Job Matching with Bayesian Learning

Robert A. Miller

Structural Econometrics

September 2017

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.

Bayesian Learning
Applying the methodology

- 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.

- The payoff from job $m \in M$ at time $t \in \{0, 1, \ldots\}$ 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 $N\left(\gamma_m, \delta_m^2\right)$;
  - 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 (Miller JPE, 1984)

Processing information

- At $t = 0$ the individual sees $\left(\gamma_m, \delta_m^2\right)$ 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 $N\left(\gamma_{mt}, \delta_{mt}^2\right)$ where:

$$
\begin{aligned}
\gamma_{mt} &= \frac{\delta_m^{-2}\gamma_m + \sigma_m^{-2}\sum_{s=0}^{t-1}\left(x_{ms} - \psi_s\right)d_{ms}}{\delta_m^{-2} + \sigma_m^{-2}\sum_{s=0}^{t-1} d_{ms}} \\
\delta_{mt}^{-2} &= \delta_m^{-2} + \sigma_m^{-2}\sum_{s=0}^{t-1} d_{ms}
\end{aligned}
$$

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

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

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

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

- 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.
- What is the optimal order for undertaking the projects?

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 $D(\cdot)$ *is the (standard) DAI for a (hypothetical) job whose match parameter* ξ *is drawn from* $N(0,1)$ *and whose payoff net of the general component is* $\sigma^2 \varepsilon_t$.

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

Optimization

Optimal turnover

- $D(\cdot)$ is a deceasing 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.)

Optimization with One Job Type

A more conventional approach to the solution

- Simplifying to a world where all jobs initially look identical:

$$V_t(\gamma_t) = \max \left\{ V_0, \gamma_t + \beta \int V_{t+1} \left(\frac{\alpha \gamma_t + (\xi + \sigma\epsilon)}{\alpha + 1} \right) f_t \left(\xi + \sigma\epsilon | \gamma_t \right) \right\}$$

where V_0 is an arbitrary constant and

$$\alpha \equiv (\sigma / \delta)^2 \qquad x_{t+1} - \psi_{t+1} = \xi + \sigma\epsilon$$

which implies:

$$\frac{\alpha \gamma_t + (\xi + \sigma\epsilon_{t+1})}{\alpha + 1} = \frac{\delta^{-2}\gamma_t + \sigma^{-2}(\xi + \sigma\epsilon_{t+1})}{\delta^{-2} + \sigma^{-2}}$$

- $\xi \sim N\left(\gamma_t, \delta_t^2\right)$ at t and $\epsilon \sim N\left(0, \sigma^2\right)$
- Therefore $f_t\left(\xi + \sigma\epsilon | \gamma_t\right) \sim N\left(\gamma_t, \delta_t^2 + \sigma^2\right)$.
- Recursively compute to obtain $V_t(\gamma)$ for different values of V_0.
- Now solve for $V_0 = V_0(\gamma)$.

Optimization with One Job Type

Simplifying the problem a little further

- The realized lifetime utility of the individual is a linear transformation of the original utility:

$$\sum_{t=0}^{\infty} \sum_{m \in M} \beta^t d_{mt} \left(\frac{x_{mt} - \gamma}{\sigma} \right)$$

- Applying the Bayesian updating formula recursively:

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

$$\rho_{t+1} \equiv \frac{\gamma_{t+1} - \gamma}{\sigma} = \frac{\alpha}{1+\alpha} \rho_t + \frac{1}{1+\alpha} \left(\frac{x_{m,t+1} - \gamma}{\sigma} \right)$$

and $\left(x_{m,t+1} - \gamma\right)/\sigma \sim N\left(0, (\alpha + t)^{-1} + 1\right)$ at t.

- We can rewrite the value function recursion with a state variable ρ_t and three parameters (W_0, α, β).

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			
		Percentage with Tenure of				Empirical Hazard		
	Number	⩾ 2	⩾ 3	⩾ 4	Number	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

Probability Distribution of Spell Lengths

Hazard rate for spell length

- We define the (discrete) hazard at t periods as the probability a spell ends after t periods conditional on surviving that long.
- In a one occupation model:

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

where from the Bayesian updating formula on Slide 5:

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

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:

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

- To help fix ideas note that $\Psi_0(\rho) = 0$ for all $\rho \leq 0$ 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

Recursively computing the distribution of normalized match qualities

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

$$\Psi_1(\rho) \equiv \Pr\left\{\frac{\gamma_1 - \gamma}{\sigma} \leq \rho\right\} = \Phi\left[\alpha^{1/2}(\alpha+1)^{1/2}\rho\right]$$

- 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:

$$h_1 = \Psi_1\left[\alpha^{-1/2} D(\alpha, \beta) - (\alpha+1)^{-1/2} D(\alpha+1, \beta)\right] > 1/2$$

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

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

Probability Distribution of Spell Lengths

The distribution of normalized match qualities

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

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

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

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

Maximum Likelihood Estimation

Complete and incomplete spells

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

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

- 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:

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

  where we now assume that $\alpha_{in} \equiv A_i X_n$ and $\beta_n \equiv B_i X_n$ 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 contacting: Pastorino (2014)
  4. Entrepreneurship: Hincape (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 necessarily solve the dynamic optimization problem for different parameter values to estimate the model;
  - 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, both in dynamic theory and in estimation.