

Best Reply Behavior

Michael Peters

June 3, 2005

1 Introduction

So far, we have concentrated on individual optimization. This unified way of thinking about individual behavior makes it possible to understand a number of interesting economic problems. Some of these will be discussed in coming chapters. However, the biggest insights in economics come from thinking about group behavior. Individual and group behavior are tied together in economics by best reply behavior.

As you have seen in prior readings, the solutions to individual maximization problems depend on characteristics of the environment. When many people make maximizing decisions at the same time, their choices may or may not be consistent with one another. Perhaps the biggest insight of all in economics is that consistency restrictions on peoples' choices will impose restrictions on the set of environments that we should ever expect to see.

The oldest application of this idea is the one that you probably saw in operation in your first year economics course. Every consumer's demand for every good depends on the price that she thinks prevails for that good. If consumers' beliefs about these prices are completely arbitrary, then some goods will be in excess demand and some in excess supply. As a consequence, all consumers should have the same belief about the price at which they can buy each good, and this belief should be that the price will be the one at which aggregate demand and aggregate supply are equal. This is where the concept of a market comes from. A market isn't a physical location where traders buy and sell, rather it is a state of beliefs that result in a kind of consistent and predictable behavior on the part of a large group of individual decision makers.

Prior to John Nash and game theory, economics couldn't really get much beyond markets in thinking about group behavior. In cases where markets obviously couldn't give a good description of behavior, very special models were proposed - for example, the monopoly model to deal with the obvious cases in which firms are not price takers. Nash showed that you can use the consistency of individual decisions in a much broader way to understand many non-market problems as well.

The basic idea behind Nash equilibrium is that instead of forming beliefs about prices, people form beliefs about what other people will do. Fixing these beliefs leads us back directly to the maximization framework that we have been studying so far. In order to think about consistency among these beliefs, we have to extend the kind of reasoning we have been using to derive the demand curve to a much more flexible description of best reply behavior.

2 Consumption Theory with Externalities

To understand this method, let's modify our description of the individual choice problem slightly. As before, let's suppose that there are two goods. The price of good y will be 1, and the price of good x will be p throughout this argument. However, let's suppose that there are two consumers instead of just 1. Each consumer has W to spend on x and y . The utility that each person gets for good x depends positively on the amount of good x that the *other* person chooses to consume. This is called a *positive consumption externality*. Such externalities are everywhere. Movies and novels are great examples - they are fine by themselves, but they are even more enjoyable if your friends have seen the same movie or read the same book, because then you can discuss it. On the margin you might be inclined to go to see a certain movie just because you know one of your friends has already seen it.

Let's call the consumers Alice and Bob, denoting Alice's consumption bundle by (x_a, y_a) and Bob's by (x_b, y_b) . Alice's utility function is $u(x_a, y_a, x_b)$ which says that her utility depends on the things that she chooses, x_a and y_a as well as the thing that Bob chooses, x_b . When we say that her *utility* is affected, what we mean is that her preferences are affected by Bob's consumption. So when Bob changes his consumption, all of Alice's indifference curves will change shape. Suppose they change in such a way that the more Bob consumes, the more Alice is willing to pay for an extra unit of good x . This kind of externality is sometimes referred to as a *positive consumption*

externality.

More formally, this means that Alice's marginal rate of substitution of good y for good x increases as Bob increases his consumption. Recall that the marginal rate of substitution is given by the ratio of Alice's *marginal* utility for good x to her marginal utility for good y . Recall as well, that the marginal rate of substitution is the same as the absolute value of the slope of Alice's indifference curve.

Now we are interested in the way that the solution to Alice's utility maximization problem changes as Bob varies his consumption. Lets analyze this a couple of different ways.

2.1 The graphical Approach

There is nothing strange at all about Alice's budget set. She can buy all the good x and y that she can afford at constant prices. So she is just going to choose the highest indifference curve that touches her budget set. The only thing that is new about this problem is the fact that the shape of Alice's indifference curves depend on Bob's consumption. We know how - as we just figured out, the amount of y she is willing to give up (or pay) for one unit of good x is increasing with Bob's consumption. That means that her indifference curves will all get steeper when Bob's consumption goes up. As a result, she will purchase more good x the more Bob purchases.

The first figure shows this simple logic.




Figure 1:

Start with a situation in which Bob's consumption of good x is x_0 . Alice's preferences in this case lead her to choose the bundle (x^*, y^*) which is

the highest indifference curve touching her budget set. If Bob increases his consumption to $x_0 + dx_0$, then Alice's indifference curves become steeper, like the red curves in the figure. The tangency then moves to the right to the point (x', y') . That's about all there is to it.

The consumer's demand function is a type of best reply function. It is common to think of the demand function in a graphical way, sloping downwards and eventually meeting some kind of supply function. We can do the same kind of thing for any best reply function. Here we don't care so much about changes in price and income, but we do care about the impact of a change in Bob's consumption.

The next Figure summarizes what we have figured out.




Figure 2:

The horizontal axis measure Bob's consumption of good x , the vertical axis measures Alice's consumption of good x . The thick upward sloping line records Alice's best reply to Bob's choice. When Bob chooses not to consume any x at all, Alice still finds that the highest indifference curve touching her budget line touches at a point whose x coordinate is $x^*(0)$. On the other hand, if Bob chooses the higher consumption x_b^e , then Alice finds that her tangency occurs at a new point given by $x^*(x_b^e)$ because her indifference curves have all become steeper.

We could do exactly the same sort of exercise for Bob, assuming that his preferences are also changed by Alice's consumption. The axis are already labelled the right way, just think of the curves the other way around. For instance, if Alice chooses to consume 0 units of good x , Bob is still like to want some. The amount that he wants is determined by his best reply function which gives an x value of $x_b^*(0)$.

With this observation, we can make the choices that Alice and Bob make consistent with one another. Alice is really guessing what consumption Bob will choose. Bob is doing the same. If Alice thinks that Bob isn't going to consume any x , her best choice is $x_a^*(0)$. But Bob never chooses 0 units of good x . You can see this by looking at the graph. No matter what consumption he thinks Alice will pick, he always chooses a positive consumption of good x . So Alice is going to see that she is wrong and start changing her beliefs about Bob. This in turn will cause her to change her consumption. It seems pretty compelling that things will settle down when Bob and Alice's beliefs about each other are correct. This happens when Alice consumes x_a^e and Bob chooses x_b^e - the point where their best reply functions meet.

This is a pretty big conceptual leap. We can't really predict what Alice is going to do in this case without knowing what Bob is going to do. We can't say anything about Bob without knowing what Alice is going to do. We seem to be at an impasse. What saves us is the best reply concept. The way to understand what Alice and Bob do is to think about the things they could have done, but didn't. In economic history, this idea is called *counterfactual analysis*. Historians love to imagine what the Canadian Economy would look like if the rail line had never been built to the west coast, or what the American Economy would look like without the American civil war. We do the same - what would Alice do in all the counterfactual situations that she might encounter with Bob. Then we can consider each of these counterfactuals and try to think through whether they make sense.

The point of this section is not so much to show you how to do this. You will learn this method when you study game theory. The point here is to show how our utility theorem is making it possible to do this. Considering all possible counterfactuals is an impossible task using words and intuition - there is just too much to think about. The maximization approach suggested by the utility theorem allows us to reduce these counterfactuals to very simple geometrical and mathematical objects that we can manipulate.

2.2 An algebraic Treatment

An alternative approach to the one described above is to provide enough structure on the preferences of Alice and Bob that we can find their best reply functions explicitly. Suppose that Alice's utility function is $u(x_a, x_b, x_b) =$

$x_a^{\alpha(x_b)} y_a^{1-\alpha(x_b)}$ where we are going to let

$$\alpha(x_b) \equiv \frac{px_b}{W} \left(\frac{3}{4} \right) + \frac{W - px_b}{W} \left(\frac{1}{4} \right) \quad (1)$$

because we need something concrete to find the best reply functions explicitly.

If you recall, the marginal rate of substitution of y for x is given by the ratio of the marginal utility of good x to the marginal utility of good y . In words, this is the amount of good y the consumer would be willing to give up in order to get 1 additional unit of good x . The MRS is then given by

$$\frac{\alpha(x_b)x^{\alpha(x_b)-1}y^{1-\alpha(x_b)}}{(1-\alpha(x_b))x^{\alpha(x_b)}y^{-\alpha(x_b)}} = \frac{\alpha(x_b)y}{(1-\alpha(x_b))x}$$

This is an increasing function of x_b , because $\frac{d\alpha}{dx_b} = \frac{1}{2} \frac{p}{W} > 0$.

From our previous work, we know that the x coordinate of the bundle that maximizes Alice's utility is

$$x_a^*(x_b) = \frac{\alpha(x_b)W}{p} = \frac{(W + 2px_b)W}{4W} \frac{W}{p} = \frac{W + 2px_b}{4p} \quad (2)$$

Now we can put some more structure on the diagram that is given above. For example it is immediate from (2) that the intercept of Alice's best reply function is $x_a^*(0) = \frac{W}{4p}$. This function is linear and has slope $\frac{1}{2}$.

If Bob has the same preferences, then his best reply function looks exactly the same. For their choices to be consistent with one another we want both Alice and Bob to guess each other's actual consumption correctly. In other words, Alice's guess about what Bob would do, x_b is equal to what Bob actually does do when he correctly predicts Alice's consumption, $x_a^*(x_b)$. That is

$$x_b = x_b^* [x_a^*(x_b)]$$

Notice how this works. We start with Bob's consumption and use this to find Alice's consumption. We take the result and use it find Bob's consumption. If we get the consumption level we started with, then we have found the consumption levels that are consistent with one another.

Lets try this in the previous diagram.




Figure 3:

Start with Bob's consumption. In the diagram this is x_0 and is labelled on the horizontal axis, where Bob's consumption is measured. We need to compute Alice's consumption. From (2), this is

$$x_a^*(x_0) = \frac{W + 2px_0}{4p}$$

Using this, we can compute Bob's consumption by measuring the (horizontal) distance over to his best reply function. This would be

$$x_b^*\left(\frac{W + 2px_0}{4p}\right) = \frac{W + 2p\left(\frac{W + 2px_0}{4p}\right)}{4p} = x_1$$

It should be clear from the diagram that this won't give us the right answer unless we choose x_0 at the point where the best reply functions cross. From the algebra above, this is simply the solution to

$$\frac{W + 2p\left(\frac{W + 2px}{4p}\right)}{4p} = x$$

or

$$\frac{W}{2} = px$$

This is a simple Nash equilibrium.