Lecture 6: Topics in Intertemporal Labor Supply

a. the extensive margin

b. "involuntary" unemployment and the separation of the supply and demand sides

c. structural models

Some References

extensive margin:

Marco Bianchi, Björn R. Gudmundsson, Gylfi Zoega. "Iceland's Natural Experiment in Supply-Side Economics." AER 91(5) 2001, pp. 1564-1579.

Raj Chetty, Adam Guren, Day Manoli, and Andrea Weber. "Does Indivisible Labor Explain the Di¤erence between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities. NBER Macroeconomics Annual 2012.

Day Manoli and Andrea Weber. "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions. July 2013. Available at http://weber.vwl.uni-mannheim.de/2476.0.html

involuntary unemployment:

John Ham and Kevin T. Reilly, "Testing Intertemporal Substitution, Implicit Contracts, and Hours Restriction Models of the Labor Market Using Micro Data," AER 2002 92(4), pp.905-927,

structural labor supply:

Pierre Olivier Gourinchas and Jonathan Parker. (2002) "Consumption Over the Life Cycle" *Econometrica* 70 (No. 1):47-89.

Michael Keane and Kenneth Wolpin. "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Calro Evidence. *Review of Economics and Statistics 76 (November 1994): 648-672.*

Michael Keane, Petra Todd and Kenneth Wolpin. "The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications. Chapter 4 in Handbook of Labor Economics Volume 4a.

The Extensive margin

A lot of the labor supply literature ignores the "extensive margin" – workers who don't work for a year are dropped. However, variation in the number of workers is potentially important for understanding aggregate movements in hours:

(a) some people do miss an entire year of work in downturns

(b) the elasticity of participation w.r.t. wages can be relatively high, even if η is small.

There is a literature in macro arguing that the extensive margin is highly elastic, and that the extensive margin needs to be taken into consideration in both tax policy analysis and in macro modeling (see Chetty et al for a discussion of this literature).

Chetty et al present a "meta analysis" of various quasi-experimental studies that measure the effects of either permanent changes in (after tax) wages, or temporary changes, on employment rates. They use the former to obtain estimates of compensated elasticities of "participation"; the latter provide estimates of the Frisch elasticities of participation. An interesting paper is the one by Bianchi et al (2001), on the effects of a "tax holiday" created in Iceland when the country switched tax systems and everyone was untaxed for a single year (1987). You may find it instructive to read the paper because it is almost impossible to understand what the original authors did (or why), despite the very clear research design.

Looking at Chetty et al's Table 1, notice that the typical compensated elasticity is around 0.25, while the typical Frish elasticity is around 0.3. These are not much different than the elasticities people have obtained for the "intensive" margin.

Manoli and Weber (2013) is a very recent attempt to look at one of the important "extensive margins" : variation in the length of time people work. This paper uses an RD design to study the effects of a benefit that is paid to workers who retire after certain tenure "milestones". Since workers start jobs at different ages, there is a smooth distribution of people across the tenure distribution at different ages, and Manoli and Weber find strong evidence that some workers appear to delay retirement to get the benefits. They use a variant of the "bunching" style estimator we discussed in Lecture 3 to relate the fraction of people who retire at the threshold point to the relative size of the extra severance payment available for those who reach the threshold. Their estimated elasticities are somewhat larger than most of those in Chetty et al's table, with a value of around 0.6.

"Involuntary" Unemployment

Ham and Reilly build on an earlier paper by Ham (ReStud, 53 (4), 1986) which asks whether "signals" from the demand side affect hours, controlling for wages and other factors. In a simple neoclassical model "market model"

$$\begin{aligned} h &= h^d(w,x) \\ &= h^s(w,y) \end{aligned}$$

where h^d and h^s are the demand and supply functions for hours (by some group of workers), and x and y represent demand and supply shocks. The effect of demand shocks on supply choices works through w: the two sides of the market both make independent decisions, taking w as given. Thus, a test of the standard model is to fit the supply function and include x directly in the estimating equation. This requires that there be instruments for w in addition to the demand shock variables - so one interpretation of their test is that they are testing whether one set of demand shock variables affect supply, when wages are instrumented with other variables.

Formally, H-R consider two specifications. Their first set of models use first differenced labor supply model of the type we presented in Lecture 5:

$$\Delta \log h_t = \Delta A_t + \eta \Delta \log w_t + \delta \xi_t - \delta (r_{t-1} - \rho) - \delta \phi_t. \tag{1}$$

Their idea is to include an extra set of explanatory variables: the changes in the unemployment rates for the industry and occupation that the agent was working in in the base year $(\Delta UR^{ind}, \Delta UR^{occ})$. These are treated as potentially "endogenous" because they may reflect the "news" shocks incorporated in the innovation in the log marginal utility of consumption, ξ_t. They also present models with future wage changes $(\Delta \log w_{t+1})$ included on the right hand side, as a potential way to incorporate non-separable preferences (basically, if people forsee high wages ahead they may work more or less this period) See Table 1 of their paper.

H-R's second specification builds on Altonji's idea of controlling directly for consumption. Recall from Lecture 5 that the baseline specification is:

$$\log h_t = (A_t - \frac{\eta}{\kappa} B_t) + (\eta - \theta \frac{\delta}{\kappa}) \log w_t + \frac{\delta}{\kappa} \log c_t + e_{1t} - \frac{\delta}{\kappa} e_{2t}.$$

In this case they augment the model with (UR^{ind}, UR^{occ}), and include specifications with future wages. See tables 2 and 3, which use PSID and CES data.

Their key finding is that predictable movements in ΔUR^{ind} and ΔUR^{occ} (or in the levels of UR^{ind}, UR^{occ}),have a lot of explanatory power. They interpret this as evidence that wages are not "fully sufficient" to translate all the necessary information about the state of the demand side to the worker.

Structural Methods

The idea of "fully structural" modelling is to estimate the parameters of the utility function that drives choices within and between period. Some advantages of this approach:

1) the model can be solved for the value of the marginal utility of wealth for an agent in a given period, conditional on the state variables he or she sees at that point. This makes it possible to assess the wealth effects of wage changes, and the net effect (via intertemporal substitution and wealth effects) on labor supply

2) the model can be used to assess "out of sample" policy changes, like a revision in social security, on outcomes at all stages of the lifecycle

There are also some costs:

3) because of computational complexity many simplifications have to be made.

4) it is often very hard to understand where identification is "coming from" - in most cases parameters are identified by a combination of functional form assumptions and general features of the data. There is rarely "local identification" based on specific design features, as occurs in IV or RD approaches to estimation of simpler 'reduced form' models

A basic example.

We will discuss a simple dynamic labor supply model that illustrates the idea of interpolation of the value function (or, actually the derivative of the value function) using a regression approximation. To keep things very simple, we will assume that wages take on only a limited set of values (say $w_1, w_2...w_J$) and $\pi_{ij} = P(w_t = w_i | w_{t-1} = w_j)$ are known. There will be two state variables: the wage, and assets. The value function at time t will be denoted $V_t(A_t, w_t)$. When the wage takes on only discrete values this is just a set of J functions $V_t(A_t, w_j)$. What is relevant for dynamic consumption and hours choices are the derivatives $\partial V_t(A, w_j)/\partial A = \lambda_t(A, w_j)$. The solution method will involve working backward from the retirement period, and at each period solving for the optimal choices of consumption and hours in that period, as a function of the wage in that period, assets, and the approximations to $\partial V_{t+1}(A, w_j)/\partial A$. With these in hand we can then compute $\partial V_t(A, w_j)/\partial A$ at each of a finite set of values for A. We will then fit a regression model to these points to get an approximating model for $\partial V_t(A, w_j)/\partial A$ at every level of A. We then continue working backward to obtain the optimal consumption and hours functions in each period for each wage and level of assets,

$$c_t^*(A_t, w_t)$$
$$h_t^*(A_t, w_t).$$

In applications these functions can be used to compute a likelihood for the observed data for a sample of people who are observed at various points in time, or to compute hours and consumption profiles that are matched to observed profiles. We defer a discussion of how to use the estimated optimal response functions till the end of the lecture.

Let's assume the within period utility function is separable:

$$U(c, h) = u(c) - d(h).$$

with $d(0) = 0$. Let's also assume that agents work until an exogenous age R, then retire. At that point the agent becomes eligible for a pension p. In addition to the pension amount, an agent with (beginning-of-period) wealth A_R buys an annuity and receives a per-period payment of rA_R for the rest of his/her life. For purposes of modeling labor supply at earlier ages we can therefore consider the value function for period R :

$$V_R(A_R) = \sum_{j=0}^{\infty} \frac{U(p + rA_R, 0)}{(1+r)^j} = \frac{1}{r} u(p + rA_R)$$

where $U(c, h)$ is the within-period utility function, and I have simplified things by assuming that the agents' discount rate and the annuity price are equivalent (with separable preferences this means that the agent wants to set consumption constant for all remaining periods). A similar setup is used by Gourinchas and Parker (2002). Note that the function $V_R(A_R)$ inherits properties from $u(.)$, so if u depends on some parameter τ then the same parameter shifts V_R.

Now let's go back to period $R-1$. In this period the agent faces a wage w_{R-1}, and has assets A_{R-1}. The value function for this period is

$$V_{R-1}(A_{R-1}, w_{R-1}) = \max_{c_{R-1}, h_{R-1}} u(c_{R-1}) - d(h_{R-1}) + \frac{1}{1+r} [\frac{1}{r} u(p + r(1+r)(A_{R-1} + w_{R-1}h_{R-1} - c_{R-1})].$$

Note that there is no uncertainty left once we get to $R-1$. So we can solve for the optimal choice in this period very easily, to get a "starting value function" for our backward recursion.

The f.o.c.'s for period $R-1$ are:

$$\begin{aligned} u'(c_{R-1}) &= \lambda_{R-1} = u'(p + r(1+r)(A_{R-1} + w_{R-1}h_{R-1} - c_{R-1})) \\ d'(h_{R-1}) &= \lambda_{R-1} w_{R-1}. \end{aligned}$$

Now lets assume

$$\begin{aligned} d(h) &= \frac{1}{1+1/\eta} h^{1+1/\eta} \\ u(c) &= \log c \end{aligned}$$

so the f.o.c. for hours implies:

$$h_{R-1} = w_{R-1}^{\eta} \ c_{R-1}^{-\eta},$$

which means optimal earnings in period $R-1$ are

$$w_{R-1}h_{R-1} = w_{R-1}^{1+\eta} \ c_{R-1}^{-\eta}$$

Now all we have to do is find an optimal choice for c_{R-1}. Equating marginal utility of consumption in period $R-1$ and R means that the levels of consumption are equal, so:

$$\begin{aligned} c &= p + r(1+r)(A_{R-1} + w_{R-1}^{1+\eta} \ c^{-\eta} - c) \\ \Rightarrow & \quad c = \frac{r(1+r)}{1+r(1+r)} A_{R-1} + \frac{1}{1+r(1+r)} p + \frac{r(1+r)}{1+r(1+r)} w_{R-1}^{1+\eta} c^{-\eta} \end{aligned}$$

This has to be solved numerically. It has the form

$$c = f(c) = k + \gamma c^{-\eta}$$

and notice that k is pretty big and γ is small. Its not hard to solve this by iterative methods.[1] With this we have now obtained numerically

$$c_{R-1}^{*}(A_{R-1}, w_{R-1})$$

(this also depends on η, p, r). We can then obtain $h_{R-1}^{*}(A_{R-1}, w_{R-1})$.

Now notice that

$$\partial V_{R-1}(A, w_{R-1})/\partial A = \lambda_{R-1}^{*}(A_{R-1}, w_{R-1}) = \frac{1}{c_{R-1}^{*}(A_{R-1}, w_{R-1})}.$$

[1] I used this method: start with the initial guess $c_1 = k$. Now $f(c) = f(c_1) + (c - c_1) f'(c_1)$, so setting $c = f(c)$ gives a new guess

$$c_2 = \frac{f(c_1) - c_1 f'(c_1)}{1 - f'(c_1)}.$$

This converges in 3-4 iterations.

This is the function we are going to need to take expectations over in solving for optimal choices at period $R-2$. In particular, if in period $R-2$ the wage is $w_{R-2}=w_i$ then we are going to need to calculate

$$E_{R-2}[\partial V_{R-1}(A,w_i)/\partial A]=\sum_j \frac{1}{c^*_{R-1}(A,w_j)}\pi_{ji},$$

treating A as an endogenous variable that depends on $c_{R-2}, w_{R-2}, h_{R-2}$, and A_{R-2}.

Our method is as follows. First, using the procedure above, we calculate $c^*_{R-1}(A,w_j)$ for a grid of values of A and each possible value of w_j. In a "test" program, I measured all monetary units in $1000's$ and assumed that the possible values for A are $1,2...1,000$ (i.e., up to a million). I assumed that w takes on values of $10,20....100$ (i.e., 10,000, 20,000... 100,000), and that $p=20$ (i.e., 20,000). Then I formed a simple $n^{th}-order$ polynomial approximation:

$$\frac{1}{c^*_{R-1}(A,w_j)}=b_{0j}+b_{1j}A+b_{2j}A^2+...b_{nj}A^n$$

For my test program I found that $n=4$ gets an extremely good fit. Now notice that once we have these coefficients, the expected derivative of the $R-1$ value function is:

$$\begin{aligned}
E_{R-2}[\partial V_{R-1}(A,w_i)/\partial A] &= \sum_j (b_{0j}+b_{1j}A+b_{2j}A^2+...b_{nj}A^n)\pi_{ji} \\
&= \sum_j b_{0j}\pi_{ji}+\sum_j b_{1j}\pi_{ji}A+...+\sum_j b_{nj}\pi_{ji}A^n \\
&= b_0^i+b_1^iA+b_2^iA^2+...+b_n^iA^n
\end{aligned}$$

where the coefficients $b_0^i, b_1^i...b_n^i$ depend on the wage in $R-2$ via the "weights" π_{ji}. Notice the benefit of having a discrete first-order process for wages: given the J approximating polynomials, all we have to do to form the expectation for a given wage in $R-2$ is weight the approximating polynomials by the appropriate transition probabilities.

Now we are ready to solve the optimal choices for c and h in $R-2$. Specifically, the Bellman equation is:

$$V_{R-2}(A_{R-2},w_{R-2})=\max_{c_{R-2},h_{R-2}} u(c_{R-2})-d(h_{R-2})+\frac{1}{1+r}E_{R-2}[V_{R-1}(A_{R-1},w_{R-1}|w_{R-2})].$$

And the f.o.c. are:

$$\begin{aligned}
u'(c_{R-2}) &= \lambda_{R-2}=E_{R-2}[\partial V_{R-1}(A_{R-1},w_{R-1}|w_{R-2})/\partial A_{R-1}] \\
d'(h_{R-2}) &= \lambda_{R-2}w_{R-2} \\
&\Rightarrow h_{R-2}=w_{R-2}^{\eta}\ c_{R-2}^{-\eta} \\
&\Rightarrow w_{R-2}h_{R-2}=w_{R-2}^{1+\eta}\ c_{R-2}^{-\eta}
\end{aligned}$$

So we need to solve

$$\frac{1}{c_{R-2}} = b_0^i + b_1^i A + b_2^i A^2 + ... + b_n^i A^n$$

where

$$A = (1+r)(A_{R-2} + w_{R-2}^{1+\eta}\ c_{R-2}^{-\eta} - c_{R-2}).$$

Thus for each value of A_{R-2} and each possible value of the wage w_i we need to solve the root of the function $g(c; A_{R-2}, w_i)$, where:

$$g(c; A_{R-2}, w_i) = \frac{1}{c} - \sum_k b_k^i((1+r)(A_{R-2} + w_i^{1+\eta}c^{-\eta} - c)^k = 0.$$

Again, a numerical solution is needed.[2] The solution is

$$c_{R-2}^*(A_{R-2}, w_{R-2})$$

(which also depends on η, p, r). We can then get $h_{R-2}^*(A_{R-2}, w_{R-2})$.

Finally, going backward one step we will need to evaluate

$$E_{R-3}[\partial V_{R-2}(A_{R-2}, w_{R-2})/\partial A_{R-2}|w_{R-3} = w_i] = \sum_j \frac{1}{c_{R-2}^*(A_{R-2}, w_j)}\pi_{ji}.$$

Thus we can proceed backwards, by estimating the approximating polynomial functions and repeating the previous steps.

Some comments:

1) Notice in this algorithm, everything is summarized by the approximating polynomial coefficients for $\lambda_t^*(A_t, w_j)$. For example, if we use a fourth order polynomial, and have 10 possible wage values, the relevant information for period t (given the transition matrix elements π_{ij}, and the parameters η, p, r) is contained in 50 numbers. The algorithm proceeds by getting the numbers sequentially from $R-2$ back to some earliest possible period ($e.g., R-40$).

2) We could introduce tastes in one of several ways. One way is to allow the marginal utilities of consumption or leisure to change with age in some way, e.g.,

$$d_t(h_t) = f(t)\frac{1}{1+1/\eta}h_t^{1+1/\eta}$$

where $f(t)$ is a simple function like $f(t) = \exp(\upsilon t)$. For a given value of υ it is possible to solve for the optimal consumption and hours functions in each

[2] The standard method is Newton-Raphson. Recall that if you are trying to find a c such that $g(c) = 0$ you can normally start with an initial guess c_1 and iterate: $c_j = c_{j-1} - g(c_{j-1})/g_c(c_{j-1})$. In the case where we are approximating the marginal utility of income with polynomials, the analytical derviative is easy.

period, and then search for a "best fitting" choice. Another way is to assume there are discrete types $\upsilon \in \{\upsilon_1, \upsilon_2, ...\upsilon_K\}$, and assume

$$d_k(h) = \exp(\upsilon_k) \frac{1}{1 + 1/\eta} h^{1+1/\eta}$$

Then we have to solve the problem for each "type",and think of how to map the behavior we see into an average across the types.

3) How do we get the π_{ij} elements?

Suppose that we want to approximate a first order serially correlated continuous process by a 1st order Markov process. G. Tauchen (1986 Economics Letters) described a simple algorithm. For example, suppose we want to approximate an AR-1 wage process:

$$w_t = a + \rho w_{t-1} + \epsilon_t$$

where $\epsilon_t \sim N(0, \sigma^2)$. Note that for this process $E[w_t] = \mu_w = a/(1-\rho)$, and $var[w_t] = \sigma_w^2 = \sigma^2/(1-\rho^2)$. To approximate this with a discrete 1st order markov model with N points of support, first find $N-1$ cut points k_j $(j = 1, ..N-1)$ such that

$$\Phi[\frac{k_{j+1} - \mu_w}{\sigma_w}] - \Phi[\frac{k_j - \mu_w}{\sigma_w}] = \frac{1}{N}$$

with $k_0 = -\infty$, and $k_N = \infty$. (This defines the boundaries so that the probability a draw from $N(\mu_w, \sigma_w^2)$ falls in each bin is 1/N). Next, find the mean value of a $N(\mu_w, \sigma_w^2)$ within each bin. These values will be the points of support for the discrete process. If $\rho = 0$ we can stop. Otherwise, the last step is to define "transition probabilities" π_{ij} such that

$$\pi_{ij} = P(k_i < w_t < k_{i+1} | k_j < w_{t-1} < k_{j+1})$$

assuming that

$$\begin{pmatrix} w_{t-1} \\ w_t \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_w \\ \mu_w \end{pmatrix}, \sigma_w^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right)$$

This can be computed using the usual formulas (e.g. in Johnson and Kotz) (or using simple simulation methods).

4) How do we use the optimal consumption and hours functions, $c_t^*(A_t, w_t), h_t^*(A_t, w_t)$?

A huge obstacle to micro research on consumption and labor supply is the absence of reliable data on assets. For example, the well known structural study of retirement by Rust and Phelan, "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets" Econometrica 65(July 1997), **assumes no savings**, in part because of the low quality of the asset information in their data set. As a result, almost no studies have tried to estimate structural labor supply models that are directly based on observed data on consumption, hours, wages, and assets. One of the few is Imai and Keane, IER 2004, which solves the problem by evaluating the value function at a discrete number of points and interpolating (rather than interpolating the marginal utility of wealth function). Imai and Keane allow for mismeasurement in assets and hours.

TABLE 1
Extensive Margin Elasticity Estimates from Quasi-Experimental Studies

Study	Elasticity	Standard Error	Population and Variation
A. Steady State (Hicksian) Elasticities			
1. Juhn, Murphy, and Topel (1991)	0.13	0.02	Men, skill-specific trends, 1971-1990
2. Eissa and Liebman (1996)	0.30	0.10	Single Mothers, U.S. 1984-1990
3. Graversen (1998)	0.24	0.04	Women, Denmark 1986 tax reform
4. Meyer and Rosenbaum (2001)	0.43	0.05	Single Women, U.S. Welfare Reforms 1985-1997
5. Devereux (2004)	0.17	0.17	Married Women, U.S. wage trends 1980-1990
6. Eissa and Hoynes (2004)	0.15	0.07	Low-Income Married Men & Women, U.S. EITC expansions 1984-1996
7. Liebman and Saez (2006)	0.15	0.30	Women Married to High Income Men, U.S. tax reforms 1991-97
8. Meghir and Phillips (2010)	0.40	0.08	Low-Education Men, U.K. wage trends, 1994-2004
9. Blundell, Bozio, and Laroque (2011)	0.30	n/a	Prime-age Men and Women, U.K., tax reforms 1978-2007
Unweighted Mean	**0.25**		
B. Intertemporal Substitution (Frisch) Elasticities			
10. Carrington (1996)	0.43	0.08	Full Population of Alaska, Trans-Alaska Pipline, 1968-83
11. Gruber and Wise (1999)	0.23	0.07	Men, Age 59, variation in social security replacement rates
12. Bianchi, Gudmunndsson, and Zoega (2001)	0.42	0.07	Iceland 1987 zero tax year
13. Card and Hyslop (2005)	0.38	0.03	Single Mothers, Canadian Self Sufficiency Project
14. Brown (2009)	0.18	0.01	Teachers Near Retirement, California Pension System Cutoffs
15. Manoli and Weber (2011)	0.25	0.01	Workers Aged 55-70, Austria severance pay discontinuities
Unweighted Mean	**0.32**		

Notes: This table reports elasticities of employment rates with respect to wages, defined as the log change in employment rates divided by the log change in net-of-tax wages. Where possible, we report elasticities from the authors' preferred specification. When estimates are available for multiple populations or for multiple specifications without a stated preference among them, we report an unweighted mean of the relevant elasticities. See Appendix B for details on sources of estimates.

Fig. 1. Payment Amounts based on Tenure at Retirement

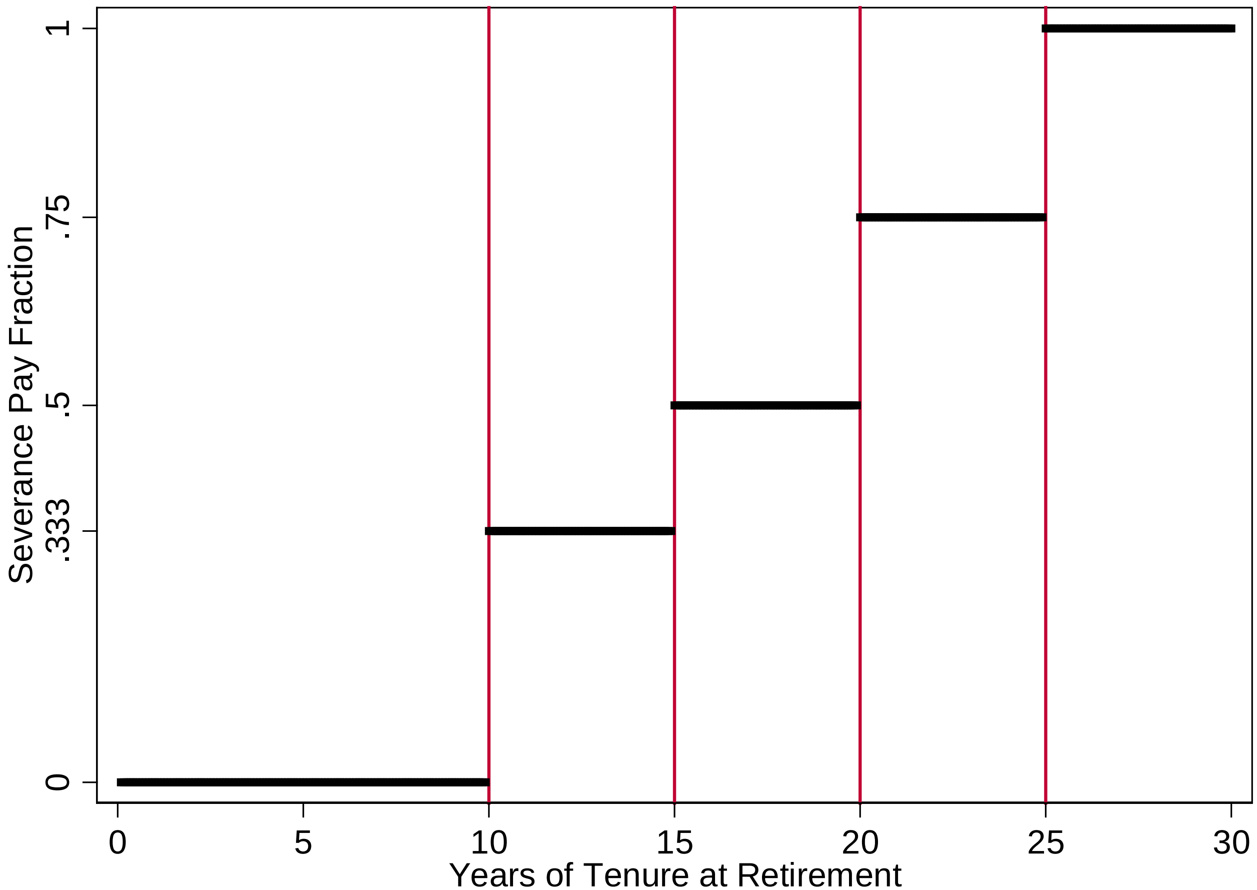

Notes: There are two forms of government-mandated retirement benefits in Austria: (1) government-provided pension benefits and (2) employer-provided severance payments. The employer-provided severance payments are made to private sector employees who have accumulated sufficient years of tenure by the time of their retirement. Tenure is defined as uninterrupted employment time with a given employer and retirement is based on claiming a government-provided pension. The payments must be made within 4 weeks of claiming a pension according to the following schedule. If an employee has accumulated at least 10 years of tenure with her employer by the time of retirement, the employer must pay one third of the worker's last year's salary. This fraction increases from one third to one half, three quarters and one at 15, 20 and 25 years of tenure respectively. Since payments are based on an employee's salary, overtime compensation and other non-salary payments are not included when determining the amounts of the payments. Provisions to make these payments come from funds that employers are mandated to hold based on the total number of employees. Severance payments are also made to individuals who are involuntarily separated (i.e. laid off) from their firms if the individuals have accumulated sufficient years of tenure prior to the separation. The only voluntary separation that leads to a severance payment, however, is retirement. Employment protection rules hinder firms from strategically laying off workers to avoid severance payments and there is no evidence on an increased frequency of layoffs before the severance pay thresholds.

Fig. 3. Distribution of Tenure at Retirement, Full Sample

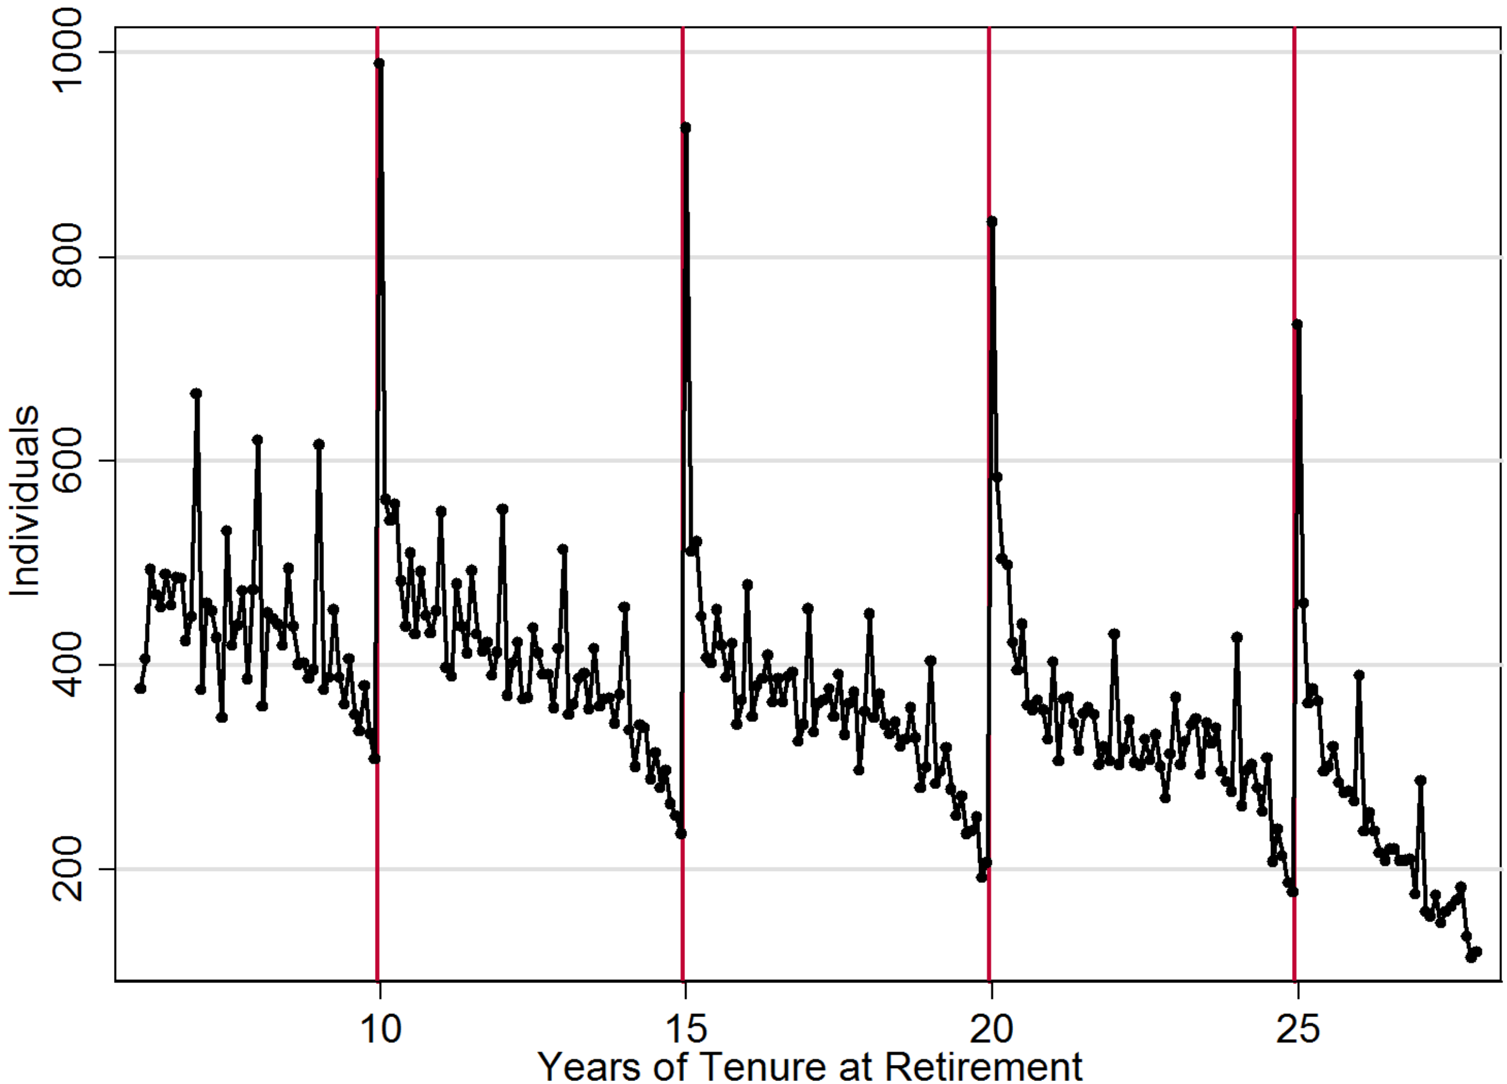

Notes: This figure plots the distribution of tenure at retirement at a monthly frequency. Each point captures the number of people that retire with tenure greater than the lower number of months, but less than the higher number of months. Tenure at retirement is computed using observed job starting and job ending dates. Since firm-level tenure is only recorded beginning in January 1972, we restrict the sample to individuals with uncensored tenure at retirement (i.e. job starting after January 1972).

Fig. 12. Estimating the Changes in Retirements

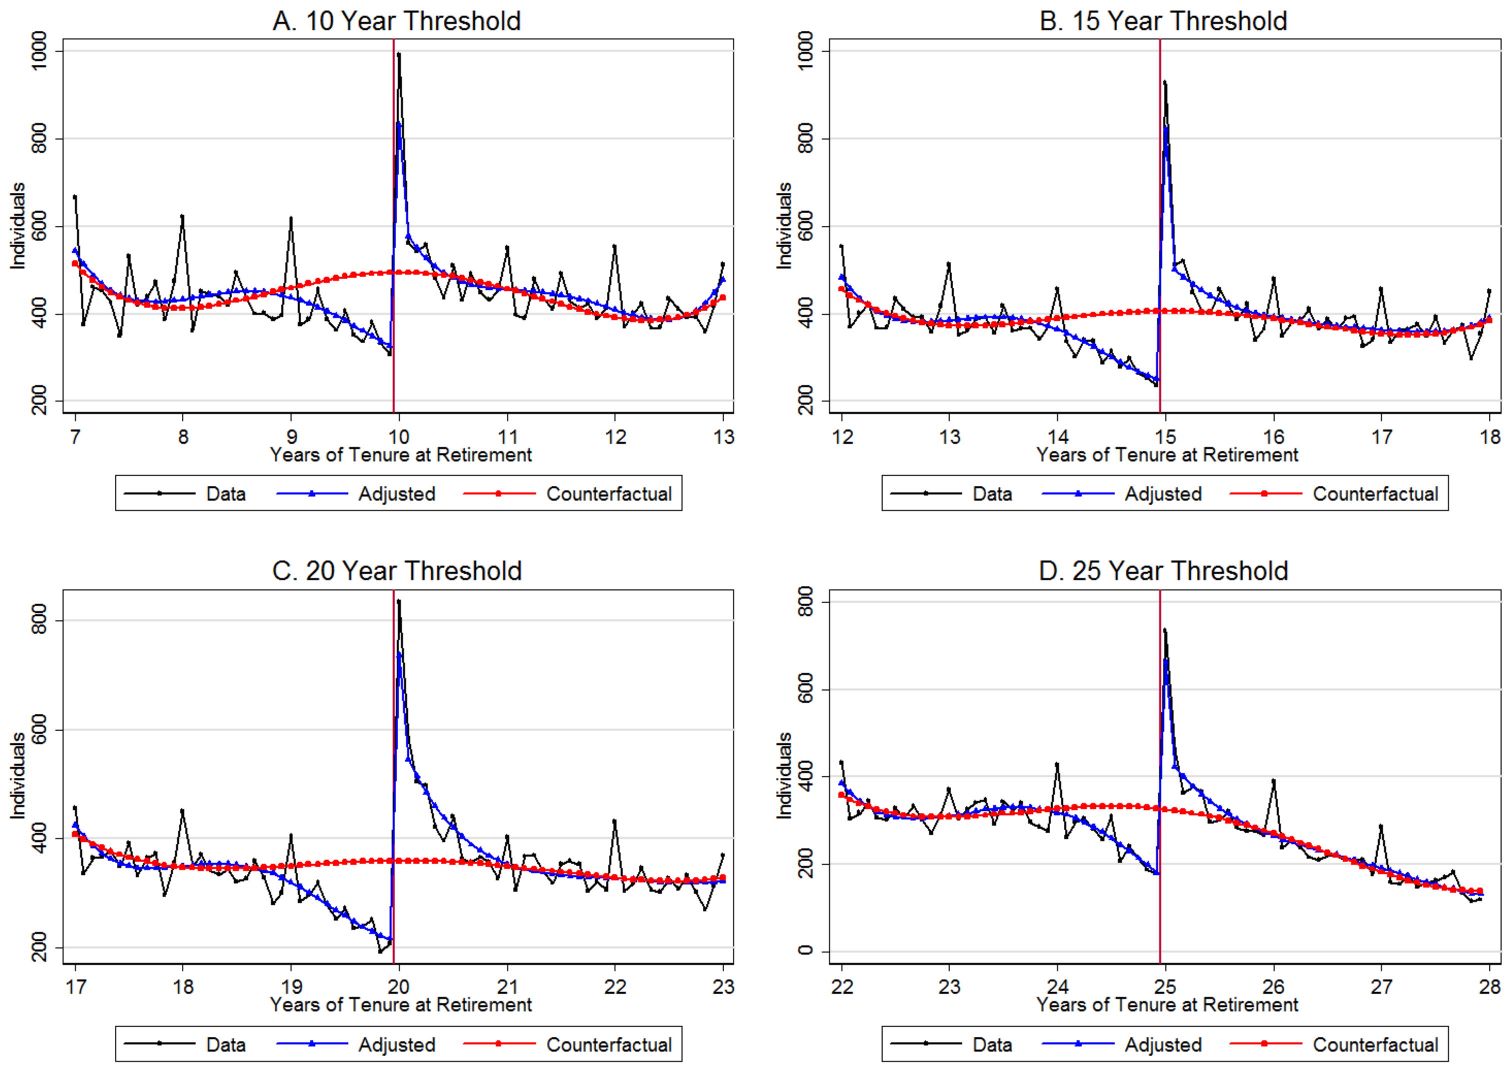

Notes: This figure combines plots for the observed retirement frequencies (black squares), the seasonally adjusted retirement frequencies (blue triangles) and the counterfactual retirement frequencies (red circles).