# Lecture Note 8, Part 1: Applying Consumer Theory to Competitive Markets – The United States Sugar Program

David Autor, MIT and NBER

14.03/14.003 Applied Microeconomics and Public Policy, Fall 2016

# 1 Applying Consumer Theory to Competitive Markets

When exchange takes place voluntarily, canonical consumer theory suggests that all participants are made better off. Otherwise, participants would not have voluntarily engaged in the exchange. (Later in the semester, we'll discuss why this may not always be true.) It's useful to have a dollar metric of the gains from transacting, similar to how we developed a dollar metric of how price changes hurt consumers in previous lectures. This measure captures consumer surplus in dollar terms. The notion of consumer surplus is critically important: although we can readily measure the direct costs of a given project or policy (i.e., building a bridge, imposing a tariff), costs are not necessarily equal to benefits. Demand curves allow us to measure these benefits. When costs exceed benefits, we call that deadweight loss. When benefits exceed costs, we call that surplus: consumer surplus, producer surplus, or more generally 'social surplus,' which does not distinguish between consumers versus producers.

As we discussed in the first lecture, we can think of the market demand curve as the set of consumers arrayed in inverse order from the person with the highest Willingness to Pay (WTP) to the person with the lowest WTP for a commodity. Similarly, we think of the market supply curve as the set of producers arrayed in order from firm willing to produce at lowest price to firm demanding the highest price to produce a good. (This is sometimes called Willingness to Accept or WTA.)

Consumer surplus is the difference between the maximum value that a consumer is willing to pay for a commodity and the market price of that commodity. In market equilibrium, the marginal consumer of a commodity gets no surplus from the purchase (he or she is indifferent) whereas the inframarginal consumers receive positive surplus. Analogous to consumer surplus, the relevant measure of surplus for firms (producers) is producer surplus. This is the area *above the supply curve*—that is, what the producer receives for goods in excess of the cost of production.

In an ideal world, the market matches consumers and producers. If in equilibrium a producer is willing to produce at a price less than or equal what a consumer is willing to pay, we expect that transaction to occur. Most consumers will be buying at a price *below* their maximal willingness to pay, and most producers will be selling at a price *above* their lowest willingness to produce. This will happen up to the point at which the marginal producer and consumer are indifferent between selling/buying and going home.

When this mechanism works correctly, it maximizes the sum of producer and consumer surplus: all gains from trade are realized; all transactions that benefit both parties occur; no transactions occur that do not benefit both parties. (These statements are equivalent.)

It is noteworthy that this metric does not place any greater weight on consumer or pro-

ducer surplus. If supply is perfectly elastic, all of the surplus is captured by consumers. If demand is perfectly elastic, all of the surplus is captured by producers. It also places the same weight on different individual consumers, and on different individual producers. Why do we want to maximize surplus without any regard for the identities of the beneficiaries? Don't we care about equity *and* efficiency? As we'll discuss shortly when we reach the topic of General Equilibrium, in a competitive market, the goals of maximizing the pie and redistributing the pie are *not* in tension. There is no intrinsic tradeoff between maximizing the pie and determining the size and allocation of the slices. This lecture focuses on the first goal – maximizing the pie – and we will address the second question soon.

To focus on the first question, you can pretend that the agents in the following examples are broadly similar to one another.

In the analysis that follows, please keep the following three points in mind:

1. A *transfer* between two agents is not necessarily a net social gain or loss

    - If I pay a worker \$5 to perform a task, that is *not* a \$5 social gain even though I have 'generated a job.' Why? This transaction is, to a first approximation, a transfer of \$5 from me to someone else – although a new task was also completed thanks to the transaction.
    - *If* there is a gain, it is because the value that the worker produces is worth more than \$5 to me (i.e., consumer surplus) or the worker's alternative use of time is worth less than \$5 to her (i.e., producer surplus). The hourly wage that a worker earns at her job is not a measure of her surplus because (a) her time has an opportunity cost—there are other valuable things she would like to be doing if she were not working; (b) work may require her to exert effort that has direct disutility (physical exertion, intellectual boredom, mental exhaustion, etc.). Thus, the social surplus could be considerably smaller than the amount of money changing hands.
    - Aside: Note that Gross Domestic Product (GDP) basically measured the amount of money that changes hands in transactions nationwide. So if you and I mow each others' lawns rather than mowing our own lawns, we've added to GDP. When you read the news, remember that GDP roughly captures the business of an economy but is *not* the number corresponding to social surplus in an economy.

2. A cost is not a benefit

    - This is a actually restatement of the above. If I raise the wages of all federal employees in the United States government by 10%, what is the social benefit of

this? Barring major changes to the behavior of federal employees or who applies for these jobs, the change in social surplus is approximately zero. I am simply transferring money from one group of citizens (taxpayers) to another (federal employees). A reduction in surplus, if present, would be due to the transfers that don't take place thanks to higher taxes (deadweight loss of taxation).

3. All costs are opportunity costs.

   - In economic reasoning, there is no such thing as *intrinsic value.* The only cost of using a given resource is the value of its alternative to which it could have been put. This is its *opportunity cost.*
   - Why is gasoline so comparatively cheap (less per gallon than mouthwash or shampoo) when it is an indispensable commodity for most U.S. households. (Which would you rather do without, mouthwash or gasoline?) The simple answer is that gasoline is relatively abundant, so the last gallon isn't worth much—its opportunity cost is low. This opportunity cost is low because the high value uses of gasoline (fueling fire trucks, ambulances, and NASCAR vehicles) will surely be satisfied even at high prices. It's only when gas gets cheap that people start taking day-long jet ski excursions, cranking up their gas-powered beer coolers, and entertaining the kids with the Acme Home Flame Thrower. When gas is scarce, prices rise substantially because the first gallon is very valuable indeed—its opportunity cost might literally mean life or death to someone whose house is on fire or who needed to get to a hospital. Put differently, the price of a gallon of gas is determined by the opportunity cost of the marginal gallon, not by the value that people place on the first gallon (or the inframarginal gallons between first and last). This observation—that prices reflect scarcity rather than intrinsic value—is often referred to the diamond-water paradox. I prefer to call it the gasoline-shampoo paradox.

# 2 The U.S. Sugar Program

As detailed in the readings, the U.S. Sugar Program is a system of import quotas administered by the U.S. Department of Agriculture that is designed to reduce the supply of imported sugar to the U.S. market with the goal of achieving a target domestic sugar price. Our objective in this exercise is to estimate the economic costs and benefits of this policy—that is, what are the gains to consumers and producers and the efficiency costs (deadweight losses), if any.

## 2.1 Analytics

Will use basic competitive theory to analyze the economic consequences of the U.S. sugar program. Working from a competitive baseline scenario of what the price and quantity of sugar sold in the U.S. market would be *absent* the Sugar Program, we can decompose the deviations from this baseline into three components:

1. *Transfers between consumers and producers.* Externally imposed changes in quantities or prices will generally induce some transfer of surplus from consumers to producers or vice versa. These transfers are *not* efficiency losses (though one may still view them as unwarranted windfalls for specific groups).

2. *Deadweight losses from foregone consumption.* A price or quantity quota will generally reduce equilibrium consumption below its competitive level. This implies that there are foregone units of the good that producers would have been willing to produce at a price that consumers would have been willing to pay, if not for the quota. These thwarted trades are therefore a deadweight loss, reflecting losses of consumer *and* producer surplus.

3. *Deadweight losses from inefficient resource allocation.* Price or quantity restrictions may also cause production distortions: low cost producers may be prevented or disincentivized from supplying goods while high cost producers take their place. In these cases, there is the usual loss in consumer surplus and/or producer surplus *and* an additional deadweight loss incurred: real resources are consumed by high-cost producers to produce goods that low-cost producers could have made absent the distortion.

To implement this analysis, we need to estimate the consumer demand curve (to assess consumer surplus) and the producer supply curve (to assess producer surplus). With these in hand, we can consider the consequences of the quota system relative to a counterfactual case in which the market is competitive.

So far we've studied consumer surplus at the level of *individual consumers* using individual demand curves stemming from utility maximization. To operationalize the notion of surplus at the *market* level, we need measures of producer costs and consumers' willingness to pay—that is, we need information on the market supply and demand curves. In point of fact, we often have pretty good data on producers' costs, which should reflect their willingness to produce at various prices (that is, their Willingness to Accept or WTA). It's much harder to measure consumer's Willingness to Pay (WTP), however, since this is a function of consumer preferences (i.e., utility) rather observable production costs in the case of producers. When it's tough to estimate a whole demand curve, we often estimate the elasticity of consumer

demand for a commodity and extrapolate what a reasonable demand curve would look like. Specifically, if we know both a price-quantity pair on the demand curve and the elasticity of demand (that is, the slope of the demand curve), we can estimate the area under a standard demand curve to calculate measures of surplus (WTP minus price).

For our analysis of the U.S. Sugar Program, we'll use the uncompensated (Marshallian) demand curve for sugar. Why not the compensated demand curve? Two explanations First, it's not really feasible to estimate a compensated demand curve at the market level—what would it mean to hold "market utility" constant? More substantively, recall from the Slutsky equation that $\partial d_x/\partial p_x = \partial h_x/\partial p_x - (\partial d_x/\partial I)\, x$. This equation says that the compensated and uncompensated demand curves don't differ by much if $(\partial d_x/\partial I) \cdot x$ is small. This will be true if the income effect is small *or* if the commodity in question (here sugar) is a small share of consumer budgets—so, there is very little change in consumers' real income when sugar prices change. The latter condition is quite likely to be true. According to the United States Department of Agriculture's Economic Research Service, U.S. sweetener deliveries (including both sugar and High Fructose Corn Syrup or HFCS) in 2010 were 131.9 pounds per capita.[1] At the historically high U.S. market price of sugar – \$0.50 per pound in 2010 – sugar expenditures would amount to only \$65 per capita (though frankly, that's still an amazing amount of sugar; on average, Americans are basically eating their weight in sugar annually). Thus, even substantial fluctuations in the price of sugar (a doubling or halving of the price, for example) would have only minimal income effects on consumers' household budgets. This suggests that the difference between compensated and uncompensated demand will be small in this setting. This distinction does sometimes matter quite a bit, but not in our example.

## Details to follow in Part 2 of this lecture note...

[1] https://www.ers.usda.gov/publications/pub-details/?pubid=39266

14.03 / 14.003 Microeconomic Theory and Public Policy
Fall 2016