# Reading on Directed Search

Michael Peters

February 23, 2021

This reading will describe a model that is used extensively in macroeconomics to understand labor markets. The model is a part of a larger literature on search and matching. I'll explain the model in its simplest form here. Richer variants of the model can be used to understand wage distributions and unemployment duration. The model has also been used to do econometric evaluation of labor markets.

## 1 Bertrand Equilibrium

The model of directed search emerged as a response to something called *Bertrand Equilibrium* - a model designed to understand price competition. So I'll describe it first. In the Bertrand model, there are two firms. They have constant production costs, say $c_1$ for firm 1 and $c_2 < c_1$for firm 2. There is a consumer who has some utility function $u(q, P)$ describing her payoff when she buys $q$ units in total and pays price $P$ for them. We want to assume $u$ is increasing in $q$ and decreasing in $P$. We'll assume that $u(Q(p), pQ(p)) \geq u(q, pq)$ for all $q$. In other words, the consumer's demand curve is $Q(p)$. Of course, the consumer can allocate her purchases between the two firms in any way that she wants - lets just use $q_1$ to be the amount she buys from firm 1 and $q_2$ as the amount she buys from firm 2.

As we will use this below, let $p_i^*$ be the *monopoly price* for firm $i$, i.e, it is the price that maximizes $pQ(p) - c_iQ(p)$

To describe the game we need to describe the strategies of each of the three players and the payoffs that are associated with each set of strategies. For strategies, firms just set prices. The consumer must choose how much to buy from each firm for each pair of prices that the firms choose. Payoffs for firm 1

$$\pi_1 (p_1, p_2, q_1 (\cdot), q_2 (\cdot)) = p_1 q_1 (p_1, p_2) - c_1 q_1 (p_1, p_2)$$

and for firm 2

$$\pi_2 (p_1, p_2, q_1 (\cdot), q_2 (\cdot)) = p_2 q_2 (p_1, p_2) - c_2 q_2 (p_1, p_2).$$

For the consumer payoffs are $U (q_1 + q_2, p_1 q_1 + p_2 q_2)$.

Lets use subgame perfect Nash equilibrium as our solution concept. Then we can use backward induction to figure out what everyone will do. Even using backward induction, to describe an equilibrium what we need to do is to specify the strategies that each of the players are using. For firms this is easy, just prices. For consumers it is more complicated.

We have to write down what the consumer is expected to do in response to any pair of prices she faces. Recall that the reason we have to do this is because the firms calculated their best price offer by considering what they would get by deviating to other offers. Subgame perfection requires the firms to believe that when they deviate, the consumer will respond in a sensible way. What that means in this context is that firms should believe that consumers will respond by buying only at the lowest price. When firms offer the same price, it doesn't matter which firm the consumer buys from. Nevertheless, to specify a strategy we have to specify exactly what consumers will do when prices are equal. We write this down formally as follows:

$$q_1 (p_1, p_2) = \begin{cases} Q (p_1) & p_1 < p_2 \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

For firm 2, the corresponding strategy is

$$q_2 (p_1, p_2) = \begin{cases} Q (p_2) & p_2 \leq p_1 \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

This relatively simple pair of expressions hides a very important restriction. The strategy specified above has the property that the consumer will buy everything from firm 2 as long as firm 2's price is *less than or equal* to firm 1's price. There are a continuum of other strategies we could have used, all of which would have specified different actions when the firms set the same price. Later you will see why the strategy specified above is the only one that supports a subgame perfect equilibrium.

Given the strategy specified above for consumers, we now want the firms to choose prices that are best replies. Start with firm 1 and suppose he is facing a price for firm 2 equal to $p_2$. There are two easy cases. If $p_2 > p_1^*$, then it should be apparent that firm 1 will want to respond by setting price equal to $p_1^*$. This has to be better than setting any price at or above firm 2's price, since 1 either won't get any demand at such a price, or will sell some amount that earns him less profit than he could get by charging what he would as a monopolist. Alternatively, if firm 2 has a price below $c_1$, then firm 1 can set any price that is no smaller than firm 2's price. This will ensure that he doesn't have to sell output at a price less than his marginal cost because of the way we specified the consumer's strategy in (1) and (2) above.

The funny stuff starts when firm 2 charges a price that is between $p_1^*$ and $c_1$. If that is the case, then firm 1 can sell some quantity at prices above his marginal cost and make a profit. In that case he won't want to set a price at or above $p_2$, for then he will get nothing. He needs to set a price below $p_2$ in order to make profitable sales. Since $p_2$ is below his monopoly price, he also wants to set his price as close to $p_1$ as possible. No matter what price below $p_2$ he sets, he can always get the price a little closer. In other words, firm 1 *doesn't have a best reply.*

This may seem a little confusing. Yet it is actually somewhat helpful. It means that we are never going to be able to find a subgame perfect equilibrium where firm 2 sets a price above $c_1$. Why? Well, a subgame perfect equilibrium is a Nash equilibrium, so both firms need to set prices that are best replies. Since we can never find a best reply for firm 1, we can rule out such an outcome.

Once firm 2's price falls to $c_1$, firm 1 has a lot of best replies - any price at or above firm 2's price will ensure that 1 doesn't make any sales. This is true even if firm 1 matches firm 2's price because of the way we specified the consumer's strategy in (1) and (2).

The outcome that is usually referred to as 'the' equilibrium in the Bertrand game is the one where both firms set the price and the consumer uses the strategy described by (1) and (2) above.

To understand why this is a Nash equilibrium you need to check unilateral deviations. First of all, you already know that the consumer's strategy is a best reply to the prices set by firms, because this strategy has the consumer doing her best no matter what the firms do. Firm 2 sets the price $c_1$ and sells $Q(c_1)$ units at a price which is strictly above his marginal cost. So he makes a profit. If he raises price, he loses all his profitable sales. So raising price is not a profitable deviation. If he lowers price he sells a bit more at a lower price. Since $c_1$ must also be below firm 2's monopoly price, this isn't profitable.

Firm 1 sells nothing in this equilibrium. If he raises price, he still sells nothing. If he lowers price, he will make sales, but at a price that is below his marginal cost. So firm 1 has no profitable deviation.

There are a bunch of other equilibrium outcomes like this one. Can you see why the same arguments can be used to support subgame perfect equilibrium with both firms charging any price between $c_1$ and $c_2$? Given the way we have specified the consumer's strategy, we can support any such price by having firm1 match firm 2's price. This keeps firm 2 from wanting to raise its price. Firm 1 is happy with this because he doesn't make any sales in equilibrium.

### 1.1 Problems to Think About

1. We could change the strategy that the consumer is using in a subgame perfect equilibrium by having her buy half of her demand from each firm when the two firms set the same price. Can you explain why it is impossible to construct a Subgame perfect equilibrium in which the consumer uses this strategy?

## 2 Directed Search

Though Bertrand equilibrium is useful in explaining how subgame perfection works in games where players have a continuum of strategies, it describes an equilibrium that just doesn't seem plausible. For example, both firms set the same price even though they have different costs. For most products there is a lot of price variability between firms that doesn't seem to be related to difference in product quality (Apple computers versus everything else for example - or Microsoft Office versus OpenOffice).

Directed search is one of the models that has been proposed to deal with this. It has some very useful characteristics, especially when applied to labor markets, where it can be used to explain unemployment. The search decisions that workers make in a directed search model lead to unemployment and unfilled vacancies even though workers can see all the firms' wages when they make their search decisions. The reason it is called directed search is that workers decisions about where to apply are guided or directed by wages. This distinguishes the model from a much bigger literature that assumes that workers and firms are matched randomly (by chance).

To see how it works, suppose that as in the story above, there are two firms. Instead of setting prices, though, they set wages. Each of the firms want to fill exactly one vacancy. Instead of a single consumer, we'll assume there are two workers who try to find jobs with these two firms. The way directed search approaches this is to assume that each of the workers makes an application to one and only one of the firms. Each of the firms then collects its applications and hires one of the workers who applied. If two workers apply to the same firm, the firm chooses one of them randomly and offers her the job. If only one worker applies, the firm just offers the job to that worker.

To keep the story simple, we will assume that firms who get no applications are just out of luck, as are workers who apply but aren't offered a job. Otherwise, we'll assume that firm 1 earns gross profit $y_1$ if it fills its vacancy, while firm 2 earns $y_2 < y_1$. Workers' payoffs are just the wages they earn.

What gives the directed search model nice properties is the assumption that workers use a *symmetric* application strategy. What that means is that each of the workers applies to firm 1 with the same probability $\pi$. This is something you are familiar with - a mixed strategy equilibrium. The extra part here is that firms will have to figure out how changes in their wages are going to affect the mixed strategy equilibrium for the workers' application game. Since workers' mixed strategies are going to mean that some worker don't get jobs, there is going to be some unemployment and some unfilled vacancies. Firms can limit the probability with which they have unfilled vacancies by raising wages, since that will increase the probability with which workers apply. We'll work out the wages that firms offer in a subgame perfect equilibrium. The outcome won't look anything like the equilibrium of the Bertrand pricing game that we studied above.

Lets do backward induction and try to figure out what the workers will do for every pair of offers by the firms. Call the wages of the two firms $w_1$ and $w_2$ for firm 1 and 2 respectively. Strictly speaking, we should model what the firms do once they receive applications, but will skip that for brevity and just assume as above that the firms mechanically select each worker with probability $\frac{1}{2}$ when it has two applications.

The normal form of the application game played among the workers now looks like the following:

| | Firm 1 | Firm 2 |
|---|---|---|
| Firm 1 | $\frac{w_1}{2}, \frac{w_1}{2}$ | $w_1, w_2$ |
| Firm 2 | $w_2, w_1$ | $\frac{w_2}{2}, \frac{w_2}{2}$ |

To understand the payoffs, just observe that if both workers apply to the same firm, they each get the job with probability $\frac{1}{2}$. That makes an expected payoff equal to $\frac{w_1}{2}$ for both of them when they both apply to firm 1.

If the other worker is applying to firm 1 with probability $\pi$, then the expected payoff to the worker if he applies there is

$$\pi \frac{w_1}{2} + (1 - \pi) w_1.$$

The explanation is that if the other worker also applies to firm 1, then there is half a chance that the worker will be hired. If the other worker applies to firm 2, then the worker is hired for sure.

Using the same reasoning to compute the expected payoff associated with an application to firm 2, the probability with which the worker expects the other worker to apply to firm 1 had better satisfy

$$\pi \frac{w_1}{2} + (1 - \pi) w_1 = \pi w_2 + (1 - \pi) \frac{w_2}{2} \tag{3}$$

or

$$\pi = \frac{2w_1 - w_2}{w_1 + w_2}.$$

Now you should recognize that in order to describe a subgame perfect equilibrium, you need to specify how workers will react to *all* pairs of wages, not just to those you think are important. In the expression above some weird stuff can happen when wages get too far apart. First, if $w_2 > 2w_1$, the solution to the equation above is negative, so something is wrong. In this case, think "one of the actions has become dominated". If you look back at the payoff matrix you can see which one - $w_2$ is so high that the worker would rather go to firm 2 than firm 1 even if he were sure that the other worker were going to apply to firm 2.

Another way to look at it is that in order to satisfy (3) it isn't enough just to maximize the probability with which the other worker applies at the same firm, you have to go even further and change the weight assigned to the good outcome so that the payoff turns negative.

Exactly the same thing occurs when $w_1 > 2w_2$ (so that the solution to (3) is greater than 1). Then applying at firm 2 is a dominated strategy.

What this algebra tells us is that the only symmetric subgame perfect equilibrium strategy looks like this:

$$\pi(w_1, w_2) = \begin{cases} \frac{2w_1 - w_2}{w_1 + w_2} & \frac{w_1}{2} \le w_2 \le 2w_1 \\ 1 & w_2 < \frac{w_1}{2} \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

This last formula says that firm 1 should expect that varying its wage will change the probability with a worker applies in the Nash equilibrium of the workers' application game. How exactly? Well, you can read this from the formula - using the quotient rule

$$\frac{d\pi(w_1, w_2)}{dw_1} =$$

$$\frac{d}{dw_1}\frac{2w_1 - w_2}{w_1 + w_2} = \frac{(w_1 + w_2)2 - (2w_1 - w_2)}{(w_1 + w_2)^2} = \frac{3w_2}{(w_1 + w_2)^2} > 0 \tag{5}$$

provided $\frac{w_1}{2} \le w_2 \le 2w_1$. Otherwise this derivative is zero.

If you write down what firm 1's expected profit is you get

$$\left(1 - (1 - \pi)^2\right)(y_1 - w_1).$$

The logic is that firm 1 is going to fill its vacancy provided at least one of the workers applies. The probability that neither of them applies is $(1 - \pi)^2$ - which gives the formula.

At this stage, lets make a guess about what equilibrium is going to look like. First, notice that there is no point for firm 1 to offer a wage more than twice $w_2$ or less than half $w_2$. In the first case, he would get applications from both workers for sure, and would still get these applications if he cut his wage a bit. In the latter case, he wouldn't get any applications at all, so he wouldn't make any profit. This means that in any subgame perfect equilibrium, the wages of the firms are going to be close enough together that the application probability will be determined by the solution to (3). Given this, it isn't too hard to see how firm 1 would choose its wage? Maximize the firm's profit by choosing the wage that makes the derivative of this profit function 0. That is, find $w_1$ by solving the equation

$$2(y_1 - w_1)\frac{d\pi}{dw_1}(1 - \pi) = \left(1 - (1 - \pi)^2\right) \tag{6}$$

where $\pi$ is given by (4) and $\frac{d\pi}{dw_1}$ is given by (5). The solution to this equation gives the best reply function for firm 1.

At this point, even the computer algebra programs are going to fail to find solutions for you, though you could try for numerical solutions. The literature on directed search has handled this by developing models with a continuum of workers and firms and using these to approximate large labor markets. Instead of studying those, lets just look at a special case that is analytically tractable (though a little too special to be of much practical use). If the profits the firms make are the same (lets say $y_1 = y_2 = y$=1), then it seems plausible that both firms would set the same wage in a subgame perfect equilibrium. If they did, the probability with which each worker would apply to them would be $\frac{1}{2}$. Here is a way to find it.

First you could write out the derivative of the profit function that applies when firms choose wages $w_1$and $w_2$ (I used a computer algebra program to find this)

$$-2\left(w_1-1\right)\left(\frac{2w_1-w_2}{w_1+w_2}-1\right)\left(\frac{2w_1-w_2}{\left(w_1+w_2\right)^2}-\frac{2}{w_1+w_2}\right)+\left(\frac{2w_1-w_2}{w_1+w_2}-1\right)^2-1$$

Now what we do is to set this derivative to 0 and solve for $w_1$. Again using computer algebra, this solution is

$$w_1 = \frac{w_2^2+4w_2}{5w_2+2} \tag{7}$$

I'll remark again, that the reason this is so simple is because I have assumed $y_1 = y_2 = 1$.

Finally, we want to find the full Nash equilibrium for this game. In other words, we want each of the two firms to set a wage that is a best reply to the other firm's wage. The two firms have exactly the same payoff function in this example, which means that both of them should want to set the same wage in equilibrium. Since the best replies are described by equation (7), you should see that we need to find a wage $w^*$ for both of them to use that is a best reply to itself, or

$$w^* = \frac{w^*\left(w^*+4\right))}{5w^*+2}$$

which has solution $\frac{1}{2}$.

The equilibrium in which the firms are identical and set the same wage isn't completely satisfactory. It hides one of the main predictions of directed search, i.e., it is harder to get a job at a higher wage. Of course, everyone knows that it is harder to get a job at a high wage firm. The advantage of the model is that you now have a precise explanation of what 'harder to get a job' means. This makes it possible to turn our intuition into formal predictions that we can check by looking at data.

**Problems:**

1. In the example above, $\frac{1}{2}$ is a Nash equilibrium. Is there another one? We didn't actually check second order conditions in this exercize. Do both or either of the potential candiates for equilibrium satisfy the second order condition?

2. Find the application probability $\pi$ to be used by both workers that maximizes the expected number of matches (which is the same as the expected level of employment). Do you see any connection with the Price of Anarchy theorem?

3. Find the application probability $\pi$ to be used by both workers that maximizes expected revenues of firms.

4. Assuming that all firms offer the same wage, write out the Nash equilibrium application probabilities for workers when there are three firms and two workers.

# 3 Directed Search with incomplete information.

When two workers apply to the same firm in the story above, one of them is chosen randomly and given the job. The unlucky worker presumably goes back to the market and tries again with another firm. This has a couple of implications that don't seem very plausible. First, the wage that a worker receives doesn't say much about the worker. If there were a distribution of wages available, then the workers who get jobs at the high wage firms are just lucky. As a result, their wages shouldn't be correlated at they move between jobs. Second, since workers just repeat their application behavior in each period that they are unemployed, the probability that they will get a job in any period shouldn't depend on the length of time they are unemployed. Neither of these predictions is right.

To deal with this an alternative model assumes that workers have different types that the firm can see when they apply and the firm interviews them. When two workers apply at the same firm, the firm just hires the best one.[1] A lottery is used to pick a worker only when they have the same type. What makes the model run is that workers don't know how good or bad their competitors are.

To see how this one works, let the types be $h$ and $l$. Suppose that the types of the two workers are independently drawn, and that each worker is believed to have type $h$ with probability $\lambda$. The payoff matrix faced by a worker then depends on his or her type. The matrix for a type $h$ worker looks like this

| | Firm 1 | Firm 2 |
|---|---|---|
| Firm 1 | $\lambda\frac{w_1}{2} + (1-\lambda)\,w_1$ | $w_1$ |
| Firm 2 | $w_2$ | $\lambda\frac{w_2}{2} + (1-\lambda)\,w_2$ |

For a low type worker the matrix looks different:

| | Firm 1 | Firm 2 |
|---|---|---|
| Firm 1 | $(1-\lambda)\,\frac{w_1}{2}$ | $w_1$ |
| Firm 2 | $w_2$ | $(1-\lambda)\frac{w_2}{2}$ |

A reasonable conjecture would seem to be that the high type worker would surely go to the high wage firm (assume $w_1 > w_2$ in what follows). Whether this conjecture is reasonable or not depends on what the other high type worker is supposed to do, and how likely it is that the other worker is high type. If the other worker is expected to apply to the high type firm for sure, then the payoff to applying to the high type firm is

$$\lambda\frac{w_1}{2} + (1-\lambda)\,w_1.$$

By applying to the low type firm, the worker would then get the job for sure, and earn $w_2$. So the high type worker can be expected to apply for sure to the high type firm provided

$$\lambda\frac{w_1}{2} + (1-\lambda)\,w_1 > w_2.$$

[1] This is the two type version of the model in [?].

This is bound to be true if $\lambda$ (the probability the other worker is a high type) is low enough. So lets just assume this inequality holds in what follows.

What may not be so obvious is that the low type worker will also apply to the high wage firm if the probability the other worker has the high type isn't too large.

To see why, observe that the payoff to the low type worker from applying to the high wage firm is

$$(1-\lambda)\left(\frac{\pi}{2}+(1-\pi)\right)w_1 \tag{8}$$

where $\pi$ is the probability that the other worker applies to the high wage firm when he or she has a low type. The payoff from applying to the low wage firm (assuming again that the high type of the other worker applies only at the high wage) is

$$w_2\left(\lambda+(1-\lambda)\left(\frac{(1-\pi)}{2}+\pi\right)\right). \tag{9}$$

If the other worker doesn't apply to the high wage firm at all when he has a low type, then the payoffs are just $(1-\lambda)\,w_1$ and $\left(\lambda+\frac{(1-\lambda)}{2}\right)w_2$. If $\lambda<\frac{2w_1-w_2}{2w_1+w_2}$ then the low type worker will prefer to take his chances with the high wage firm unless the low type of the other worker also applies with some probability.

As before, we can find a fixed point by setting (8) equal to (9) and solving for $\pi$. The solution is

$$\pi^*=\frac{\lambda\,(w_2+2w_1)-(2w_1-w_2)}{-\,(1-\lambda)\,(w_2+w_1)}. \tag{10}$$

Notice that the condition that ensures that this is positive and less than one is the same as the condition above that determines when the low type worker will want to apply to the high type firm.

Once again, lets defer wondering why firms wages might differ and just take it for granted that they do. We can now do the same basic calculations what we did in the previous model to see if this change in the modeling procedure improves things at all.

It is now possible to do all kinds of probability calculations using this information. If you aren't used to doing this, it might help to look at the following picture:

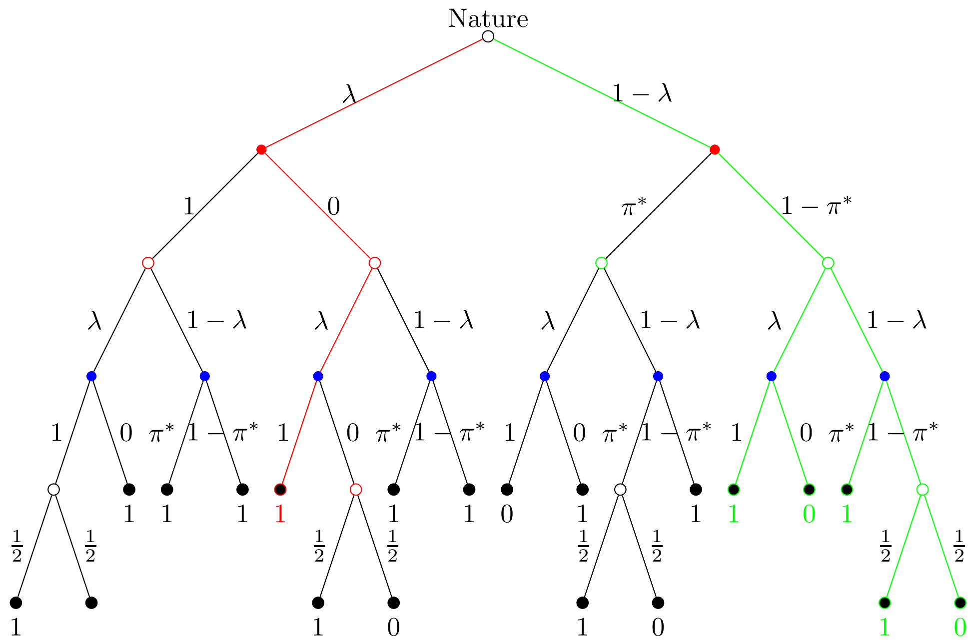


From an outsider's perspective there are many things happening in this model. First, one might say that nature determines whether each of the workers is high or low. In the figure above, the nodes that represent this are the open nodes. For example, at the very top of the tree there is an open node that represents nature choosing whether worker 1 will be a high type, which occurs with probability $\lambda$, or a low type, $1-\lambda$. The two edges or branches that extend below this node represent these different choices, and represent these choices, and are labeled with the appropriate probabilities.

A bit lower down the tree, there are some more open circles, for example where nature chooses whether worker 2 is high or low, or a couple at the very bottom where nature decides which of the two workers gets the job in case they both apply to the same firm. We'll get back to that momentarily.

At the end of the first two edges in the diagram, there are two red nodes. These are places where worker 1 makes a decision about where to apply - the edge that leads down to the left of these red nodes represents a decision to apply at firm 1, while the edge that leads down to the right means apply at firm 2. These edges are labeled with the probabilities with which the worker actually makes these choices. As we calculated above, the high type worker will apply to firm 1 (the high wage firm) with probability 1. So we labelled the two edges with a 1 and a 0 to indicate that this is what we believe will happen.

At the rightmost red node, we are thinking about the worker having a low type. As we calculated above, the low type worker should apply to firm 1 with probability $\pi^*$, so we label the edge leading down to the left with $\pi^*$ to indicate this, similarly for the edge leading to the right.

Then, as we move down the tree, we have four more open nodes, indicating nature's choice for worker 2's type. Finally, at the blue nodes, worker 2 makes a choice about whether to apply to firm 1 or firm 2.

If you follow the branches down through the tree, you will eventually end up at the black nodes, called terminal nodes. Each of these represents a complete history of play. For example, the path colored red has nature choosing high for worker 1, worker 1 choosing to apply to firm 2, nature making player 2 a high type as well, and worker 2 applying to firm 1.

At the very bottom of the tree, I labeled each of the terminal nodes with a 1 if worker 1 gets the job, and with a 0 otherwise. Since this path tells me everything that happened, I could label the terminal node with anything, for example, the payoffs of both players. Here we'll focus on a simpler thing and just keep track of whether or not worker 1 gets the job.

We don't know whether or not the red path will be followed. To find the probability that this particular history will occur we just multiply together the probabilities that are listed along the edges that make up the path. Again, along the red path we would calculate $\lambda * 0 * \lambda * 1$ to be the probability that this path will be followed (this probability is 0 here because a high type worker 1 would never apply to firm 2).

An *event* is a collection of paths. For example, the event that the worker is a low type and applies to firm2 is colored green in the diagram. To find the probability of an event, you find the probability of each of the paths in the event, then add them all together. For example, there is a total of 5 paths in the event where the worker is low and applies to firm 2. The probabilities of each of the paths, from left to right, are

$$(1-\lambda) * (1-\pi^*) * \lambda * 1$$

then

$$(1-\lambda) * (1-\pi^*) * \lambda * 0$$

then

$$(1-\lambda) * (1-\pi^*) * \lambda * \pi^*$$

then

$$(1-\lambda) * (1-\pi^*) * \lambda * (1-\pi^*) * \frac{1}{2}$$

then

$$(1-\lambda) * (1-\pi^*) * \lambda * (1-\pi^*) * \frac{1}{2}.$$

Summing these last five lines gives the probability of the event.

We are almost ready to do some reasoning. We need one last concept - *Bayes Rule.* I copied the following diagram from oscarbonilla.com:

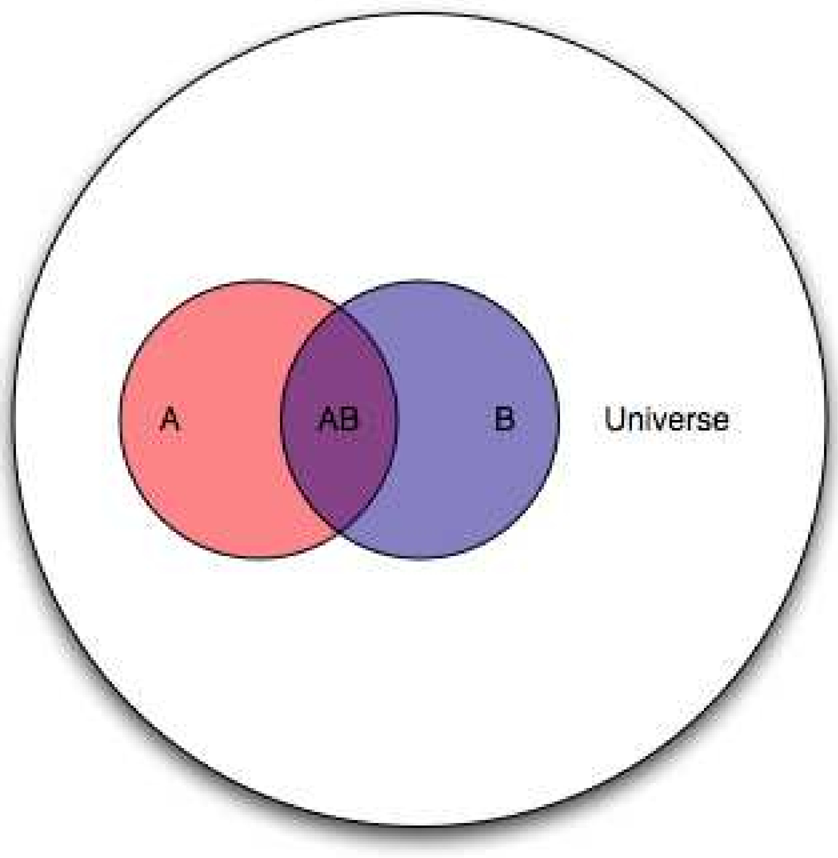


This is the standard way to describe conditional probability. Lets relate it back to the tree diagram that represents our equilibrium. The big outer (white) circle - labelled Universe - is one way to represent the collection of all the paths that appear in our diagram. All the paths together form an event. Since we have to follow one of the paths in the tree diagram, the event "something happens" occurs with probability 1. In other words, the Universe in the diagram above has probability 1.

The colored circles in the diagram represent other events. For example, the circle labelled $A$ in the figure might represent the event in which the first worker is a high type worker. That event would be all of the histories that follow the left branch out of the node marked 'Nature' at the very top. You might try to write down the sum of all the probabilities for those paths as we did above. If you did, you would find that each string of multiplied probabilities begins with $\lambda$. If you factored out the $\lambda$, you would have a sum inside your brackets which would be equal to one (make sure to verify that for yourself). Then we would say, the probability of the event in which the first worker is a high type is $\lambda$ and the area of the colored circle (it looks orange in my browser) marked $A$ in the figure above would be $\lambda$ to represent that.

A second event might be the one where worker 1 finds a job. In the tree diagram, this is the collection of all the paths that lead to terminal nodes that have the label 1 underneath them. Again we would multiply out the probabilities along each path and sum them up. The total probability we get from that sum could then be represented by the area of the blue circle labeled B above.

Notice that the circles overlap, but aren't the same. Thats because there are paths along which even a high type worker *doesn't* get a job, and paths along which a low type worker *does* get a job. Naturally, the area in the intersection of the two circles represents the event in which worker 1 is both a high type worker, and that he does get a job. We'd find that probability by collecting all the paths that start out down the leftward edge at the very top and also end at a terminal node that has a 1 beneath it.

So the probability of events can be found by summing up the probability of all the histories that are included in the event, where the event itself is defined by a bunch of things that have to occur. The event $\{A \cap B\}$ is just a smaller number of branches than in the event A or the event B. Now we can define *conditional probability.*

Formally, the probability of the event $A$ conditional on the event B is

$$\Pr\{A|B\} = \frac{\Pr\{A \cap B\}}{\Pr\{B\}}.$$

Since joint probabilities like $\Pr\{A \cap B\}$ can always be written as products of a conditional and marginal as in $\Pr\{A \cap B\} = \Pr\{B|A\}\Pr\{A\}$ we get the Bayes rule version

$$\Pr\{A|B\} = \frac{\Pr\{B|A\}\Pr\{A\}}{\Pr\{B\}}.$$

Sometimes it is easier to calculate conditional probabilities and multiply them by marginals than it is to find all the branches then sum. This is particularly true when you condition on a type after having found a Bayesian equilibrium. For example, suppose $A$ is the event "gets a job at the high wage firm". The event $B$ is "has a high type". Since $\Pr\{A|B\}\Pr(B) = \Pr\{A \cap B\}$ and $\Pr\{B\} = \lambda$, $\Pr(A|B)$ is just the probability calculation you get by summing brances *after you have already gone down the first branch.* This calculation is quite intuitive and you have been doing it implicitly from the start. It is

$$\Pr(A|B) = Q^h = \frac{\lambda}{2} + (1-\lambda). \tag{11}$$

Notice something about this formula - we used the equilibrium that we found above to compute this probability because we wrote it as if a high type worker applies to the high wage firm for sure. The same computation for the low type gives

$$Q^l = \pi(1-\lambda)\left(\frac{\pi}{2} + (1-\pi)\right).$$

Now, we can find the probabilty the worker has a high type and gets a job at the high wage firm, or the worker is a low type and gets a job at the high wage firm. This would be given by

$$\lambda Q^h + (1-\lambda) Q_l.$$

As before remember that $Q^h$ is computed assuming that high type workers apply for sure to the high wage, while low type workers apply to the high wage firm with a probability given by (10). The formula for $\pi^*$ depends on wages and $\lambda$.

Before we go on observe something about this whole process. Our model is a game which is defined by the two wage offers which we can see, and $\lambda$, which we can't. Suppose we play this game over and over (in a lab for example), and observe that after many plays of the game the proportion of workers who get jobs at the high wage firm is $\hat{Q}$. Assuming $\lambda$ is the same in all of them, then we should have

$$\lambda Q^h + (1-\lambda) Q_l = \hat{Q}.$$

Since the left hand side of this equation involves $\lambda$, which we can't see, while the right hand side is just a proportion that we measure in our experimental data, we could actually solve this equation to find a particular value$\hat{\lambda}$. This process is called *identifying* $\lambda$ *from the data.* The solution $\hat{\lambda}$ would be called our *estimate* of $\lambda$.

Now (as analysts) we are in business. Suppose we want to try a policy that will work if and only if the worker is a high type. For example, suppose we want to spend money training the worker, and that the worker will only benefit from the training if they are sufficiently educated. We can now estimate the probability with which our policy will succeed.

Now conditional probability isn't as simple as it was when conditioning on type. Conditional on getting a job at the high wage firm, the probability with which the worker has a high type, by Bayes rule, is

$$\Pr\{\text{type is high}|\text{job at high wage}\} =$$

$$\frac{\Pr\{\text{worker gets a job at the high wage}|\text{type is high}\} * \Pr\{\text{type is high}\}}{\Pr\{\text{worker gets a high wage job}\}}$$

$$= \frac{\left(\frac{\lambda}{2} + \left(1-\hat{\lambda}\right)\right)\hat{\lambda}}{\hat{\lambda}\left(\frac{\hat{\lambda}}{2} + \left(1-\hat{\lambda}\right)\right) + \left(1-\hat{\lambda}\right)\left(\hat{\pi}\left(1-\hat{\lambda}\right)\left(\frac{\hat{\pi}}{2} + (1-\hat{\pi})\right)\right)}. \tag{12}$$

In this expression $\hat{\pi}$ is the expression given by substituting $\hat{\lambda}$ into (10). We have estimated that our policy will succeed with probability given by (12).

## Problems:

1. Find the equilibrium in the problem above when $y_1 = y_2$. How does it compare to the model presented in the first section above.

2. What is the symmetric application strategy that maximizes the expected number or matches? the expected revenue for both firms?

3. Suppose that $\lambda \frac{w_1}{2} + (1-\lambda)\, w_1 < w_2$. Write down the conditions that describe the Nash equilibrium for the workers' application game. Can you find the two application probabilities? Numerically?

4. Can you provide conditions under which $Q_l < Q_h$? As a hint, note that this will depend on $\lambda$, $w_1$ and $w_2$. To make this manageable suppose that $w_1 = \alpha w_2$ and let $w_2$ be some constant $w$ so that your answer depends only on $\lambda$ and $\alpha$. Try to draw a diagram with $\alpha$ on one axis and $\lambda$ on the other, dividing the diagram into regions where $Q_1 > Q_2$ (and remember that there are going to be some regions where the high type worker won't apply for sure to the high wage firm).

5. In the tree diagram, label all the paths in which the worker gets a job. Use Bayes Rule to calculate the probability that the worker has a high type conditional on him finding a job.

6. Use Bayes Rule to find the probabiity that a worker gets a job at each of the two different wages conditional on the two events finding a job and being a high type worker, then finding a job and being a low type worker. Use these conditional probability distributions to calculate the expected wage of low and high type workers conditional on finding jobs.

# 4 A continuous version

The idea that players can have either a good or bad type is actually pretty awkward. It works for something like a drug test, which an applicant can either pass or fail. Yet firms are interested in much more. As we tried to find the equilibrium we had to worry about things like whether there would be an equilibrium in which good workers would apply only at the high wage firm. Ultimately, the algebra we had to do became very complicated.

Another approach is to assume that workers 'types' are taken from a very large set - a good example here would be the set $[0, 1]$. It might seem that would make it much harder to find an equilibrium. Often that isn't the case.

What we'll do here is to assume that when workers apply, firms interview them and place them on a scale between 0 and 1. For example, they might test their math skill, their writing skill, and various personality attributes. Their score becomes their 'type' and the firm just offers the job to the worker with the highest type.[2]

Our second assumption will be that each worker knows his or her own type, but doesn't know the type of the other worker. For that reason, we need to allow for the possibility that each different type of worker will act differently. The way this is done is to imagine that each worker uses a *strategy rule* $\pi$ that specifies for each value of their type, the probability with which they will apply to the high wage firm. Formally we would write $\sigma : [0,1] \to [0,1]$. That just says that for each type $t$ in $[0,1]$ we specify a $\pi(t)$ which is the probability the worker applies to the high wage firm.

This is really no different from what we did in the previous section where the good worker used a different probability from the bad worker.

To write the payoffs, we have to specify what each worker believes about the type of the other worker. For this, lets assume that each worker's type is independently drawn from a fixed distribution $F$ whose support is $[0,1]$. The support is the smallest closed set that has probability 1.

As before, if the worker gets the job at the high wage firm, he or she earns $w_1$, similarly for firm 2. The worker gets the job if the other worker doesn't apply, or if the other worker has a lower type. To work out the probability of that event, we'll use the strategy rule $\pi$. Suppose the worker has type $t$. The probability the other worker has a higher type is $1 - F(t)$. Just because the worker has a higher type, that doesn't necessarily mean they'll actually apply to the wage $w_1$. The probability of that event is

$$\int_t^1 \pi(t)\, dF(t)$$

which means that the expected payoff the worker of type $t$ gets from applying to firm 1 is

$$w_1 \left(1 - \int_t^1 \pi(t)\, dF(t)\right)$$

A 'Bayesian' equilibrium of this game is a strategy rule that has the

[2] Many companies who are hiring programmers use online testing sites to check skills of potential employees. (search for online testing) Passing the online test then results in one or more interviews. In economics, a relatively new innovation is something called a 'pre-doc', in which students with masters or honors degrees can work as (paid) research assistants for professors in well known universities for a year as a way of showing their suitability for the ph'd program. Applications for pre-docs typically involve a test.

property that $\pi(t) >$ only if

$$w_1 \left(1 - \int_t^1 \pi\left(\tilde{t}\right) dF\left(\tilde{t}\right)\right) \geq w_2 \left(F(t) + \int_t^1 \pi\left(\tilde{t}\right) dF\left(\tilde{t}\right)\right)$$

I will show you that there is a Bayesian equilibrium in which each worker uses a very simple rule. This rule is

$$\pi(t) = \begin{cases} 1 & t \geq t^* \\ \overline{\pi} & \text{otherwise.} \end{cases}$$

Here $t^*$ and $\overline{\pi}$ are fixed constants between 0 and 1. I'll tell you exactly what they are shortly.

To show that these rules can be part of a Bayesian equilibrium, I have to show you that if a worker uses them, he or she will never be able to do better by deviating. To do this, I have to substitute these rules into the payoff functions I described above.

If worker 2 is expected to use this rule, then the payoff to worker 1 when he applies to firm 1 will be given by

$$\begin{cases} F(t)\, w_1 + \left(F(t^*) - F(t)\right)(1 - \overline{\pi})\, w_1 & t < t^* \\ F(t)\, w_1 & \text{otherwise.} \end{cases}$$

If, on the other hand, worker 1 applies to firm 2, the same calculation gives expected payoff

$$\begin{cases} F(t)\, w_2 + \left(F(t^*) - F(t)\right) \overline{\pi} w_2 + \left(1 - F(t^*)\right) w_2 & t < t^* \\ w_2 & \text{otherwise.} \end{cases}$$

Now choose $t^*$ to satisfy

$$F(t^*)\, w_1 = w_2.$$

Its easy to verify that when $t^*$ is chosen this way, both workers will strictly prefer to apply to firm 1 if their types are larger than $t^*$ (because they have a better chance of getting the job that a worker of type $t^*$does, and that worker is indifferent).

On the other hand, if a worker has a type below $t^*$ and he/she is supposed to randomize with a probability between 0 and 1, then as you have learned, he/she will only be willing to do so if the payoffs to the two choices are exactly the same, i.e.,

$$F(t)\,w_1 + \left(F(t^*) - F(t)\right)(1-\overline{\pi})\,w_1 =$$

$$F(t)\,w_2 + \left(F(t^*) - F(t)\right)\overline{\pi} w_2 + \left(1 - F(t^*)\right) w_2.$$

This equality has to be true uniformly for types less than $t^*$, which means that the derivatives of the functions with respect to $t$ have to be the same, i.e.,

$$w_1\left(f(t) - f(t)(1-\overline{\pi})\right) =$$

$$w_2\left(f(t) - \overline{\pi} f(t)\right).$$

You can see that the $f(t)$ cancels out, allowing you to solve for

$$\overline{\pi} = \frac{w_2}{w_1 + w_2}.$$

The payoffs from apply to both firms have the same slope when $\overline{\pi}$ is chosen this way. As $t \to t^*$, the payoff at firm 1 converges to $w_1 F(t^*)$ while the payoff at firm 2 converges to $w_2$, and these two payoffs are equal by the choice of $t^*$. So the payoffs must be equal to each other uniformly in $t$.

This means that, as promised, workers whose types are below $t^*$ are indifferent, so that the simple strategy is always a best reply for them, as it is for workers whose types are above $t^*$ who strictly prefer to apply to firm 1 as the strategy specifies. Since no worker of either type can find a profitable deviation, our rule is a Bayesian equilibrium strategy rule.

$$(y_1 - w_1)\left(1 - \left(F(t^*)(1-\overline{\pi})\right)^2\right)$$