Fixed Point Theorems

Definition: Let X be a set and let $f : X \to X$ be a function that maps X into itself. (Such a function is often called an **operator**, a **transformation**, or a **transform** on X, and the notation $T(x)$ or even Tx is often used.) A **fixed point** of f is an element $x \in X$ for which $f(x) = x$.

Example 1: Let X be the two-element set $\{a, b\}$. The function $f : X \to X$ defined by $f(a) = b$ and $f(b) = a$ has no fixed point, but the other three functions that map X into itself each have one or two fixed points. More generally, let X be an arbitrary set; every constant function $f : X \to X$ mapping X into itself has a unique fixed point; and for the identity function $f(x) = x$, every point in X is a fixed point.

Remark: Note that the definition of a fixed point requires no structure on either the set X or the function f.

Example 2: Let X be the unit interval $[0, 1]$ in $\mathbb{R}$. The graph of a function $f : X \to X$ is a subset of the unit square $X \times X$. If f is continuous, then its graph is a curve from the left edge of the square to the right edge (see Figure 1). A fixed point of f is an element of $[0, 1]$ at which the graph of f intersects the 45°-line. Intuitively, it seems clear that if f is continuous then it must have a fixed point (its graph must cross or touch the 45°-line), and also that discontinuous functions f may not have a fixed point.

Fixed points show up in a number of contexts, but most prominently in the notion of equilibrium. We typically represent a "system" by specifying the set of "states" the system can be in, and also specifying how the system moves, or "transitions," from one state to another. For example, we can represent the *state* of a system of markets as a list $\mathbf{p}$ of prices and a list $\mathbf{z}$ of net demand quantities. If we want to say this is the state at time t, we could write $\mathbf{s}_t = (\mathbf{p}_t, \mathbf{z}_t)$. We sometimes also include some kind of specification of how the economy moves from state to state; let's say $\mathbf{s}_{t+1} = f(\mathbf{s}_t)$ for some *transition function* $f : S \to S$, where S is the set of all possible states. A *stationary state* would be defined as a state $\mathbf{s}$ for which $f(\mathbf{s}) = \mathbf{s}$, so that if $\mathbf{s}_t = \mathbf{s}$ then $\mathbf{s}_{t+1} = \mathbf{s}_t$. We would also call this an *equilibrium* of the system described by $f : S \to S$ — an equilibrium is a stationary state. This is exactly the motivation for our definition of Walrasian equilibrium as a pair $(\mathbf{p}, \mathbf{z})$ at which $\mathbf{z} = \mathbf{0}$ (markets clear): while we don't specify a function f, we believe that whatever the true f is, it satisfies $f(\mathbf{s}) \neq \mathbf{s}$ when $\mathbf{z} \neq \mathbf{0}$ and $f(\mathbf{s}) = \mathbf{s}$ when $\mathbf{z} = \mathbf{0}$ — that the fixed points of f are the states at which markets clear. Similarly, we define a game-theoretic equilibrium to be a list $\mathbf{s}$ of strategies or actions by the players that is a stationary state of any transition function f that captures the idea that $f(\mathbf{s}) = \mathbf{s}$ for strategy-lists $\mathbf{s}$ in which no player wants to change his strategy.

Contractions and The Banach Fixed Point Theorem

Definition: Let (X, d) be a metric space. A **contraction** of X (also called a **contraction mapping** on X) is a function $f : X \to X$ that satisfies

$$\forall x, x' \in X : d\big(f(x'), f(x)\big) \leqq \beta d(x', x)$$

for some real number $\beta < 1$. Such a β is called a **contraction modulus** of f. (Note that if β is a contraction modulus of f and $\beta < \beta' < 1$, then β' is also a contraction modulus of f.)

In other words, a transformation is a contraction if the images of any pair of points are always closer together than the points themselves, and if the ratio of these two distances is bounded away from 1. In Example 3 their ratio is not bounded away from 1.

Example 3: Let $I = (0, 1)$, the open unit interval, and let $f : I \to I$ be the function $f(x) = \frac{1}{2} + \frac{1}{2}x^2$. Then $|f(x') - f(x)| = \frac{1}{2}(x' + x)|x' - x| < |x' - x|$ for all $x, x' \in I$, but f is not a contraction because if $\beta < 1$ and $x, x' > \beta$, then $|f(x') - f(x)| > \beta|x' - x|$. There is no $\beta < 1$ that will satisfy the inequality in the definition of a contraction.

Example 4: Let $f : \mathbb{R} \to \mathbb{R}$ be a differentiable real function. If there is a real number $\beta < 1$ for which the derivative f' satisfies $|f'(x)| \leqq \beta$ for all $x \in \mathbb{R}$, then f is a contraction with respect to the usual metric on $\mathbb{R}$ and β is a modulus of contraction for f. This is a straightforward consequence of the Mean Value Theorem: let $x, x' \in \mathbb{R}$ and wlog assume $x < x'$; the MVT tells us there is a number $\xi \in (x, x')$ such that $f(x') - f(x) = f'(\xi)(x' - x)$ and therefore $|f(x') - f(x)| = |f'(\xi)||x' - x| \leqq \beta|x' - x|$. The same MVT argument establishes that if $\beta < 1$ and $f : (a, b) \to (a, b)$ satisfies $|f'(x)| \leqq \beta$ for all $x \in (a, b)$, then f is a contraction of (a, b).

Theorem: Every contraction mapping is continuous.

Proof: Let $T : X \to X$ be a contraction on a metric space (X, d), with modulus β, and let $\overline{x} \in X$. Let $\epsilon > 0$, and let $\delta = \epsilon$. Then $d(x, \overline{x}) < \delta \Rightarrow d(Tx, T\overline{x}) \leqq \beta\delta < \epsilon$. Therefore T is continuous at $\overline{x}$. Since $\overline{x}$ was arbitrary, T is continuous on X. ■

The above proof actually establishes that a contraction mapping is *uniformly* continuous:

Definition: Let (X, d_X) and (Y, d_Y) be metric spaces. A function $f : X \to Y$ is **uniformly continuous** if for every $\epsilon > 0$ there is a $\delta > 0$ such that

$$\forall x, x' \in X : d_X(x, x') < \delta \Rightarrow d_Y\big(f(x), f(x')\big) < \epsilon.$$

Notice how this definition differs from the definition of continuity: uniform continuity requires that, for a given ϵ, a single δ will work across the entire domain of f, but continuity allows that the δ may depend upon the point x at which continuity of f is being evaluated.

Theorem: Every contraction mapping is uniformly continuous.

Banach Fixed Point Theorem: Every contraction mapping on a complete metric space has a unique fixed point. (This is also called the Contraction Mapping Theorem.)

Proof: Let $T : X \to X$ be a contraction on the complete metric space (X, d), and let β be a contraction modulus of T. First we show that T can have at most one fixed point. Then we construct a sequence which converges and we show that its limit is a fixed point of T.

(a) Suppose x and x' are fixed points of T. Then $d(x, x') = d(Tx, Tx') \leqq \beta d(x, x')$; since $\beta < 1$, this implies that $d(x, x') = 0$, *i.e.*, $x = x'$.

(b) Let $x_0 \in X$, and define a sequence $\{x_n\}$ as follows:

$$x_1 = Tx_0, \quad x_2 = Tx_1 = T^2 x_0, \quad \ldots \quad , x_n = Tx_{n-1} = T^n x_0, \quad \ldots$$

We first show that adjacent terms of $\{x_n\}$ grow arbitrarily close to one another — specifically, that $d(x_n, x_{n+1}) \leqq \beta^n d(x_0, x_1)$:

$$\begin{aligned} d(x_1, x_2) &\leqq \beta d(x_0, x_1) \\ d(x_2, x_3) &\leqq \beta d(x_1, x_2) \leqq \beta^2 d(x_0, x_1) \\ &\cdots \\ d(x_n, x_{n+1}) &\leqq \beta d(x_{n-1}, x_n) \leqq \beta^n d(x_0, x_1). \end{aligned}$$

Next we show that if $n < m$ then $d(x_n, x_m) < \beta^n \frac{1}{1-\beta} d(x_0, x_1)$:

$$\begin{aligned} d(x_n, x_{n+1}) &\leqq \beta^n d(x_0, x_1) \\ d(x_n, x_{n+2}) &\leqq d(x_n, x_{n+1}) + d(x_{n+1}, x_{n+2}) \\ &\leqq \beta^n d(x_0, x_1) + \beta^{n+1} d(x_0, x_1) = (\beta^n + \beta^{n+1}) d(x_0, x_1) \\ &\cdots \\ d(x_n, x_m) &\leqq (\beta^n + \beta^{n+1} + \cdots + \beta^{m-1}) d(x_0, x_1) \\ &= \beta^n (1 + \beta + \beta^2 + \cdots + \beta^{m-1-n}) d(x_0, x_1) \\ &< \beta^n (1 + \beta + \beta^2 + \cdots) d(x_0, x_1) \\ &= \beta^n \frac{1}{1-\beta} d(x_0, x_1). \end{aligned}$$

Therefore $\{x_n\}$ is Cauchy: for $\epsilon > 0$, let N be large enough that $\beta^N \frac{1}{1-\beta} d(x_0, x_1) < \epsilon$, which ensures that $n, m > N \Rightarrow d(x_n, x_m) < \epsilon$. Since the metric space (X, d) is complete, the Cauchy sequence $\{x_n\}$ converges to a point $x^* \in X$. We show that x^* is a fixed point of T: since $x_n \to x^*$ and T is continuous, we have $Tx_n \to Tx^*$ — *i.e.*, $x_{n+1} \to Tx^*$. Since $x_{n+1} \to x^*$ and $x_{n+1} \to Tx^*$, we have $Tx^* = x^*$. ■

Note that the proof of uniqueness did not require that the space be complete.

First Cournot Equilibrium Example: Two firms compete in a market, producing at output levels q_1 and q_2. Each firm responds to the other firm's production level when choosing its own level of output. Specifically (with a_1, a_2, b_1, b_2 all positive),

$$q_1 = r_1(q_2) = a_1 - b_1 q_2$$
$$q_2 = r_2(q_1) = a_2 - b_2 q_1$$

but $q_i = 0$ if the above expression for q_i is negative. See Figure 2. The function $r_i : \mathbb{R}_+ \to \mathbb{R}_+$ is firm i's *reaction function*. Define $\overline{r} : \mathbb{R}^2_+ \to \mathbb{R}^2_+$ by $\overline{r}(q_1, q_2) = \big(r_1(q_2), r_2(q_1)\big)$. The function $\overline{r}$ is a contraction with respect to the city-block metric if $b_1, b_2 < 1$:

$$\begin{aligned}
d\big(\overline{r}(\mathbf{q}), \overline{r}(\mathbf{q}')\big) &= |\overline{r}_1(\mathbf{q}) - \overline{r}_1(\mathbf{q}')| + |\overline{r}_2(\mathbf{q}) - \overline{r}_2(\mathbf{q}')| \\
&= |(a_1 - b_1 q_2) - (a_1 - b_1 q_2')| + |(a_2 - b_2 q_1) - (a_2 - b_2 q_1')| \\
&= b_1 |q_2' - q_2| + b_2 |q_1' - q_1| \\
&\leqq \max\{b_1, b_2\}(|q_1 - q_1'| + |q_2 - q_2'|) \\
&= \max\{b_1, b_2\}\, d(\mathbf{q}, \mathbf{q}').
\end{aligned}$$

Therefore we have an "existence and uniqueness result" for Cournot equilibrium in this example: $\overline{r}$ has a unique fixed point $\mathbf{q}^*$ — a unique Cournot equilibrium — if each $b_i < 1$.

There are several things to note about this example. First, note that while the condition $b_1, b_2 < 1$ is *sufficient* to guarantee the existence of an equilibrium, it is not necessary. Second, note that we could have obtained the same result, and actually calculated the equilibrium production levels, by simply solving the two "response" equations simultaneously. The condition $b_1, b_2 < 1$ is easily seen to be sufficient (and again, not necessary) to guarantee that the two-equation system has a solution. (Note that $b_1 b_2 \neq 1$ is in fact sufficient.) However, things might not be so simple if the response functions are not linear, or if there are more firms (and therefore more equations). We'll consider a second, nonlinear example shortly.

A third thing to note about the example is that we used the city-block metric instead of the Euclidean metric. This highlights an important and useful fact: a given function mapping a set X into itself may be a contraction according to one metric on X but not be a contraction according to other metrics. Recall that the definition of a fixed point, and therefore whether a given point in the function's domain *is* a fixed point, does not depend on the metric we're using, and in fact does not even require that X be endowed with any metric structure. We've earlier seen that a judicious choice of metric can make a proof easier; here we see that a judicious choice of metric can make a method of proof available that would not work with a different metric.

Exercise: Let A be the matrix

$$A = \begin{bmatrix} \frac{1}{4} & \frac{1}{2} \\ \frac{1}{6} & \frac{2}{3} \end{bmatrix},$$

and let $T : \mathbb{R}^2 \to \mathbb{R}^2$ be the transformation defined by $T(x) = Ax$. Let e_1 and e_2 denote the unit vectors in $\mathbb{R}^2$,

$$e_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad e_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

(a) Plot the locus of all points $x \in \mathbb{R}^2$ that satisfy $\|x\|_\infty = 1$ and the locus of all points that satisfy $\|x\|_1 = 1$. In the same diagram, plot the points $T(e_1)$ and $T(e_2)$.

(b) Is the transformation T a contraction? Does this depend on which norm you're using? (The geometry in (a) should provide you with some help in answering this question.) For each of the norms $\|\cdot\|_\infty$ and $\|\cdot\|_1$, determine whether T is or is not a contraction with respect to the norm.

The Brouwer Fixed Point Theorem

In Example 2 in the preceding section the Banach Theorem seems somewhat limited: it seems intuitively clear that any continuous function mapping the unit interval into itself will have a fixed point, but the Banach Theorem applies only to functions that are contractions. An elementary example is the function $f(x) = 1 - x$, which has an obvious fixed point at $x = 1/2$. But for every x and x' in the interval $[0, 1]$, $d\big(f(x), f(x')\big) = d(x, x')$, so f is not a contraction and the Banach Fixed Point Theorem doesn't apply to f. The fixed point theorem due to Brouwer covers this case as well as a great many others that the Banach Theorem fails to cover because the relevant functions aren't contractions.

Brouwer Fixed Point Theorem: Let S be a nonempty, compact, convex subset of $\mathbb{R}^n$. Every continuous function $f : S \to S$ mapping S into itself has a fixed point.

The Brouwer Theorem requires only that f be continuous, not that it be a contraction, so there are lots of situations in which the Brouwer Theorem applies but the Banach Theorem doesn't. In particular, Brouwer's Theorem confirms our intuition that any continuous function mapping $[0, 1]$ into itself has a fixed point, not just the functions that satisfy $|f'(x)| \leqq \beta$ for some $\beta < 1$. But conversely, the Banach Theorem doesn't require compactness or convexity — in fact, it doesn't require that the domain of f be a subset of a vector space, as this version of Brouwer's Theorem does. So there are also lots of situations where Banach's Theorem applies and Brouwer's doesn't.

Proofs of Brouwer's Theorem require some highly specialized mathematical ideas. For us, the benefit of developing these ideas in order to work through a proof of the theorem doesn't come close to justifying the time cost it would require, so we won't go there.

Here's a generalization of Brouwer's Theorem to normed vector spaces (which, in particular, don't have to be finite-dimensional, as Brouwer's Theorem requires):

Schauder Fixed Point Theorem: Let S be a nonempty, compact, convex subset of a normed vector space. Every continuous function $f : S \to S$ mapping S into itself has a fixed point.

This theorem would apply, for example, to any compact convex subset of $C[0, 1]$, the vector space of continuous functions on the unit interval, with the max norm.

The Method of Successive Approximations

The Banach and Brouwer Theorems are *existence theorems*: when a function satisfies the assumptions of one of the theorems, the theorem tells us that the function has a fixed point. We've described how economic and game theoretic equilibria can generally be represented as fixed points; therefore fixed point theorems can tell us when an economic or strategic model has an equilibrium, which is important information.

Often we also want to *find* the equilibrium, or equilibria. For example, if we want to know what will be the result of some policy change or other exogenous change in the economy, we typically want to find the new equilibrium for the new economic parameter values. So what we need are methods for *computing* equilibria, *i.e.*, fixed points. Here we'll present only the most elementary version of this problem. In subsequent courses you'll study more powerful methods for dealing with particular classes of problems.

Let's start by going back and taking a look at our proof of the Banach Fixed Point Theorem. We began by using the given function f to recursively define a sequence of points in the space X via the *recursion formula* $x_{n+1} = f(x_n)$. Then everything we did in the proof prior to the last two sentences was in the service of proving that the sequence converges: in order to get that result we used the facts that the space is complete and that the function f is a contraction. Then, once we knew that the sequence converges, in the last two sentences of the proof we simply showed that the sequence's limit is a fixed point of f. The argument in these last two sentences used only the fact that f is continuous and that $\lim x_n$ exists — we no longer needed to use either completeness of the space or the fact that f is a contraction.

The following theorem and its proof repeat the result and the argument in those last two sentences of the Banach Theorem's proof: if we have a sequence that's defined recursively from a continuous function f and the sequence converges, then the sequence's limit is a fixed point of f.

Theorem: Let (X, d) be a metric space, let $f : X \to X$, let $x_0 \in X$, and let $\{x_n\}$ be the sequence defined recursively from f and x_0 by $x_{n+1} = f(x_n)$. If f is continuous and $\{x_n\}$ converges, then $\lim x_n$ is a fixed point of f.

Proof: Let $x^* = \lim x_n$. Then

$$f(x^*) = f(\lim x_n) = \lim f(x_n) = \lim x_{n+1} = \lim x_n = x^*,$$

where $f(\lim x_n) = \lim f(x_n)$ follows from continutiy of f. ■

This theorem is an example of *The Method of Successive Approximations* — recursively constructing a sequence that will converge to the fixed point (or other value) we're trying to find. The sequence's terms can be thought of as approximations to the fixed point, and — if the sequence converges! — the terms eventually become arbitrarily close approximations to the fixed point. The Banach Fixed Point Theorem tells us that a contraction mapping defined on a complete metric space will have a fixed point, because in that case any sequence defined by recursively applying the function will in fact converge.

Ideally, the sequence we construct in implementing the method of successive approximations would be one in which successive terms of the sequence are *successively closer* approximations to the fixed point. The above theorem doesn't guarantee that. But the proof of the Banach Theorem works precisely because the terms *are* successively closer to the the function's fixed point, as the following theorem guarantees for any sequence defined recursively from a contraction mapping. (Note that the theorem does not assume that the space is complete, and it does not guarantee the *existence* of a fixed point. We've shown earlier that if a contraction *does* have a fixed point, it will be unique.)

Theorem: If T is a contraction with contraction modulus β on a metric space (X, d), and if T has a fixed point x^*, then for any $x \in X$,

$$\forall n \in \mathbb{N} : d(T^n x, x^*) \leqq \beta^n d(x, x^*).$$

Proof:

$$\begin{aligned} d(T^n x, x^*) &= d(TT^{n-1}x, Tx^*), \quad \text{because } x^* \text{ is a fixed point of } T \\ &\leqq \beta d(T^{n-1}x, x^*), \quad \text{because } T \text{ is a contraction} \\ &\leqq \beta^2 d(T^{n-2}x, x^*) \\ &\leqq \quad \cdots \quad \leqq \beta^{n-1} d(Tx, x^*) \\ &= \beta^n d(x, x^*). \quad \blacksquare \end{aligned}$$

The Cournot Equilibrium Example Again: Suppose the current "state" of the market is $\mathbf{q}(0) = (q_1(0), q_2(0))$ — Firm i is producing $q_i(0)$ units. Suppose further that "tomorrow" (at time $t = 1$) each firm chooses its production level $q_i(1)$ by responding to the amount its rival firm produced today (at time $t = 0$), and similarly at each subsequent date:

$$q_1(t+1) = r_1(q_2(t)) = a_1 - b_1 q_2(t) \qquad \text{and} \qquad q_2(t+1) = r_2(q_1(t)) = a_2 - b_2 q_1(t),$$

or more concisely, $\mathbf{q}(t+1) = \overline{r}(\mathbf{q}(t))$. We've already proved that the function $\overline{r}(\cdot)$ is a contraction if $b_1, b_2 < 1$, in which case the above theorem guarantees that the sequence $\{\mathbf{q}(t)\}$ will converge monotonically (in distance) to the Cournot equilibrium $\mathbf{q}^*$. Try it yourself by choosing values for the parameters a_1, a_2, b_1, b_2 — for example, $a_1 = 25, a_2 = 30, b_1 = 3/5, b_2 = 1/2$ — and any starting state $\mathbf{q}(0)$. Plot the sequence along with the reaction curves in the q_1-q_2-space.

Second Cournot Equilibrium Example: Suppose the response functions in our Cournot example are

$$q_1 = r_1(q_2) = \frac{1}{2(1+q_2)} \qquad \text{and} \qquad q_2 = r_2(q_1) = \frac{1}{2}e^{-q_1}.$$

Now it's not so easy to calculate the equilibrium, or even to tell whether an equilibrium exists, as it was in the linear example, where we could simply solve the two response equations simultaneously. But it's possible to show that the function $\overline{r} : \mathbb{R}^2_+ \to \mathbb{R}^2_+$ defined as before is a contraction (so a Cournot equilibrium exists and is unique), and one can use the Method of Successive Approximations to compute an approximation to the equilibrium.

Exercise: Use Excel (or any other computational program) and the Method of Successive Approximations to compute the Cournot equilibrium to three decimal places in the Second Cournot Example.

Example: Let a be a positive real number and define the function $f : \mathbb{R}_{++} \to \mathbb{R}_{++}$ as

$$f(x) = \frac{1}{2}x + \frac{a}{2x}.$$

Does this function have a fixed point? Neither the Banach Theorem nor the Brouwer Theorem gives us the answer, because the space $\mathbb{R}_{++}$ is neither complete nor compact. And we can't extend the domain of f to all of $\mathbb{R}_+$ (which would make the space complete) because $\lim_{x\to 0} f(x) = \infty$. But the method of successive approximations leads us to the (unique) fixed point of f, as follows.

Let's choose some (arbitrary) positive real number as x_0 and then recursively define the sequence $\{x_n\}$ by

$$x_{n+1} = f(x_n) = \frac{1}{2}x_n + \frac{a}{2x_n}.$$

We've seen this sequence before: we showed that it converges to $\sqrt{a}$. Therefore, since f is continuous, we know from the above theorem that $\sqrt{a}$ is a fixed point of f.

It's also straightforward to show that $\sqrt{a}$ is the *unique* fixed point of f: it's easy to show that

$$(1) \text{ if } x < \sqrt{a} \text{ then } f(x) > \sqrt{a}, \quad \text{and} \quad (2) \text{ if } x > \sqrt{a} \text{ then } \sqrt{a} < f(x) < x;$$

therefore no $x \neq \sqrt{a}$ can be a fixed point of f.

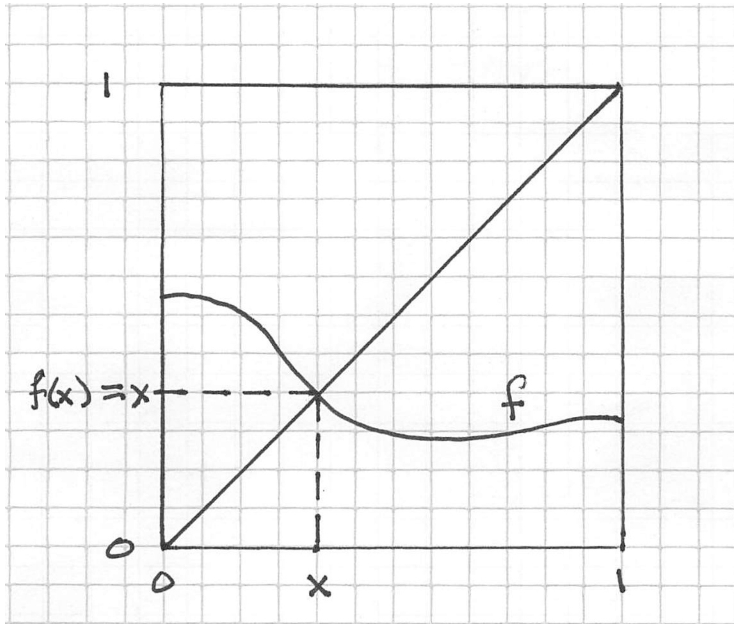

Figure 1

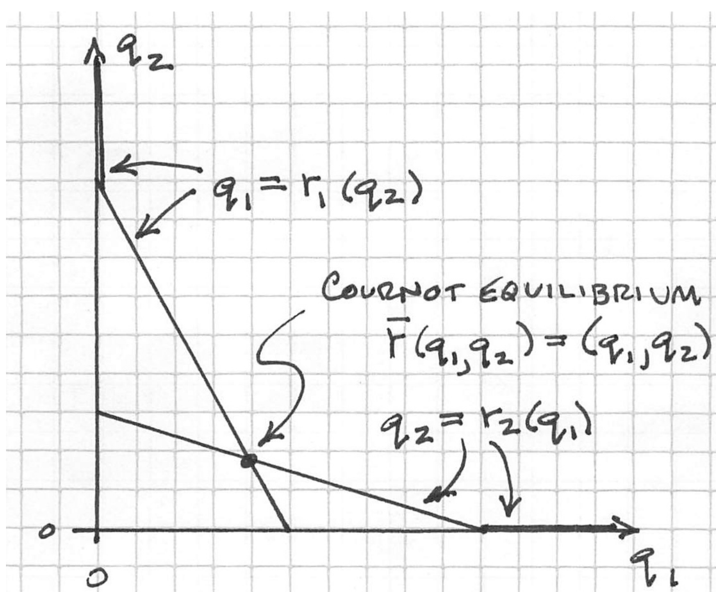

Figure 2