

Job Matching with Bayesian Learning

Robert A. Miller

Structural Econometrics

September 2020

Bayesian Learning

Motivation

- Adam Smith, and many others, including perhaps your parents, have commented on "the hasty, fond, and foolish intimacies of young people" (Smith, page 395, volume 1, 1812).
- One approach to explaining such behavior is to argue that some people are not rational all the time.
- A challenge for this approach is to develop an axiomatic theory for irrational agents that has refutable predictions.
- There is ongoing research in behavioral economics and economic theory in this direction.
- Another approach, embraced by many labor economists, is that by repeatedly sampling experiences from an unfamiliar environment, rational Bayesians update their prior beliefs as they sequentially solve their lifecycle problem.

- This issue seems like a candidate for applying the methodology described in the previous slides:
 - 1 Write down a dynamic discrete choice model of Bayesian updating and sequential optimization problem;
 - 2 Solve the individual's optimization problem (for all possible parameterizations of the primitives);
 - 3 Treat important factors to the decision maker that are not reported in the sample population as unobserved variables to the econometrician;
 - 4 Integrating over the probability distribution of unobserved random variables, form the likelihood of observing the sample;
 - 5 Maximize the likelihood to obtain the structural parameters that characterize the dynamic discrete choice problem;
 - 6 Predict how the individual would adjust her behavior if she was confronted with new opportunities to learn or different payoffs.

Job Matching and Occupational Choice (Miller JPE, 1984)

Individual payoffs and choices

- The payoff from job $m \in M$ at time $t \in \{0, 1, \dots\}$ is:

$$x_{mt} \equiv \psi_t + \xi_m + \sigma_m \epsilon_{mt}$$

where:

- ψ_t is a lifecycle trend shaping term that plays no role in the analysis;
 - ξ_m is a job match parameter drawn from $N(\gamma_m, \delta_m^2)$;
 - ϵ_{mt} is an idiosyncratic *iid* disturbance drawn from $N(0, 1)$
- Every period t the individual chooses a job m to work in. The choice at t is denoted by $d_{mt} \in \{0, 1\}$ for each $m \in M$ where:

$$\sum_{m \in M} d_{mt} = 1$$

- The realized lifetime utility of the individual is:

$$\sum_{t=0}^{\infty} \sum_{m \in M} \beta^t d_{mt} x_{mt}$$

Job Matching and Occupational Choice

Processing information

- At $t = 0$ the individual sees (γ_m, δ_m^2) for all $m \in M$.
- At every t , after making her choice, she also sees ψ_t , and $d_{mt}x_{mt}$ for all $m \in M$.
- Following Degroot (*Optimal Statistical Decisions 1970, McGraw Hill*) the posterior beliefs of an individual for job $m \in M$ at time $t \in \{0, 1, \dots\}$ are $N(\gamma_{mt}, \delta_{mt}^2)$ where:

$$\gamma_{mt} = \frac{\delta_m^{-2} \gamma_m + \sigma_m^{-2} \sum_{s=0}^{t-1} (x_{ms} - \psi_s) d_{ms}}{\delta_m^{-2} + \sigma_m^{-2} \sum_{s=0}^{t-1} d_{ms}}$$

$$\delta_{mt}^{-2} = \delta_m^{-2} + \sigma_m^{-2} \sum_{s=0}^{t-1} d_{ms}$$

- She maximizes the sum of expected payoffs, sequentially choosing d_{mt} for each $m \in M$ at t given her beliefs $N(\gamma_{mt}, \delta_{mt}^2)$.

Optimization

Maximization using Dynamic Allocation Indices (DAIs)

Corollary (from Theorem 2 in Gittens and Jones, 1974)

At each $t \in \{1, 2, \dots\}$ it is optimal to select the $m \in M$ maximizing:

$$DAI_m(\gamma_{mt}, \delta_{mt}) \equiv \sup_{\tau \geq t} \left\{ \frac{E \left[\sum_{r=t}^{\tau} \beta^{r-t} (x_{mr} - \psi_r) \mid \gamma_{mt}, \delta_{mt} \right]}{E \left[\sum_{r=t}^{\tau} \beta^{r-t} \mid \gamma_{mt}, \delta_{mt} \right]} \right\}$$

- To understand the intuition for this rule, consider two projects, m' taking 4 periods with payoffs $\{1, 8, 7, x'\}$ and another m'' taking 2 periods with payoffs $\{6, x''\}$.
- Suppose m' can be split into a 3 period project with payoffs $\{1, 8, 7\}$ and an additional 1 period project with payoff $\{x'\}$ that cannot be undertaken before the 3 period project is completed, but does not have to be undertaken immediately afterwards.
- Prove the DAI rule optimally schedules the projects.

Optimization

An interpretation of the DAI

- Consider a project with payoffs $\{x_{mt}\}_{t=0}^{\infty}$ and form the value function for the following renewal problem:

$$\begin{aligned} V_{mt} &\equiv \sup_{\tau \geq t} E_t \left[\sum_{r=t}^{\tau} \beta^{r-t} x_{mr} + \beta^{\tau+1-t} V_{mt} \right] \\ &\equiv E_t \left[\sum_{r=t}^{\tau^o} \beta^{r-t} x_{mr} + \beta^{\tau^o+1-t} V_{mt} \right] \end{aligned} \quad (1)$$

- Thus V_{mt} is the maximal value from continuing with project m until some nonanticipating stopping time τ and then restarting from t , drawing a new path of rewards, optimally stopping again, and so on.
- Now define the certainty renewal flow equivalent $D_m(z_{mt})$ as:

$$D_{mt} \equiv V_{mt} \Big/ \sum_{r=t}^{\infty} \beta^{r-t}$$

Optimization

Proof sketch for optimality of DAI rule

- Substituting for $V_{mt}(z_{mt})$ in (1) yields:

$$D_{mt} \sum_{r=t}^{\infty} \beta^{r-t} = E_t \left[\sum_{r=t}^{\tau^o} \beta^{r-t} x_{mr} + \beta^{\tau^o+1-t} D_{mt} \sum_{r=t}^{\infty} \beta^{r-t} \right]$$

$$D_{mt} \left\{ \sum_{r=t}^{\infty} \beta^{r-t} - E_t \left[\beta^{\tau^o+1-t} \sum_{r=t}^{\infty} \beta^{r-t} \right] \right\} = E_t \left[\sum_{r=t}^{\tau^o} \beta^{r-t} x_{mr} \right]$$

and rearranging gives:

$$D_{mt} = E_t \left[\sum_{r=t}^{\tau^o} \beta^{r-t} x_{mr} \right] / E_t \left[\sum_{r=t}^{\tau^o} \beta^{r-t} \right]$$

- The next slide shows that for a specialization of the general framework it is optimal to undertake action m instead of another action m' with (independent) payoff structure $\{x_{m't}\}_{t=0}^{\infty}$ iff $V_{mt} \geq V_{m't}$.
- Since $V_{mt} \geq V_{m't} \Leftrightarrow D_{mt} \geq D_{m't}$ the optimality of the DAI rule follows immediately (in this special case).

Optimization

Proof in a simple case

- Suppose project m lasts τ_m periods and yields a present value reward of R_m and m' lasts τ'_m periods and yields a present value reward of R'_m . It is optimal to start with m instead of m' iff:

$$\begin{aligned} R_m + \beta^{\tau_m+1} R'_m &> R'_m + \beta^{\tau'_m+1} R_m \\ \iff R_m (1 - \beta^{\tau'_m+1}) &> R'_m (1 - \beta^{\tau_m+1}) \\ \iff R_m / (1 - \beta^{\tau_m+1}) &> R'_m / (1 - \beta^{\tau'_m+1}) \\ \iff V_m &> V'_m \\ \iff V_m / \sum_{r=t}^{\infty} \beta^{r-t} &> V'_m / \sum_{r=t}^{\infty} \beta^{r-t} \end{aligned}$$

the second last line following the fact that in this simple case:

$$V_m = R_m + \beta^{\tau_m+1} R_m + \dots = (1 - \beta^{\tau_m+1})^{-1} R_m$$

and similarly for V'_m .

Corollary (Proposition 4 of Miller, 1984)

In this model:

$$DAI_m(\gamma_{mt}, \delta_{mt}) = \gamma_{mt} + \delta_{mt} D \left[\left(\frac{\sigma_m}{\delta_m} \right)^2 + \sum_{s=0}^{t-1} d_{ms} \right]$$

where $D(\sigma)$ is the (standard) DAI for a (hypothetical) job whose fixed match parameter ξ is drawn from $N(0, 1)$ and whose random component in the payoff is $\sigma \varepsilon_t$.

- $D(\cdot)$ can be numerically computed by solving for the fixed point of a contraction mapping. (See Proposition 5 of Miller, 1984.)

Optimization

Occupations and optimal turnover

- We define an occupation as a set of jobs with the same $(\gamma_m, \delta_m, \sigma_m)$ value.
- Thus $(\gamma_m, \delta_m, \sigma_m)$ differs across jobs in a multi-occupational world.
- Now $D(\cdot)$ is a decreasing function. Thus $DAI_m(\gamma_{mt}, \delta_{mt}) \uparrow$ as:
 - γ_{mt}, δ_{mt} and $\delta_m \uparrow$
 - σ_m and $\sum_{s=0}^{t-1} d_{ms} \downarrow$.
- Given γ_m :
 - Occupations with high δ_m and low σ_m are experimented with first;
 - Matches with low σ_m are resolved for better or worse relatively quickly;
 - Turnover declines with tenure. (See also Jovanovic, 1979.)
- Lastly, β also affects the DAI because this parameter indexes how much future payoffs are discounted.

Empirical Application

A world with only one occupation

- It is just as easy to compute the DAIs for an economy with many occupations as a world with only one.
- However (we shall see that) the multiple integration required for a more complex world is essentially unmanageable if the econometrician does not observe $d_{mt}x_{mt}$, the payoff from choosing job $m \in M$ at time $t \in \{0, 1, \dots\}$ net of the lifecycle trend shaping term.
- It is entirely reasonable to think this is indeed the case, because match quality specific factors often revolve around nonpecuniary intangibles that are only partly reflected in wages (in a possibly nonmonotone way).
- Consider the a discrete time analogue to Jovanovic's (1979) one occupation job matching economy.
- The limited objective in this study was to seek evidence against this economy, as a way of empirically motivating why a multi-occupational world seems plausible.

The Colman-Rossi Data Set

Tenure and turnover by employment and profession

TABLE 1
TENURE AND TURNOVER BY EMPLOYMENT AND EDUCATION

	CURRENT POSITION				PAST SPELLS			
	Number	Percentage with Tenure of			Number	Empirical Hazard		
		≥ 2	≥ 3	≥ 4		1	2	3
Employment:								
Professional	67	76	65	31	183	61	49	65
Farm owner	22	95	90	9	44	55	50	30
Manager	80	80	73	33	128	60	55	61
Clerk	40	82	67	35	175	69	55	44
Salesman	27	77	62	29	138	64	51	54
Craftsman	107	81	65	25	379	61	53	59
Operative	84	80	78	39	553	68	59	53
Serviceman	13	92	61	46	60	73	63	33
Farm laborer	6	83	83	33	144	72	54	63
Nonfarm laborer	21	76	57	33	281	78	55	39
Education:								
Grade school	177	84	75	28	779	70	55	64
High school	113	81	67	33	566	68	58	42
College	84	76	67	35	463	61	50	50

The Colman-Rossi Data Set

Transitions with and between employment groups

TABLE 2
TRANSITIONS WITHIN AND BETWEEN EMPLOYMENT GROUPS

	Professional	Farm Owner	Manager	Clerk	Salesman	Craftsman	Operative	Serviceman	Farm Laborer	Nonfarm Laborer
Professional (183)	67	1	11	4	4	5	5	1	0	1
Farm owner (44)	0	25	2	2	2	9	39	2	14	5
Manager (128)	11	2	39	4	20	10	9	1	1	3
Clerk (175)	10	0	14	33	7	11	15	2	0	7
Salesman (138)	1	1	27	6	30	9	17	4	0	5
Craftsman (379)	5	0	7	6	5	48	18	2	2	7
Operative (553)	4	3	5	6	4	19	38	3	4	14
Serviceman (60)	3	0	5	8	7	10	30	18	3	15
Farm laborer (144)	2	8	1	1	2	8	28	2	31	16
Nonfarm laborer (281)	1	2	2	8	2	18	40	3	1	22

Empirical Application

Hazard rate for spell length

- Define h_t as the (discrete) hazard at t periods as the probability a spell ends after t periods conditional on surviving that long.
- In a one occupation economy with an infinite number of jobs, it suffices to only keep track of the current job match. (Why?)
- Appealing to the corollary above:

$$\begin{aligned}h_t &\equiv \Pr \left\{ \gamma_t + \delta_t D \left[\left(\frac{\sigma}{\delta} \right)^2 + t, \beta \right] \leq \gamma + \delta D \left[\left(\frac{\sigma}{\delta} \right)^2, \beta \right] \right\} \\&= \Pr \left\{ \frac{\gamma_t - \gamma}{\sigma} \leq \frac{\delta}{\sigma} D \left[\left(\frac{\sigma}{\delta} \right)^2, \beta \right] - \frac{\delta_t}{\sigma} D \left[\left(\frac{\sigma}{\delta} \right)^2 + t, \beta \right] \right\} \\&= \Pr \left\{ \rho_t \leq \alpha^{-1/2} D(\alpha, \beta) - (\alpha + t)^{-1/2} D(\alpha + t, \beta) \right\}\end{aligned}$$

where $\rho_t \equiv (\gamma_t - \gamma) / \sigma$ and $\alpha \equiv (\sigma / \delta)^2$ which implies:

$$\frac{\delta_t}{\sigma} = \frac{[\delta^{-2} + t\sigma^{-2}]^{-1/2}}{\sigma} = \left[\left(\frac{\delta}{\sigma} \right)^{-2} + t \right]^{-1/2} = (\alpha + t)^{-1/2}$$

Probability Distribution of Spell Lengths

Relating the hazard rate to the distribution of normalized match qualities

- Define the probability distribution of transformed means of spells surviving at least t periods as:

$$\Psi_t(\rho) \equiv \Pr\{\rho_t \leq \rho\} = \Pr\{\sigma^{-1}(\gamma_t - \gamma) \leq \rho\} = \Pr\{\gamma_t \leq \gamma + \rho\sigma\}$$

- To help fix ideas note that $\Psi_0(\rho) = 0$ for all $\rho < 0$ and $\Psi_0(0) = 1$.
- From the definition of h_t and $\Psi_t(\rho)$:

$$\begin{aligned} h_t &= \Pr\left\{\rho_t \leq \alpha^{-1/2} D(\alpha, \beta) - (\alpha + t)^{-1/2} D(\alpha + t, \beta)\right\} \\ &= \Psi_t\left[\alpha^{-1/2} D(\alpha, \beta) - (\alpha + t)^{-1/2} D(\alpha + t, \beta)\right] \end{aligned}$$

- To derive the discrete hazard, we recursively compute $\Psi_t(\rho)$.

Probability Distribution of Spell Lengths

Inequalities relating to normalized match qualities after one period

- By definition every match survives at least one period, and hence:

$$\Psi_1(\rho) = \Pr\{\gamma_1 \leq \gamma + \rho\sigma\}$$

- From the Bayesian updating rule for γ_t :

$$\begin{aligned}\gamma_1 &\leq \gamma + \rho\sigma \\ \Leftrightarrow \frac{\delta^{-2}\gamma + \sigma^{-2}(x_1 - \psi_1)}{\delta^{-2} + \sigma^{-2}} &\leq \gamma + \rho\sigma \\ \Leftrightarrow \delta^{-2}\gamma + \sigma^{-2}(\xi + \sigma\epsilon) &\leq (\gamma + \rho\sigma)(\delta^{-2} + \sigma^{-2}) \\ \Leftrightarrow \alpha\gamma + \xi + \sigma\epsilon &\leq (\gamma + \rho\sigma)(\alpha + 1) \\ \Leftrightarrow (\xi - \gamma) + \sigma\epsilon &\leq \sigma(\alpha + 1)\rho \\ \Leftrightarrow \delta^{-1}(\xi - \gamma) + \alpha^{1/2}\epsilon &\leq \alpha^{1/2}(\alpha + 1)\rho\end{aligned}$$

Probability Distribution of Spell Lengths

Computing the distribution of normalized match qualities after one period

- Since every match survives at least one period, we can calculate $\Psi_1(\rho)$ for all matches:

$$\Psi_1(\rho) \equiv \Pr\{\gamma_1 \leq \gamma + \rho\sigma\}$$

- Appealing to the inequalities from the previous slide:

$$\begin{aligned}\Psi_1(\rho) &= \Pr\{\gamma_1 \leq \gamma + \rho\sigma\} \\ &= \Pr\{\delta^{-1}(\xi - \gamma) + \alpha^{1/2}\epsilon \leq \alpha^{1/2}(\alpha + 1)\rho\} \\ &= \Pr\{\epsilon' + \alpha^{1/2}\epsilon \leq \alpha^{1/2}(\alpha + 1)\rho\} \\ &= \Pr\{(\alpha + 1)^{1/2}\epsilon'' \leq \alpha^{1/2}(\alpha + 1)\rho\} \\ &= \Phi\left[\alpha^{1/2}(\alpha + 1)^{1/2}\rho\right]\end{aligned}$$

where ϵ , ϵ' and ϵ'' are independent standard normal random variables.

Probability Distribution of Spell Lengths

Solving for the one period hazard rate and the probability distribution of survivors

- The spell ends if:

$$\rho_1 < \alpha^{-1/2} D(\alpha, \beta) - (\alpha + 1)^{-1/2} D(\alpha + 1, \beta)$$

- Therefore the proportion of spells ending after one period is:

$$\begin{aligned} h_1 &= \Psi_1 \left[\alpha^{-1/2} D(\alpha, \beta) - (\alpha + 1)^{-1/2} D(\alpha + 1, \beta) \right] \\ &= \Phi \left\{ \begin{array}{l} \left[\alpha^{1/2} (\alpha + 1)^{1/2} \right] \\ \times \left[\alpha^{-1/2} D(\alpha, \beta) - (\alpha + 1)^{-1/2} D(\alpha + 1, \beta) \right] \end{array} \right\} \\ &> 1/2 \text{ (because } D(\cdot) \text{ in } \alpha) \end{aligned}$$

- So the truncated distribution of ρ for survivors after one draw is:

$$\tilde{\Psi}_1(\rho) \equiv (1 - h_1)^{-1} [\Psi_1(\rho) - h_1]$$

Probability Distribution of Spell Lengths

Recursively computing the distribution of normalized match qualities

- To derive $\Psi_2(\rho)$ from $\tilde{\Psi}_1(\rho)$ the worker takes another draw, and appealing to Bayes rule one more time:

$$\begin{aligned}\Psi_2(\rho) &\equiv \frac{\int_{-\infty}^{\infty} \Psi_1\left(\rho - \epsilon [(\alpha + 1)(\alpha + 2)]^{-1/2}\right) d\Phi(\epsilon) - h_1}{1 - h_1} \\ &= \frac{\int_{-\infty}^{\infty} \Phi\left[\alpha^{1/2}(\alpha + 1)^{1/2} \times \left(\rho - \epsilon [(\alpha + 1)(\alpha + 2)]^{-1/2}\right)\right] d\Phi(\epsilon) - h_1}{1 - h_1}\end{aligned}$$

- More generally (from page 1112 of Miller, 1984):

$$\Psi_{t+1}(\rho) \equiv \frac{\int_{-\infty}^{\infty} \Psi_t\left(\rho - \epsilon [(\alpha + t)(\alpha + t + 1)]^{-1/2}\right) d\Phi(\epsilon) - h_t}{1 - h_t}$$

Maximum Likelihood Estimation

Complete and incomplete spells

- Suppose the sample comprises a cross section of spells $n \in \{1, \dots, N\}$, some of which are completed after τ_n periods, and some of which are incomplete lasting at least τ_n periods. Let:

$$\rho(n) \equiv \begin{cases} \tau_n & \text{if spell is complete} \\ \{\tau_n, \tau_{n+1}, \dots\} & \text{if spell is incomplete} \end{cases}$$

- Let $p_\tau(\alpha_n, \beta_n)$ denote the unconditional probability of individual n with discount factor β_n working τ periods in a new job with information factor α_n before switching to another new job in the same occupation:

$$p_\tau(\alpha_n, \beta_n) \equiv h_\tau(\alpha_n, \beta_n) \prod_{s=1}^{\tau-1} [1 - h_s(\alpha_n, \beta_n)]$$

- Then the joint probability of spell duration times observed in the sample is:

$$\prod_{n=1}^N \sum_{\tau \in \rho(n)} p_\tau(\alpha_n, \beta_n)$$

Maximum Likelihood Estimation

The likelihood function and structural estimates

- We could allow for an additional source of unobserved heterogeneity by writing the likelihood as:

$$L_N(A_1, B_1, A_2, B_2, \lambda) \equiv \prod_{n=1}^N \sum_{\tau \in \rho(n)} \left[\begin{array}{l} \lambda p_{\tau}(\alpha_{1n}, \beta_{1n}) \\ + (1 - \lambda) p_{\tau}(\alpha_{2n}, \beta_{2n}) \end{array} \right]$$

where we now assume that $\alpha_{in} \equiv A_i X_n$ and $\beta_{in} \equiv B_i X_n$ for $i \in \{1, 2\}$ and the parameter space is $(A_1, B_1, A_2, B_2, \lambda)$.

- Briefly, the structural estimates show that:
 - 1 individuals care about the future and value on job experimentation;
 - 2 the occupational dummy variables are significant, suggesting that the choice of different occupations is not random;
 - 3 educational groups have different beliefs and learning rates;
 - 4 these three results are not sensitive to whether the additional unobserved heterogeneity is incorporated or not.

Recent Work

Recent studies estimating dynamic discrete choice models with Bayesian learning

- There is renewed interest within structural estimation for modeling Bayesian learning as the Markov process driving the state variables:
 - ① Pharmaceuticals: Crawford and Shum (2005)
 - ② Occupational choice: James (2011)
 - ③ Wage contracting: Pastorino (2014)
 - ④ Entrepreneurship: Hincapie (2016), Dillon and Stanton (2016).
 - ⑤ College choices: Arcidiacono, Aucejo, Maurel and Ransom (2016)
- Compared to earlier work, recent studies:
 - draw upon larger samples;
 - focus more closely on wages and less on nonpecuniary characteristics;
 - do not solve the dynamic optimization problem to estimate the model;
 - use simulation methods instead of directly integrating;
 - predict the outcomes of counterfactual regimes induced by hypothetical technical change and alternative public policies;
 - use similar numerical techniques to this study when solving optimization problems to conduct counterfactuals.