Job Matching and Turnover

Robert A. Miller

Structural Econometrics

October 2025

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

A structural approach

- How could we formulate this as an economic model and estimate it?
  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 $m \in \{1, 2, \ldots\}$ at time $t \in \{0, 1, \ldots\}$ is:

$$x_{mt} \equiv \psi_t + \xi_m + \sigma \epsilon_{mt} \tag{1}$$

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 $N\left(\gamma, \delta^2\right)$;
  - ε_{mt} is an idiosyncratic *iid* disturbance drawn from $N\left(0, 1\right)$

- 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 $m \in \{1, 2, \ldots\}$ where:

$$\textstyle\sum_{m=1}^{\infty} d_{mt} = 1$$

- The realized lifetime utility of the individual is:

$$\textstyle\sum_{t=0}^{\infty} \sum_{m=1}^{\infty} \beta^t d_{mt} x_{mt}$$

Job Matching and Occupational Choice

Processing information

- At $t = 0$ the individual sees $\left(\gamma, \delta^2\right)$, the same for all m.
- At every t, after making her choice, she also sees ψ_t, and $d_{mt}x_{mt}$ for all m.
- Following DeGrroot *(Optimal Statistical Decisions 1970, McGraw Hill*) the posterior beliefs of an individual for job m at time $t \in \{0, 1, \ldots\}$ are $N\left(\gamma_{mt}, \delta^2_{mt}\right)$ where:

$$\begin{aligned} \gamma_{mt} &= \frac{\delta^{-2}\gamma + \sigma^{-2}\sum_{s=0}^{t-1}\left(x_{ms} - \psi_s\right)d_{ms}}{\delta^{-2} + \sigma^{-2}\sum_{s=0}^{t-1} d_{ms}} \\ \delta^{-2}_{mt} &= \delta^{-2} + \sigma^{-2}\sum_{s=0}^{t-1} d_{ms} \end{aligned} \tag{2}$$

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

Optimization

Renewal problem

- Let $\{d_{mt}\}_{m=1}^{\infty}$ denote the period t choice.
- Also denote by V_0 the exante value function,defined as:

$$V_0 = \max_{\{d_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \sum_{m=1}^{\infty} \beta^t d_{mt} x_{mt} \right] \equiv E_0 \left[\sum_{t=0}^{\infty} \sum_{m=1}^{\infty} \beta^t d_{mt}^o x_{mt} \right]$$

- A simple contradiction argument proves that after leaving a job, it is never optimal to return to it:
    - Intuitively the first time you quit one job for another, the value of staying is less than V_0, and starting a new job is always an option here.
- Optimization problems with this feature (of always having the choice to restart), are called renewal problems.

Optimization

A recursive representation

- Suppose the current job m has a match distribution of $(\gamma_{mt}, \delta_{mt})$.
- Note distributions of all previous matches jobs are irrelevant.
- Let $V(\gamma_{mt}, \delta_{mt})$ denote the value of optimally solving the worker's problem from this point forwards:

$$V(\gamma_{mt}, \delta_{mt}) = \max\left\{V_0, E\left[x_{mt} + V\left(\gamma_{m,t+1}, \delta_{m,t+1}\right) | \gamma_{mt}, \delta_{mt}\right]\right\}$$

- Then $V_0 = V(\gamma, \delta)$, and appealing (2):

$$\begin{aligned} E\left[x_{mt} | \gamma_{mt}, \delta_{mt}\right] &\equiv \psi_t + \gamma_{mt} \\ \gamma_{m,t+1} &= \gamma_{mt} + \frac{x_{mt} - \psi_t}{\sigma^2 \delta_{mt}^{-2} + 1} \\ \delta_{m,t+1}^{-2} &= \delta_{mt}^{-2} + \sigma^{-2} \end{aligned}$$

and hence $E\left[x_{mt} + V\left(\gamma_{m,t+1}, \delta_{m,t+1}\right) | \gamma_{mt}, \delta_{mt}\right] =$

$$\psi_t + \gamma_{mt} + E\left[V\left(\gamma_{mt} + \frac{\xi_m + \sigma\epsilon_{mt}}{\sigma^2 \delta_{mt}^{-2} + 1}, \left[\delta_{mt}^{-2} + \sigma^{-2}\right]^2\right) | \gamma_{mt}, \delta_{mt}\right]$$

A Generalization

Individual payoffs and choices

- We can:
  - generalize this model by distinguishing between jobs and occupations;
  - reduce the complexity of the numerical algorithm solving the model.
- Suppose the payoff from job $m \in M \leq \infty$ at time $t \in \{0, 1, \ldots\}$ is:

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

where ξ_m is drawn from $N\left(\gamma_m, \delta_m^2\right)$, and as before:

  - the individual sees $\left(\gamma_m, \delta_m^2\right)$ for all $m \in M$ at $t = 0$.
  - she maximizes the sum of expected payoffs, sequentially choosing d_{mt} for each $m \in M$ at t given her beliefs $N\left(\gamma_{mt}, \delta_{mt}^2\right)$.
- Note that if:
  - $\left(\gamma_k, \delta_k^2\right) \neq \left(\gamma_m, \delta_m^2\right)$ then we say that k and m belong to different occupations.
  - $M < \infty$ then a worker might return to a job she quit.

A Generalization

Maximizing using Dynamic Allocation Indices (DAIs)

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:

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

- If τ is fixed and there is perfect foresight, the fundamental ratio is:
  - the discounted sum of benefits $\sum_{r=t}^{\tau} \beta^{r-t}\left(x_{mr} - \psi_r\right)$
  - divided by the discounted sum of time $\sum_{r=t}^{\tau} \beta^{r-t}$.
- For example if project A yields 5 and takes 2 periods to complete, and B yields 3 but only takes 1 period, do A first if and only if:

$$\begin{aligned} & 5 + 3\beta^2 > 3 + 5\beta \\ \iff \quad & 5\left(1 - \beta\right) > 3\left(1 - \beta\right)\left(1 + \beta\right) \\ \iff \quad & DAI_A \equiv 5 \big/ \left(1 + \beta\right) > 3 \equiv DAI_B \end{aligned}$$

A Generalization

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:

$$\begin{aligned} V_{mt} &\equiv \sup_{\tau \geq t} E_t \left[\sum_{r=t}^{\tau} \beta^{r-t} x_{mr} + \beta^{\tau+1-t} V_{mt} \right] \\ &\equiv E_t \left[\sum_{r=t}^{\tau^o} \beta^{r-t} x_{mr} + \beta^{\tau^o+1-t} V_{mt} \right] \end{aligned} \qquad (3)$$

- 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_{mt} as:

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

Optimization

Proof sketch for optimality of DAI rule

- Substituting for $V_{mt}\left(z_{mt}\right)$ in (3) 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:

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

- 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:

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

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

$$V_m = R_m + \beta^{\tau_m+1} R_m + \ldots = \left(1 - \beta^{\tau_m+1}\right)^{-1} R_m$$

  and similarly for V'_m.

Optimization

Bayesian learning with a normal distribution

Corollary (Proposition 4 of Miller, 1984)

In this model:

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

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* $\sigma\varepsilon_t$*.*

Optimization

Occupations and optimal turnover

- Define an occupation as jobs with the same initial $(\gamma_m, \delta_m, \sigma_m)$.
- In a multi-occupational world $(\gamma_m, \delta_m, \sigma_m)$ differs across jobs.
- We can prove $D(\cdot)$ is a decreasing function.
- Consequently $DAI_m(\gamma_{mt}, \delta_{mt}) \uparrow$ as:
  - γ_{mt} and $\delta_{mt} \uparrow$
  - 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 (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 the multiple integration required for a more complex world is essentially unmanageable if $d_{mt}x_{mt}$ is not not observed for $m \in M$ at time $t \in \{0, 1, \ldots\}$.
- Yet match quality specific factors often revolve around nonpecuniary intangibles that are only partly reflected in wages (in a possibly nonmonotone way).
- 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.
- More specifically: could risky behavior be rational?
- We return to the single occupation we started the lecture with.

The Colman-Rossi Data Set

Tenure and turnover by employment and profession

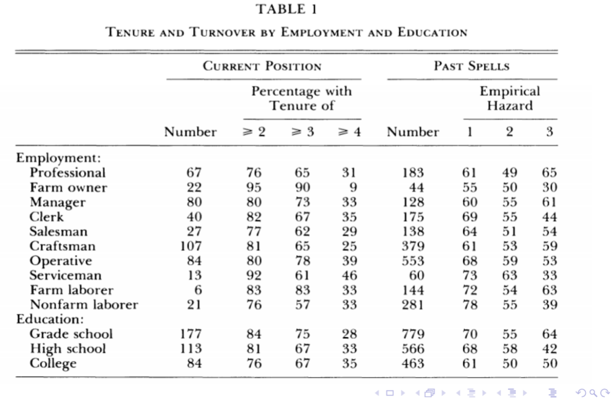

TABLE 1

TENURE AND TURNOVER BY EMPLOYMENT AND EDUCATION

	CURRENT POSITION				PAST SPELLS			
		Percentage with Tenure of				Empirical Hazard		
	Number	$\geqslant 2$	$\geqslant 3$	$\geqslant 4$	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

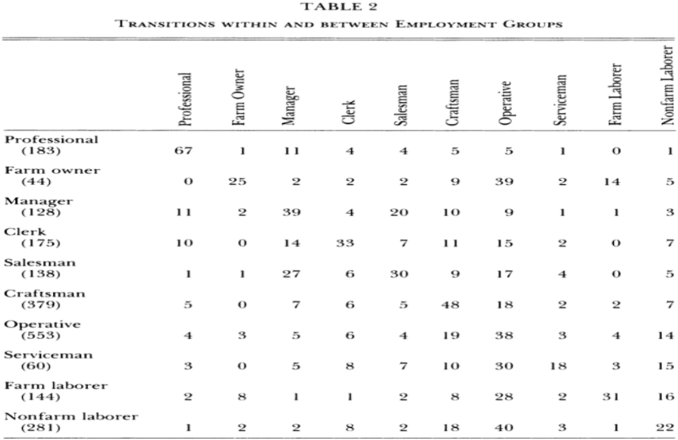

TABLE 2

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

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:

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

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:

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

- 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:

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

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:

$$\Psi_1(\rho) \equiv \Pr\{\gamma_1 \leq \gamma + \rho\sigma\} \equiv \Pr\{\rho_1 \leq \rho\}$$

- Appealing to the inequalities from the previous slide:

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

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:

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

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

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

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

$$\widetilde{\Psi}_1\left(\rho\right) \equiv \left\{ \begin{array}{l} \left(1 - h_1\right)^{-1} \left[\Psi_1\left(\rho\right) - h_1\right] \text{ if } \rho > \rho_1^* \\ 0 \text{ if } \rho \leq \rho_1^* \end{array} \right.$$

Probability Distribution of Spell Lengths

The distribution of (standardized) mean beliefs after a second draw

- Appealing to (1) and (2), for workers taking another draw:

$$\begin{aligned}\gamma_{m2} &= (\alpha+1)(\alpha+2)^{-1}\gamma_{m1} + (\alpha+1)^{-1}(\xi_m + \sigma\epsilon_{mt}) \\ &= \gamma_{m1} + \sigma(\alpha+1)^{-1/2}(\alpha+2)^{-1/2}\epsilon'''\end{aligned}$$

  where ϵ''' is standard normal, and the second line follows the same logic as in slide 21.

- Hence $\Pr\{\rho_2 \leq \rho \,|\epsilon'''\}$, the probability distribution of ρ_2 of one-period survivors conditional on ϵ'''. is:

$$\begin{aligned} & \Pr\left\{\gamma_{m2} \leq \gamma + \sigma\rho \,\middle|\epsilon''', \right\} \\ = \; & \Pr\left\{\gamma_{m1} + \sigma(\alpha+1)^{-1/2}(\alpha+2)^{-1/2}\epsilon''' < \gamma + \sigma\rho \,\middle|\epsilon''', \right\} \\ = \; & \Pr\left\{\rho_1 < \rho - \sigma(\alpha+1)^{-1/2}(\alpha+2)^{-1/2}\epsilon''' \,\middle|\epsilon''', \right\} \\ = \; & \widetilde{\Psi}_1\left[\rho - \sigma(\alpha+1)^{-1/2}(\alpha+2)^{-1/2}\epsilon'''\right]\end{aligned}$$

Probability Distribution of Spell Lengths

Recursively computing the distribution of normalized match qualities

- Margining over ϵ''' and appealing to the definition of $\widetilde{\Psi}_1(\rho)$ now yields:

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

- 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

- Suppose the information and discount factors depend on X_n, some individual socio-economic factors;

$$\begin{aligned} \alpha_n &\equiv AX_n \\ \beta_n &\equiv BX_n \end{aligned}$$

  where A and B are the structural parameters to be estimated. Then the likelihood is:

$$L_N(A,B) \equiv \prod_{n=1}^{N} \sum_{\tau \in \rho(n)} p_\tau(AX_n, BX_n)$$

- Briefly, the structural estimates show that:
  1. individuals care about the future and the information value from 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.

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. Wage contacting: Pastorino (2014)
    3. College attrition: Arcidiacono, Aucejo, Maurel and Ransom (2016)
    4. Entrepreneurship: Hincapie (2020)
    5. Task assignment: Golan, James and Sanders (2021)
- 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.