LARGE GAMES - DIRECTED SEARCH

MICHAEL PETERS

One of the interesting properties of directed search is that with finite numbers, when a firm changes the wage it offers, it changes the outside option available to all the other workers. The reason is that when a firm raises its wage, all workers are more likely to apply. This increases the chance that any worker will be hired at all the other firms. It would seem reasonable to expect this effect would disappear as the number of workers and firms become large. This reasoning leads to a theory of directed search that that begins to explain wage distributions. For the purposes of this course, this also makes it possible to introduce one of the more useful ideas in economics, large games.[1]

A large game consists of a measurable set T of players, often taken to be an interval in $\mathbb{R}$, a measurable space of actions X and a measure F on T describing the size of the set of players in any measurable subset of T. The distribution of actions in a large game is a measure on $T \times X$ having the property that for every measurable subset $B \subset T$, $G(B, X) = F(B)$ (i.e, the marginal measure for Y associated with G coincides with F).

The payoff to a player of type t in a large game is given by some function

$$u_t(x_t, G).$$

A large game is a collection consisting of the type space T, the distribution F on T, the action space X, and the family of payoff functions u_t.

Define the best reply distribution in the following way: For each measurable subset B of X define the set

$$H(B) \equiv \{t \in T : \arg\max \{u_t(x', G)\} \cap B \neq \emptyset\}. \tag{0.1}$$

Definition 1. G is a Nash equilibrium distribution if for every measurable subset B of X, $F(H(B)) \geq G(T, B)$.

In words, what this condition says is that the measure of the set of players who have a best reply to the distribution G that lies in the set B is always at least as large as the measure of the set of players who are supposed to take actions in B under the distribution G.

Notice that the distribution $G(T, B)$ is the marginal measure for X evaluated on the subset B of X. This means that $G(T, B)$ is the measure of the set of players who choose actions in B. In equilibrium, each of the players who chooses an action in B should find that that action to be a best reply. This explains why every player who is expected to choose an action in B must be in $H(B)$. However, not all the players who can find best replies in B will necessarily choose actions in B in equilibrium, because they may also have best replies that lie outside of B. That is why equilibrium only requires $F(H(B)) \geq G(T, B)$.

[1]See (Ali-Khan and Y.Sun 2002) for the mathematics of large games.

An easier heuristic you might use to understand this stuff is to imagine that all players have unique best replies. In that case the definition would require that for every collection of actions B, the measure of players assigned to choose actions in B by the distribution G should be exactly equal to the measure of the set of players who have best replies in B. This is the reason for the heuristic short hand - a Nash equilibrium is a distribution of actions that is a best reply to itself.

Its pretty hard to see from this abstract definition how a large game might be a useful thing to think about. One advantage to this approach is that the marginal measure $G(T, B)$ is often something observable in a market. For example, in the directed search problem to be described below this marginal measure might be a wage distribution. If we can work out from equilibrium conditions what the joint distribution G looks like, it might be possible to recover the unobservable distribution of types F that most closely supports the observed distribution.

Return to the simple directed search game in which firms start by committing to wages, following which workers submit applications.. Assume there is a continuum of workers distributed uniformly on the interval $\left[0, \overline{k}\right]$. (The measure of the entire set of workers is $\overline{k}$). As it was described earlier, workers have types t in the interval $\left[0, \overline{k}\right]$. The exogenously given measure on this set is just the uniform measure on $\left[0, \overline{k}\right]$.

Firms constitute a second group of players. Suppose the set of firm types is distributed on the interval $Y \equiv \left[\underline{y}, \overline{y}\right]$ according to some differentiable distribution function F, where y is interpreted as the revenue a firm earns when it succeeds in hiring a worker. In other words, firm types are just given by their revenue. To make life easy, suppose that the measure of the set Y of firms is 1, and interpret $f(y)$ as the density of the set of firms.

As before, when a firm of type y offers a wage w and succeeds in hiring a worker, the firm gets payoff $y - w$ while the worker who is hired gets payoff w. Firms and workers maximize their expected payoff. Firms are all different - a fact we will use to get a wage distribution - while workers are all identical.

Now firms offer wages in some interval $W \equiv [\underline{w}, \overline{w}]$. Since the action and type spaces are are completely ordered sets we don't really have to think about measures, we can just use distribution functions. Let G be a distribution of wage offers. The *support* of G (the smallest closed set that has full measure under G) is denoted by $\overline{G}$. We need to try to turn this data into some kind of useful payoff function. There are a variety of ways to do this. One of the oldest suggestions is based on the following logic: Go back to the finite model and instead of just 2 workers and 2 firms, suppose there are n workers and m firms. Suppose that all the firms offer the same wage w, say, as they do in (Burdett, Shi, and Wright 2001). In response to this, each worker applies to each firm with probability $\frac{1}{n}$. Then the probability that a worker gets a job at that firm when he applies is given by

$$\sum_{j=0}^{n-1} \frac{(n-1)!}{j!\,(n-j-1)!} \left(\frac{1}{m}\right)^j \left(1 - \frac{1}{m}\right)^{n-1-j} \frac{1}{j+1} \tag{0.2}$$

$$= \frac{1 - \left[1 - \frac{1}{m}\right]^n}{n/m}$$

To understand this complicated formula start from the right, $\frac{1}{j+1}$ is the probability that our worker will be selected as the one to get the job when j of the other workers apply. The probability that exactly j of the other workers apply when there are m firms is given by $\left(\frac{1}{n}\right)^j\left(1-\frac{1}{n}\right)^{n-1-j}$ while the combinatorial formula $\frac{n!}{j!(n-j)!}$ is the number of different groups of j workers that can be selected from n of them. Let $k \equiv \frac{n}{m}$.

If you want to understand where the right hand side of this equation comes from, it uses the following transformation:

$$\sum_{j=0}^{n-1} \frac{(n-1)!}{j!\,(n-j-1)!}\left(\frac{1}{m}\right)^j\left(1-\frac{1}{m}\right)^{n-1-j}\frac{1}{j+1} =$$

$$\frac{m}{n}\sum_{j=0}^{n-1} \frac{n!}{(j+1)!\,(n-(j+1))!}\left(\frac{1}{m}\right)^{j+1}\left(1-\frac{1}{m}\right)^{n-(1+j)} =$$

$$\frac{m}{n}\sum_{j=1}^{n} \frac{n!}{j!\,(n-j)!}\left(\frac{1}{m}\right)^{j}\left(1-\frac{1}{m}\right)^{n-j} =$$

$$\frac{m}{n}\left\{1-\left(1-\frac{1}{m}\right)^n\right\}.$$

Imagine letting the number of workers and firms get large together in such a way the ratio of the number of workers to the number of firms is always k. It turns out that the expression (0.2) has a limit, and it is

$$\frac{1-e^{-k}}{k}.$$

The important part about this formula is that the probability depends only on the ratio of the number of buyers to sellers. Of course, this is an obvious consequence of the fact we took a limit. In a large market, we don't have numbers n and m, we have measures. What this formula implicitly says is that provided firms are identical, the probability of matching with one of them will be given by the formula above no matter what the measure is of the set of sellers.

If you frig around with the algebra enough, you can show that this expression is a decreasing function of k.[2] The important point is that the limit matching probability is a function of the ratio of the number of buyers to sellers. Doing the corresponding limit operation for firms gives the probability that they fill their position as $\left(1-e^{-k}\right)$.

These things have a very natural interpretation in the large game setting. Begin by creating a distribution - let $P(w)$ be the measure of the set of workers who make applications at wages at or below w. Since firms offer wages first, it is clearly necessary to require that the support of P be a subset of the support of G . The distributions of actions of the players are now given by the two distributions P and G. If $P(W) = k$ and $G(W) = 1$ then we can always write the distributions as marginals associated with the type distributions - the uniform on worker types and F on firm types. The limit formula above now supports some natural probabilities.

For example, suppose that the distribution function G has an *atom* at w of size dG. This means that w lies in the support of G and $\lim_{\tilde{w}\downarrow w} G(\tilde{w}) - \lim_{\tilde{w}\uparrow w} G(\tilde{w}) > 0$ (by convention, distribution functions are always taken to be right continuous). If P also has an atom at w of size dP, then the natural interpretation there will be

[2] The derivative can be written as $\frac{k+1}{e^k} - 1$. Just draw the graph of the two functions in the fraction to see that this is negative.

on average $\frac{dP}{dG}$ workers applying to each firm offering the wage w. Then using the reasoning above, the probability that a worker who applies at w gets a job at wage w should be

$$\frac{1-e^{-\frac{dP}{dG}}}{\frac{dP}{dG}}$$

while the probability that a firm offering w fills its vacancy is

$$1-e^{-\frac{dP}{dG}}.$$

Should both G and P be differentiable at w, then we might as well use the same interpretation, and imagine that the matching probability should be computed using the ratio of the densities. Atoms divided by densities should be interpreted as $\frac{dP}{dG}=\infty$ (yielding matching probabilities 0 and 1 for workers and firms respectively). Densities over atoms are interpreted as $\frac{dP}{dG}=0$ with corresponding probabilities.

The payoff to a worker who applies to a any wage w in the support of G is then given by

$$u_i\left(w,G,P\right)=\frac{1-e^{-\frac{dP(w)}{dG(w)}}}{\frac{dP(w)}{dG(w)}}w \tag{0.3}$$

while for firms, the corresponding payoff when a firm offers a wage in the support of the distribution G is

$$v_y\left(w,G,P\right)=1-e^{-\frac{dP(w)}{dG(w)}}\left(y-w\right).$$

Finally, there is the problem of how to assign payoffs to firms who offer wages outside the support of G. Compute

$$v^{\max}\left(P,G\right)=\max_{w\in\overline{G}}w\frac{1-e^{-\frac{dP(w)}{dG(w)}}}{\frac{dP(w)}{dG(w)}}.$$

This is the maximum payoff a worker can attain by applying to some wage in the support of the distribution G. Now for any firm who offers a wage w' outside $\overline{G}$, define the payoff to be

$$w'\left(1-e^{-k'}\right)$$

where k' satisfies

$$\frac{1-e^{-k'}}{e^{-k'}}w'=v^{\max}\left(P,G\right)$$

if this equation has a positive solution, and $k'=0$ otherwise.

The logic is simply that when some firm offers a wage that isn't offered by any other firms, then workers will apply up to the point where workers are indifferent between this new wage, and the wage offered by all the other firms.

Now, in order to find a Nash equilibrium, we need to find distributions P and G that support underlying joint distributions that are best replies to themselves. To start getting to this, notice that the probability term in the expression (0.3) could be written

$$Q\left(w,P,G\right)=\frac{1-e^{-\frac{dP(w)}{dG(w)}}}{\frac{dP(w)}{dG(w)}}.$$

It would seem sensible to expect that the probability of finding a job would be a decreasing function of the wage, at least using the logic from the 2 firm 2 worker example we discussed before. However, at this point, this function is just something that is derived from the distribution functions G and F. Here is a picture illustrating what this tradeoff might look like for an arbitrary pair of distributions G and P.

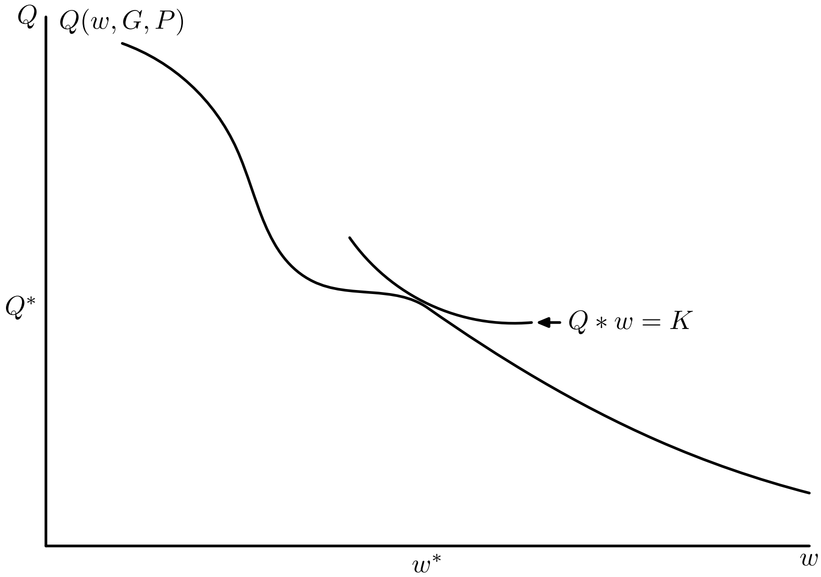

Implicit in this diagram is the assumption that the density of both P and G are strictly positive over the entire interval where $Q(w, G, P)$ is defined. However, all the workers have the same preferences. We could draw families of iso expected payoff curves for workers. The highest expected payoff workers can attain is indicated by the point (w^*, Q^*) where some iso-expected utility curve is tangent to the function $Q(w, G, P)$. The implication is that *all workers* would apply at wage w^* given G and P. In words, $P(w) = 0$ for any $w < w^*$. In terms of our original definition of Nash equilibrium, for any interval of wages $[\underline{w}, w]$ for which $w < w^*$, the measure of the set of workers who have best replies in the interval is zero. On the other hand, since the density $p(w') > 0$ for each wage below w^*, the distribution of applications assigns a strictly positive measure of applicants to this interval.

This picture makes it pretty clear that if we want to support a distribution of wages in a Nash equilibrium, then $Q(w, G, P)$ is going to have to coincide with some iso-expected utility curve for workers. For the moment, let V be the expected utility level associated with curve. Define the function $k(w)$ such that

$$\frac{1 - e^{-k(w)}}{k(w)} w = V. \tag{0.4}$$

To see how to fit firms into all this, we just change the label on the diagram above from Q to k, and draw the function $k(w)$ defined by (0.4) above. The new picture looks like this:

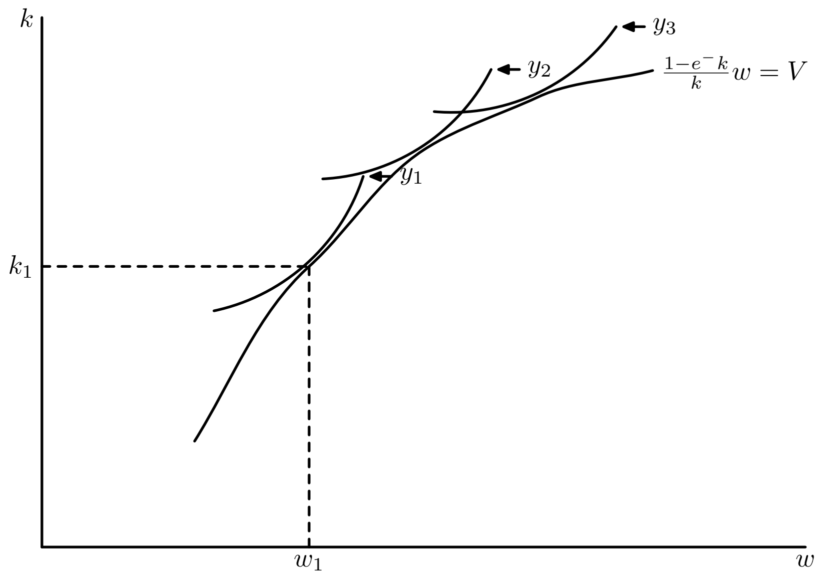

What this picture means is that if we expect to have an equilibrium in which different firms offer different wages, then the wages the firms offer are going to have to lie somewhere along a curve that looks like the one above.

A firm of type y earns expected payoff

$$(y - w)\left(1 - e^{-k}\right).$$

If we tried to draw iso-expected profit curves for firms, they would be convex shaped, and at any point, the slope of such a curve would become flatter as y increases. Since $k(w)$ is concave, this means that the higher a firm's type, the higher the wage it will choose. We are now ready to construct a Nash equilibrium.

To do it, fix v then for each firm type y find the point where the iso-profit curve for firm y is just tangent to the curve above. The picture above shows what the picture looks like with the tangencies drawn in for a couple of firms. There are a continuum of different firms like this - the tangency for each of them is at a slightly different point. When we match firms with revenue y_1 with the point (w_1, k_1) lying on the worker's market payoff line, we know that the density of firms with revenue y_1 is given by $f(y_1)$. The way the curve is constructed, we then need to have the density of workers applying at wage w_1 to be equal to $k_1 f(y_1)$.

This gives a method to find the fixed point - start with a value of v, then construct the market payoff line and locate each firm along this line using the tangency. That is, define

$$k^*(y, v) = \arg\max_k \left\{(y - w)\left(1 - e^{-k}\right) : \frac{1 - e^{-k}}{k} w = v\right\}$$

In words, what this expression gives is the point on the market value line where a firm of type y finds its tangency, For example, in the picture above $k^*(y_1, V) = k_1$. Now integrate across all the values of y as follows

$$\int_{\underline{y}}^{\overline{y}} k^*(y, v) f(y)\, dy.$$

This expression gives the total measure of workers we would need to ensure that every worker was indifferent between all the wages firms offer, and every firms' wage puts it at a tangency point on the market value line. We are going to make the argument that we can define a Nash equilibrium provided that we can find a value for v so that this integral is equal $\overline{k}$.

Theorem 2. *If the equation*

$$\int_{\underline{y}}^{\overline{y}} k^*(y,v)\, f(y)\, dy = \overline{k} \tag{0.5}$$

has a solution, then there is a Nash equilibrium for the large directed search game.

Proof. Let v^* be the solution to equation (0.5) above. Define the wage distribution as follows: for each $y \in \left[\underline{y}, \overline{y}\right]$, let

$$w^*(y) = \left\{ w : \frac{1 - e^{-k^*(y,v^*)}}{k^*(y,v^*)} w = v^* \right\}.$$

Notice that in order to prove that some distribution of wages and applications is part of a Nash equilibrium (see Definition 1), we don't actually use the joint measures on actions and types, we only need marginal measures on the set of actions. These marginal measure are, of course, the wage distribution G for firms and the application distribution P for workers. The distributions we want are given by

$$G^*(w) = F\{y : w^*(y) \le w\}$$

and

$$P^*(w) = \int_{\underline{y}}^{y:w^*(y)=w} k^*(y,v^*)\, dF(y)$$

(If you want to see the connection back to Definition 1, note that $G^*(w)$ is equivalent to $G^*(w,Y)$).)

For each interval $[\underline{w}, w]$ of wages, the set of firm types who have best replies in the interval is equal to the set of firms for whom $w^*(y) \le w$ (since the best reply function $w^*(\cdot)$ is monotonically increasing). The measure of this set is $F\{y : w^*(y) \le w\}$, and that is exactly the mass that has been assigned by $G^*(w)$. In words, the measure of the set of firms who are assumed to offer wages in any interval is exactly equal to the measure of firms who have best replies in that interval.

One the other hand, by the construction of each pair $w^*(y)$ and $k^*(y,v^*)$, every worker gets exactly the same expected payoff from applying at any of the wages in the support of G^*. It is then immediate that for any interval of wages, the measure of the set of workers who have best replies in the interval is equal to $\overline{k}$, the size of the entire set of workers. Since $\int_{\underline{y}}^{\overline{y}} k^*(y,v^*)\, f(y)\, dy = \overline{k}$ by assumption, the measure of the set of workers who apply in any interval is never larger than k and the distribution satisfies the requirement for a Nash equilibrium. Again, by construction, the measures $G(\overline{w}) = F(Y)$ while $P(\overline{w}) = \overline{k}$, so by Definition 1, G^* and P^* constitute a Nash equilibrium for the directed search game. □

To see how this kind of model might be used, suppose we start with a wage (offer) distribution G and we know the unemployment rate U in some labor market as well as the measure $\overline{k}$ of workers who are searching for jobs. Here, what I mean is that these objects are given to us in some kind of data set. The job is to try to figure out something about the underlying conditions in the market - for example, do firms differ in productivity. This problem is made somewhat more difficult by the fact that the distribution of applications P is data that probably isn't available. We seem to need that, along with the market value v^* to say something about the market.

So to begin, lets just guess the market payoff v and see what can be recovered. Suppose the support of the empirical wage distribution is $[w_1, w_2]$. From this we can construct the curve $k^*(w, v)$ on the interval $[w_1, w_2]$ using the equation

$$k^*(w,v) = \left\{ k : \frac{1-e^{-k}}{k} w = v \right\}.$$

For the sake of argument, suppose the estimated wage distribution G is differentiable. Then, we could recover the application distribution by solving the equation

$$dP(w) = \left\{ p : \frac{p}{dG(w)} = k^*(w,v) \right\}.$$

It is easy enough to show that this equation always has a unique solution. The recovered distribution of applications is $\int_{w_1}^{w} dP(w')$.

The next step is to try to compute what the unemployment rate would be if this recovered distribution were the correct distribution. Since $dP(w)$ is our current guess about how many applications are submitted at wage w we can calculate this. The probability that a worker fails to get the job when he or she applies at wage w, according to our current guess, is

$$\frac{k^*(w,v) - 1 + e^{-k^*(w,v)}}{k^*(w,v)},$$

the measure of unemployed workers at the end of the process would be

$$\int_{w_1}^{w_2} \frac{k^*(w,v) - 1 + e^{-k^*(w,v)}}{k^*(w,v)} dP(w).$$

We can actually recover v in that case by asking that this integral divided by $\overline{k}$ be equal to the empirically observed measure of workers in the market.

In words, we can find a market payoff v^* that has the properties that the recovered distribution P along with the curve $k^*(w, v^*)$ actually support the unemployment rate we observe in the market. Now we are in business. Choose some wage w in the support of the empirically observed wage distribution. This firm should have a revenue y that makes its iso-profit line tangent to $k^*(\cdot, v^*)$ at w. From the empirically observed wage distribution, we recover the fact that the density of firms with this revenue is $dG(w)$ and we have now recovered the distribution of firm types that best supports the data available.

Why would we want to know this distribution of firm types? The answer is that we can now change the large game by implementing some kind of new policy, for example, imposing a minimum wage wage w_1. Since we think we know the distribution of firm types, and we think the firms are playing our large game, we can then estimate the impact of the policy by finding the new equilibrium after the imposition of the minimum wage. This would involve solving for the equilibrium wage distribution and application distribution after imposition of the new policy.

References

Ali-Khan, M., and Y.Sun (2002): *Non-Cooperative Games with Many Players*Elsevier.
Burdett, K., S. Shi, and R. Wright (2001): "Pricing with Frictions," *Journal of Political Economy*, 109(5), 1060–1085.