Directed Search

Michael Peters

October 23, 2017

One of the problems with a standard Arrow Debreu market clearing approach in practice that that market outcomes reveal a lot of heterogeneity. The same good trades at many different prices, brands proliferate, and, perhaps worst of all, markets don't seem to clear - the labor market being a classic example. Search theory has provided one class of models that help to resolve some of these problems. Directed search is one of the models that has been proposed to deal with this.[1] 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.

To see how it works, suppose there are two firms who set wage in an attempt to maximize expected profits. Each of the firms want to fill exactly one vacancy. 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 apply. If two workers apply to the same firm, the firm flips a coin and chooses each of them with equal probability 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,

[1] The are also random matching models. ((Burdett and Mortensen 1998), (Mortensen 1986)

while firm 2 earns $y_2 < y_1$. Workers' payoffs are just the wages they earn. Each worker is assumed to be risk neutral, and to maximize their expected wage.

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 π. 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.

Lets start by figuring out how workers apply. Call the wages of the two firms w_1 and w_2 for firm 1 and 2 respectively.

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 π, 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

$$\pi\frac{w_1}{2} + (1-\pi)\,w_1 = \pi w_2 + (1-\pi)\,\frac{w_2}{2} \quad (1)$$

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 (1) 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 (1) 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:

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

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

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

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

provided $\frac{w_1}{2} \leq w_2 \leq 2w_1$. 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 (1). 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

$$2\left(y_1 - w_1\right)\frac{d\pi}{dw_1}\left(1-\pi\right) = \left(1-\left(1-\pi\right)^2\right) \tag{4}$$

where π is given by (2) and $\frac{d\pi}{dw_1}$ is given by (3). 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 $y_1 = y_2 = y$), 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 $\frac{1}{2}$. Whatever common wage they set should satisfy the first order condition (4) for both firms. This first order condition would then simplify to

$$2\left(y_1 - w\right)\frac{1}{2}\frac{3}{4w} = \frac{3}{4}$$

This has a very simple solution $w = \frac{y}{2}$.

What this means is that there is only one wage, $\frac{y}{2}$ that could satisfy the first order condition that is required to hold when firms wages are best replies to one another. This means either that the wages both firms offer in the subgame perfect equilibrium are equal to $\frac{y}{2}$, or that there is no pure strategy equilibrium at all. There are some second order conditions that need to hold. Lets leave those for some future headache and just take it for granted that the subgame perfect equilibrium has both firms offering the wage $\frac{y}{2}$ while the workers both use the strategy described in (2). [2]

Though this seems very simple minded, it has begun to lay the framework for an empirically rich model. Generally speaking, if the firms have different revenues, i.e., $y_1 \neq y_2$, then they will set different wages. So this model will ultimately be consistent with wage dispersion. Instead of worrying about how exactly to explain the wage dispersion, lets take it for granted and make some observations that come from this.

First go back to the very first observation

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

Since π (the probability both workers apply to firm 1 - the high wage firm) is greater than $\frac{1}{2}$, workers are less likely to succeed when they apply to firm 1, and firm one is more likely to fill its vacancy. Imagine that they way this works is that when a worker fails to find a job, he or she simply returns to the market and tries again. Though this isn't plausible in this two firm example, it seems reasonable that in some steady state version of this theory, that which ever firm succeeded in hiring would be replaced by a clone (we'll talk about this more in the large game version), so that the worker would again face the same two wages. Since the probability of matching is just the inverse of the expected time it takes to find a job, we suddenly have a model that predicts duration of unemployment.

Workers who are suddenly displaced from their jobs come to this market then try the strategy π. Since there is nothing to distinguish the workers, the expected duration of their unemployment spells is independent of their employment history. In particular, workers who are laid off from high wage jobs will take the same amount of time to find a job as workers who are laid off low wage jobs. This is a prediction that seems potentially testable.

Secondly, observe that a worker might get lucky and find a job right away. Or, he or she might remain unemployed for a number of periods. If that latter event happens, then there is no reason for them not to use the Nash equilibrium strategy π. So the second prediction of this model is that

[2] If you want to see the calculus for the symmetric case worked out for many workers and firms, you can read (Burdett, Shi, and Wright 2001)

the expected duration of unemployment should be independent of how long the worker has already been unemployed.

You have no doubt heard of discouraged workers, so you know that expected duration tends to rise with duration of unemployment. (See for example (Machin and Manning 1999)). On the second point, here is a regression that shows the relationship between wages earned before and after a transition:

Table 1: Regression Results

Variables	(1) $w_{i,t}$	(2) $\mu^2_{i,t}$	(3) $w_{i,t}$
$w_{i,t-1}$	0.41***	-0.05***	0.36***
	(0.114)	(0.016)	(0.007)
$w^2_{i,t-1}$	-0.06		-0.04***
	(0.055)		(0.004)
Constant	0.54***	0.04***	0.57***
	(0.049)	(0.011)	(0.005)
Covariates	Yes	Yes	Yes
nob	61,572	61,572	61,572
R^2	0.854	0.050	0.853
	OLS	OLS	GLS

You will notice that there is a significant positive correlation between the wage at a job an executive leaves, and the wage he or she receives in their new job. This correlation is inconsistent with the basic directed search model we have just discussed. One might also observe that the correlation is far from perfect - if the current wage goes up by 10% the future wage only goes up by 4%. So workers who are employed at high wage firms will tend to be employed at high wage firms throughout their career. Nonetheless a significant fraction of workers will lose their high wage status once they move.

0.1 A Second Approach

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.[3] 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 λ. 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, w_2
Firm 2	w_2, w_1	$(1-\lambda) \frac{w_2}{2}$

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).

[3]This is the two type version of the model in (Peters 2010).

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

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

This is bound to be true if λ (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

$$(1-\lambda)\left(\frac{\pi}{2} + (1-\pi)\right) w_1 \tag{5}$$

where π 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

$$w_2\left(\lambda + (1-\lambda)\left(\frac{(1-\pi)}{2} + \pi\right)\right). \tag{6}$$

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 (5) equal to (6) and solving for π. The solution is

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

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.

First, note that the probability with which a high type worker matches in any period is

$$Q^h = \frac{\lambda}{2} + (1-\lambda)\,. \tag{7}$$

The only thing that might prevent him from getting a job is the possibility that he runs into another high type worker.

On the other hand, the probability that the low type worker gets a job is

$$Q^l = \lambda\,(1-\pi^*) + (1-\lambda)\left\{\pi^*\left(\frac{\pi^*}{2} + (1-\pi^*)\right) + (1-\pi^*)\left(\pi^* + \frac{(1-\pi^*)}{2}\right)\right\} \tag{8}$$

Notice that the term multiplying $(1-\lambda)$ in this expression must be less than 1. Provided that π^* is near $\frac{1}{2}$, this means that the probability with which the low type worker matches is smaller than that for a high type worker. This seems reasonable since the low type worker faced competition that the high type worker doesn't from the other low type worker. The low type worker also uses a 'riskier' application strategy since they deliberately try to compete against the high type worker as well.

Now when a worker appears in the unemployment queue, the worker has a high type with probability λ (at least this is true provided that job separations are independent of type). His or her expected duration of unemployment is then the reciprocal of (7). If the worker has been unemployed for a period already, then we can use Bayes rule to calculate the probability that the worker is a high type worker.

Bayes rule says that the probability that the worker has a high type conditional on observing that he was unemployed last period is equal to the probability that he is unemployed conditional on being a high type times the probability he has a high type divided by the probability that he is unemployed. The formula is then

$$\frac{\lambda\left(1-Q^h\right)}{\lambda\,(1-Q^h) + (1-\lambda)\,(1-Q^l)} < \lambda$$

because of the fact that $\left(1 - Q^l\right) > \left(1 - Q^h\right)$. So a worker who has been unemployed for a period has a longer expected duration of unemployment than a worker who has just entered the queue of unemployed workers.

A second straightforward calculation provides some additional insight. Suppose that at the end of a period, we see a worker employed at the high wage, and another worker employed at the low wage. At the end of the same period, a shock occurs and both workers are laid off after working for a single period. What wage do we expect each of the workers to get at their next job?

Suppose we don't have any other information on their work or unemployment histories. The model says that high wage workers don't apply at low wage firms. So the worker at the low wage job must be a low type. On the other hand, low type workers do apply at high wage jobs, so the high wage worker might also be a low type. At a coarse level, this suggests a couple of things. First, the low wage worker will tend to have highly variable wages, sometimes working at the low wage, sometimes, the high wage. The high type worker will have a very steady wage.

If you look at the second column of the table with regressions above, it shows a regression in which the residuals (squared) associated with the regression in the first column are regressed on the past wage. The residuals proxy the unpredictability of the future wage. The regression says that this unpredictability falls, the higher was the worker's wage at his last job. This is now consistent with the way the model above actually works. The impact of this past wage, though significant in a statistical sense is small in an economic sense. The reason, as just described, is that the past wage is a very poor signal of worker type.

Additional Reading:

1. (Virag 2011)
2. (Galenianos and Kircher 2012)
3. (Camera and Kim 2013)

References

Burdett, K., and D. Mortensen (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39(2), 257–273.

Burdett, K., S. Shi, and R. Wright (2001): "Pricing with Frictions," *Journal of Political Economy*, 109(5), 1060–1085.

Camera, G., and J. Kim (2013): "Uniqueness of Equilibrium in Directed Search Models," Discussion paper, Purdue University, to appear in Journal of Economic Theory.

Galenianos, M., and P. Kircher (2012): "On the game theoretic foundations of competitive search equilibrium," *International Economic Review*, 53, 1–21.

Machin, S., and A. Manning (1999): *The Causes and Consequences of Long Term Unemployment in Europe*Elsevier.

Mortensen, D. T. (1986): *Job Search and Labour Market Analysis*North Holland, Amsterdam.

Peters, M. (2010): "Unobservable Heterogeneity in Directed Search," *Econometrica*, 78(4), 1173–1200.

Virag, G. (2011): "High Profit Equilibria in Directed Search Models," *Games and Economic Behavior*, 71, 224–234.