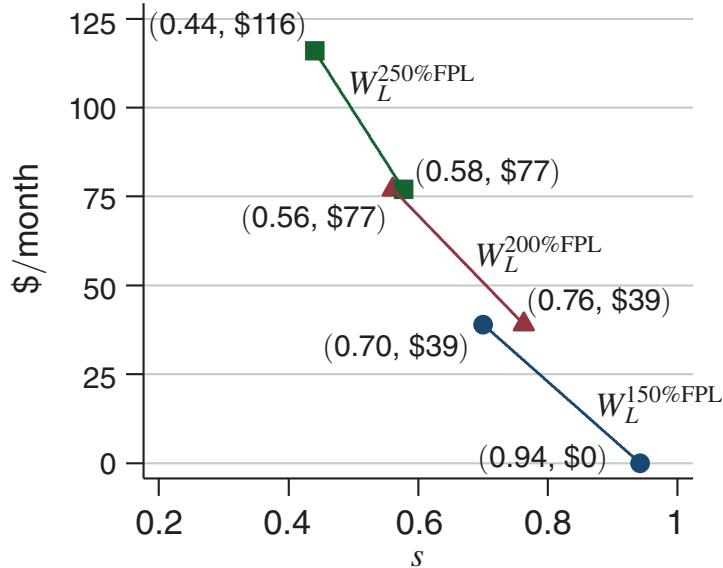
Building on this evidence, as well as a number of other empirical insights and plausible assumptions, FHS construct an estimate of the demand curve for health insurance in the Commonwealth Care eligible population (Figure 10).³ Panel A plots six points corresponding to the Y_1 and Y_0 for the three different regression discontinuity (RD) estimates. Because those three RDs are estimated for potential enrollees with different incomes, Panel C translates them into a single demand curve. (In particular, the authors are adjusting for the fact that insurance demand may change with income itself. We saw in Figure 5 that there's only a modest downward trend over income levels between the discontinuity thresholds, so this translation is small.)

These estimates imply that 94% of eligible consumers would ‘buy’ the policy if it were free (perhaps 6% cannot get organized to do even that) but only 36% would purchase it if the cost was

³For a professional economist, this paper is highly readable and admirably simple—given the complexity of the subject matter. But the paper is far more technically complex than, say, Card and Krueger’s 1994 article, and I would not expect 14.03/003 students to master this paper without guidance.

116. This evidence suggests that *almost no one* in the eligible population is willing to pay even close to the full policy cost.⁴

Panel A. W_L (based on $1 - D_U$)



Panel C. Adjusted W_L

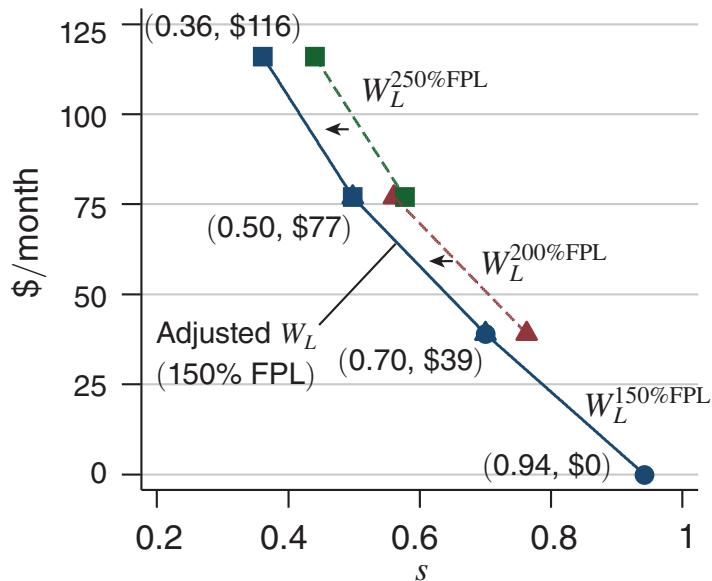


FIGURE 10. WILLINGNESS TO PAY CURVES: EMPIRICAL

Like SNAP, the Commonwealth Care subsidy is an *in-kind* transfer: the government does not give consumers cash but merely pays a part of their insurance premium. This subsidy loosens the

⁴ Additionally, FHS separately considering demand for a high-quality and low-quality policy, both of which were offered by Commonwealth care—but this distinction turns out to be not very important in practice.

budget constraint for consumers who would otherwise have purchased health insurance. But for consumers who would not have otherwise purchased insurance, the subsidy actually requires them to spend money to receive the benefit. (If they purchase the subsidized policy, we can infer that they are better off as a result, even though they have to increase out-of-pocket expenditures to obtain it. Conversely, if they don't take up the subsidy, they are no worse off, and the government faces no direct cost.) This points to an important difference between health insurance subsidies and more conventional subsidies. Unlike say gallons of gasoline or pounds of cheddar cheese, health insurance is an indivisible good—you either buy it or not. Consumers cannot purchase insurance until the marginal utility of the last dollar of insurance is equal to the marginal utility of the last dollar spent on any/all other goods. This indivisibility potentially makes the efficiency costs of health insurance subsidies quite large. (Note that there are some gradations of quality and price in the health insurance market place, but the range of price/quality is small compared to the price difference between any insurance and no insurance.)

Take a stark case. Recall that the health insurance policy cost was roughly \$400 per month in Massachusetts between 2009 and 2013. The FHS evidence demonstrates that many low-income consumers would not buy health insurance at full price. (Indeed, Figure 10 suggests that essentially no one would buy at this price.) For consumers with incomes between 150% and 00% of poverty, the government offers a \$360 subsidy, while the consumer's out of pocket cost is \$40 *if* she chooses to buy the policy. Subsidized consumers will therefore purchase the policy if they value it at $\geq \$40.01$. In the extreme case where the consumer values the policy at exactly \$40.01, she is close to indifferent after the \$360 subsidy is paid (the consumer's surplus is \$0.01). Arguably, the government is “wasting” the entire subsidy, since in the end, the consumer does not perceive herself to be better off. Of course, “wasting” is too strong a term—though useful for our *Carte Blanche* discussion—since the government may value citizens having insurance policies even if those citizens don't really want to pay much for them. We'll talk about this in class.

Compare this health insurance subsidy to the case of an in-kind transfer like SNAP. Food is a divisible commodity and pretty much everyone spends money on it. If the government provides the consumer with \$360 in monthly SNAP benefits that can be spent *only* on food, a lower bound on the benefit the consumer receives from that in-kind transfer is the cash amount she would have spent on food absent SNAP. This is surely a meaningful number, probably a couple of hundred dollars a month (at a minimum). Hence, it's very hard to imagine a case where a consumer is essentially indifferent between receiving and non receiving SNAP. This is an economically consequential difference between SNAP and the Commonwealth Care subsidy.

5 The Regression Discontinuity Design

Finkelstein, Hendren, and Shepard exploit discontinuities in the CommCare subsidy schedule to estimate the willingness to pay (WTP) for insurance among low-income Massachusetts residents. Specifically, they employ a Regression Discontinuity (RD) design. This section explains the RD

design in more detail. You can add RD to your toolbox along with Difference-in-Differences (DD) and Randomized Control Trials (RCT).

We seek to estimate the causal effect of a treatment. We posit that for each individual i , there exists a pair of potential outcomes: Y_{i1} for what would occur if i were exposed to the treatment and Y_{i0} if i were not exposed. The causal effect of the treatment is represented by the difference $T = Y_{i1} - Y_{i0}$. The fundamental problem of causal inference (FPCI) is that we cannot observe both Y_{i1} and Y_{i0} .

We have so far handled the FPCI using two techniques: randomization into treatment and control groups (RCT), and difference-in-difference estimation. Both methods attempt to find treated and control units that are in expectation comparable—that is, their potential outcomes if treated (or if untreated) are expected to be the same—and then contrasts outcomes among those treated relative to those not treated to estimate the average effect of treatment on the treated (ATT).

The Regression Discontinuity (RD) estimator takes a fresh approach to identifying a causal relationship when the treatment and control groups do *not* have potential outcomes that are identical in expectation. It instead looks for units that are *arbitrarily close* in terms of their potential outcomes and yet are treated differently (one assigned to treatment, the other assigned to control) due to some bright line rule that determines assignment. This situation occurs more commonly than one might expect. For example, the result of an election can be decided by a single vote, or the cutoff for which children are allowed to enter 1st grade in a given year may depend on whether they were born before or after midnight on September 1 six years earlier. Arbitrary cutoffs are inevitable for administrative purposes. A driver either is or is not speeding. A potential candidate for office either does or does not have the requisite number of signatures to get on the ballot. A library book is not overdue until the moment that it is.

While arbitrary cutoffs are necessary for administration, why are they useful for economists? Define a variable X that is used to determine the cutoff above/below which a person (or unit) i is or is not assigned to treatment. For example, X could be the percentage of voters for candidate A or X could be the exact hour/minute/second of birth. We will refer to X as the *running variable*, and we'd like that variable to be continuous.

Imagine there are two underlying relationships between potential outcomes and treatment, represented by $E[Y_{i1}|X_i]$ and $E[Y_{i0}|X_i]$. Thus at each value of X_i , the causal effect of treatment is $E[T|X_i = x] = E[Y_{i1}|X_i = x] - E[Y_{i0}|X_i = x]$. Let's say that individuals to the right of a cutoff c (e.g., $X_i \geq 0.5$) are exposed to treatment, while those to the left ($X_i < 0.5$) are denied treatment. We therefore observe $E[Y_{i1}|X_i]$ to the right of the cutoff and $E[Y_{i0}|X_i]$ to the left of the cutoff.

As we consider units i that are arbitrarily close (within ε) to the threshold, it may be reasonable to assume that:

$$\begin{aligned}\lim_{\varepsilon \downarrow 0} E[Y_{i1}|X_i = c + \varepsilon] &= \lim_{\varepsilon \uparrow 0} E[Y_{i1}|X_i = c + \varepsilon], \\ \lim_{\varepsilon \downarrow 0} E[Y_{i0}|X_i = c + \varepsilon] &= \lim_{\varepsilon \uparrow 0} E[Y_{i0}|X_i = c + \varepsilon].\end{aligned}$$

That is, for units that are *almost identical*, we may be willing to assume that had both been treated

(or not treated), their outcomes would have been arbitrarily similar. If this assumption is plausible, we can form a Regression Discontinuity estimate of the causal effect of treatment on outcome \mathbf{Y} using the contrast:

$$\hat{T} = \lim_{\varepsilon \downarrow 0} E[Y_i | X_i = c + \varepsilon] - \lim_{\varepsilon \uparrow 0} E[Y_i | X_i = c + \varepsilon],$$

which in the limit is equal to:

$$T = E[Y_{i1} - Y_{i0} | X_i = c].$$

The RD estimator estimates the causal effect of a treatment as the “jump” in an outcome variable, \mathbf{Y} , as near-identical units on one side of a discontinuity, c , are allocated to treatment while those on the other side are allocated to non-treatment. Note that while RD estimation does estimate the treatment effect given that $x_i = c$, if the treatment effect is not the same for everyone, it will not give you the average treatment effect on the treated. For example, imagine you are studying the effect of a scholarship on student grades. If you randomly assign scholarships, you would get the average treatment effect for the entire sample of students (i.e the average treatment effect on the treated). If, instead, scholarships are given to students with SAT scores above 2100, and you use an RD design, you will get the treatment effect on those students with SAT scores of 2100, but not the average treatment effect for all students who received the scholarship.

It’s important to note that RD only works under a few conditions:

1. *There is no manipulation of the running variable.* If people know about the rule *and* they are able to manipulate \mathbf{X} , then they may selectively change the recorded value of \mathbf{X} to ensure that a given individual does or does not receive the treatment. We can check this in part by comparing the observable characteristics of individuals just on either side of the discontinuity – if they appear to be similar, then we are less worried about manipulation.
2. *In the absence of the discontinuity, the outcome \mathbf{Y} is changing smoothly as a function of the running variable.* If the running variable is causing large, discontinuous changes in \mathbf{Y} , then that will confound any effect of the discontinuity. As an extreme example, consider a school that grades its students on a scale of A, B, C, D, F (without any underlying numerical scores) and gives all of its A students a scholarship for college. Imagine that we want to study the effects of this scholarship on the likelihood that a student goes to college. If we compare A students to B students, A students probably differ greatly from B students, so we cannot attribute the effects of being an A student on college-going to the scholarship alone. However, if we had the raw numerical scores and we could compare students with a 90.1 grade point average (barely an A) with those who have an 89.9 grade point average (barely a B), then the assumption of smooth changes is much more likely to be met.
3. *Nothing else changes at the discontinuity (besides the policy of interest).* If many things change at a single discontinuity, then we cannot use an RD design to study the effects of one policy in particular. For example, imagine that we are interested in the persistent effects on income of having been part of one medieval kingdom versus another. We know where the border between

the kingdoms is, but it now is also the same border as the border that separates France and Spain. We can look at the effects of the border on income, but we cannot attribute that difference purely to the historical effects of the kingdoms, since the modern border also likely matters.

The 2022 paper by Bleemer and Mehta that we discussed in class (“Will studying economics make you rich? A regression discontinuity analysis of the returns to college major”) provides a particularly crisp RD example. Here’s the authors’ summary of what their paper does:

The specific case we analyze is the economics department at the University of California, Santa Cruz (UCSC). UCSC Economics imposed a grade point average (GPA) restriction policy in 2008: students with a GPA below 2.8 in Economics 1 and 2 were generally prevented from declaring an economics major. Students who just met the GPA threshold were 36 percentage points more likely to declare the economics major than those who just failed to meet it. Most of these students would have otherwise earned degrees in other social sciences.

Below are two key figures from their paper that tell the story. Figure 1 documents the sharp discontinuity in the probability of majoring in economics at the GPA threshold. During the years 2008–2012, about 41% of graduating UCSC students who took Economics 1 and 2 and who were just below the GPA cutoff of 2.75, ultimately majored in economics. Among their peers who were just above the 2.75 GPA cutoff, about 77% majored in economics. Thus, the RD estimate of the causal effect of passing the GPA threshold on the probability of majoring in economics is 36.1%.

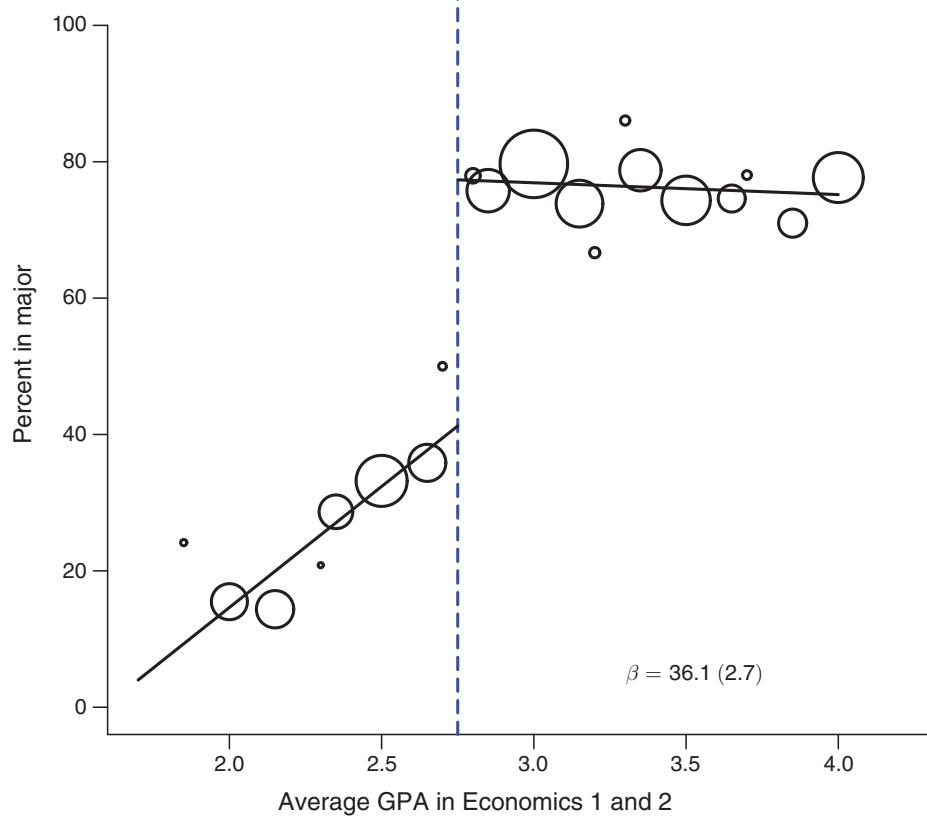


FIGURE 1. THE EFFECT OF THE UCSC ECONOMICS GPA THRESHOLD ON MAJORING IN ECONOMICS

Notes: Each circle represents the percent of economics majors (y-axis) among 2008–2012 UCSC students who earned a given *EGPA* in Economics 1 and 2 (x-axis). The size of each circle corresponds to the proportion of students who earned that *EGPA*. *EGPAs* below 1.8 are omitted, leaving 2,839 students in the sample. Fit lines and beta estimate (at the 2.8 GPA threshold) from linear RD specification; standard error (clustered by *EGPA*) in parentheses.

Figure 2 documents the corresponding discontinuity in average annual earnings among these same students during the years 2017–2018 (averaging about 8 years after graduation). The earnings jump at that GPA threshold is \$7,989 — from approximately \$47.5K to \$55.5K — an increment of about 18% relative to average annual earnings of students immediately above the cutoff.

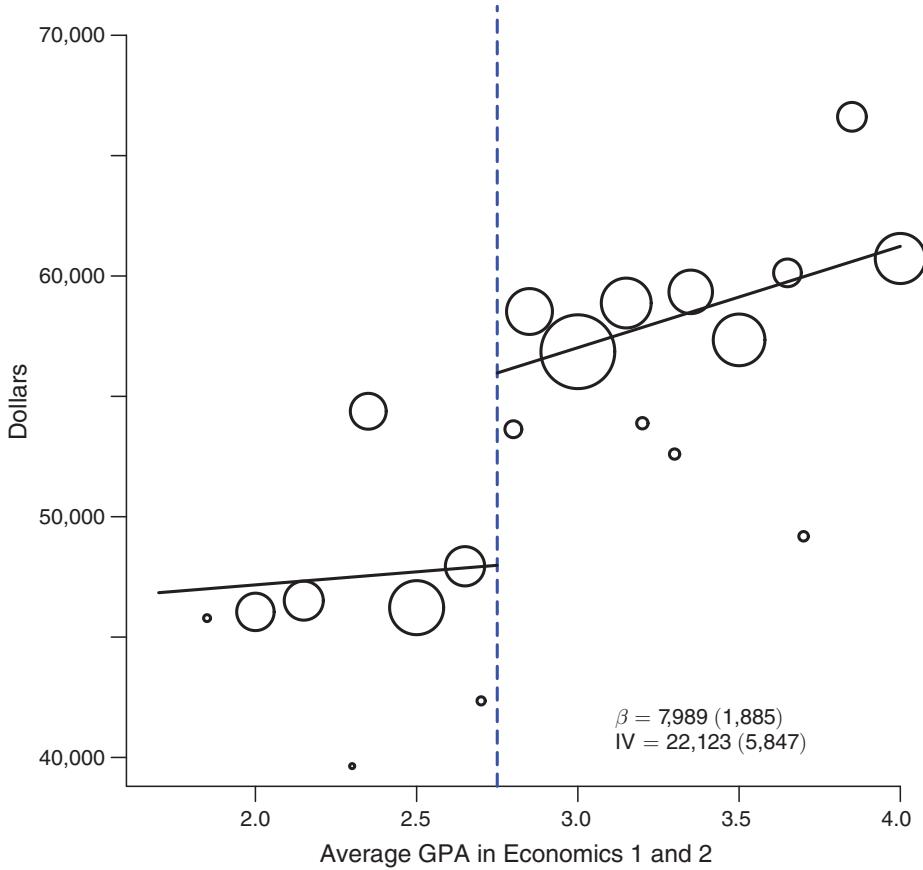


FIGURE 2. THE EFFECT OF THE UCSC ECONOMICS GPA THRESHOLD ON ANNUAL WAGES

Notes: Each circle represents the mean 2017–2018 wages (y-axis) among 2008–2012 UCSC students who earned a given *EGPA* in Economics 1 and 2 (x-axis). The size of each circle corresponds to the proportion of students who earned that *EGPA*. 2017–2018 wages are the mean EDD-covered California wages in those years, omitting zeroes. Wages are CPI adjusted to 2018 and winsorized at 2 percent above and below. *EGPAs* below 1.8 are omitted, leaving 2,446 students with observed wages. Fit lines and beta estimate (at the 2.8 GPA threshold) from linear RD specification and instrumental variable specification (with majoring in economics as the endogenous variable); standard errors (clustered by *EGPA*) in parentheses.

As we briefly mentioned in class, and as we'll discuss further soon, this \$7,989 jump is (under our maintained assumptions) a credible estimate of the causal effect of *being permitted to major* in economics on adult earnings, but it is *not* an estimate of the causal effect of *majoring* in economics on earnings. The reason is that at the GPA threshold, the probability of majoring in economics does not rise from 0% to 100% but, instead, from 41% to 77%. This implies that the earnings gap at the cutoff *understates* the effect of economics degree vs. no economics degree. Why? As we cross the threshold, we are contrasting two mixtures of economics and non-economics majors: one group with 41% economics majors, the other with 77% economics majors. If we want to estimate the pure effect of *majoring in economics* — rather than being allowed to major in economics — on earnings, we need to take one more step: rescaling the earnings estimate to account for these differing proportions of majors at the threshold. You may already see how to do this rescaling. We will discuss it further soon.

6 The Expenditure Function

So far, we've analyzed problems where income was held constant and prices changed. This gave us the Indirect Utility Function. Now, we want to analyze problems where utility is held constant and expenditures change. This gives us the Expenditure Function.

These two problems are closely related—in fact, they are ‘*duals*.’ Most economic problems have a *dual problem*, which means an inverse problem. For example, the dual of choosing output in order to maximize profits is minimizing costs at a given output level; cost minimization is the dual of profit maximization. Similarly, the dual of maximizing utility subject to a budget constraint is the problem of minimizing expenditures subject to a utility constraint. Minimizing costs subject to a minimum utility constraint is the dual of maximizing utility subject to a (maximum) budget constraint.

6.1 Setup of expenditure function

Consumer's primal problem: maximize utility subject to a budget constraint. Consumer's dual problem: minimizing expenditure subject to a utility constraint (i.e. a level of utility the consumer must achieve). The dual problem yields the “expenditure function,” the minimum expenditure required to attain a given utility level.

1. Start with:

$$\begin{aligned} & \max U(x, y) \\ & \text{s.t. } p_x x + p_y y \leq I \end{aligned}$$

2. Solve for $x^*, y^* \Rightarrow u^* = U(x^*, y^*)$ given p_x, p_y, I .

$$V = V(p_x, p_y, I)$$

V is the indirect utility function, and its value is equal to u^*

3. Now solve the following problem:

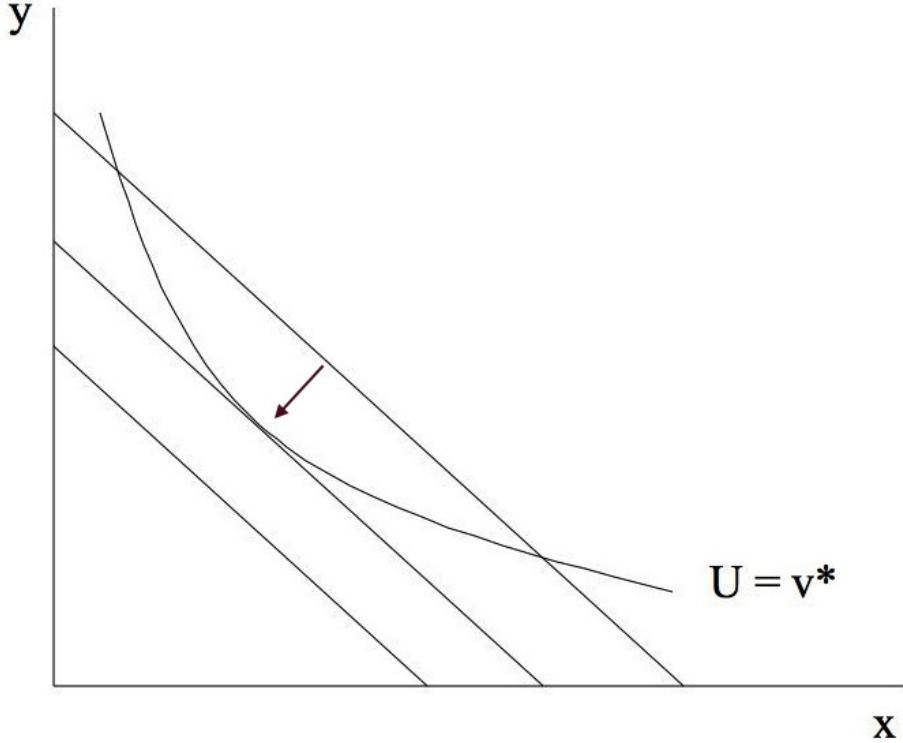
$$\begin{aligned} & \min p_x x + p_y y \\ & \text{s.t. } U(x, y) \geq u^* \end{aligned}$$

which gives $E = p_x x^* + p_y y^*$ for $U(x^*, y^*) = u^*$.

$$E = E(p_x, p_y, V^*)$$

E is the expenditure function, and its value is equal to $p_x x^* + p_y y^*$

6.2 Graphical representation of the dual problem



- The dual problem consists of choosing the lowest budget set tangent to a given indifference curve. Example:

$$\begin{aligned} \min E &= p_x x + p_y y \\ \text{s.t. } x^{0.5} y^{0.5} &\geq U_p \end{aligned}$$

where U_p comes from the primal problem.

$$\mathcal{L} = p_x x + p_y y + \lambda (U_p - x^{0.5} y^{0.5})$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= p_x - \lambda \cdot 0.5 x^{-0.5} y^{0.5} = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= p_y - \lambda \cdot 0.5 x^{0.5} y^{-0.5} = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= U_p - x^{0.5} y^{0.5} = 0 \end{aligned}$$

- The first two of these equations simplify to:

$$x = \frac{p_y y}{p_x}$$

We substitute into the constraint $U_p = x^{0.5}y^{0.5}$ to get

$$\begin{aligned} U_p &= \left(\frac{p_y y}{p_x}\right)^{0.5} y^{0.5} \\ x^* &= \left(\frac{p_y}{p_x}\right)^{0.5} U_p, \quad y^* = \left(\frac{p_x}{p_y}\right)^{0.5} U_p \\ E^* &= p_x \left(\frac{p_y}{p_x}\right)^{0.5} U_p + p_y \left(\frac{p_x}{p_y}\right)^{0.5} U_p \\ &= 2p_x^{0.5} p_y^{0.5} U_p \end{aligned}$$

6.3 Relationship between the Expenditure function and the Indirect Utility function

How do the solutions to the Dual and Primal problems compare?

- Examining the relationship between the expenditure and indirect utility functions:

$$\begin{aligned} V(p_x, p_y, I_0) &= U_0 \\ E(p_x, p_y, U_0) &= I_0 \\ V(p_x, p_y, E(p_x, p_y, U_0)) &= U_0 \\ E(p_x, p_y, V(p_x, p_y, I_0)) &= I_0 \end{aligned}$$

- The Expenditure function and Indirect Utility function are *inverses* one of the other.
- Let's verify this in the example we saw above. Recall that the primal problem gave us factor demands x_p^* , y_p^* as a function of prices and income (not utility).
- The dual problem gave us expenditures (budget requirement) as a function of utility and prices.

$$x_p^* = \frac{I}{2p_x}, \quad y_p^* = \frac{I}{2p_y}, \quad U^* = \left(\frac{I}{2p_x}\right)^{0.5} \left(\frac{I}{2p_y}\right)^{0.5}$$

Now plug these into expenditure function:

$$E^* = 2U_p p_x^{0.5} p_y^{0.5} = 2 \left(\frac{I}{2p_x}\right)^{0.5} \left(\frac{I}{2p_y}\right)^{0.5} p_x^{0.5} p_y^{0.5} = I$$

Finally notice that the multipliers are such that the multiplier in the dual problem is the inverse of the multiplier in the primal problem.

$$\begin{aligned} \lambda_P &= \frac{U_x}{p_x} = \frac{U_y}{p_y} \\ \lambda_D &= \frac{p_x}{U_x} = \frac{p_y}{U_y} \end{aligned}$$

6.4 Expenditure function: What is it good for?

The expenditure function is an essential tool for making consumer theory operational for public policy analysis. Using the expenditure function, we can ‘monetize’ otherwise incommensurate trade-offs to evaluate costs and benefits. The need for this type of calculation arises frequently in policy analysis and is the basis for most cost-benefit analyses.

As we have stressed earlier the semester, we don’t know what ‘utils’ are. This presents a problem if we want to determine *how much* harm or benefit a certain policy imposes on an individual. The expenditure function gives us a convenient way to potentially circumvent this problem. Using the expenditure function, we can figure out *how much money* a consumer would have to be compensated (which could be a positive or negative number) to leave her equally well off after a policy is implemented as she was initially. So, the expenditure function permits us to calculate a ‘money metric.’

Let’s say we were considering a policy that raised prices for some consumers, perhaps by raising the cost of gasoline. Policymakers might be legitimately concerned that this policy change would adversely affect low income consumers. To offset this effect, they might provide cash compensation to offset their loss. How large should this transfer be?

A typical policy response would be to set compensation equal to the full amount of the price increase multiplied by the consumer’s initial expenditure on gasoline. Let C equal the compensation amount, with

$$C = \Delta P_g \times Q_{g,0}.$$

Here, $\Delta P_g = P_{g,1} - P_{g,0}$ is the policy-induced price change and $Q_{g,0}$ is the quantity that the consumer was purchasing initially (i.e., at time $t = 0$).

Is C the right amount of compensation—this is, neither too much or too little? If you knew the consumer’s utility function (a tall order, of course), you could calculate the exact answer as

$$C^* = E(P_{g1}, P_a, V(P_{g0}, P_a, I_0)) - E(P_{g0}, P_a, V(P_{g0}, P_a, I_0)),$$

where I_0 is the consumer’s initial budget, P_a are the prices of all other goods (assumed constant over time), and $V(\cdot)$ is the indirect utility function. You would then directly compare $C^* \leq C$ to see if C is above or below the exact compensation required. Absent knowledge of each consumer’s utility function, can we say anything more?

The answer is yes. A bit of thought should convince you that it must be the case that

$$C \geq C^*.$$

That is, the simple compensation scheme $C = \Delta P \times Q_0$ *always* weakly overestimates the actual compensation required. Why? As a starting point, note that C *must* be an upper bound on C^* . Clearly, if we compensate the consumer the full amount of money required to buy her initial bundle, she must be at least as well off as before; she can have the original bundle *or* she can choose many others that were not initially affordable. (These other bundles would have less of the good

whose prices has increased but more of other goods. You should demonstrate to yourself that some previously infeasible bundles are now feasible with prices P_{g1}, P_a and income $I_0 + C$.)

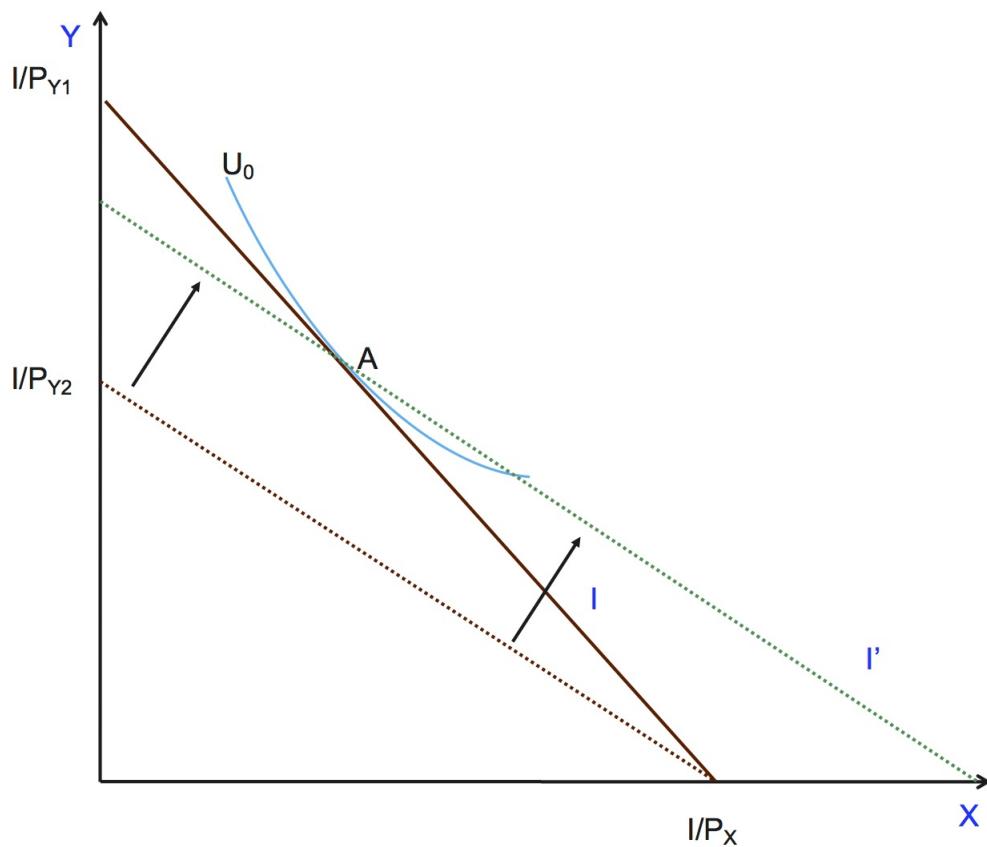
But we can say more? Again, yes. If the consumer has standard indifference curves that are bowed towards the origin (diminishing MRS), a change in the price of one good (e.g., gasoline) will cause the consumer to partially substitute towards other goods. This substitution partly blunts the effect of the price increase as the consumer re-optimizes her bundle given the new prices. If we raise the price of one good but hold the consumer's utility at its initial level, her optimally chosen bundle will rotate along the original indifference curve to a new location where the new price ratio is tangent to the initial indifference curve. This new bundle will cost more at the new prices than the original bundle at the old prices (unless there exists a perfect substitute for the original good at the original price⁵). But this new bundle (which holds utility constant) will cost strictly less than $I_0 + C$. The difference between the cost of the new bundle at the new prices and the cost of the old bundle at the old prices (both lying on the initial indifference curve) is equal to C^* .

Compensation and Over-Compensation: Graphical Illustration

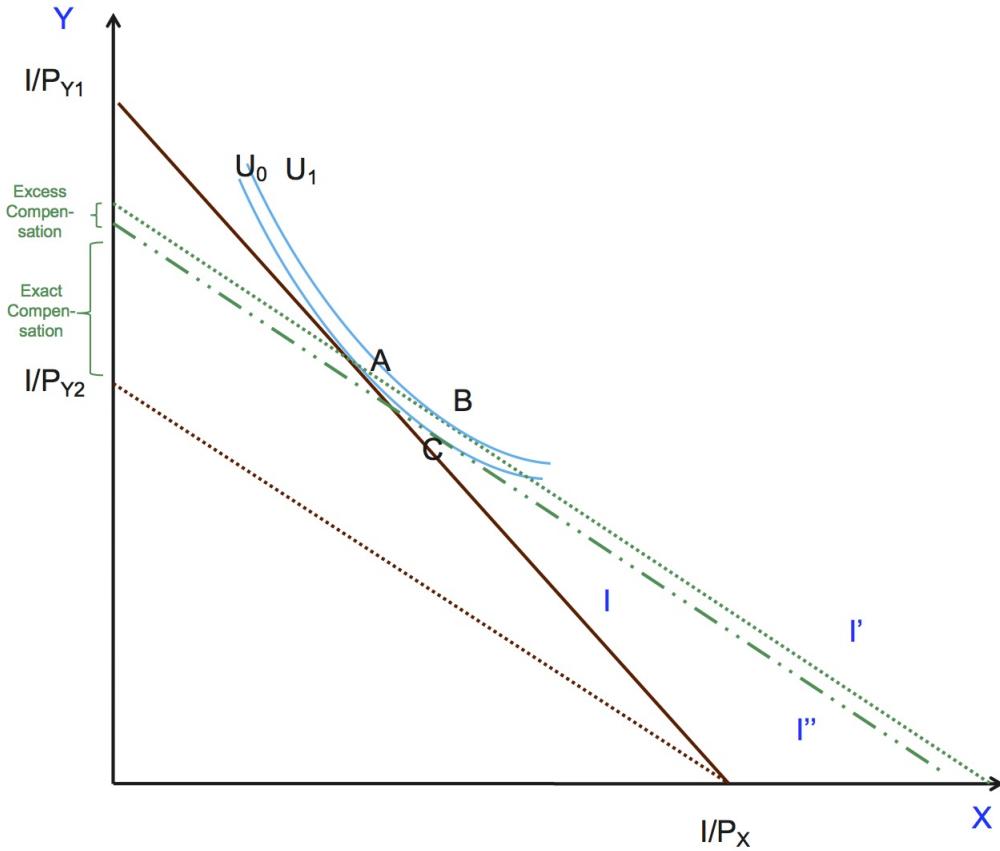
You can see the operation of this logic in the two following two figures. In the first figure, the consumer has budget set I and faces prices (P_X, P_{Y1}) and chooses point A on the budget set, generating utility U_o . Notice that the points of intersection of the budget set with the axes correspond to I/P_X and I/P_{Y1} ; these are the points at which the entire budget is spent on one good or the other.

Imagine the price of Y rises from P_{Y1} to P_{Y2} , rotating the budget set counterclockwise from the x -intercept so that the new point of intersection with the y -axis is given by I/P_{Y2} . Clearly, the consumer can no longer afford bundle A . How much compensation (in budget terms) do we need to give the consumer to make her as well off as she was initially at point A ?

⁵Consider an extreme case where gasoline and kerosene are *perfect* substitutes and have identical initial prices. In that case, a rise in the price of gas would have no effect on consumer welfare since consumers would simply switch to kerosene. If consumers received compensation C along with the price change, they would be strictly better off than before the price change.



An intuitive answer is that we would need to increase her budget from I to I' , thus permitting the consumer to afford her old bundle at the new prices. But if you stare at the diagram for a moment, you can see that this is *excess* compensation. On the budget set I' , the agent can certainly afford A , but that budget will not be tangent to the original indifference curve; indeed it cuts right through it!



Facing budget set I' , the consumer would select a point like B that is on a higher indifference curve. How do we know it will be on a higher indifference curve? So long as the consumer has preferences that are at least weakly convex, she will adjust her bundle as relative prices change, with the degree of adjustment depending on the curvature of her indifference curves. If her indifference curves were linear (implying that X and Y are perfect substitutes), the agent would simply consume exclusively Y when the price of X rose and the price change would have no impact on her utility. With conventional strongly convex indifference curves as pictured above, the consumer would substitute towards Y as P_X rises, but she would not substitute all the way.

So, how much would we need to compensate the agent to make her as well off as she was at A ? We need to find the minimum expenditure that allows her to reach utility level U_0 at the new prices. This budget corresponds to I'' . Facing that budget, the consumer would pick point C , which lies on the same indifference curve as A . The amount of compensation required is depicted in terms of the good Y . The additional income required is the gap between the y-intercept of I and I'' multiplied by P_{Y2} . (This converts quantities of Y to dollars. You could of course do the same calculation in terms of X .) If we had compensated the consumer the difference between I and I' , she would be overcompensated.

This figure illustrates a key principle: the expenditure function optimizes quantities in response to prices. If consumers have convex preferences, the expenditure adjustment required to fully offset

a large price change is less than the initial quantity consumed times the change in price. (For an infinitesimal price change, however, the envelope theorem applies; the consumer would not re-optimize quantity for an infinitesimal price change.)

Aside: Is it valid to measure utility in dollar equivalent terms given that utility functions are ordinal?

- Since utility functions are only defined up to a monotone transformation, doesn't this mean that welfare loss/gain calculations based on the expenditure function computed *in dollars* (a cardinal measure) are not uniquely defined for a given utility function? Actually, it does not.
- Consider the following thought experiment. Utility functions $U_1(\cdot)$ and $U_2(\cdot)$ are identical for consumer theory; $U_2(\cdot)$ is a monotone transformation of $U_1(\cdot)$. Hence, these two utility functions have identical preference rankings and choose the same bundles of goods for given income and prices. If we gave $U_1(\cdot)$ and $U_2(\cdot)$ each \$100 in cash, they would consume identical bundles to one another. Likewise, if we gave them \$100 in monthly food stamps (meaning that they must spend *at least* \$100 per month on food—though they could spend more), they would consume identically.
- Imagine that $U_1(\cdot)$ and $U_2(\cdot)$ are distorted by food stamps so that they are forced to consume more food using \$100 in stamps than they would if given \$100 cash. How much additional cash (in addition to food stamps) would it take to make $U_1(\cdot)$ and $U_2(\cdot)$ indifferent between \$100 in cash versus \$100 in stamps plus additional cash?
- We don't know the numerical answer without an explicit functional form. But we do know that the answer must be the same for $U_1(\cdot)$ and $U_2(\cdot)$. Why? Both $U_1(\cdot)$ and $U_2(\cdot)$ would choose to buy the same bundles using the extra cash to get back on the original indifference curve associated with receiving \$100 in cash—and of course those bundles would cost the same since all consumers face the same prices.
- Hence, the DWL associated with food stamps (in dollars, not utils) is identical for both utility functions, despite the fact that the functions are not identical. If you want to demonstrate this to yourself, work an example with $U_1(X, Y) = X^{1/2}Y^{1/2}$ and $U_2(X, Y) = 1/2 \ln X + 1/2 \ln Y$.]