Lecture Note 9 — Applying the GE Framework to Consumer Markets: Fishing in the state of Kerala, India

David Autor, MIT Economics and NBER

14.03/14.003, Microeconomic Theory and Public Policy, Fall 2025

Revised October 20, 2025

Applying the GE Framework to Consumer Markets: Fishing in Kerala (“The Digital Provide,” Robert Jensen, 2007)

We will discuss the empirical details of the “Digital Provide” paper in class. There are three substantive points related to the paper that I explore in these lecture notes:

1. How do we illustrate the gains from trade in Kerala using the Edgeworth box?

2. How can we be sure that the integration of markets in Kerala is welfare enhancing?

3. What is “the law of one price” and why is it relevant here?

1 Gains from trade in Kerala in the Edgeworth box

Consider the Edgeworth box below with two consumers, A and B, who have two sources of nourishment, rice and fish. Each day, A and B harvest a fixed quantity of rice, depicted in the figure by the dashed line, and they also go fishing. On some days, the fish schools are primarily in A’s fishing area and on some days, the fish schools are primarily in B’s fishing area. Let’s assume for simplicity that their total catch (the sum of their two catches) is identical on each day. That is, their catches are inversely correlated so that when A has a big catch, B has a small catch and vice versa.

In this diagram, the curve connecting the vertices is the Contract Curve, i.e., the set of Pareto efficient allocations. In autarky (that is, no trade), A and B each consume their own endowments each day. This is clearly not Pareto efficient. The lens shaped regions extending from the two different endowments (big catches for A and B, respectively) show the unexploited gains from trade.

If A and B could trade fish for rice on any day, this would improve the welfare for both parties under both sets of circumstances (i.e., A or B has a big day). Even though A is relatively wealthy in one state and B is relatively wealthy in the other, both benefit from the opportunity to trade on either occasion. The reason is that neither party prefers to consume either a *lot* of fish or a *little* fish with the average daily allotment of rice. On a day when A has a big catch, he’d prefer to trade some of that catch for additional rice; on a day when A has a small catch, he’d prefer to trade some of his rice for additional fish—and vice versa for B. Thus, there are potential gains from trade no matter which endowment prevails on a given day.

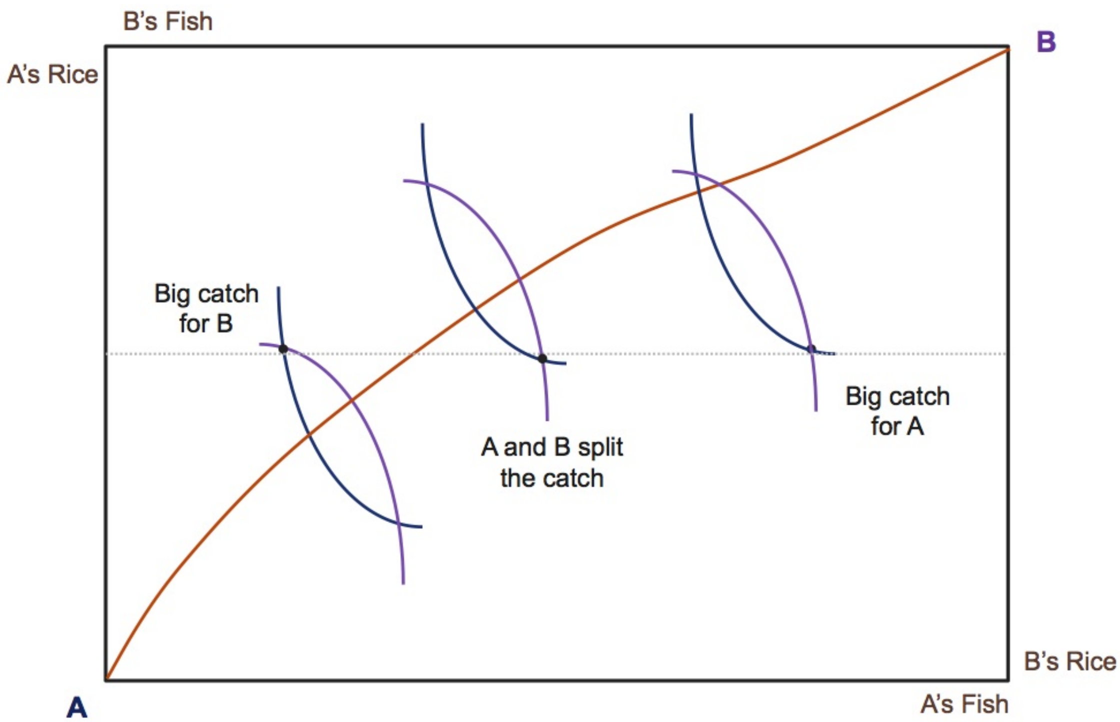

2 Further gains from market integration when endowments vary: consumption smoothing

When we first discussed the Edgeworth box, we gave each person a fixed endowment. In this setting, however, people's endowments vary from one day to the next. If fish were storable, this distinction theoretically wouldn't matter: A could simply pool (pun intended) her catch across good and bad days. This would give her the option to consume the same amount of fish each day. Of course, B could do similarly. Yet fish are not storable in a setting with no giant freezers (don't test this), so consumers' endowments do vary over time in this setting. With an integrated market, A and B could agree to share their fish so that on any given day, they pool and split their catch.

The possibility of pooling the catch is illustrated in the diagram above as the midpoint between the two possible endowments (corresponding to good days for A and and good days for B). Notice that even at this mid-point, A and B would want to trade further to reach the contract curve. This reflects differences in their preferences. Starting from identical endowments (the point in the middle), A would happily give up fish for a little more rice and B would happily give up rice for a little more fish.

Would A and B prefer to "pool" their fish catch each day and then trade, or would they prefer to trade from their randomly drawn, more volatile endowments? Or is the answer indeterminate? In a world of risk and uncertainty, in which A or B could have lots of "bad" days in a row, this question is actually difficult to address using the tools currently at our disposal. We'll come back to the question later in the semester when we study risk and insurance. We can, however, address a slightly simpler question right now: would A and B prefer to alternate "good" and "bad" days or keep the same consumption every day?

We'll use a simple parametric example (i.e. an example with specific functional forms) to answer this question, but the lesson is in fact a general one. You can put aside the Edgeworth box for a few minutes, and pretend that A and B have a constant amount of wealth each day but don't have any alternative suppliers of fish. Assume that both A and B have standard preferences that satisfy our five axioms, and that their Marshallian demand curves for fish are unambiguously downward sloping. Assume that fish is a small part of their consumption bundles, so that income effects can be ignored (Hicksian and Marshallian demands don't differ by much).

Holding wealth and the price of all other goods constant, assume that both A and B have personal demand curves of the form:[1]

$$Q = 60 - P.$$

For example, if the price of fish is 20, A and B would each buy $60 - 20 = 40$ fish.

On a good day, A catches 40 fish, and on a bad day, she catches 20. $B's$ catches are just the opposite: when A has a good day, B catches only 20 fish, and when A has a bad day, B catches 40 fish.

Would A and B prefer to consume 40 and 20 fish each on alternating dates, or would they prefer to consume 30 every day? To answer this question, we need to calculate the utility A gets from these bundles. Although we cannot answer this question in terms of "utils," we can answer it in terms of surplus, which is simply the area under the demand curve. (While surplus isn't exactly the same as welfare, we know that situations with more surplus for a given consumer also produce more welfare for that consumer.)

Inverting the demand curve above gives us the price that A and B would be willing to pay

[1] Our substantive conclusions do not depend on their demand curves being identical. They just need to be downward sloping.

for fish as a function of quantity:

$$P = 60 - Q.$$

They are willing to pay $59 for the first fish, $58 for the second, etc. So, how much are they willing to pay for a bundle of Q fish? The answer is the area under the demand curve:

$$\begin{aligned} WTP(Q) &= \int_0^Q P(Q)\, dQ \\ &= \int_0^Q (60 - Q)\, dQ \\ &= 60Q - \frac{Q^2}{2} + c\Big|_0^Q \end{aligned}$$

Hence, on high and low days, their total WTP for the bundle of fish they consume is:

$$\begin{aligned} WTP(40) &= 60Q - \frac{Q^2}{2} + c\Big|_0^{40} = 2400 - 800 + c - c = 1,600 \\ WTP(20) &= 1200 - 200 = 1,000. \end{aligned}$$

Hence, total WTP across the two days is $2,600$ and average WTP is $1,300.$

Now, if A and B split their catch each day, they'd each get 30 fish. How much would they be willing to pay to consume this quantity daily?

$$WTP(30) = 1800 - 450 = 1,350.$$

In this case, total WTP across two days is $2,700$ and average WTP is $1,350.$ So A and B would prefer to split their catch every day rather than eat more fish on high days and less fish on low days (that is, $2 \times WTP(30) > WTP(40) + WTP(20)$).

We can also illustrate this result graphically. Consider the diagram below

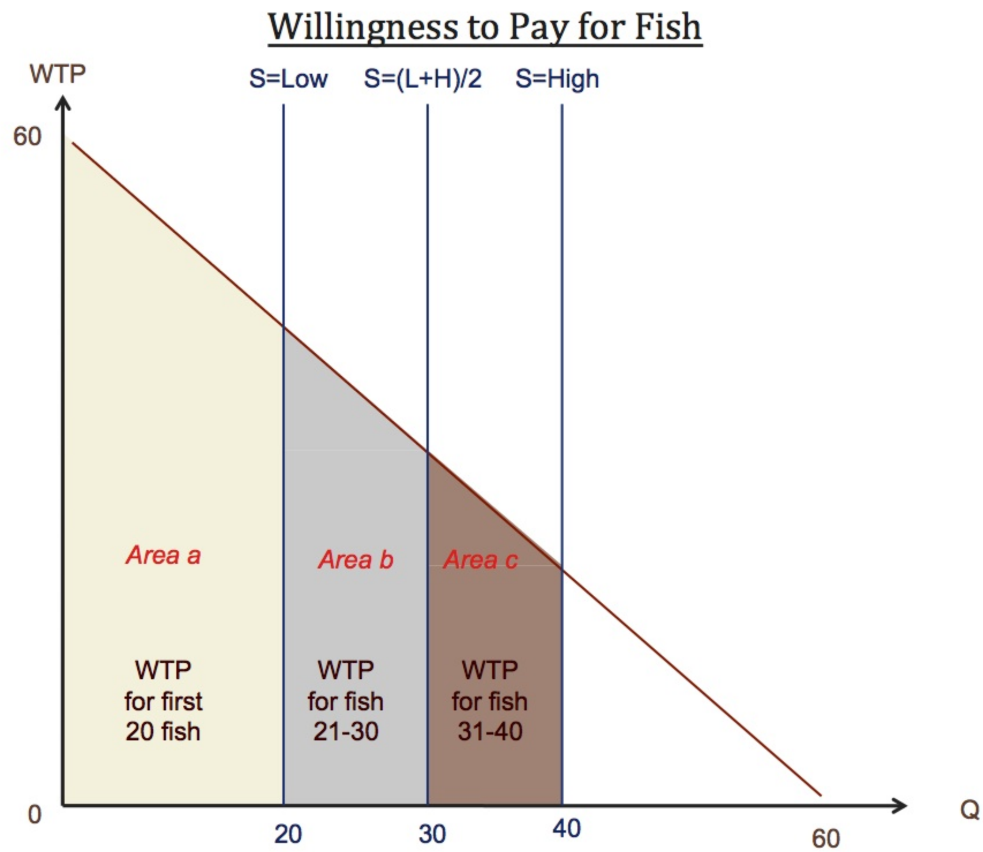

The three regions of this figure correspond to A or $B's$ willingness to pay for the first 20, next 10, and subsequent 10 fish. Call these areas a, b and c. To consume 40 fish one day and 20 the next, A or $B's$ total willingness to pay is:

$$WTP\,(40) + WTP\,(20) = a + (a + b + c) = 2a + b + c$$

To instead consume 30 fish on both days, A or $B's$ total WTP would be:

$$2 \times WTP\,(30) = 2a + 2b$$

Since area b is greater than area c, the consumer prefers to consume a consistent quantity rather than experiencing feast and famine. This idea of evening out consumption swings is often referred to as "consumption smoothing".

While we used a specific example here, this general principal follows directly from consumers' diminishing marginal rate of substitution. If the marginal utility of fish is declining in consumption, the consumer obtains less utility from consuming two extreme bundles (one high in fish, one low in fish) than from consuming the average of these two bundles twice. Notice that this result would not necessarily hold absent diminishing MRS for fish. In that case, the demand curve for fish would be perfectly elastic (meaning a straight horizontal line at a fixed height), and the consumer would be indifferent between the average bundle and the two extreme bundles.

Returning to the gains from trade, this example suggests that A and B can improve welfare not merely by trading from their autarkic position each day but also from trading *over time* to pool resources (you can envision this as a new Edgeworth box for each discrete time period). Although A and B cannot store fish, they could agree to a binding contract to share their catches, which would have the effect of reducing the swings in consumption over time for each party.

3 Applying "the law of one price"

Now let's return to our standard Edgeworth box, and discuss *the law of one price*. The law of one price states that in competitive equilibrium, prices for a given commodity should be equalized across markets. This is a necessary condition for Pareto efficiency: if prices differ across consumers for the same commodity, they will not be able to trade until they equate their MRS's across goods. Their psychic rates of tradeoff will differ *at the margin*, which leaves unexploited gains from trade. A situation with unexploited gains from trade is often called an *arbitrage* opportunity.

Definition 1 *Arbitrage. (1) Taking advantage of a price difference between two or more markets. (2) Striking a combination of matching deals that capitalize upon the imbalance between prices.*

The law of one price implies that arbitrage opportunities should not exist in equilibrium. You may ask: who enforces this "law"?

Let's say that you observe that rice sells at one price in Northern China and another price in Southern China. Is there any limit on how much we might expect these prices to differ? Yes. If rice can be transported between North and South—and if buyers and sellers are aware of this price difference—then the price difference between markets should not be greater than the transport cost. If it were, a trader would find it profitable to transport rice between North and South to arbitrage the price difference. This would lower the price in the more expensive market (by increasing supply) and raise it in the less expensive market (by reducing supply). This trade would be expected to continue until prices in these markets differ *by no more than the transport cost.* (I this transport cost is non-zero, we would not expect prices to fully equalize across markets.) At this point, there would be no further arbitrage opportunities. In short, *arbitrage enforces the law of one price.* You may hear of similar scenarios in finance, where professionals are constantly looking for opportunities to "buy low" in one market and "sell high" in another.

Jensen tests the Law of One Price in Table VII of the "Digital Provide." Does the law of one price hold across fish markets in Kerala? The answer appears to be that *before* the introduction of cell phones, it did not. After the introduction of cell phones, the law of one price appears to hold almost all of the time. This example highlights that for arbitrage to function effectively, two preconditions are needed: (1) There is a cost-effective means to transport fish between local markets in Kerala; (2) There is a mechanism that allows sellers to learn *about* the price differentials across markets. Accurate and inexpensive information is crucial to the efficient operation of markets. In Kerala, mobile phones provide that information transmission mechanism.

To test your understanding a bit: assume now that fishermen and consumers are not the same people. In particular, fisherman C and fisherman D face the demand curves above from A and B. On alternating days, C catches 20 fish and D catches 40 fish (and vice versa). C and D can choose to pool their catches and sell 30 fish each to A and B each day. Or, they can keep their catches separate and sell 20 and 40 fish to consumers A and B, respectively, on alternating days. (For simplicity, assume that C sells to A exclusively and that D sells to B exclusively.) Which approach is profit maximizing for C and D? Which maximizes consumer surplus for A and B? Which maximizes social welfare (producer plus consumer surplus)? Let's say that each fisherman recognized that he was 'large' relative to the market (since each seller faces exactly one buyer) and could affect the market price by restricting the quantity of fish available. How would the market equilibrium look different in this case?