Hotelling's game

Hotelling's game represents a form of competition. In its two interpretations, two entrepreneurs are choosing either the location of their business (fixing market prices) or the price of their goods (fixing their locations). For the sake of argument, we assume their goods are perfect substitutes so that this is competition in the most straightforward sense.

Although Hotelling's initial paper in 1929 did not specify this, most modern takes on the game represent it as two merchants trying to attract customers from along a beach; we'll assume they're selling ice cream, but could just as well be trashy t-shirts or various tchotchkes. Firms face 0 costs of production. The beach is represented as the real numbers from 0 to L — that is, the beach is of length L — and consumers are spread evenly along it. Each consumer desires one serving of ice cream, and doesn't care which store he gets it from. Importantly, customers dislike having to walk along the beach (lazy Americans!) and if they have to walk distance d to get to the ice cream shop, they resent it to the tune of dt where t is the cost per unit walked (here, we assume that $t > 0$ and that the cost is subtracted from utility; we could also assume that $t < 0$ and the cost is *added* to utility — the math is identical — but I consider it simpler to think about subtracting costs as positive quantities; it is truly a matter of personal preference).

To clarify what is meant when we say, "consumers are spread evenly along [the beach]," let's use an example: suppose all consumers louging between 0 and $\frac{L}{3}$ decide to get their ice cream from firm A, while all consumers lounging between $\frac{L}{3}$ and L decide to get their ice cream from firm B. We then say that firm A sells $\frac{L}{3} - 0 = \frac{L}{3}$ units of ice cream, while firm B sells $\frac{L}{-}\frac{L}{3} = \frac{2L}{3}$ units. Mathematically, each so-called "consumer" is infinitesimal, and we integrate the mass of consumers going to a particular ice cream shop to determine that firm's market demand. However, since we assume that customers are evenly distributed we can avoid any idea of integration and just look at intervals containing customers which go to one firm or another.

Since customers don't care one way or another which firm they get ice cream from, and they really want the one unit they demand, their choice of supplier will revolve around two things: the price the firms are charging, and the distance to walk to one firm or another. At this point, it will be useful to assume that firm A is located at point a, and firm B is located at point b along the boardwalk. To keep analysis simple, we'll assume that $a < b$.

Suppose that some consumer is located at point x along the beach. If this customer buys ice cream from firm A, she will have to pay price p_A for the ice cream cone and will also have to walk $|x - a|$ to get to the shop. Her transportation cost is then $|x - a|t$, giving her a total cost of $p_A + |x - a|t$. Similarly, if she gets ice cream from firm B her total cost will be $p_B + |x - b|t$. Since she wants ice cream no matter what, she will look only to minimize her cost and will buy from firm A if $p_A + |x - a|t < p_B + |x - b|t$, from firm B if $p_B + |x - b|t < p_A + |x - a|t$, and we will allow her behaviour to vary as fits our needs if she is indifferent between the two.

Although this may seem simplistic, firms competing in prices can obtain three possible outcomes:

(a) The firm can be outpriced by its rival, and obtain 0 profits.

(b) The firm can outprice its rival and receive the entire market. Its profits are then Lp_i.

(c) The firm can share the market with its rival and receive m_i as a fraction of the market. Its profits are then $m_i p_i$.

Since those facts are fairly obvious, a good question is when one situation or another might obtain. We'll assume that we're talking about firm A for now; hopefully it's obvious how this discussion will generalize to firm B's problem.

(a) Suppose firm A is outpriced by firm B, and $\pi_A = 0$. Then by assumption, no customer wants to purchase from firm A; formally, for all x, $p_A + |x - a|t > p_B + |x - b|t$. Suppose that $x > b$; then we have

$$p_A + (x-a)t = p_A + |x-a|t > p_B + |x-b|t = p_B + (x-b)t$$

This tells us

$$p_A > p_B + (a-b)t$$

Similarly, suppose $x < a$; then we have

$$p_A + (a-x)t = p_A + |x-a|t > p_B + |x-b|t = p_B + (b-x)t$$

This tells us

$$p_A > p_B + (b-a)t$$

Lastly, suppose $a < x < b$; then we have

$$p_A + (x-a)t = p_A + |x-a|t > p_B + |x-b|t = p_B + (b-x)t$$

This tells us

$$p_A > p_B + (a+b-2x)t$$

We can use algebra to see that

$$p_B + (b-a)t > p_B + (a+b-2x)t > p_B + (a-b)t$$

So for this zero-market condition to hold, we need

$$p_A > p_B + (b-a)t$$

(b) Suppose firm A outprices firm B, and $\pi_A = Lp_A$. This is simply the converse of the first situation above — in which firm B outprices firm A — and so by analogy this will happen when

$$p_A < p_B - (b-a)t$$

(c) Suppose the firms split the market. From the previous two analyses, this will happen when

$$p_A \in (p_B - (b-a)\,t, p_B + (b-a)\,t)$$

Firstly, notice that since the market is split there must be *some* customer x who is willing to buy from firm A, $p_A + |x-a|t < p_B + |x-b|t$. Suppose that this customer is such that $x < b$. Two cases arise (if you like, skip past this since it's just a mess of algebra):

- $x \in [a, b)$. Pick some x', $a \leq x' < x$. Then $|x'-a| = x'-a < x-a = |x-a|$, and $|x'-b| = b-x' > b-x = |x-b|$. So we have

$$p_A + |x'-a|t < p_A + |x-a|t < p_B + |x-b|t < p_B + |x'-b|t$$

  and the customer at x' prefers firm A to firm B on a price basis.
- $x \leq a$. Pick some $x' < x$. Then

$$|x'-a| = a-x' = a-x-x'+x = |a-x| + |x-x'|$$

  and

$$|x'-b| = b-x' = b-x-x'+x = |b-x| + |x-x'|$$

  Then we see

$$p_A + |x'-a|t = p_A + |a-x|t + |x-x'|t < p_B + |b-x|t + |x-x'|t = p_B + |x'-b|t$$

and so any consumer to the left of x also prefers firm A to firm B.

Pick some x', $x < x' \leq a$. Then

$$|x' - a| = a - x' = a - x - x' + x = |a - x| - |x - x'|$$

and

$$|x' - b| = b - x' = b - x - x' + x = |b - x| - |x - x'|$$

Then we see

$$p_A + |x' - a|t = p_A + |a - x|t - |x - x'|t < p_B + |b - x|t - |x - x'|t = p_B + |x' - b|t$$

and so any consumer between x and a also prefers firm A to firm B.

The moral of the story is that if the consumer at x prefers firm A to firm B, all consumers to the left of x also prefer firm A to firm B. Similar logic will apply to those consumers who prefer firm B continuing to the right of any one consumer who prefers firm B. This moral is derived from this: if the consumer at x prefers firm A to firm B, then the consumer at a will prefer firm A to firm B. Then all customer to the left of a prefer firm A to firm B. The analogous result for firm B will hold, noting that if there is a consumer with $x > b$ who prefers firm A to firm B then no consumers with $x' > b$ will prefer firm B (as if there was such a customer x', all customers between b and L would prefer B, contradicting the fact that $x > b$ prefers A). So in order to split the market, there must be some consumer x and some consumer y, both between a and b, one of whom prefers firm A and one of whom prefers firm B.

If this is the case, then there is a consumer $x_\sim$ who is perfectly indifferent between purchasing ice cream from firm A and purchasing ice cream from firm B. This customer's costs must be

$$p_A + (x_\sim - a)t = p_B + (b - x_\sim)t \qquad \Longrightarrow \qquad x_\sim = \frac{1}{2}\left(b + a + \frac{p_B - p_A}{t}\right)$$

where we were able to do away with absolute values since $a < x < b$. From the previous discussion, all consumers to the left of x must prefer firm A to firm B, and all consumers to the right of x must prefer firm B to firm A. Firm A's profits are then

$$\pi_A = p_A x_\sim = \frac{1}{2} p_A \left(b + a + \frac{p_B - p_A}{t}\right)$$

So we now know profits in the three cases, in which firm A obtains none of the market, some of the market (particularly the left portion of the market), and all of the market. Suppose for now that the two firms share the market. We can compute best responses in the usual way, by taking derivatives.

$$\begin{aligned}
& \max_{p_A} \pi_A(p_A, p_B) = \max_{p_A} \frac{1}{2} p_A \left(b + a + \frac{p_B - p_A}{t}\right) \\
\Longrightarrow \quad & 0 = b + a + \frac{p_B - p_A}{t} - \frac{p_A}{t} \\
\Longrightarrow \quad & p_A = \frac{1}{2}\left((b + a)t + p_B\right)
\end{aligned}$$

Similarly, for firm B,

$$\begin{aligned}
& \max_{p_B} \pi_B(p_B, p_A) = \max_{p_B} p_B \left(L - \frac{1}{2}\left(b + a + \frac{p_B - p_A}{t}\right)\right) \\
\Longrightarrow \quad & 0 = L - \frac{1}{2}\left(b + a + \frac{p_B - p_A}{t}\right) - \frac{p_B}{2t} \\
\Longrightarrow \quad & p_B = Lt - \frac{1}{2}\left((b + a)t - p_A\right)
\end{aligned}$$

To find a candidate Nash equilibrium, we substitute in to find the point (in prices) at which the two best-response functions cross.

$$\begin{aligned} p_A^* &= \frac{1}{2}\left((b+a)\,t + Lt - \frac{1}{2}\left((b+a)\,t - p_A^*\right)\right) \\ \iff \quad \frac{3}{4}p_A^* &= \frac{1}{2}\left(\frac{1}{2}(b+a)\,t + Lt\right) \\ \iff \quad p_A^* &= \frac{1}{3}(2L+b+a)\,t \end{aligned}$$

To find candidate p_B, we can either apply the same method to obtain the crossing point or just substitute in for what we now know p_A to be.

$$\begin{aligned} p_B^* &= Lt - \frac{1}{2}\left((b+a)\,t - \frac{1}{3}(2Lt+b+a)\right) \\ \iff \quad p_B^* &= \frac{4}{3}Lt - \frac{1}{3}((b+a))\,t \\ \iff \quad p_B^* &= \frac{1}{3}(4L-(b+a))\,t \end{aligned}$$

So we now have the prices which will be used if firms A and B are forced/decide to share the market. This allows us to explicitly compute the location of the indifferent consumer,

$$\begin{aligned} x_\sim &= \frac{1}{2}\left(b+a+\frac{p_B^*-p_A^*}{t}\right) \\ \iff \quad x_\sim &= \frac{1}{2}\left(b+a+\frac{1}{3}(4L-(b+a)-(2L+b+a))\right) \\ \iff \quad x_\sim &= \frac{1}{2}\left(b+a+\frac{1}{3}(2L-2(b+a))\right) \\ \iff \quad x_\sim &= \frac{1}{6}(2L+b+a) \end{aligned}$$

From this, we obain

$$\begin{aligned} \pi_A\left(p_A^*,p_B^*\right) &= p_A^* x_\sim \\ &= \left[\frac{1}{3}(2L+b+a)\,t\right]\left[\frac{1}{6}(2L+b+a)\right] \\ \pi_A^* &= \frac{t}{18}(2L+b+a)^2 \end{aligned}$$

$$\begin{aligned} \pi_B\left(p_A^*,p_B^*\right) &= p_B^*\left(L-x_\sim\right) \\ &= \left[\frac{1}{3}(4L-(b+a))\,t\right]\left[L-\frac{1}{6}(2L+b+a)\right] \\ &= \frac{t}{3}(4L-(b+a))\left(\frac{2}{3}L-\frac{1}{6}(b+a)\right) \\ \pi_B^* &= \frac{t}{18}(4L-(b+a))^2 \end{aligned}$$

Now we must address the question of when the firms prefer to share the market. Since $\frac{t}{18} > 0$, profits at the optimum are always positive (this is obvious for firm A; you can check algebraically that $a < b \leq L$ implies $4L-(b+a) > 0$). So we know that neither firm has incentive to deviate so that it loses all of the market to its rival and obtains 0 profits.

However, will a firm want to "undercut" its rival to obtain the entire market? We can check this for firm A alone, as the logic will hold identically for firm B. Recall that firm A obtains the entire market if it sets $p_A < p_B - (b-a)t$. When we setup the problem, we allowed for some fudging of consumer indifference conditions, so let's say this: if the consumer is indifferent between firm A and firm B, he will buy ice cream from the firm with the lower price for its good (not considering transportation costs). Then firm A can obtain the entire market by setting $p_A \leq p_B - (b-a)t$; since any price under this boundary does not affect its demand curve, firm A will choose the largest price allowable and will set $p_A = p_B - (b-a)t$. At this price, profits are

$$\begin{aligned}\pi_A &= L\left(p_B - (b-a)\,t\right)\\ &= \frac{Lt}{3}\left(4L - (b+a) - 3\,(b-a)\right)\\ &= \frac{2Lt}{3}\left(2L - 2b + a\right)\end{aligned}$$

With these profits in mind, will firm A ever want to outprice firm B? We check these profits against its profits from sharing the market.

$$\begin{aligned}& \frac{2Lt}{3}\left(2L - 2b + a\right) \circ \frac{t}{18}\left(2L + b + a\right)^2\\ \iff\quad & 12L\left(2L - 2b + a\right) \circ \left(2L + b + a\right)^2\\ \iff\quad & 24L^2 - 24Lb + 12La \circ 4L^2 + 4Lb + 4La + b^2 + 2ba + a^2\\ \iff\quad & 20L^2 - 28Lb + 8La \circ b^2 + 2ba + a^2\end{aligned}$$

If the left-hand side is bigger, firm A prefers to steal the market from B; if the right-hand side is bigger, firm A prefers to share the market. Since this math is ugly, let's try a few cases to see what's going on.

- $b = L$. Then the relation above is

  $$20L^2 - 28L^2 + 8La = -8L^2 + 8La \circ L^2 + 2La + a^2$$

  Since $a < b = L$, the left-hand side is always smaller than the right-hand side and firm A prefers to share the market regardless. That is, if firm B's position is disadvantageous enough firm A is willing to share rather than undercut, ostensibly because firm A will obtain very few customers from undercutting firm B. Note that this will hold if b is sufficiently close to L.

- $a = 0$. Then the relation above is

  $$20L^2 - 28Lb + 8La = 20L^2 - 28Lb \circ b^2 = b^2 + 2ab + a^2$$

  Then the left-hand side will be larger so long as

  $$b^2 + 28Lb - 20L^2 < 0$$

  According to the quadratic formula, the roots for this form are

  $$\frac{-28L \pm \sqrt{(28L)^2 + 80L^2}}{2} = \left(-14 \pm \sqrt{216}\right) L$$

  Since this is an upward-facing quadratic in b, we see that the left-hand side is larger when $b \in (-14L - \sqrt{216}L, -14L + \sqrt{216}L)$. Clearly, the left-hand bound of this interval is negative and out of the range of this question. So according to the right-hand bound, if $b < 0.697L$ the left-hand side of the profit comparison inequality is larger, and firm A will prefer to steal the market from firm B.

Here's the kicker: suppose that we have a situation in which firm A *does* prefer to steal the market from firm B. Firm B has a well-defined best response to this action by firm A: since it chooses

$$p_B = Lt - \frac{1}{2}(b+a)t + \frac{1}{2}p_A$$

its best response will fall when p_A drops. Notably, since $b + a < 2L$ its price will be positive so long as p_A is positive. *So there is no point at which A can steal the entire market and earn positive profits!* This is a case in which Nash equilibrium does not necessarily exist; eventually, to continue undercutting firm B firm A will need to set prices to be negative, which is clearly a stupid decision. Its option then is to share the market, but if the two firms choose to share the market firm A wants to undercut! There is no equilibrium this parameterization of this game.

So in order to obtain Nash equilibrium, we need b to be sufficiently large compared to a. If this condition holds (as represented in the profit comparison inequality above) then the Nash equilibrium in prices will be

$$p_A^* = \frac{1}{2}\left(2L + (b+a)\right)$$
$$p_B^* = \frac{1}{2}\left(4L - (b+a)\right)$$

Note that in class we did not have this issue: we made the seemingly innocuous assumption that every consumer to the left of firm A will buy from firm A (since she must walk past firm A to get to firm B). While this assumption isn't unreasonable on its face, we can now see that it changes the outcome of the game measurably — notably, it makes Nash equilibrium in prices exist where it may not without the assumption! We could, of course, introduce mechanisms to support this decision (in a fit of intervention, the government mandates that you cannot walk past an ice cream shop without buying some ice cream) but that was not the spirit in which the assumption was introduced in section.

Importantly, if the existence of Nash equilibrium requires b to be significantly large relative to a, all previous arguments (from section) about making this a two-stage game in which location is chosen first fall apart. As you recall, we decided that it was only reasonable for the two firms to locate at the same point; but the small-deviations argument assumed price competition had a well-defined Nash equilibrium! Since we cannot rely on this anymore, the location game argument becomes somewhat nonsensical.

What causes Nash equilibrium to not exist here? In a word, discontinuity. The individual firms' best-response functions are discontinuous in the actions of their opponents: if p_B is high enough, firm A steals the entire market. At the point at which p_B transitions from firm A wanting to share to firm A wanting to steal, firm A's profits jump from the sharing level to the stealing level since it captures a nonnegligible market segment from an infinitesimal change in price. This jump causes issues with the existence of Nash equilibrium.

Hotelling's location game

To wrap up, let's discuss a simpler version of Hotelling's game. In this model, market prices are fixed at p and the firms are choosing a and b, the location of their ice cream stands. Since prices are identical, consumers will walk to the firm which is closer. If firms are equidistant, we can think of a consumer as having a 50% chance of going to either location (this will make more sense when we apply it later).

Assume that firms choose locations a and b, $a < b$. Since prices are identical, we see (from above) that

$$x_\sim = \frac{1}{2}\left(a + b + \frac{p_B - p_A}{t}\right) = \frac{1}{2}(a+b)$$

That is, the indifferent consumer is located exactly halfway between firms A and B (this holds regardless of whether $a < b$ or $b < a$, but let's keep the convention going). Then firm A's profits are

$$\pi_A = px_\sim = \frac{p}{2}(a+b)$$

Since $a < b$, there is some a' such that $a < a' < b$. Then we have

$$\frac{p}{2}(a+b) < \frac{p}{2}(a'+b)$$

so firm A prefers to deviate to location a' rather than remain at a. This argument holds for all $a < b$, so we cannot have $a < b$ in a Nash equilibrium.

Notice that we could say the same thing about supporting a Nash equilibrium with $b < a$. So our only candidate for Nash equilibrium is $a = b$, both firms locating at the same point.

If this is the case, each firm is equidistant to *all* consumers, so each consumer has a 50% chance of going to either firm. Think of it this way: you walk down the beach boardwalk to get ice cream, and there are two no-name ice cream shacks right next to each other charging the same price. Absent any other information (Yelp is not allowed), you'll just randomly choose which one you go to. With this construction, each firm receives 50% of the market, or $\frac{L}{2}$ consumers. Profits to each firm are $\frac{Lp}{2}$.

Suppose that $a = b > \frac{L}{2}$. Firm A can choose a' such that $\frac{L}{2} < a' < a$, and

$$\frac{p}{2}(a'+b) > \frac{Lp}{2}$$

So by shifting a little closer to the center of the market firm A is able to earn higher profits. Firm B faces the same decision, so $a = b > \frac{L}{2}$ cannot constitute a Nash equilibrium. Similarly, if $a = b < \frac{L}{2}$ both firms also face an incentive to shift closer to the center of the market, although in this case they are deviating rightward. This tells us that our only candidate equilibrium is $a = b = \frac{L}{2}$.

Suppose firm A deviates to some $a' < a$ from this strategy. Then its profits are

$$\frac{p}{2}(a'+b) = \frac{a'p}{2} + \frac{Lp}{4} < \frac{ap}{2} + \frac{Lp}{4} = \frac{Lp}{2}$$

So by deviating it reduces its own payoff. It follows that Nash equilibrium in the Hotelling location game — where prices are exogenously fixed by the market! — is both firms locating at the same point, each precisely halfway along the beach.

Now, many people consider this to be a reason for why gas stations often crowd the same corner, or you find auto-repair shops all along the same drag in a town. There is some merit to this argument, but there are additional complexities regarding zoning and other concerns which start confounding the issue. Hotelling's analysis has found a much different home, however, in the world of political theory. Consider the electorate of a country as existing along a one-dimensional continuum of beliefs (the political beach, if you will); a candidate who wants to win office must select a location along this continuum — find a political position to campaign for — and then attract voters. If voters are fixed in their locations along the continuum, it turns out to be optimal for political candidates to locate precisely along the middle of the political spectrum; since "middle" is a little bit of a weasel word, it will help to mention that this phenomenon is referred to as the Median Voter Theorem (and is well-described on Wikipedia). Of course, it makes some silly assumptions — political beliefs are not one-dimensional, although in a two-party system they can appear to be — but we can actually see this intuitively in American elections: in the primaries, candidates tend to stake out more extreme positions than in the general elections. This is because in the primaries, the "middle" of the political spectrum is extremely skewed (Republican candidates face only, roughly, the right half of the spectrum while Democratic candidates face only the left half) while in the general election a candidate will face the electorate as a whole. Political commentators often bring this up when discussing how a candidate's ability to appeal to her base may negatively reflect on her chances of winning the overall election.