Matching

Christopher Taber

Department of Economics
University of Wisconsin-Madison

November 18, 2015

Outline

1 Jovanavic

2 Nagypal

3 Neal

Jovanavic, *Journal of Political Economy*,1979

We use a simpler version of Jovanavic's model using discrete time and following Chapter 6 in Sargent and Ljungqvist's textbook (second edition, 2000) with my own notation.

- Worker's are matched with firm
- They draw a match distribution, θ
- They don't know its value exactly but learn it over time
- Must "experience" the job to learn the match value
- They are risk neutral so maximize expected wages

Can explain why

- Wages rise with tenure
- Separations fall with tenure
- Probability of a separation negatively correlated with wage

Simple Model with three stages

Period 0

- Worker starts at a new job
- V_0 is the value function of any job at this point
- Worker/firm match drawn $\theta \sim N\left(\mu, \sigma_0^2\right)$
- This match is not observed
- Firm offers wage contract and worker will accept some job (indifferent between all).
- After the first period the output

$$y_1 = \theta + u_1$$

is observed where $u_1 \sim N\left(0, \sigma_u^2\right)$.

Period 1

- Based on previous output, firm makes wage offer (dictated by contract)
- Worker decided whether to take it or not.
- At the end of this period the actual output θ is observed

Period 2

- Firm makes wage offer (dictated by contract)
- Worker decides whether to take it or not
- If they take it they keep it forever.

Also, If a worker chooses to switch jobs they must wait out a period before starting next one

Wages

In the model

- Firms offer the wages
  - μ in the first period
  - $E(\theta \mid y_1)$ in the second period
  - θ every year after that
- Jovanavic thinks of this process as an implicit contract that will have zero profits on average so it is an equilibrium in a competitive model

Solving Wage

To solve for the wage in the second period we will use the Kalman filter

After observing y_1 it is straight forward to show that the posterior distribution of θ is Normal with mean and variance

$$
\begin{aligned}
m_0 \equiv E(\theta \mid y_1) =& (1 - K_0)\mu + K_0 y_1 \\
\sigma_1^2 \equiv Var(\theta \mid y_1) =& K_0 \sigma_u^2
\end{aligned}
$$

where

$$K_0 \equiv \frac{\sigma_0^2}{\sigma_0^2 + \sigma_u^2}$$

σ_1^2 is between initial and last period

Next consider the distribution of m_0.

$$y_1 \sim N\left(\mu, \sigma_0^2 + \sigma_u^2\right)$$

so

$$\begin{aligned} m_0 &\sim N\left(\mu, K_0^2\left[\sigma_0^2 + \sigma_u^2\right]\right) \\ &\sim N\left(\mu, K_0\sigma_0^2\right) \end{aligned}$$

OK lets solve the model

Let

- V_0 be the value function of starting a new job
- $V_1(m_0)$ the value of continuing in the job given m_0
- $V_2(\theta)$ the value function of continuing in the job given θ

Solving backward if I decide to continue on the job given my θ there is not reason to change after that so letting β be the discount rate

$$\begin{aligned} V_2(\theta) =& \theta + \beta V_2(\theta) \\ =& \frac{\theta}{1-\beta} \\ V_1(m_0) =& m_0 + \beta E\left(\max\left\{V_2(\theta), \beta V_0\right\} \mid m_0\right) \\ V_0 =& \mu + \beta E \max\left\{V_1(m_0), \beta V_0\right\} \end{aligned}$$

Clearly there are cutoff points for θ and m_0

Define $\overline{\theta}$ so that

$$V_2\left(\overline{\theta}\right) = \beta V_0$$

if $\theta \geq \overline{\theta}$ we keep the job, otherwise we draw a new job

Similarly define $\overline{m}_0$ so that

$$V_1\left(\overline{m}_0\right) = \beta V_0$$

So the model is characterized by $\overline{m}_0$ and $\overline{\theta}$.

Lets try to understand this further.

By definition of $\overline{\theta}$

$$\frac{\overline{\theta}}{1-\beta} = \beta V_0$$

and by definition of $\overline{m}_0$

$$\overline{m}_0 + \beta \int_{-\infty}^{\overline{\theta}} \beta V_0 d\Phi\left(\theta; \overline{m}_0, \sigma_1^2\right) + \beta \int_{\overline{\theta}}^{\infty} \frac{\theta}{1-\beta} d\Phi\left(\theta; \overline{m}_0, K_0 \sigma_u^2\right)$$
$$= \beta V_0$$

Combining these and solving gives

$$\overline{\theta} - \overline{m}_0 = \int_{\overline{\theta}}^{\infty} \frac{\beta\left(\theta - \overline{\theta}\right)}{1-\beta} d\Phi\left(\theta; \overline{m}_0, K_0\sigma_u^2\right)$$

so one can see that $\overline{\theta} > \overline{m}_0$ because there is an option value of staying after the first period that is no longer there after the second period.

Now lets see if the model gives us the results

Do wages rise with tenure?

The key thing is we are conditioning on staying so we need to deal with the selection

$$\begin{aligned}\overline{w}_0 &= \mu \\ \overline{w}_1 &= E\left(m_0 \mid m_0 > \overline{m}_0\right) \\ \overline{w}_2 &= E\left(\theta \mid \theta > \overline{\theta}, m_0 > \overline{m}_0\right)\end{aligned}$$

As long as $\overline{m}_0 > -\infty$, since

$$\begin{aligned}\overline{w}_0 &= \mu \\ &= E\left(m_0\right) \\ &< E\left(m_0 \mid m_0 > \overline{m}_0\right) \\ &= \overline{w}_1\end{aligned}$$

Similarly using the law of iterated expectations

$$\begin{aligned}
\overline{w}_1 =& E\left(m_0 \mid m_0 > \overline{m}_0\right) \\
=& E\left(E\left(\theta \mid m_0\right) \mid m_0 > \overline{m}_0\right) \\
=& E\left(\theta \mid m_0 > \overline{m}_0\right) \\
<& E\left(\theta \mid \theta > \overline{\theta}, m_0 > \overline{m}_0\right) \\
=& \overline{w}_2
\end{aligned}$$

To see that separations fall with tenure note that in this simple model after the second period, turnover is zero.

(we can't sign turnover after period 1 from after period 2 because it depends on σ_u^2 but this is more an artifact of the fact that you learn everything at the end of the first period)

It is also clear that quits will be negatively correlated with wages since $\overline{w}_1 = m_0$ and

$$Pr\left(\theta < \overline{\theta} \mid m_0\right) = \Phi\left(\frac{\overline{\theta} - m_0}{\sigma_1}\right)$$

is declining in m_0

More Time Periods

Extending this to more time periods is straight forward. Assume that θ is revealed at time $T+1$ (sending $T \to \infty$) you can approximate an infinite time model.

Each period prior to $T+1$ the firm observes

$$y_t = \theta + u_t$$

where u_t is i.i.d. with

$$u_t \sim N\left(0, \sigma_u^2\right)$$

What makes this easy to solve is that with Normality we can use the Kalman filter so the only state variable is

$$m_t = E(\theta \mid y_0, ..., y_t)$$

because our posterior distribution of θ at time t is normal with mean m_t and variance σ_t^2 where

$$m_t = (1 - K_t)\, m_{t-1} + K_t y_t$$

$$K_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_u^2}$$

$$\sigma_{t+1}^2 = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_u^2}\sigma_u^2$$

and σ_0^2 is the unconditional variance.

The Bellman equations are as before

$$
\begin{aligned}
V_{T+1}(\theta) &= \theta + \beta V_{T+1}(\theta) \\
&= \frac{\theta}{1-\beta}
\end{aligned}
$$

$$
V_t(m_{t-1}) = m_{t-1} + \beta E\left(\max\left\{V_{t+1}(m_t), \beta V_0\right\} \mid m_{t-1}\right)
$$

All of the intuition we got from the simpler model goes through

- There are cutoffs $\overline{m}_t$ each period
- $\overline{m}_t$ is increasing with t
- Mean wages for the stayers are increasing with t
- Quitting declines with t
- There is a negative correlation between wages and subsequent quits

Outline

1 Jovanavic

2 Nagypal

3 Neal

Nagypal, *Review of Economic Studies*, 2007

Nagypal is worried about distinguishing returns to tenure from learning about the value of the match.

In Jovanavic we see that wages rise with seniority and the separation rate declines with seniority

Specific human capital will lead to exactly the same implications

Nagypal tries to sort these things out.

The key to identification is that the implications of a firm productivity shock are different for workers of different tenure,

First note that without any productivity shock, once the productivity of the worker has been revealed you will not fire him/her. Thus you only fire the newer guys.

However, when you are hit by a shock this will no longer be the case. You may well want to fire a worker who has been with you a long time, but has been revealed to be mediocre.

You want to hang on to the young guys because they still have high option value.

This is the basic intuition

Lets look at the details of the model

Environment:

- Continuum of ex-ante identical individual lived agents
- Continuum of firms
- Many workers per firm

Production

Match quality

$$\mu \sim N\left(\bar{\mu},\sigma_{\mu}^{2}\right)$$

and is unknown at the start of the match.

x_{τ} is worker productivity at tenure τ and

$$x_{\tau} \sim N\left(\mu, \sigma_{x}^{2}\right)$$

Speed of learning depends on σ_{μ}^{2} relative to σ_{x}^{2}.

She also allows for learning by doing based on Jovanovic and Nyarko so that output

$$q_t = x_\tau h\left(\varepsilon_\tau\right)$$

where

$$h\left(\varepsilon_\tau\right) = \prod_{i=1}^{N}\left(A - \varepsilon_{i\tau}^2\right)$$
$$\varepsilon_{i\tau} \sim N\left(0, \Sigma(\tau)\right)$$
$$\Sigma(\tau) = \frac{\sigma_\gamma^2\sigma_y^2}{\left(\tau-1\right)\sigma_\gamma^2+\sigma_y^2} + \sigma_y^2.$$

This means that

$$E_{\tau-1}\left[h\left(\varepsilon_t\right)\right] = \left(A - \Sigma(\tau)\right)^N$$

Expected output increases with tenure and is concave.

Macro shocks are driven by price differences.

The output is sold at price p_ℓ for $\ell = 1, ..., M$ with markov transition matrix Π.

A match may dissolve exogenously during any period with probability δ

Timing within each period:

1. Production of the good
2. Sale price, output, and ε_τ are observed
3. Match may dissolve exogenously
4. If not, worker decides whether to remain at firm (will stay if indifferent)

Evolution of beliefs

Everyone is going to have rational expectations

Thus everything will be updated based on Bayes rule

Preferences

Workers and firms are risk neutral with discount factor β

Hiring

Firm n has vacancies v_{nt} at the end of period t

There are u_t unemployed workers

Matching function is just

$$m_t = min(v_t, u_t)$$

Vacancies are costly for two reasons:

- Pay c_0 per open vacancy
- Pay $c(e_{nt})$ for hiring e_{nt} new workers

Equilibrium

- Agents in period t in existing matches make continuation decisions to maximize the surplus of the relationship
- Agents have rational expectations
- Firms choose vacancies to maximize discounted sum of revenue
- The distribution of workers across price and belief states at the end of the period and the state of unemployment is consistent with the optimal decisions of the agents in the model and is constant

Separation Decisions

Now we can figure out the separation decision.

It is assumed that separations are efficient so we can write down the Bellman equation for the joint decision of the worker and the firms

$$
\begin{aligned}
W(p_t, \tilde{u}_t, \tau) =& max\{U + V, \\
& \sum_{j=1}^{M} \pi(p_j \mid p_t)[p_j \tilde{\mu}_\tau \left(A - \Sigma(\tau)\right)^N \\
& + \beta(\delta(U + V) + (1 - \delta)E_\tau W(p_j, \tilde{\mu}_{\tau+1}, \tau + 1))]\}
\end{aligned}
$$

where $\tilde{\mu}_\tau$ is the posterior belief about the match and $\tilde{\mu}_\tau \to \mu$ as $\tau \to \infty$.

Then asymptotically

$$
\begin{aligned}
W(p_t, \mu) =& max\{U + V, \\
& \sum_{j=1}^{M} \pi(p_j \mid p_t)[p_j\mu \left(A - \sigma_y^2\right)^N \\
& + \beta(\delta(U + V) + (1-\delta)E_\tau W(p_j, \mu))]\}
\end{aligned}
$$

Finally we need to worry about the vacancy behavior of firms

Firm n chooses vacancy level v_n to maximize

$$\sum_{e_n=0}^{v_n} \left(\begin{array}{c} v_n \\ e_n \end{array} \right) \lambda^{e_n} \left(1-\lambda\right)^{v_n-e_n} \left[e_n(W(p_t\bar{u},0)-U-V)-c(e_n)\right]-c_0 v_n]$$

where e_n is the number of new workers and λ is the probability that a particular vacancy is filled so that

$$\lambda = min\left(\frac{u}{v},1\right).$$

She will assume that parameters are such that $\lambda = 1$.

She finally shows that

$$\begin{aligned} V &= 0 \\ U &= \frac{w}{1-\beta} \end{aligned}$$

the first comes from the fact that competition bid this to zero.

The second because firms have all the bargaining power so the wage is constant (at w).

Thats the model.

First she simulates it to show the difference between learning by doing and learning about the value of the match.

She assumes only two different values of the prices.

The base set of parameters are:

TABLE 1

Parameter values used in representative simulations

Parameter	Interpretation	Value
A	No-error output on a task	1·0
N	Number of tasks	5·0
ρ	Persistence of price shocks	0·95
β	Discount factor	0·99
$\bar{\mu}$	Average match quality	1·0
p_{l}	Low price	1·0
p_{h}	High price	2·0
δ	Exogenous separation rate	0·003

She then considers two cases

Case 1: only learning by doing

- $\sigma_\mu = \sigma_\gamma = \sigma_y = 0.4$
- $N = 5$
- There is still variation in the match component, but it is observed instantly

Here is what the cutoffs look like in the two states of the world and how they vary with tenure.

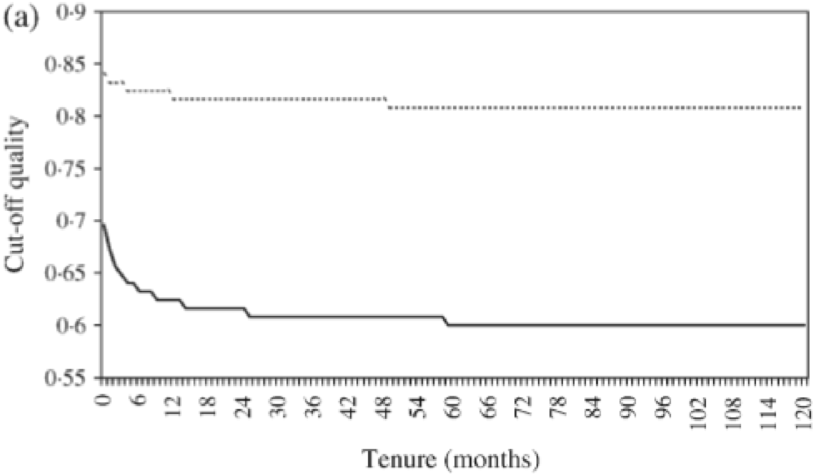
(a)
Cut-off quality
0·9
0·85
0·8
0·75
0·7
0·65
0·6
0·55
0
6
12
18
24
30
36
42
48
54
60
66
72
78
84
90
96
102
108
114
120
Tenure (months)

Case 2: only Learning about the value of the match

- $q_t = x_t$
- $\sigma_\mu = 0.4, \sigma_x = 0.6$

Now we can look at how the cutoff value varies with the belief of the cutoff

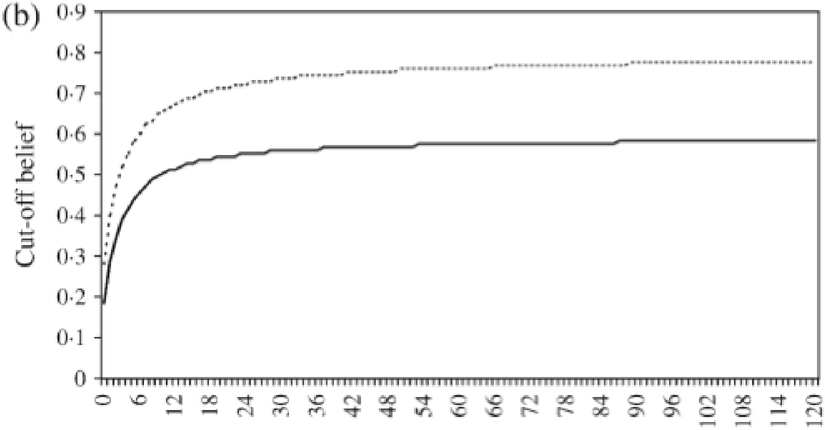

These are quite different

Now look at the hazard rates in the two different cases, but for two different types of firms:

- low endogenous separation rate
- High endogenous separation rate

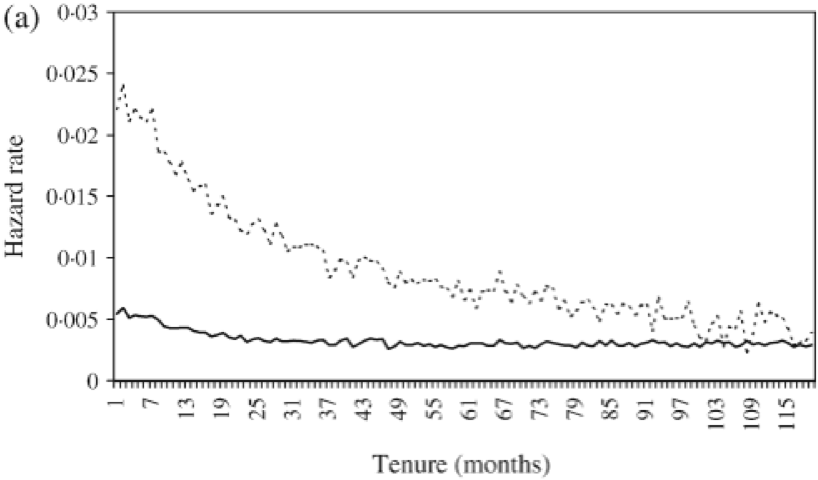
(a)
Hazard rate
0·03
0·025
0·02
0·015
0·01
0·005
0
1
7
13
19
25
31
37
43
49
55
61
67
73
79
85
91
97
103
109
115
Tenure (months)

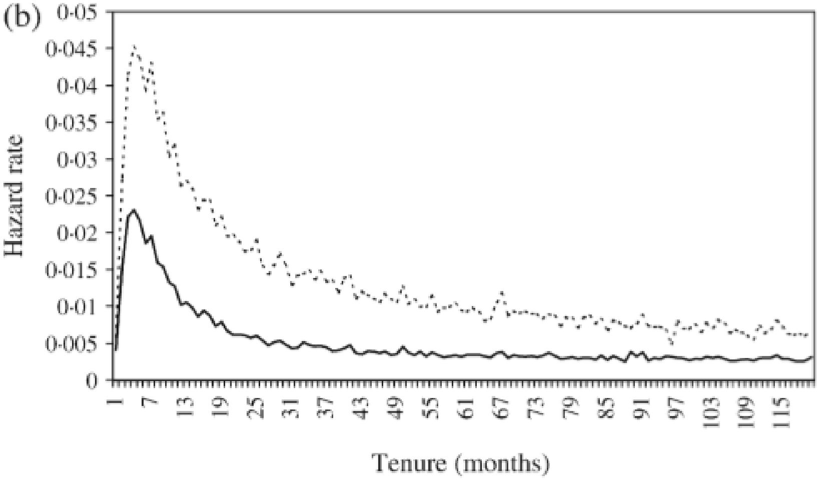
(b)
Hazard rate
0·05
0·045
0·04
0·035
0·03
0·025
0·02
0·015
0·01
0·005
0
1
7
13
19
25
31
37
43
49
55
61
67
73
79
85
91
97
103
109
115
Tenure (months)

Estimation

Model is estimated using Efficient Method of Moments (although is not necessarily efficient in this case)

Basic idea is

1. Propose and Auxiliary Model
2. Estimate Parameters of auxiliary model
3. Define objective function $f(\theta)$ in the following way:
    1. For a given θ simulate data from the model
    2. Estimate the parameters of the auxiliary model from this simulated data
    3. Define $f(\theta)$ to be the distance between the parameters of the auxiliary model estimated from the simulated data relative to the parameters coming from the actual data
4. Minimize $f(\theta)$

She uses a discrete time hazard model so that for $\tau = \tau_0, ..., \tau_n$ on the interval m $\{\tau_{m-1}+1, \tau_m\}$ the hazard rate function is

$$h(\tau, s; \eta) = \frac{exp(\eta_m + \eta_{n+m}s)}{1 + exp(\eta_m + \eta_{n+m}s)}$$

Here s is the endogenous separation rate from the employing firm

This is defined as

$$\frac{\text{The number of workers who are laid off or quit during a quarter}}{\text{Total number of workers at firm}}$$

This is estimated using Maximum Likelihood

She uses the score of the likelihood function as the measure of distance

That is the score on the actual data is zero using the parameters of the actual data

$f(\theta)$ is a function of the score evaluated at the parameters estimated from the simulated data

Data

The data comes from two sources

the first is the French Labor Force Survey Enquete Emploi

It has

- about 60,000 households
- Each household is in for 3 years

Second data set is Declaration Mensuelle des Mouvements de Main-d'Oeuvre

It is

- a survey of firms that employ at least 50 workers
- Contains monthly data on exits and entrants
- She tries to match the two data sets together as well as possible

She first estimates the auxiliary model

TABLE 3

Auxiliary model estimation results (standard errors in parentheses)

Variable	Coefficients	Average hazard rate (%)
η_1	−5·7787 (0·2724)	0·414
η_2	−4·7664 (0·1576)	1·084
η_3	−4·7397 (0·1322)	1·027
η_4	−4·8932 (0·1377)	0·892
η_5	−5·3637 (0·1089)	0·528
η_6	−5·5251 (0·1001)	0·452
η_7	−5·8743 (0·1185)	0·311
Variable	**Coefficients**	**Marginal effect (%)**
η_8	0·1394 (0·0405)	0·029
η_9	0·1223 (0·0243)	0·066
η_{10}	0·1136 (0·0353)	0·059
η_{11}	0·1167 (0·0249)	0·052
η_{12}	0·0953 (0·0139)	0·026
η_{13}	0·1139 (0·0164)	0·026
η_{14}	0·1002 (0·0234)	0·016
Log likelihood	−2998·6	

And then the structural model

TABLE 4

Structural model estimation results (standard errors in parentheses)

Variable	Interpretation	Coefficient
δ	Exogenous separation rate	0·00322 (0·00169)
w	Wage	0·5189 (0·4546)
σ_μ	Dispersion of initial match quality	0·6261 (0·2652)
σ_x	Dispersion of productivity around match quality	1·0283 (0·3740)
σ_γ	Dispersion of initial uncertainty about tasks	0·6016 (7·1750)
σ_y	Dispersion of signals about tasks	0·3075 (0·1667)
N	Number of tasks	5·0901 (2·1846)
EMM criterion function		$8{\cdot}17 \times 10^{-5}$

Lets look at some stuff to see what the parameters of the model mean

Here are the cutoff values from the model

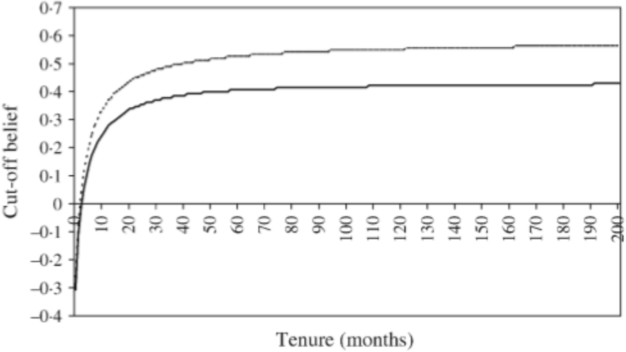
Cut-off belief
0·7
0·6
0·5
0·4
0·3
0·2
0·1
0
–0·1
–0·2
–0·3
–0·4
0
10
20
30
40
50
60
70
80
90
100
110
120
130
140
150
160
170
180
190
200
Tenure (months)

Figure 7 shows the importance of learning about the match versus learning by doing

First you can see that it takes a long time to learn the quality of the match

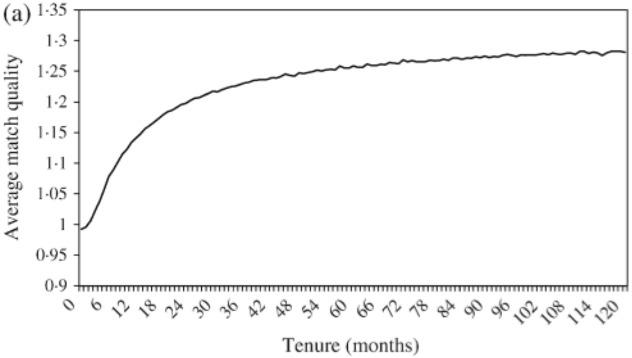
(a)
Average match quality
1·35
1·3
1·25
1·2
1·15
1·1
1·05
1
0·95
0·9
0
6
12
18
24
30
36
42
48
54
60
66
72
78
84
90
96
102
108
114
120
Tenure (months)

The learning by doing model is imprecise, but learning by doing does not appear to be important

It appears to happen very quickly

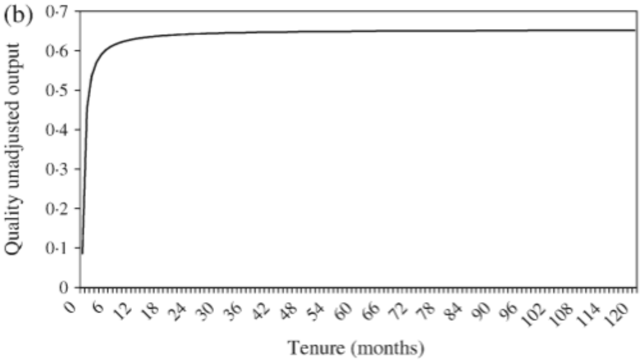

Putting them together you get this:

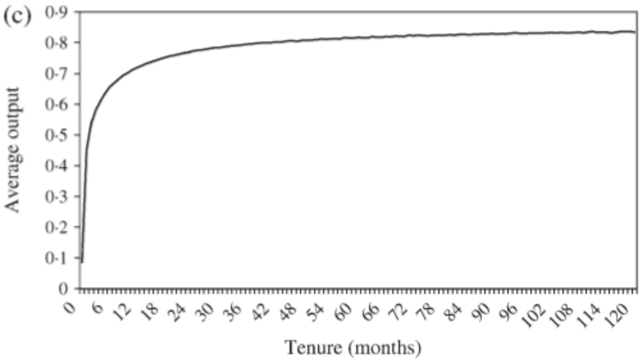

Outline

1. Jovanavic
2. Nagypal
3. Neal

Neal, *Journal of Labor Economics*, 1999

So far you either switch jobs or you stay at the same job: human capital is either completely job specific or completely tenure

Neal will be the first paper to relax this

He wants to distinguish between "complex" job switches in which workers switch careers from simple job shifts in which workers switch firms but do not switch careers

He develops a simple model of this and shows that the data is consistent with the basic predictions of the model: workers first shop for a career and then shop for a firm within the career

The key components of the model are:

- Career match θ distributed $F(\theta)$
- Job match ξ distributed $G(\xi)$

The key restriction of the model is that

- to switch careers, you must switch firms, but
- to switch firms, you do not have to switch careers

He is going to abstract from everything but the most necessary components-clearly one could make this model more complicated if you want.

Assuming that people are paid $\theta + \xi$ and that there are no search costs in the sense that you can always find a new job of the type you want-but you don't observe the match component until you start working there

You can write the Belman equation as

$$\begin{aligned} V(\theta, \xi) &= \theta + \xi + \beta max \left\{ V(\theta, \xi), \int V(\theta, s) dG(s), \right. \\ & \left. \int \int V(x, s) dF(x) dG(s) \right\} \end{aligned}$$

where V is the value function and β is the discount factor.

You can think of it in terms of the following figure

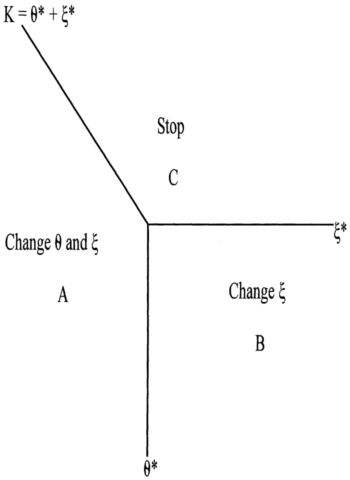

Note that once you get to region B, you will never go back to A

Once you get to C, you will stay

This has the implication that as workers age, the fraction of job changes that are complex should fall

Note also that if we condition on people who have never made a simple job change, the probability that the next job change will be simple does not depend on age

Neal looks for these implications in the data

Data

He uses the NLSY which is great for constructing data on job changes and how they vary with occupation and industry

He looks at Males only

The question here is what represents a career

Neal assumes that a complex job change represents both an occupation change and an industry change

Lets look at the first piece of evidence.

Each observations is a sequence of job changes.

He groups by the total number of job changes and documents the fraction consistent with the pure model (i.e. no complex changes following simple changes)

Table 1
The Prevalence of the Two-Stage Search

Number of Employer Changes in the Job History (1)	Percentage of Actual Job Histories That Satisfy the Two-Stage Search (2)	Percentage of Job Histories That Satisfy the Two-Stage Search Given Random Behavior (3)
2	83.6 (318)	75.3
3	63.3 (283)	50.5
4	52.8 (235)	31.8
5	42.1 (209)	19.4
6	30.7 (163)	11.4
7	27.5 (120)	7
8	16.1 (93)	3.8
9	22 (59)	2.2
10	12.5 (48)	1.2

You can see that the results are not precisely the two stage model, but they are much closer than you would expect by chance

Next an observation is a single job change

He groups by the number of simple changes since working in the current career (and by education)

Table 2
The Frequency of Career Changes among Workers Who Are Changing Employers (in %)

	Dropout	High School Graduate	College Graduate
Leaving first job in a given career	70.9 (2,270)	68.5 (3,942)	54.0 (658)
Prior employer changes while working in current career = 1	23.8 (361)	22.1 (700)	16.5 (176)
Prior employer changes while working in current career = 2	22.5 (178)	15.6 (283)	10.5 (67)
Prior employer changes while working in current career >2	14.6 (192)	16.0 (331)	7.8 (51)

Key thing is that (for example) for high school graduates for whom this is there first firm in the career, the chances that the next switch is complex is 69%

However, for those who underwent a previous job switch in this career, it is only 22%

The next tables are similar, but we group by experience

Table 3
The Frequency of Career Changes as a Function of Work Experience (in %)

Months of Work Experience	Education Level: Dropout	High School Graduate	College Graduate
$\exp \leq 12$	64.2	64.0	49.5
	(509)	(888)	(99)
$12 < \exp \leq 24$	63.4	64.5	52.3
	(535)	(965)	(172)
$24 < \exp \leq 36$	59.4	56.8	44.7
	(441)	(759)	(141)
$36 < \exp \leq 60$	56.5	55.4	36.8
	(616)	(1,038)	(242)
$60 < \exp \leq 84$	56.5	50.3	36.4
	(455)	(737)	(154)
$\exp > 84$	51.9	44.5	33.3
	(445)	(869)	(144)

Table 4
The Frequency of Career Changes: The Roles of Work Experience and Prior Employer Changes within Career (in %)

	Months of Work Experience					
	$\text{exp} \leq 12$	$12 \leq \text{exp} < 24$	$24 \leq \text{exp} < 36$	$36 \leq \text{exp} < 60$	$60 \leq \text{exp} < 84$	$\text{exp} \geq 84$
Dropouts:						
1. Prior simple changes = 0	69.1	72.7	68.5	73.0	71.9	70.2
	(469)	(440)	(352)	(414)	(306)	(289)
2. Prior simple changes > 0	7.5	20.0	23.6	22.8	24.8	18.0
	(40)	(95)	(89)	(202)	(149)	(156)
High school graduates:						
3. Prior simple changes = 0	67.7	73.4	66.6	70.5	68.3	62.0
	(838)	(812)	(571)	(748)	(457)	(516)
4. Prior simple changes > 0	2.0	17.0	27.1	16.6	21.1	19.0
	(50)	(153)	(188)	(290)	(280)	(353)
College graduates:						
5. Prior simple changes = 0	53.9	56.3	54.6	52.0	56.3	50.0
	(91)	(151)	(108)	(152)	(80)	(76)
6. Prior simple changes > 0	.0	23.8	12.1	11.1	14.9	14.7
	(8)	(21)	(33)	(90)	(74)	(68)

NOTE.—The numbers in parentheses give the number of job changes in each category. The percentages give the fraction of job changes that involve a career change. Work experience (exp) is measured in months beginning with the worker's transition to full-time employment. For all education categories, I test the null hypothesis that conditional on no prior simple changes the probability of a career change is constant across experience categories. For high school graduates, the data easily reject this null at a significance level of .001. But, for dropouts and college graduates, the test statistics, which are distributed $\chi^2(5)$, are only 3.52 and 1.24.

One concern is that this could be about career specific human capital rather than about search.

Neal addresses this with the following Table

Table 5
The Frequency of Career Changes: The Roles of Career-Specific Experience and Prior Employer Changes within Career (in %)

	Months of Experience in Current Career					
	exp ≤ 12	12 ≤ exp < 24	24 ≤ exp < 36	36 ≤ exp < 60	60 ≤ exp < 84	exp ≥ 84
Dropouts:						
1. Prior simple changes = 0	71.5	71.8	68.1	65.4	62.2	76.2
	(1,491)	(454)	(163)	(104)	(37)	(21)
2. Prior simple changes > 0	8.8	21.0	29.0	19.0	32.9	20.4
	(114)	(200)	(124)	(174)	(70)	(49)
High school graduates:						
3. Prior simple changes = 0	71.3	65.8	65.2	63.0	67.7	51.1
	(2,369)	(783)	(302)	(292)	(102)	(94)
4. Prior simple changes > 0	6.9	19.1	24.6	20.3	21.6	18.3
	(159)	(299)	(252)	(291)	(176)	(137)
College graduates:						
5. Prior simple changes = 0	62.0	49.4	50.0	42.4	53.1	31.3
	(295)	(166)	(64)	(85)	(32)	(16)
6. Prior simple changes > 0	4.8	14.0	8.6	13.7	29.0	9.4
	(21)	(50)	(58)	(95)	(38)	(32)

NOTE.—The numbers in parentheses give the number of job changes in each category. The percentages give the fraction of job changes that involve a career change. Career-specific experience is defined as the months of work experience in the career associated with the job that the worker is leaving.

While the strict version of the model is not precisely true, the data is broadly consistent with the idea.