Class notes. Econ 741b.

François Gourio*

Part I
Lecture 1

Outline: (1) Organization, (2) Facts on Firm Dynamics, (3) Simple Viner model, (4) "Dynamic Viner", (5) Adjustment cost model, (6) Hopenhayn 1992, starting with a static example, (7) Hopenhayn-Rogerson 1993.

1. Organization

(0) Syllabus, office hours, books (Ljungqvist-Sargent, Stokey Lucas Prescott).

(1) Speak up in class. Don't hesitate to make comments, ask questions, ...

(2) Accent, writing - just have me repeat.

(3) Techniques: dynamic programming, stationary distributions, numerical computation. Hopefully in some cases we'll be able to study models with aggregate shocks.

(3) My specialties (business cycles, investment, asset pricing) vs. the class: I am trying to learn this material because I think there will be some interesting cross-breeding.

(4) Grading: replication and/or modification of a paper, on the reading list below, or approved by me. This could be either empirical or numerical. You will need to write report explaining what you did and your results. You will also have to provide me with your code/data. You can work in groups, but the size of a group is limited to 2. **It is required that you come discuss your project**

*Boston University, Department of Economics. Email fgourio@bu.edu, Phone 617-353-4534, office #400.

with me before November 17. No need for a long talk, just to make sure this is doable, if there are some not-too-complicated extensions, and some hints perhaps on how to do it. For now, let's say that you will be able to turn in these reports on 12-17 at the latest.

(5) General motivation: the typical macro model looks at production either through a representative production function $Y = AF(K, L)$, or through Dixit-Stiglitz monopolistic competition. Sometimes researchers add adjustment costs.

We'll try to dig a bit deeper and look at the individual producers, i.e. go beyond the representative agent assumption. Hopefully this will give us a better idea of how production and productivity change over time or across countries.

Lots of heterogeneity in productivity as we will see. Hence reallocation may be important. And reallocation varies over the business cycle.

(6) We're going to cover a lot of models. I'd like that we try and think, for each model: (1) what is the definition of the firm (or plant) in this model? (2) What kind of cross-sectional distributions does this model generate? (3) Where does the heterogeneity come from? (4) What kind of interesting comparative statics does the model deliver? (5) Are there interesting testable empirical implications? Time-series or cross-section? Aggregate, industry or micro-level? (6) How is this interesting for macroeconomics? In some cases our models will have constant aggregates, which is why the last question seems appropriate.

(7) These notes: probably contain errors or typos. Please raise "issues" if you see any. Some notes in brackets are really reminders for myself.

2. Facts on Firm and Plant Dynamics

Later in the class we'll talk about job flows and worker flows.

Sources of data that I know: Compustat, Census bureau (LRD), Industry-level data (NIPA and BLS, NBER manufacturing database), Annual survey of manufactures from Census.

Some empirical papers: Bartelsman-Doms (JEL 2000), Dunne-Roberts-Samuelson (QJE 1989, Rand 1989), Doms-Dunne (on investment, RED 1998). See also Davis and Haltiwanger later.

Facts regarding firm dynamics

Fact #1: The distribution of the size of firms is skewed to the right, with fat tails. In particular it is well approximated by a Pareto distribution: $\Pr(size \leq x) =$

$1 - kx^{-\gamma}$ for $x \geq$some threshold. This can be rewritten as:

$$\log \Pr(size \geq x) = \log k - \gamma \log x.$$

If $\gamma = 1$ this is called "Zipf Law". The lognormal is a less precise approximation. Plot of a firm's log rank against its log size (in sales, or employment): you get a straight line. There is a large literature asking why this distribution arises. See below for a short explanation. I attach below a graph I did with Compustat (It seems that this law doesn't work very well in Compustat, so this example is actually not the most convincing one.) It's unclear to me whether this law works on the whole distribution or only for firms large enough - Rossi-Hansberg and Wright have a plot which covers a large range of the distribution.

Fact #2: Small firms tend to grow faster than large firms. (The usual references are Hall 1987 and Evans 1987 JPE.)

Fact #3: Small firms have a low probability of survival.

Note that Fact 2 is a mixture of two things: Fact 3 says small firms die more than large firms, but the ones that do survive grow a lot faster than large firms, hence on average small firms do tend to grow faster.

Fact #4: Entry and exit rates are large. There is a lot of simultaneous entry and exit. These entry and exit rates are correlated across industry: there are high and low turnover industries.

This fact is very important. It contrasts sharply with the simple Viner model below, where there is either entry or exit in an industry at a given point in time, depending on the state of demand as compared to supply. In the data there is a large, continuous restructuring with many firms appearing, growing and dying - a form of "creative destruction". This suggests we need to think of firm-level shocks as more important than industry-level shocks.

Hopenhayn (1992): "Over a five-year intervals, 30% of jobs and 40% of firms disappear in manufacturing."

Facts regarding the productivity of plants

[See survey by Bartelsman and Doms in JEL 2000.]

Fact #6: there is a large heterogeneity in measured productivity between plants.

A typical number: within the same 4-digit industry (i.e., a rather well-defined product), the producer at the 75th percentile has a labor productivity which is twice the one of the producer at the 25th percentile. Looking at the tails yields bigger numbers, e.g. 90/10: 4 to 1. Looking at TFP yields smaller numbers, e.g.

the 75/25 ratio is around 1.4. It is tempting to think that these differences are due to differences in (mismeasured) capital intensity or vintages, but:

Fact #7: a large share of these productivity differentials is not explained by observable characteristics.

To understand why this is "shocking", remember that in your basic micro model, price equals marginal cost, and thus marginal cost must be the same in all firms. High costs producers are not supposed to survive in equilibrium - they will be undercut (in prices) by the more productive producers. But (1) these studies measure *average,* not marginal, productivity or cost; and (2) one can easily imagine an equilibrium where firms make some investment, which increases their productivity up to some output level. Then we would see ex-post some firms having higher productivity, but this higher productivity would compensate for an initial investment. (It may be a good exercise to write down such a model.) The problem with this story is that we should find some trace of this investment, and I do not know that anybody has found this. Finally, adding a bit to the mystery:

Fact #8: these productivity differentials are persistent.

There remains the possibility that some of these productivity differentials are just measurement error.

Derivation of Gibrat's law⇒Zipf Law:

Assume $X_{t+1} = X_t \varepsilon_{t+1}$ with ε_{t+1} iid with c.d.f. F and mean 1 (so that there is no trend to growth, otherwise the distribution will keep shifting to the right). Let $G_t(s) = \Pr(X_t \leq s)$. This cross-sectional c.d.f. satisfies the law of motion:

$$G_{t+1}(s) = \int_0^\infty F\left(\frac{s}{x}\right) dG_t(x).$$

This comes from:

$$\begin{aligned} G_{t+1}(s) &= \Pr(X_{t+1} \leq s) \\ &= \int_0^\infty \Pr(X_{t+1} \leq s \mid X_t = x) \Pr(X_t = x)\, dx \\ &= \int_0^\infty \Pr\left(\varepsilon_{t+1} \leq \frac{s}{x}\right) \Pr(X_t = x)\, dx \\ &= \int_0^\infty F\left(\frac{s}{x}\right) dG_t(x). \end{aligned}$$

A stationary cross-sectional distribution thus satisfies:

$$G(s) = \int_0^\infty F\left(\frac{s}{x}\right) dG(x). \tag{2.1}$$

It is possible to show that the only solution of this equation is the Pareto distribution with $\gamma = 1$ (Zipf Law). What I want to show is that the Pareto distribution is at least one solution. The reason I had some difficulty deriving this in class is due to the fact that it's not exactly true. More precisely, the Pareto distribution is $G(x) = 1 - \kappa x^{-\gamma}$ for $x \in [\kappa^{1/\gamma}, +\infty[$. This means that there will always be some probability that you leave the support of the distribution (i.e. go below $\kappa^{1/\gamma}$) if you have enough bad shocks ε. Hence what we show instead is that the Pareto distribution with $\gamma = 1$ is one solution in the tail, i.e. for large s, $G(s) = 1 - c/s$ is a solution for some c.

It seems simpler to rewrite equation (2.1) by integrating by parts first:

$$\begin{aligned}
G(s) &= \left[F\left(\frac{s}{x}\right)G(x)\right]_0^\infty + \int_0^\infty G(x)\frac{s}{x^2}f\left(\frac{s}{x}\right)dx \\
&= \int_0^\infty G(x)\frac{s}{x^2}f\left(\frac{s}{x}\right)dx \\
&= \int_0^\infty G\left(\frac{s}{v}\right)f(v)\,dv.
\end{aligned}$$

Now let's check that $G(s) = 1 - c/s$ works (not taking into account the range problems). The right-hand side is

$$\begin{aligned}
\int_0^\infty \left(1 - \frac{cv}{s}\right)f(v)\,dv &= \int_0^\infty f(v)\,dv - \frac{c}{s}\int_0^\infty vf(v)dv \\
&= 1 - \frac{c}{s},
\end{aligned}$$

which equals the left-hand side. A reference for more on this is Xavier Gabaix's work (MIT). The Pareto distribution has fat tails so big firms matter a lot. In particular with fat tails it may be impossible to apply a law of large numbers to say that idiosyncratic shocks cancel out. Gabaix works out some interesting consequences (see his paper on "Granular effects..."). There have also been some papers which try to write down richer models of firm dynamics which generate the Zipf Law and the small deviations from it. (Lutmer, Rossi-Hansberg and Wright are recent examples, Lucas 78 Bell Journal Econ: an older one.)

The bottom line is that Zipf holds approximately because Gibrat holds approximately. Neither really works fully. The way people "test" Zipf by running an OLS regression of log size on log rank and saying that the coefficient seems statistically confused (what's H_0 and what's the alternative?).

3. The static competitive model (Marshall-Viner)

Useful framework to keep in mind while discussing production theory and industry equilibrium.

The industry has an aggregate demand curve $Q^d = D(P)$. Assume firms are defined by a cost function $C(q)$ which is increasing and convex, and which has some fixed costs.[1] The typical example, to keep in mind, is $C(q) = \alpha + \beta q + \frac{\gamma}{2}q^2$. Given the price P, each firms solves the concave program

$$\pi(P) = \max_{q \geq 0} \{Pq - C(q)\},$$

yielding the output supply q defined implicitly by $C'(q) = P$.

The *short-run equilibrium* is defined, given the number N of firms in the industry, as a triplet (Q, q, P) of total industry output, output per firm and output price satisfying:

$$\begin{aligned} Q &= Nq, \\ Q &= D(P), \\ C'(q) &= P. \end{aligned}$$

The *long-run equilibrium* is defined as a quadruplet (N, Q, q, P) satisfying the short-run equilibrium conditions, plus the zero-profit condition:

$$\pi(P) = Pq - C(q) = 0.$$

One can draw the usual picture of a U-shaped average cost curve intersecting the marginal cost curve at its minimum. From $P = C(q)/q = C'(q)$, one gets that the equilibrium price is the minimum average cost. This determines output per firm q and total demand $Q = D(P)$. The number of firms then adjusts to make the two equal: $N = Q/q$. In other words, aggregate supply in this economy is perfectly elastic at price P. The size of a firm is fixed at its efficiency scale q given by the minimum average cost.

The effect of an increase in demand D:

- In the short-run, we have a rising supply curve so P increases, profits increase, and the average output of each firm increases as well.

[1] Fixed costs will generate a well-defined size. If there are only increasing marginal costs, firms would be infinitely small.

- In the long run, N adjusts, and the output per firm falls back to its efficient level, profits decrease back to zero, P falls and total quantity produced increases further ("long-run supply more elastic than short-run supply"). Hence, all the long-run response is along the extensive margin (more firms, not more output per firm).

[See Lucas 1978 Bell for an interesting discussion and the first span of control model.]

[Extension: heterogeneity in costs.]

4. A dynamic extension of the competitive model

The static model makes a simple distinction between the short-run and the long-run. However it is a limitation that firms, in this model, do not take into account the future when making decisions. Moreover we cannot talk about the transition between the two static equilibria. For these reasons, in the 1960s were developed adjustment costs model which allow to analyze the firm's dynamic problem in a consistent way.

We now assume that there is a second industry producing firms for our industry. (Since our firms are really production units or machines, one can think of the other industry as a capital-good producing sector.) This industry can produce in each small interval of time of length dt a flow $I_t dt$ of new firms at a cost $g(I_t)\, dt$. I assume that $g' > 0$ and $g'' > 0$. This reflects external adjustment costs: there is an increasing marginal cost in the capital goods sector.

To generate ongoing replacement of firms I assume that firms disappear at rate δ. Hence the law of motion for the number of firms:

$$\frac{dN_t}{dt} = I_t - \delta N_t.$$

The firm's static problem is to choose the production scale q_t :

$$\pi_t = \pi(P_t) = \max_{q_t \geq 0} \{P_t q_t - C(q_t)\}.$$

The firm's value is the present discounted value of its profits, or assuming a constant discount rate:

$$V_t = \int_t^\infty e^{-(r+\delta)(s-t)} \pi(P_s) ds.$$

Differentiating this equation with respect to t yields:

$$\frac{dV_t}{dt} = (r+\delta)V_t - \pi(P_t).$$

Now consider the firm-producing sector. At time t it can produce a flow I_t of new firms, and sell them at price V_t. Its problem is thus

$$\max_{I_t} \left\{ V_t I_t - g(I_t) \right\},$$

leading to the first-order condition $g'(I_t) = V_t$: this is an upward-sloping supply of new firms or machines.

The equilibrium is thus determined by:

$$\begin{aligned}
P_t &= C'(q_t), \\
Q_t &= N_t q_t, \\
Q_t &= D(P_t), \\
\frac{dN_t}{dt} &= I_t - \delta N_t, \\
g'(I_t) &= V_t, \\
\frac{dV_t}{dt} &= (r+\delta)V_t - \pi(P_t),
\end{aligned}$$

plus an initial condition N_0.

Note that the instantaneous equilibrium is given by:

$$q_t = D(P_t)/N_t = C'^{-1}(P_t) \rightarrow P_t = \phi\left(N_t\right)$$

Substituting yields a system differential equations:

$$\begin{aligned}
\frac{dN_t}{dt} &= g'^{-1}(V_t) - \delta N_t, \\
\frac{dV_t}{dt} &= (r+\delta)V_t - \pi(\phi\left(N_t\right)).
\end{aligned}$$

The only state variable in this economy is N_t. We thus have the usual saddle-path stable system with V_0 adjusting to keep the system on the saddle arm. [You can linearize if you want to compute the speed of converge to the steady-state, etc.]

Draw pictures.

This model delivers a smooth adjustment in the number of firms in the industry in response to a shock. Work out the dynamics of profits and value. Note that long-run supply is not infinitely elastic supply in this case.

What's the dynamic effect of a shift in demand? The effect of a change in r or δ?

5. An adjustment cost model with a representative firm

I assume that you are familiar with this, both in the deterministic and in the stochastic case. You're encouraged to look at the recursive version of Ljungqvist and Sargent (chapter 6 of first edition, first section) which explains the "small k, Big K" representation. Lucas and Prescott 1971 first studied a model with adjustment costs, rational expectations, but no heterogeneity and no entry or exit.

6. Hopenhayn (Econometrica 1992)

6.1. A quasi-static model with one-time entry

The process of entry: there is an infinite number of potential entrants. Pay cost c_{in} to enter the industry. Next draw a productivity level $\varphi \geq 0$ from a c.d.f. $G(.)$. Next you decide whether to stay in the industry or exit. Exiting costs c_{out}, which may be negative (some recoverable value).

If you stay you can produce using the technology $y = \varphi n^{\alpha}$ with $0 < \alpha < 1$. Can hire labor on a competitive market at exogenous price w. Must pay a fixed cost c_f each period while in operation.

Discount profit at rate β. No exit once the initial stage is played.

Demand side is standard: $Q = D(P)$.

Solve the model backward: first, if you decide to stay, profits are

$$\pi(P, \varphi) = \max_{n \geq 0} \left(P\varphi n^{\alpha} - wn - c_f \right),$$

which is increasing in φ. Hence there is a threshold $\varphi^*(P)$ such that a firm will decide to stay if $\varphi \geq \varphi^*(P)$, with φ^* defined implicitly by:

$$\pi(P, \varphi^*) = 0.$$

Total supply in the industry is:

$$Q^s(P) = \int_{\varphi^*}^{+\infty} y^*(P, \varphi) dG(\varphi),$$

where $y^*(P, \varphi) = \varphi n^*(P, \varphi)^{\alpha}$ and $\alpha\varphi n^*(P, \varphi)^{\alpha-1} = w$.

Next we examine the entry decision. The cost of entry must be compensated by ex-post profits. This yields the condition:

$$\sum_{t\geq 0}\beta^t\int_{\varphi^*}^{\infty}\pi(P,\varphi)dG(\varphi)+\int_0^{\varphi^*}-c_{out}dG(\varphi)=c_{in}$$

$$\int_{\varphi^*}^{\infty}\pi(P,\varphi)dG(\varphi)=(1-\beta)\left(c_{in}+c_{out}G\left(\varphi^*\right)\right).$$

An equilibrium is thus given by φ^*, P, Q, M (= number of entrants) such that:

$$\pi(P,\varphi^*)=0,$$

$$\int_{\varphi^*}^{\infty}\pi(P,\varphi)dG(\varphi)=(1-\beta)\left(c_{in}+G\left(\varphi^*\right)c_{out}\right).$$

$$Q=D(P)=M\int_{\varphi^*}^{\infty}y^*(P,\varphi)dG(\varphi).$$

Note that the first two equations determine P and φ^* independently of the rest. P and φ^* are determined as a function of the entry and exit costs, the distribution G, the production function, the wage and the fixed cost. The price is thus determined by supply conditions, independently of the demand curve $D(P)$: in other words, supply is infinitely elastic at price P.

This model generates a selection effect: only the most productive firms survive. Hence the surviving firms are more productive than the pool of potential entrants. Exercise: plot the cross-sectional distribution of productivity, size, profits, employment, given a form for $G(.)$.

If there is more selection, average productivity increases. A theory of TFP in a way. What determines the extent of selection, i.e. φ^*? the entry and exit costs, the fixed costs, and the shape of $G(.)$.

Comparative statics: an increase in the entry/exit cost leads to a higher φ^*, i.e. less selection; whereas an increase in the fixed cost c_f may increase or decrease φ^* - selection effect vs. profitability effect as pointed out in class). I think it's interesting that for given total costs, a higher fixed cost will lead to more selection and a higher entry cost to less selection. (Unfortunately in practice it seems hard to measure and distinguish these two costs!) More competitive markets, i.e. with lower entry costs, have higher productivity, everything else equal.

Note 1: relaxing the infinitely elastic supply result: fixed number of entrants.

The infinitely elastic supply curve follows from the fact that there is a large number of potential entrants who all draw their productivity from the same distribution G. Alternatively we could assume that the entrants know their productivity and that the price will induce the ones with high productivity to enter. (Try and write the equation for this model!)

Note 2: extension to general equilibrium

One can easily extent this framework to a static general equilibrium model. Assume a representative consumer with utility over consumption and leisure $U(C, 1-N)$. The output price is then normalized to one, the wage is then given by the marginal rate of substitution $w = U_2(C, 1-N)/U_1(C, 1-N)$ and the demand curve is replaced by the resource constraint

$$C = Y = wN.$$

(Note average profits are zero since ex-post profits compensate for the initial investment, which is why they do not show up in our resource constraint.)

The unknowns are now C, w, N, φ^*, M, such that

$$\begin{aligned} C &= M \int_{\varphi^*}^{\infty} \left(y^*(\varphi, w) - c_f\right) dG(\varphi), \\ N &= M \int_{\varphi^*}^{\infty} n^*(\varphi, w) dG(\varphi), \end{aligned}$$

$$\pi(\varphi^*, w) = 0,$$

$$\begin{aligned} \int_{\varphi^*}^{\infty} \pi(\varphi, w) dG(\varphi) &= (1-\beta)\left(c_{in} + G\left(\varphi^*\right) c_{out}\right), \\ \frac{u_2(C, 1-N)}{u_1(C, 1-N)} &= w. \end{aligned}$$

Note 3: Imperfect competition

Melitz (Econometrica) has extended this framework to allow for monopolistic competition. He used this framework to study the impact of trade on reallocation within an industry. Syverson (JPE 2002) has studied how the extent of competition shapes the cross-sectional distribution of the survivors.

Allowing for continuous exogenous exit and entry

Building up progressively to a more complicated model: I now make the following modification. Firms die at rate δ, which we take to be independent of productivity. Each period new firms can enter. Clearly we can extend the previous equilibrium with $E = \delta N$ firms entering each period, compensating for the same number of firms leaving the industry. The condition for zero-profit has to be changed since now firms discount the future faster:

$$\begin{aligned} \sum_{t\geq 0} \beta^t (1-\delta)^t \int_{\varphi^*}^{\infty} \pi(P,\varphi) dG(\varphi) &= (c_{in} + G(\varphi^*) c_{out}) \\ &\Rightarrow \\ \int_{\varphi^*}^{\infty} \pi(P,\varphi) dG(\varphi) &= (1-\beta(1-\delta)) (c_{in} + G(\varphi^*) c_{out}) \end{aligned}$$

Note that the assumption that the death rate is independent of productivity is counterfactual since it implies that firms with different size (or productivity) will have same survival rate.

6.2. Idiosyncratic shocks every period (the Hopenhayn model)

The firm-level dynamics of the previous models are too simplistic since size (as measured by sales, or employment) and productivity are constant through time. To allow for a more complex pattern of growth, survival and death we now add idiosyncratic shocks every period: φ follows a first-order Markov process, i.e. φ_{t+1} is draw from the c.d.f. $F(. \mid \varphi_t)$:

$$\Pr\left(\varphi_{t+1} \leq x \mid \varphi_t = \varphi\right) = F(x \mid \varphi).$$

The timing is the following:

- re the incumbents: first firms learn their new shock; they produce this period; next they decide whether or not to leave the industry; and those who stay draw a new shock next period, and so on.

- re the entrants: they draw a shock from G, they produce and next period they join the entrants.

We look for a stationary equilibrium i.e. one where prices are constant. In Hopenhayn's 1992 paper the exposition is made a bit complicated by the fact that he defines the equilibrium generally (even if prices depend on time).[2]

[2] [Check results of convergence to stationary equilibrium]

Bellman equation for an incumbent with productivity:

$$v(\varphi; p) = \pi(\varphi, p) + \beta \max\left(0, \int_0^\infty v(\varphi'; p) dF(\varphi' \mid \varphi)\right).$$

Value of an entrant:

$$v_e(p) = \int_0^\infty v(\varphi; p) dG(\varphi).$$

Note the implicit dependence of value functions on p (and could add an input market which price is a function of total industry demand).

The key step is to prove that v is increasing in φ. Then rule of decision is a threshold rule: stay if $\varphi \geq \varphi^*$.

Let's now determine the law of motion of the cross-sectional distribution. If incumbents are distributed at the beginning of the period (after the new shocks have been drawn) according to the c.d.f. H_t, then the c.d.f. of the distribution next period is:

$$H_{t+1}(\varphi) = M_t.G(\varphi) + \int_{\varphi \geq \varphi_t^*} F(\varphi \mid u) dH_t(u),$$

A stationary equilibrium will be such that this distribution reproduces itself:

$$H(\varphi) = M.G(\varphi) + \int_{\varphi \geq \varphi^*} F(\varphi \mid u) dH(u),$$

Note that M, φ^* are endogenous and F, G are exogenous. The equilibrium cross-sectional distribution of productivity is determined both by the productivity of entrants, the stochastic process of productivity, the extent of selection, and the number of entrants. In this model once we have figured out the distribution of productivity one can easily infer the distribution of output (sales), employment, profits, etc. since they are all simple functions of productivity.

The free entry condition is

$$v_e \leq c_e,$$

where c_e is the entry cost, with equality if $M > 0$. The market-clearing conditions is:

$$Q = D(p) = \int y^*(\varphi, p) dH(\varphi).$$

The unknowns of a stationary equilibrium are Q, p, φ^*, H, and M. See below for an algorithm.

Hopenhayn gives conditions under which this model will generate that older firms are more productive and less likely to die. (And since they are more productive, they are bigger along all dimensions.)

This model also generates that ex-post value is greater than the physical creation cost of a business, because of selection. Part of the value stems from the φ which it is not possible to reproduce. Hence the firms that do survive would have high Tobin's q.

[Develop this using the JEDC paper?]

A possible algorithm to solve this model (as we did in Problem set 1):

(a) guess a price.

(b) given the price, solve the Bellman equation for incumbents. Find the value function.

(c) check whether the free entry condition holds (approximately). if not, adjust the price and go back to (b).

(d) given the correct price and the correct value and policy functions, solve for the stationary distribution, up to the scale M. compute $Q = D(p)$ and deduce the required scale M.

7. Hopenhayn-Rogerson (JPE 1993)

This is really an application of the Hopenhayn 1992 model. The two main differences are (1) it is general equilibrium and (2) there are costs to adjusting labor downward ("firing cost") within one firm. These costs turn out to have large effects on average productivity and welfare. This is because firing costs impede the reallocation of workers to the most productive jobs. There is however a debate as to how to model the firing costs (are they rebated to consumers?; etc.). There is also an interesting discussion in Atkeson, Cole and Ohanian (Carnegie Rochester Series) about the fact that despite large differences in firing costs across countries, the flows of jobs seem similar. This suggests that the benefits to reallocation are very large, so that they swamp the costs.

8. Lecture 2

Outline of lecture: (1) A guide to chapter 4 of SLP, with the main results (principle of optimality, existence and uniqueness of solution, construction of solution, properties of value function; see handout or SLP); (2) Hopenhayn with aggregate shocks, some methodological issues; (3) discussion of problem set 1; (4) a first stab at Cooley-Quadrini.

9. Hopenhayn with aggregate shocks: motivation

The Hopenhayn 1992 model and other interesting models of firm dynamics have constant aggregates. To understand the role that these firm-level dynamics play in the business cycle or in transitional dynamics, we need models where aggregates vary over time. The phenomena we can study are: the behavior of entry/exit over the business cycle, the pattern of growth and death of firms over the business cycle, which firms are most exposed to the business cycle, the source of changes in productivity over the business cycle - does productivity change within firms, is it driven by entry or exit of firms, or finally by reallocation of factors between firms?

As you may know and will see, models with heterogeneous agents and shocks are often intractable. Even computing numerically an equilibrium is complicated. Simplifying assumptions can be crucial. We will see some special cases here. The problem is that the state space of the recursive problem is large, since in general the cross-sectional distribution of firms must be included in it. Most numerical methods suffer from a "curse of dimensionality" that makes high-dimensional problems very costly (in computer time) to solve. The main exception is linear rational expectations methods of the type used to solve RBC models, which can easily handle a large state space. Unfortunately these linear methods are sometimes not very good approximations when the problems are very nonlinear. As you will see below this problem arises with Hopenhayn's model because the exit decision is non-linear.

See Ljungqvist and Sargent (LS) p. 39 (of the first edition) for some discussion of the 'curse of dimensionality'. An important example where such a problem was solved approximately is Krusell and Smith 1998 JPE (discussed p. 393 of LS).

9.1. Hopenhayn with an aggregate shock: recursive competitive equilibrium, social planner problem, curse of dimensionality

See Ljungqvist and Sargent chap. 6 for the definition of the recursive competitive equilibrium. Here we do a more complicated application of the same concept.

Assume a constant wage w and discount factor β. We now introduce an aggregate shock, which I take to be a demand shock: the demand curve is $Q_t = D(P_t, z_t)$ with z_t a Markov process with c.d.f. $U(z' \mid z)$. An increase in z increases the quantity demanded at any price level, e.g. $D(P, z) = z/P^{\varepsilon}$.

We know that the equilibrium is characterized by the following conditions:

$$
\begin{aligned}
v_t(\varphi) &= \pi_t(\varphi) + \beta \max\left(0, \int_0^{\infty} v_{t+1}(\varphi') dF(\varphi' \mid \varphi)\right), \\
v_{e,t} &= \int_0^{\infty} v_t(\varphi) dG(\varphi) \leq c_e, \text{ with equality if } M_t > 0, \\
\pi_t(\varphi) &= \max_{n_t(\varphi) \geq 0} \left\{P_t \varphi n_t(\varphi)^{\alpha} - w n_t(\varphi) - c_f\right\},
\end{aligned}
$$

$$
H_{t+1}(\varphi) = M_t G(\varphi) + \int_{s \geq \varphi_t^*} F(\varphi \mid s) H_t(ds),
$$

$$
\int_0^{\infty} v_{t+1}(\varphi') dF(\varphi' \mid \varphi_t^*) = 0,
$$

$$
Q_t = \int_0^{\infty} \varphi n_t(\varphi)^{\alpha} dH_t(\varphi) = D(P_t, z_t),
$$

given $H_0(\varphi)$.

The *recursive competitive equilibrium* is defined as follows: a value function $V(\varphi, z, H(.))$, equilibrium outcome functions $P(z, H(.))$, and $M(z, H(.))$, $\varphi^*(z, H(.))$, such that:

[1] The value function and policy function φ^* satisfy each firm's Bellman equation, given the price function P, entry function M, and law of motion for the

$$
V(\varphi, z, H(.)) = \pi(\varphi, z, H(.)) + \beta \max\left(0, \int_0^{\infty} \int_0^{\infty} V(\varphi', z', H'(.)) dF(\varphi' \mid \varphi) dU(z' \mid z)\right),
$$

$$
\int_0^{\infty} V(\varphi', z', H'(.)) dF(\varphi' \mid \varphi^*(z, H(.)) = 0,
$$

where the profit function comes from the static choice, and [2] the functions $H(.)$, $P(.)$ are in turn determined by equilibrium behavior:

$$
\begin{aligned}
\pi(\varphi, z, H(.)) &= \max_{n \geq 0} \left\{P\left(H(.), z\right) \varphi n^{\alpha} - wn - c_f\right\} \rightarrow n = n(\varphi P), \\
D(P, z) &= \int_0^{\infty} \varphi n(\varphi P)^{\alpha} dH(\varphi) \rightarrow P = P(H(.), z)
\end{aligned}
$$

$$H'(\varphi) = M\left(z, H(.)\right) G(\varphi) + \int_{s \geq \varphi^*(z,H(.))} F(\varphi \mid s) H(ds),$$

The standard way to solve such a problem (due to Lucas and Prescott 1971) is to write the *social planner problem* corresponding to the previous equilibrium. This problem is maximize consumer surplus, subject to the technology constraints:

$$\max_{Q_t, H_{t+1}(\varphi), n_t(\varphi), \varphi_t^*} E_0 \sum_{t \geq 0} \beta^t \left(S(Q_t, z_t) - M_t c_e - w_t N_t - c_f \int_0^\infty dH_t(\varphi) \right)$$

$$S(Q, z) = \int_0^Q P(s, z) ds,$$

where $P(s, z)$ is the inverse demand curve $[P(D(p, z), z) = p]$.

$$\begin{aligned}
Q_t &= \int_0^\infty \varphi n_t(\varphi)^\alpha dH_t(\varphi), \\
N_t &= \int_0^\infty n_t(\varphi) dH_t(\varphi), \\
H_{t+1}(\varphi) &= M_t G(\varphi) + \int_{s \geq \varphi_t^*} F(\varphi \mid s) H_t(ds), \\
& \quad H_0 \text{ given.}
\end{aligned}$$

[The simple way to see that these two problems are equivalent is to note that they are concave (?) problems with the same first-order conditions, or to invoke directly an infinite-dimensional version of the first welfare theorem.]

Next, note that the social planner problem has the following recursive formulation:

$$\begin{aligned}
W(H(.), z) &= \max_{Q, M, H'(.), \varphi^*} \left\{ S(Q, z) - Mc_e - wN - c_f \int_0^\infty dH(\varphi) + \beta E_{z'/z} W(H'(.), z') \right\} \\
Q &= \int_0^\infty \varphi n(\varphi)^\alpha dH(\varphi), \\
N &= \int_0^\infty n(\varphi) dH(\varphi), \\
H'(\varphi) &= MG(\varphi) + \int_{s \geq \varphi^*} F(\varphi \mid s) H(ds).
\end{aligned}$$

The key point is that the distribution $H(.)$ is a state variable. This is a "high-dimensional object" - even if φ can take only a few values this will require several additional state variables. This makes this problem hard to handle with standard numerical methods.

Note on a general equilibrium variant

What we do now can be applied directly to a general equilibrium version of the Hopenhayn model. More precisely, the social planner problem is:

$$\max_{C_t, M_t, \varphi_t^*, N_t, Y_t, n_t(\varphi), H_{t+1}(\varphi)} E \sum_{t \geq 0} \beta^t U(C_t, 1 - N_t)$$

$$C_t + M_t e \leq Y_t$$

$$H_{t+1}(\varphi) = M_t G(\varphi) + \int_{s \geq \varphi_t^*} F(\varphi \mid s) H_t(ds),$$

$$Y_t = A_t \int_0^\infty \varphi n_t(\varphi)^\alpha dH_t(\varphi) - c_f \int_0^\infty dH_t(\varphi),$$

$$N_t = \int_0^\infty n_t(\varphi) dH_t(\varphi),$$

given $H_0(\varphi)$. We now have shocks to productivity A_t, e.g. an AR(1).

9.2. Taking first-order conditions of the sequence problem

First, rewrite the law of motion for $H(.)$ as the law of motion of the pdfs:

$$h_{t+1}(\varphi) = M_t g(\varphi) + \int_{s \geq \varphi_t^*} f(\varphi \mid s) h_t(s) ds \tag{9.1}$$

Next, write the Lagrangian corresponding to the sequence problem of the planner

$$L = E \sum_{t \geq 0} \beta^t \left[\begin{array}{c} S\left(\int_0^\infty \varphi n_t(\varphi)^\alpha h_t(\varphi) d\varphi, z_t\right) - M_t c_e - w \int_0^\infty n_t(\varphi) h_t(\varphi) d\varphi - c_f \int_0^\infty h_t(\varphi) d\varphi \\ + \int_0^\infty \mu_t(\varphi) \left(M_t g(\varphi) + \int_{s \geq \varphi_t^*} f(\varphi \mid s) h_t(s) ds - h_{t+1}(\varphi)\right) d\varphi \end{array} \right],$$

where $\beta^t \mu_t(\varphi)$ is the Lagrange multiplier on 9.1 for each φ.

The first-order conditions are:

$$/M_t \quad : \quad : c_e = \int_0^\infty \mu_t(\varphi) g(\varphi) d\varphi,$$
$$/\varphi_t^* \quad : \quad : h_t(\varphi_t^*) \int_0^\infty \mu_t(\varphi) f(\varphi \mid \varphi_t^*) d\varphi = 0,$$

$$/n_t(\varphi) \quad : \quad : S_1(Q_t, z_t) \alpha \varphi n_t(\varphi)^{\alpha-1} = w,$$
$$/h_{t+1}(u) \quad : \quad : \mu_t(u) = \beta E_t \left(1_{u \geq \varphi_{t+1}^*} \int_0^\infty f(\varphi \mid u) \mu_{t+1}(\varphi) d\varphi + S_1(Q_{t+1}, z_{t+1}) u n_{t+1}(u)^\alpha - w n_{t+1}(u) - c\right.$$

$\mu_t(u) =$ value of one unit of u at the beginning of time $t+1$, after the idiosyncratic shock has been revealed.

Problem with this approach: the step function on the right-hand side of the last equation is not differentiable, and so the last equation cannot be linearized.

Gilchrist and Williams (JPE 2000) have a model with similar heterogeneity but (i) no endogenous exit (firms can be switched on or off in any period), and (ii) log-normal distributions which simplify the analysis. That model can be log-linearized. (My dissertation is really a simplified version of Gilchrist and Williams which is even simpler to handle - in the end you do not need to worry about the distribution when computing the equilibrium, which makes this very tractable.)

Jeff Campbell (RED 1998) used a model similar to Hopenhayn but with added features (time-to-build, capital) to study the behavior of entry and exit over the business cycles. He developed some numerical methods to linearize it (you can find the appendix on his webpage). This also would be an interesting project to do, though more complicated.

Marcelo Veracierto (AER 2001, WP 2005) also solved a similar problem, but with different methods. We will study that paper later in the class.

One possibility I looked at in the last few days is to change the timing of the model or change the variables to make it "smooth" and be able to linearize it. In particular, define the value of a firm at the end of the previous period:

$$\phi_t(s) \stackrel{def}{=} \int_0^\infty f(u \mid s) \mu_t(u) du.$$

Then you can rewrite the model and substitute the ϕ's for the μ's. The resulting model is smooth so it seems one should be able to linearize it. The only problem is that $\phi(\varphi) = 0$ for $\varphi \leq \varphi^*$. This cannot be linearized either.I need to spend more time looking at it...

Next I turn to some simplifications which allow to solve for this problem more easily.

9.3. Exogenous exit and aggregate shocks

To summarize: because the cross-sectional distribution is a state variable, we need to linearize. But the truncation of the distribution means it will be hard to linearize. Hence, a natural simplifying assumption is to avoid the truncation and say that firms die at an exogenous rate δ independent of productivity.[3] Then the LOM for the cross-sectional distribution is

$$H_{t+1}(\varphi) = M_t G(\varphi) + (1-\delta)\int_0^\infty F(\varphi \mid s) H_t(ds),$$

or in terms of the p.d.f.:

$$h_{t+1}(\varphi) = M_t g(\varphi) + (1-\delta)\int_0^\infty f(\varphi \mid s) h_t(s) ds,$$

and the last first-order condition of the problem of the previous section is now changed to

$$\mu_t(u) = \beta E_t \left((1-\delta)\int_0^\infty f(\varphi \mid u)\mu_{t+1}(\varphi)d\varphi + S_1(Q_{t+1}, z_{t+1}) u n_{t+1}(u)^\alpha - w n_{t+1}(u) - c_f \right).$$

Of course there is the risk that the value $\mu_t(u)$ could become negative if the profit $S_1(Q_{t+1}, z_{t+1}) u n_{t+1}(u)^\alpha - w n_{t+1}(u)$ becomes negative. This will depend on the parameter values especially the shock transitions $f(. \mid .)$.

Given our simplifying assumption, one can now solve for this model with linearization techniques. Remember that this entails the following steps:

[Assume $D(P,z) = D(P)e^z$ where z follows an ARMA process, e.g. an AR(1): $z_t = \rho z_{t-1} + \varepsilon_t$.]

(1) Find the nonstochastic steady-state, i.e. the quantities $\mu^*(\varphi), Q^*, P^*, n^*(\varphi), h^*(\varphi), M^*$ that satisfy the first-order conditions and feasibility when $z_t = 0$:

$$h^*(\varphi) = M g(\varphi) + (1-\delta)\int_0^\infty f(\varphi \mid s) h^*(s) ds,$$

[3] I believe this is what Ghirano and Melitz (forthcoming QJE, plus another WP) do, but I have yet to read their paper. They have monopolistic competition but this doesn't change anything.

$$\mu^*(u) = \beta(1-\delta)\int_0^\infty f(\varphi \mid u)\mu^*(\varphi)d\varphi + P^* u n^*(u)^\alpha - w n^*(u) - c_f,$$

$$\begin{aligned} c_e &= \int_0^\infty \mu^*(\varphi)g(\varphi)d\varphi, \\ Q^* &= D(P^*, 0) = \int_0^\infty \varphi n^*(\varphi)^\alpha h^*(\varphi)d\varphi, \\ w &= P^*\alpha n^*(\varphi)^{\alpha-1}. \end{aligned}$$

(2) Log-linearize the first-order and feasibility conditions around this nonstochastic steady-state.

(3) Solve the linear rational expectation model on the computer after feeding in your steady-state calculations and log-linearizations (either do it yourself, or use a solver like the ones from Paul Klein (U W Ontario), Uhlig, or Chris Sims which you can all find on the web).

(4) Play with it: compute impulse responses, second moments, implied spectral density if you like, etc.

It would be a reasonable project for the class to solve this, and look perhaps at some extensions.

[I assume you're familiar with these methods. If not you can find detailed explanations in several places: Walsh's monetary theory textbook; John Campbell's 1994 JME paper; Uhlig's or Klein's explanations on their websites; lots of lecture notes on the web.]

9.4. Imposing an exogenous aggregate price ("Partial Equilibrium")

First we start by looking at a model without an aggregate equilibrium condition "supply equals demand" i.e. we forget about the demand curve. Instead just assume that the output price follows a Markov process with transition c.d.f. $U(p' \mid p)$. I assume that this price process is independent of all the individual idiosyncratic shocks to productivity φ.

The individual problem is given by:

$$v(p,\varphi) = \pi(p,\varphi) + \beta \max\left(0, \int\int v(p',\varphi')dF(\varphi' \mid \varphi)dU(p' \mid p)\right),$$

or in the case where φ and p follows simple Markov chains:

$$v(p,\varphi) = \pi(p,\varphi) + \beta \max\left(0, \sum_{p',\varphi'} v(p',\varphi') f(\varphi' \mid \varphi) u(p' \mid p)\right).$$

One can iterate on this Bellman equation to find the value v and associated threshold φ^*. Note that φ^* will now depend on p.

Of course by "forgetting" about equilibrium we lose many implications. For instance, there is no guarantee that the free-entry condition will hold [$c_e = \int v(p,\varphi) dG(\varphi)$]; and we cannot simply adjust the price to make it hold since we care about the whole stochastic process for the price. Moreover, the level of entry M_t is similarly undetermined.. Finally, we don't know what aggregate quantities are.

Despite all these limitations, in some case it can be useful, at least as a first step, to solve the model with some exogenous, reasonable-looking aggregate price process. This allows to analyze the impact of aggregate shocks on the firm's decision rules.

To go a bit beyond that, you can make some additional assumptions - for instance in this case, if you assume that entry $M_t = M$ is constant, you will see how the cross-sectional distribution of productivity changes over the business cycle. In particular we can study if recessions kill the inefficient firms and so increase average productivity; or if entry breeds "excessive" entry. More generally, we can analyze the effect of a "shock" to the cross-sectional distribution of productivity.

Part II
Lecture 3

Announcement: distribute PS2, remind students of 11-17 deadline.

(1) Finish Cooley-Quadrini, (2) Review of cross-sectional distributions, (3) Facts from Davis and Haltiwanger, (4) Caballero-Hammour, (5) Campbell-Fischer.

10. Cooley and Quadrini (AER 2001)

Big points: on top of the regularities re. firm growth/size dynamics, there are empirical regularities re the financial structure of firms: small, young firms pay less dividends, have more debt, invest more, and have high Tobin's q. Their investment is also more sensitive to cash flows. This is widely interpreted as evidence of financing constraints for small firms.

Prelim remark: with complete market the cash-flow effect is zero and the value function is linear in wealth. This is independent of the technology:

$$\begin{aligned} V(W) &: \max \sum_{t\geq 0} \sum_{s^t \in S^t} q_t(s^t) \left[F(k_t(s^{t-1})) - \left[k_{t+1}(s^t) - (1-\delta) k_t(s^{t-1}) \right] \right] \\ s.t. &: k_{-1} = W. \end{aligned}$$

With (binding) incomplete markets there will be a concavity, which will depend on the curvature of the profit function in capital (DRS). (Math to add.)

Model: Exogenous one-period financial contracts: debt, with a cost of bankruptcy, and equity, with a cost of raising equity. For the "right" parameter values, this delivers the so-called pecking order, with firm using first retained earnings, then debt, and finally equity to finance investment. This justifies the financial facts enumerated above. The debt-equity trade-off is the bankruptcy cost vs. the cost of raising equity.

Related work: Gomes AER 2001 used a similar model to study the impact of financing constraints on cash flows. A next step is to make the contracts endogenous and dynamic perhaps (Hopenhayn-Albuquerque, Hopenhayn-Clementi, etc.)

11. Some notions on Markov Processes and Cross-Sectional Distributions

The following comes from the handout on Markov processes I distributed you, but there is an extra section at the end.

Let $(S, \Im, \mathbb{P})$ be a probability space.[4]

Definition 11.1. *A transition function is a function Q from $S \times B(\Im) \to \mathbb{R}$ such that*

(1) for any $s \in S$, $Q(s, .)$ is a probability measure on $(S, \Im)$.

(2) for any $A \in \Im$, $Q(., A)$ is a measurable function from $(S, \Im)$ into $\mathbb{R}$.

$Q(s, A)$ should be interpreted as $\Pr(X_{t+1} \in A \mid X_t = s)$.

Example 11.2. *In the case of a Markov chain with transition matrix P, $Q(i, \{j\}) = P(i, j)$. Hence we define $Q(s, A)$ for any $A \subset \{1, ...S\}$ as $Q(s, A) = \sum_{i \in A} P(s, i)$.*

Example 11.3. *Suppose that $X_t \in \mathbb{R}$ and we define $F(x \mid s) = \Pr(X_{t+1} \leq x \mid X_t = s)$. Then $Q(s, A) = \int_A dF(x \mid s)$.*

Example 11.4. *A normal AR(1) process: $X_{t+1} = \rho X_t + \varepsilon_{t+1}$ with ε_t iid $N(0, \sigma^2)$. We have*

$$\begin{aligned} Q(s, A) &= \Pr(X_{t+1} \in A \mid X_t = s) \\ &= \Pr(\rho s + \varepsilon_{t+1} \in A) \end{aligned}$$

Hence we can define for any b :

$$Q(s,]-\infty, b]) = \Pr(\rho s + \varepsilon \leq b) = \Pr(\varepsilon \leq b - \rho s) = \Phi_{0,1}\left(\frac{b - \rho s}{\sigma}\right),$$

where $\Phi_{0,1}$ is the standard normal cdf. [Having defined Q for any interval $]-\infty, b]$ is enough...]

---

[4] Recall this means that S is a set, $\Im$ is a sigma-algebra over S, and $\mathbb{P}$ is a probability measure over $(S, \Im)$. A sigma-algebra is a set of parts of S satisfying (i) $\emptyset \in \Im$, (ii) $A \in \Im \Rightarrow A^c \in \Im$, and (iii) if A_n is a countable collection of sets in $\Im$, then $\cup_{n \in \mathbb{N}} A_n$ is in $\Im$. A probability measure $\mathbb{P}$ on $(S, \Im)$ is a mapping from $\Im$ into $[0, 1]$ such that: (i) $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(S) = 1$, (ii) if the A_n are countable and disjoint, then $\mathbb{P}(\cup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} \mathbb{P}(A_n)$.

Let $\Lambda(S, \Im)$ the set of probability measures over $(S, \Im)$ and $M(\Im)$ the set of measurable functions from $(S, \Im)$ into $\mathbb{R}$. We now define two operators associated with the Markov chain Q.

Definition 11.5. *The Markov operator T is a mapping from the set of measurable functions $(S, \Im) \to \mathbb{R}$ into itself defined as:*

$$\forall f \in M(S, \Im), \forall s \in S : (Tf)(s) = \int f(s')Q(s, ds').$$

Interpretation: $(Tf)(s) = \mathbb{E}\left(f(s_{t+1}) \mid s_t = s\right)$.

Remark 1. *T is linear: $T(\alpha f + \beta g) = \alpha Tf + \beta Tg$ for any $f, g : (S, \Im) \to \mathbb{R}$ and $\alpha, \beta \in \mathbb{R}$.*

Remark 2. *In the case of Markov chains, $(Tf)(s) = \sum_{s' \in S} Q(s, s')f(s')$ which we can write in vector terms as $Tf = Qf$, where f is the column vector $\{f(s)\}_{s=1...S}$ and Tf the vector $\{(Tf)(s)\}_{s=1...S}$.*

SLP show in theorem 8.1 and its corollary that T maps positive measurable functions into positive measurable functions, and bounded measurable functions into bounded measurable functions: $T : M^+(S, \Im) \to M^+(S, \Im)$, and $T : B(S, \Im) \to B(S, \Im)$.

Example 11.6. *Let's write down T for the example of an AR(1). We have*

$$\begin{aligned} (Tf)(s) &= \int f(s')Q(s, ds') \\ &= \int_{-\infty}^{+\infty} f(s') \frac{1}{\sigma} \phi_{0,1} \left(\frac{s' - \rho s}{\sigma} \right) ds'. \end{aligned}$$

Definition 11.7. *The adjoint of T is T^*, a mapping from the set of probability measures over $(S, \Im)$ into itself, defined as:*

$$\forall \lambda \in \Lambda(S, \Im), \forall A \in \Im : (T^*\lambda)(A) = \int Q(s, A)\lambda(ds).$$

Interpretation: $T^*\lambda$ is next period's probability distribution over the states if λ is today's probability distribution over the states: $(T^*\lambda)(A) = \Pr(s_{t+1} \in A \mid s_t \rightsquigarrow \lambda)$.

Remark 3. *T^* is satisfies $T^*(\alpha\lambda + (1-\alpha)\mu) = \alpha T^*\lambda + (1-\alpha)T^*\mu$ for any $\alpha \in [0,1], \lambda, \mu \in \Lambda(S, \Im)$.*

Remark 4. *In the case of Markov chains, $T^*\mu = \mu Q$.*

Example 11.8. *In the case of an AR(1),*

$$\begin{aligned}(T^*\lambda)(A) &= \int Q(s,A)\lambda(ds) \\ &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} 1_{u\in A}\frac{1}{\sigma}\phi_{0,1}\left(\frac{u-\rho s}{\sigma}\right) du\, \lambda(ds)\end{aligned}$$

Or, for some $b \in \mathbb{R}$ and $A =]-\infty, b]$:

$$\begin{aligned}(T^*\lambda)(]-\infty,b]) &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} 1_{u\leq b}\frac{1}{\sigma}\phi_{0,1}\left(\frac{u-\rho s}{\sigma}\right) du\, \lambda(ds) \\ &= \int_{-\infty}^{+\infty} \Phi_{0,1}\left(\frac{b-\rho s}{\sigma}\right)\lambda(ds).\end{aligned}$$

Finally if λ has a density i.e. $\lambda(ds) = g(s)ds$, then

$$(T^*\lambda)(]-\infty,b]) = \int_{-\infty}^{+\infty} \Phi_{0,1}\left(\frac{b-\rho s}{\sigma}\right) g(s)ds$$

Definition 11.9. *We can define the $n-$period probability ahead recursively. For any $s \in S, A \in \Im$, let $Q^1(s,A) = Q(s,A)$, and $\forall n \geq 1, Q^{n+1}(s,A) = \int Q^n(u,A)Q(s,du)$.*

Proposition 11.10. *For all $n \geq 1$, Q^n is a transition function, and the operators of the iterate are given as the iterates of the operator of the one-step ahead transition function: $T^{(n)} = T^n, T^{*(n)} = T^{*n}$.*

Proof: see SLP.

Remark 5. *Hence, given an initial distribution $\lambda_0, \lambda_n = T^{*n}\lambda_0$ is the distribution over the states after n periods. Moreover $T^n f$ is the expected value of f n-periods in advance: $(T^n f)(s) = \mathbb{E}(f(s_{t+n}) \mid s_t = s)$.*

Example 11.11. *Consider again the case of an AR(1). Let G_{t+1} be the c.d.f. and g_{t+1} the p.d.f. of X_{t+1}. Then we have the recursion, $\forall x \in \mathbb{R}$:*

$$\begin{aligned} G_{t+1}(x) &= \Pr(X_{t+1} \leq x \mid X_t \rightsquigarrow G_t) \\ G_{t+1}(x) &= \int_{-\infty}^{+\infty} \Phi_{0,1}\left(\frac{x-\rho s}{\sigma}\right) G_t(ds) \\ G_{t+1}(x) &= \int_{-\infty}^{+\infty} \Phi_{0,1}\left(\frac{x-\rho s}{\sigma}\right) g_t(s)ds \end{aligned}$$

Hence:

$$\forall x \in \mathbb{R} : g_{t+1}(x) = \frac{1}{\sigma}\int_{-\infty}^{+\infty} \phi_{0,1}\left(\frac{x-\rho s}{\sigma}\right) g_t(s)ds,$$

and a stationary distribution must satisfy the equation $T^*\lambda = \lambda$, with we can write in p.d.f. terms as

$$\forall x \in \mathbb{R} : g^*(x) = \frac{1}{\sigma}\int_{-\infty}^{+\infty} \phi_{0,1}\left(\frac{x-\rho s}{\sigma}\right) g^*(s)ds.$$

This is a functional equation in g^*. In our case, we can guess and verify that the stationary distribution of X will be normal (μ, Σ^2). We see that μ, Σ^2 have to satisfy: $\mathbb{E}X_{t+1} = \rho\mathbb{E}X_t + \mathbb{E}\varepsilon_{t+1}, \mathbb{V}X_{t+1} = \rho^2\mathbb{V}X_t + \sigma^2 \Rightarrow \mu = 0, \Sigma^2 = \sigma^2/(1-\rho^2)$. To check this guess, we need to see that $g^*(x) = \phi_{0,1}\left(\frac{x-\mu}{\Sigma}\right)/\Sigma$ satisfies the equation above, i.e. we need to check that

$$\phi_{0,1}\left(\frac{x-\mu}{\Sigma}\right) = \frac{1}{\sigma}\int_{-\infty}^{+\infty} \phi_{0,1}\left(\frac{x-\rho s}{\sigma}\right)\phi_{0,1}\left(\frac{s-\mu}{\Sigma}\right) ds,$$

which is possible to do by hand. In this special case, there is also another way, which is simpler: given the properties on the sum and product by a scalar of normal distributions, we see that if X_t is normal (μ, Σ), then X_{t+1} is normal (μ, Σ). This suffices to prove that this is one invariant distribution.

More generally, the definition of an invariant distribution is still $T^*\mu = \mu$, i.e.:

$$\forall A \in \Im : \int Q(s, A)\mu(ds) = \mu(A),$$

which we can first rewrite by requiring this to hold only for the subsets $A =]-\infty, b]$, denoting $F(b \mid s) = Q(s,]-\infty, b]) = \Pr(X_{t+1} \leq b \mid X_t = s)$:

$$\forall b \in \mathbb{R} : \int F(b \mid s)\mu(ds) = \mu(]-\infty, b]),$$

and in the special case when (i) F has a continuous density i.e. $F(b \mid s) = f(b \mid s)ds$, and (ii) μ has a continuous density too, i.e. $\mu(ds) = g(s)ds$, we can rewrite this equality as

$$\forall b \in \mathbb{R} : \int f(b \mid s)g(s)ds = g(b).$$

Example 11.12. *There is a simple formula for* $\mathbb{E}\left(\sum_{t\geq 0} \beta^t f(x_t) \mid x_0\right) = \sum_{t\geq 0} \beta^t \left(T^t f\right)(x_0)$ *where* T^t *denotes the* $t-th$ *iterate of* T *in the case of Markov chains:* $\sum_{t\geq 0} \beta^t \left(T^t f\right)(x_0) = \left(\left(\sum_{t\geq 0} \beta^t Q^t\right) f\right)(x_0) = e(x_0)(I - \beta Q)^{-1} f$ *where* $e(x_0)$ *is the* $1 \times S$ *vector with* 1 *in position* x_0 *and* 0 *elsewhere.*

Remark 6. *We can obtain the expression for* $\mathbb{E}\left(\sum_{t\geq 0} \beta^t f(x_t) \mid x_0\right)$ *by a slightly different road. The distribution of* x_t *given* x_0 *is* $\left(T^{*t}\right)(x_0) = e(x_0)Q^t$. *Next,* $\mathbb{E}\left(f(x_t) \mid x_0\right) = e(x_0)Q^t f$ *where* f *is the* $S \times 1$ *vector of values* $f(s)$. *Finally,*

$$\mathbb{E}\left(\sum_{t\geq 0} \beta^t f(x_t) \mid x_0\right) = \sum_{t\geq 0} \beta^t e(x_0)Q^t f = e(x_0)(I - \beta Q)^{-1} f.$$

What we do in this example is compute of the future distribution using T^* instead of computing the expected value of f using T. This is an instance of the following theorem.

Proposition 11.13. $\forall f \in M^+(S, \Im), \forall \lambda \in \Lambda(S, \Im)$:

$$\int (Tf)(s)\lambda(ds) = \int f(s)(T^*\lambda)(ds).$$

Proof: see SLP (Theorem 8.3).

This simply says that if today's state is drawn according to λ, the expected value of f tomorrow can be found by first finding the probability distribution tomorrow if today's state is drawn according to λ, and then integrating f with respect to this distribution. In the case of Markov chains we already know this since the equality can be written: for any $f \in \mathbb{R}^S$ (a $S \times 1$ vector) for any $\lambda \in \Delta^S$ (a $S \times 1$ vector):

$$\lambda' (Qf) = (\lambda' Q) f,$$

which is obviously true.

Remark 7. *At the risk of beating a dead horse, there is a third way to compute the discounted sum of the earlier example. Let* $V(x_0) = \mathbb{E}\left(\sum_{t\geq 0}\beta^t f(x_t) \mid x_0\right)$*, then*

$$\begin{aligned} V(x_0) &= f(x_0) + \beta\mathbb{E}\left(V(x_1) \mid x_0\right) \\ V(x_0) &= f(x_0) + \beta\sum_{s\in S} Q(x_0, s)V(s) \end{aligned}$$

which we can write in vector terms as

$$V = f + \beta QV \Rightarrow V = (I - \beta Q)^{-1} f.$$

Example 11.14. *(Campbell and Cochrane, JPE 1999). Assume log consumption is a random walk:* $\Delta \log C_t = \mu + \varepsilon_t$ *with* ε_t *iid* $N(0, \sigma^2)$*. CC specify a utility function* $\frac{1}{1-\sigma}\mathbb{E}\sum_{t\geq 0}\beta^t (C_t - X_t)^{1-\sigma}$ *where* X_t *depends implicitly on past consumption and today's consumption as follows. Let* $S_t = (C_t - X_t)/C_t$ *and let* $s_t = \log S_t$ *satisfy for some chosen parameter values* $\phi, \overline{s}$ *:*

$$\begin{aligned} s_{t+1} &= (1-\phi)\overline{s} + \phi s_t + \lambda(s_t)\left(\Delta \log C_{t+1} - \mu\right), \\ \lambda(s) &= \frac{1}{e^{\overline{s}}}\sqrt{1 - 2\left(s - \overline{s}\right)} - 1, \;\; s \leq s_{\max} \\ &= 0, \;\; s \geq s_{\max} \\ s_{\max} &= \overline{s} + \frac{1}{2}\left(1 - e^{2\overline{s}}\right). \end{aligned}$$

Write the transition function associated with this Markov process. Find the stationary distribution numerically. What is interesting about this example is that it is nonlinear, and the nonlinearities are important.

We finish with two important definitions:

Definition 11.15. Q *has the Feller property if its Markov operator* T *maps* $C(S)$*, the set of continuous function of* S *into* $\mathbb{R}$*, into itself, i.e. if the conditional expectations are continuous in today's state:* $(Tf)(s) = \int f(s')Q(s, ds')$ *is continuous in* s *for any* f *continuous.*

Small variations in the state today move only a little the conditional expectation tomorrow of any continuous function; intuitively this means that the probability distribution over states tomorrow moves only a little, and since f *is continuous, the conditional expectation moves only a little too.*

Definition 11.16. *Q is monotone if its associated Markov operator T maps nondecreasing functions into nondecreasing functions.*

This means that the higher the state today, the higher the conditional expectation of tomorrow's f, for any f. This is equivalent to saying that for any $s_2 \geq s_1$, $Q(s_2, .)$ first-order stochastically dominates $Q(s_1, .)$.

[Recall that a c.d.f. F first-order stochastically dominates G iff any of the following equivalent conditions are satisfied: (i) $\forall x \in \mathbb{R}, F(x) \leq G(x)$, (ii) for any h weakly increasing, $\int h(s)dF(s) \geq \int h(s)dG(s)$, (iii) X drawn from F equals Y drawn from G plus a random variable $Z \geq 0$; of course a consequence that F FOSD G is that the mean of F is higher than the mean of G.]

Convergence results for Markov Processes

Chapters 11 and 12 in SLP study the convergence properties of the more general Markov processes in detail. Here I only quote two important results.

Proposition 11.17. *If $S \subset \mathbb{R}^l$ is compact and Q has the Feller property, then there exists an invariant distribution.*

Proposition 11.18. *Assume that $S = [a, b]$, Q is monotone, has the Feller property, and satisfies the following mixing condition:*

$$\exists c \in S, \exists \varepsilon > 0, \exists N \geq 1, P^N(a, [c, b]) \geq \varepsilon \text{ and } P^N(b, [a, c]) \geq \varepsilon.$$

Then there is a unique invariant distribution μ^* and for any initial distribution μ_0, the sequence $\mu_n = (T^*)^n \mu_0$ converges weakly (i.e. in distribution) to μ^*.

12. General discussion of methods to solve models with heterogeneity using dynamic programming

(1) Computing a stationary equilibrium with exogenous prices

This section gives "the big picture", i.e. the steps to take to solve a stationary equilibrium with exogenous prices. Start with the stochastic Bellman equation:

$$v(x, z) = \max_{y \in \Gamma(x,z)} \left\{ F(x, y, z) + \beta \int v(y, z')Q(z, dz') \right\}.$$

Under certain assumptions on F, Q, Γ, β, there exists a unique continuous bounded value function satisfying this equation; and moreover under additional assumptions there is a policy *function* $y^*(x,z)$. (More generally this policy function is actually a correspondence.)

In that case, $X=(x,z)$ is a Markov process with transition Q^X given by:

$$\begin{aligned} Q^X\left((x,z),(X,Z)\right) &= Q(z,Z)\times \Pr\left(x_{t+1}\in X \mid x_t=x, z_t=z\right) \\ &= Q(z,Z)\times 1_{y^*(x,z)\in X}. \end{aligned}$$

Note that it is very important that X is Markov, but x may not be.

Given this transition function, I can compute the evolution of the cross-sectional distribution over x and z :

$$\mu_{t+1}=T^*\mu_t,$$

where T^* is the Markov adjoint operator associated with Q^X :

$$\left(T^*\mu\right)(A)=\int Q(s,A)\mu(ds).$$

See examples in the section above for Markov chains or AR(1).

A stationary distribution is a measure μ satisfying $T^*\mu=\mu$.

There are some technical conditions for (1) existence of this stationary distribution, (2) uniqueness, (3) convergence to it. [See handout based on SLP.] The analysis for the case of a finite state space is simpler than for a general process. In particular, a sufficient condition for existence, uniqueness and convergence to the invariant distribution is : $\exists N\geq 1$ s.t. the N-period ahead transition matrix has only positive elements.

In practice, once you have found v and y^* you can compute the mapping T^*, and then finding the invariant distribution involves simply iterating on T^* : given an arbitrary μ_0, $T^{*t}\mu_0\to\mu^*$ if the conditions of the above theorem **??** are satisfied.

Example: Campbell-Fisher (see below).

(2) Computing a transition between two stationary equilibria with exogenous prices

We consider the following experiment: before time 0, agents where optimizing given the exogenous parameter vector θ, and they were assuming this would go on forever. At time 0, this parameter vector changes unexpectedly to θ' and agents expect it will remain equal to θ' forever.

Step 1: compute the stationary equilibrium for θ, i.e. the value and policy functions and the invariant distribution μ_θ.

Step 2: compute the stationary equilibrium for θ', in particular compute the operator $T^{*\prime}$ associated with the policy function and shocks for θ'.

Step 3: to compute the evolution from time 0 onward, start from μ_θ and obtain the cross-sectional distribution at time $t \geq 1$ by applying $T^{*\prime}$:

$$\mu_t = T^{*\prime t} \mu_\theta.$$

Given this, you can compute aggregates and microeconomic variables for each unit. For $t \to \infty$, μ_t will converge (under the assumptions given above in **??**) to $\mu_{\theta'}$, the invariant distribution of the new steady-state.

Example: take the Campbell-Fisher model from below and change the process of firm-level uncertainty. This would allow to evaluate the effects of the increase in firm-level risk that some recent work has examined. (See Comin and Philippon NBER Macro 2005, Campbell et al. JF 2001, Campbell and Fisher RED 2004, and their references.)

(3) Computing a stationary equilibrium with endogenous prices

In this case the firms' behavior depends on a price that needs to be determined in equilibrium. I write this as

$$v(x, z; P) = \max_{y \in \Gamma(x,z)} \left\{ F(x, y, z) + \beta \int v(y, z'; P) Q(z, dz') \right\}.$$

For given P, and again under some assumptions on F, Γ, β, X, there is a unique value function satisfying this Bellman equation, call it $v(.,.; P)$, and there is a policy function $y^*(x, z; P)$ which also depends on P. We can thus define the Markov operator which we also index by $P : T^*_P$. We then obtain the invariant distribution μ^*_P.

To obtain P, we now compute the aggregate supply and equate it with demand:

$$Q^s(P) = \int q(x, z) d\mu^*_P(x, z) = D(P),$$

which gives one equation in P, which finally determines the equilibrium price P^*. Then we know that the invariant distribution is μ^*_{P*}, with which we can compute everything.

Note 1: to prove that aggregate supply is continuous and upward-sloping in P, which we require in principle to show that we can choose P to equate supply and demand, requires some work. (See SLP for an example in the case of the savings problem.)

Note 2: in Hopenhayn 1992 (PS1) this is simpler because we can find the price from the free-entry condition, without having to compute the aggregate demand.

(4) Computing a transition between two stationary equilibria with endogenous prices

We do again the same experiment of a one-time, unexpected, permanent change in parameters. If prices are endogenous however the dynamics will be more complicated, because prices will feed back on decisions in general (i.e., tomorrow's price will have an impact on today's decision).

This suggests the following approach.

Step 1: compute the stationary equilibrium with endogenous prices for parameter θ. The price will be P^*.

Step 2: do the same for parameter θ'. The price will be $P^{*\prime}$.

Step 3: to compute the transition, choose a period T by which we assume that the new equilibrium has been reached, so $P_t = P^{*'}$ for $t \geq T$.

Next, make a guess for the price path $\{P_t\}_{t=0}^T$ with $P_T = P^{*'}$ and $P_0 = P'$.

Given this guess, solve the Bellman equation with parameter θ', but the value is indexed by t, with $0 \leq t \leq T$:

$$v(x,z,t) = \max_{y \in \Gamma(x,z)} \left\{ F(x,y,z,P_t) + \beta \int v(y,z',t+1) Q(z,dz') \right\},$$

which will yield a time-dependent value function (to reflect the fact that the problem is not stationary, since prices are not constant.)

Given this solution, compute for each time t the cross-sectional distribution and the aggregate supply $Q_t^s(P)$. Compute the equilibrium price at each date : $D(\widetilde{P}_t) = Q_t^s(\widetilde{P}_t)$. If this new equilibrium price is close to the guess, stop; otherwise update your guess using the new equilibrium price and recompute the value functions, etc.

Last, check that there is indeed convergence at time T i.e. $P_{T-1} \simeq P^{*'}$.

[You can adapt this method to the precise problem of course...]

Example 1: take the Hopenhayn-Rogerson model from above and compute the impact of a decrease in the firing cost - this will allow you to evalute the welfare benefit of such a policy reform. (See Veracierto IER for a similar computation, with capital on top.)

Example 2: Lucas-Golosov.

(5) Aggregate shocks with exogenous prices

In this case, solving the Bellman equation is not especially complicated, since we only need to add the aggregate shock in z^a in the basic Bellman equation to the idiosyncratic shock z^i:

$$v(x, z^i, z^a) = \max_{y \in \Gamma(x, z^i, z^a)} \left\{ F(x, y, z^i, z^a) + \beta \int v(y, z^{i\prime}, z^{a\prime}) P(z^a, dz^{a'}) Q(z^i, dz^{i\prime}) \right\}.$$

(Of course writing the shock as a vector $z = (z^i, z^a)$ would allow to use the basic framework, but it's useful for cross-sectional distributions to make the distinction.)

Now we have again a transition function induced by the optimal policy function $y^*(x, z^i, z^a)$ and the shock processes P and Q : the vector $X = (x, z^i, z^a)$ is Markov. Hence one can apply the results from before to obtain an invariant distribution which will give the average across both idiosyncratic and aggregate shocks of the variables in x.

One can also condition on z^a : let $T^*(z^a)$ be the operator mapping measures over (x, z^i) into themselves, then I can simulate the evolution of the cross-sectional distribution by applying this operator to an initial cross-sectional distribution. This allows to simulate the model and to examine its aggregate dynamics.

Example: Zhang JF 2005 studies a cross-section of firms facing adjustment costs [but there is a twist]; Imrohoglu JPE 1989 studies a savings problem; etc.

(6) Aggregate shocks with endogenous prices

This is the complicated case. In general firms need future prices to make decisions. But future prices are impacted by tomorrow's cross-sectional distribution. Hence the firm needs to forecast tomorrow's cross-sectional distribution, and for this it needs today's cross-sectional distribution. This is why a "curse of dimensionality" arises.

Krusell and Smith's solution (JPE 1998) is to use good simple forecasts of future prices instead of the "exact" theoretical one. If the loss in forecasting due to the use of a simple predictor rather than the best one is very small, this is likely to deliver us the correct equilibrium approximately.

13. Davis-Haltiwanger (QJE 1992, MIT Press 1996)

Very influential empirical work documenting the large flows of jobs across plants in the manufacturing sector. This handout is based on the summary of the book, chapter 2.

(See also "The flow approach to labor market...", a recent paper by Davis, Faberman and Haltiwanger published in Journal of Economic Perspectives which presents datasets and results in an updated form. Available on Steve Davis's webpage at Chicago GSB.)

Dataset used: LRD.

Definitions:

$\Delta n_{it} =$ Change in #jobs in plant i during year $t =$ # jobs at end of year minus minus #jobs at beginning of year.

Gross Job Creation = sum of jobs created across all the plants which have been expanding or have been created (i.e. $\Delta n_{it} > 0$): $GJC_t = \sum_{i\ s.t. \Delta n_{it} \geq 0} \Delta n_{it}$.

Gross Job Destruction = sum of jobs destroyed across all the plants where there has been a job loss (or where the plant has closed): $GJD_t = \sum_{i\ s.t. \Delta n_{it} < 0} -\Delta n_{it}$.

Net employment change: change in total employment between $t-1$ and t : $N_t = GJC_t - GJD_t$.

Gross Job Reallocation: sum of all gains and losses between $t-1$ and t : $GJR_t = GJC_t + GJD_t$.

Excess Reallocation = $GJR_t - |N_t|$. This is a measure of how large the flows are relative to the minimum required to do the net change.

These jobs concept need to be distinguished from the worker concepts, such that:

Gross Worker Reallocation: number of workers who changed jobs or employment status (unemployed, employed, out of labor force) between $t-1$ and t.

Facts:

(1) Job destruction and job creation are large: in a typical year in manufacturing roughly 10% of jobs are created and about as much are destroyed.

(2) Job creation and job destruction are persistent: a job that is created (resp. destroyed) is likely to exist (resp. not exist) a few years later.

(3) Concentration: a large share of job creation occurs at plants that adjust their employment a lot. Two thirds of all JC or JD takes place at plants that adjust by more than 25% their labor force. In particular, shutdowns account for about 23% of job destruction, start-ups for about 15% of job creation.

(4) GJD fluctuates more over the business cycle than GJC. Of course GJC is procyclical and GJD is countercyclical, but it appears that GJD is more volatile than GJC.

Note that fact (1) means that even when there is little change in aggregate employment, there is a lot of reallocation of jobs between plants expanding and contracting: GJR and ER are large. This again suggests important idiosyncratic shocks at the firm- or plant-level.

(5) A possibility is that job flows reflect systematic departure of workers from certain industries, or certain type of establishments (e.g. bigger, more capital

intensive, etc.). However this again doesn't seem to be the case - there are large flows in and out of plants which look alike "on paper".

(6) More work in the book on how the flows vary with plants characteristics - e.g., old, bigger plants have less reallocation.

One limitations of this study is that it covers only manufacturing - an interesting sector but increasingly unrepresentative of US economy (10% of jobs today). Keep in mind this is jobs flows, not worker flows. We'll talk about worker flows next week.

Data: some available on Haltiwanger's website.

14. Caballero-Hammour (AER 1994)

Deterministic model to study the impact of business cycles on the structure of industry and on the creation and destruction of production units. Nice story about the Great Depression. Motivated by Davis-Haltiwanger finding of cyclical GJD.

1 unit = 1 worker with some capital (Leontief).

Frontier technology $A(t)$. Depreciate at rate δ.

wage = 1 each period (normalization).

Let $f(a,t) =$mass of units at time t of age a. Note $f(0,t)$ is creation of new units today.

$$\rightarrow f(a,t) = f(0,t-a)e^{-\delta a}.$$

Let $\overline{a}(t) =$ max age existing at time t.

Output:

$$Q(t) = \int_0^{\overline{a}(t)} A(t-a)f(a,t)da.$$

Employment (and proportional to capital stock):

$$N(t) = \int_0^{\overline{a}(t)} f(a,t)da.$$

Thus:

$$\frac{dQ}{dt} = A(t)f(0,t) - \left\{ A(t-\overline{a}(t))f(\overline{a}(t),t)\left(1 - \frac{d\overline{a}(t)}{dt}\right) + \delta Q(t) \right\},$$

$$\frac{dN}{dt} = f(0,t) - \left\{ f(\overline{a}(t),t) \left(1 - \frac{d\overline{a}(t)}{dt}\right) + \delta N(t) \right\}.$$

Interpretation: new entrants, units dying, units retired, change in margin of destruction.

Unit elastic demand $P(t)Y(t) = \overline{D}(t)$, with $\overline{D}(t)$ exogenous.

Creation decision: unit built today will live for $T(t)$ periods where

$$\overline{a}(t + T(t)) = T(t).$$

Note: perfect foresight.

Free entry condition implies that cost of construction = present value of profits:

$$c\left(f(0,t)\right) = \int_t^{t+T(t)} e^{-(r+\delta)(s-t)} \left(P(s)A(t) - 1\right) ds.$$

$c(x)$ = unit cost of creation if flow is x.

Destruction decision: $P(t)A(t - \overline{a}(t)) = 1$ i.e. at the margin of profitability.

First, we can look for a steady-state when demand is constant $\overline{D}(t) = D^*$ constant, and all quantities are independent of t. Then $f^*(a) = f^*(0)e^{-\delta a}$ for $a \in [0, \overline{a}^*]$, and one can solve for everything. In the special case where the creation cost is proportional, $c\left(f(0,t)\right) = c$, then $\overline{a}^*$ is determined by

$$c = \frac{e^{\gamma \overline{a}^*} - e^{-(r+\delta)\overline{a}^*}}{\gamma + r + \delta} - \frac{1 - e^{-(r+\delta)\overline{a}^*}}{r + \delta},$$

where γ = rate of growth of A, and then $f^*(0) = \frac{(\gamma+\delta)\overline{D}^*}{e^{\gamma\overline{a}^*} - e^{-\delta\overline{a}^*}}$.

Using this expression, I can compute the steady-state effect of an increase in $\overline{D}^*$.

It will be to increase $f^*(0)$ since $\overline{a}^*$ does not change. Hence there is an increase in creation and an increase in destruction $= f^*(0)e^{-\delta\overline{a}^*}$ in the new steady-state.

Business cycle effects: assume a deterministic pattern for $\overline{D}(t)$. Then can solve numerically the system of ODEs.

C-H also prove a general result, the *insulation effect*: if c is constant independent of $f(0,t)$, then all adjustment takes places at the entry margin. To prove this, it is sufficient to note that the equilibrium of the steady-state continues to work, with only $f(0,t)$ changing over time to match the changes in demand. If the constraint $f(0,t)$ is not hit, then we can keep the same equilibrium for all the other variables as in the steady-state. This construction does not work however if c depends on $f(0,t)$. They offer a numerical example in which for enough "curvature" (adjustment costs) in terms of c, the result is reversed and job creation is less volatile than job creation.

One problem that is a bit left out of this nice model is, what is the big convexity in job creation that is necessary to get volatile job destruction? In light of the Hall-Shimer work discussed in the next lecture, it is tempting to use their model to affirm the "insulation effect". There remains the evidence on the Depression that they quote...

Question for study: why did C-H only examine to make this model stochastic? Can you figure a way to make it stochastic?

15. Campbell-Fisher (AER 2001)

Motivated mainly by Davis and Haltiwanger, Campbell and Fisher study the problem of a firm facing idiosyncratic shocks to productivity or demand as well as aggregate shocks, and adjusting labor subject to linear costs of adjustment. Prices are exogenous.

Here I follow their simple example of section 1. Consider the problem without aggregate shocks. Let τ_c = creation cost per job and τ_d = destruction cost per job. Then the corresponding Bellman equation is:

$$V(n_{-1},z)=\max_{n\geq 0}\left\{zn^{\alpha}-wn-\tau_c\left(n-n_{-1}\right)1_{n\geq n_{-1}}-\tau_d\left(n_{-1}-n\right)1_{n\leq n_{-1}}+\beta\sum_{z'\in Z}\pi_{z,z'}V(n,z')\right\},$$

which leads to the FOCs,

$$\begin{aligned}\text{If } n &> n_{-1}:\alpha z n^{\alpha-1}=w+\tau_c-\beta\sum_{z'\in Z}\pi_{z,z'}V_n(n,z'),\\ \text{If } n &< n_{-1}:\alpha z n^{\alpha-1}=w-\tau_d-\beta\sum_{z'\in Z}\pi_{z,z'}V_n(n,z').\end{aligned}$$

Let's further simplify by assuming that there are two possible values for z only: $z\in\{z_l,z_h\}$, and the transition matrix is:

$$P=\begin{pmatrix}1-p & p\\ p & 1-p\end{pmatrix}.$$

Clearly, the optimal policy if adjustment costs are small enough will be to adjust employment every time you change state, and only when you change state.

Hence upon switching from z_l to z_h, employment will shift from n_l to n_h where these values are determined by:

$$\alpha z_h n_h^{\alpha-1} = w + \tau_c - \beta \left\{ (1-p) V_n(n_h, z_h) + pV_n(n_h, z_l) \right\},$$

$$\alpha z_l n_l^{\alpha-1} = w - \tau_d - \beta \left\{ (1-p) V_n(n_l, z_l) + pV_n(n_l, z_h) \right\}.$$

Apply the envelope theorem to the value function to get

$$\begin{aligned} V_n(n_{-1}, z) &= \tau_c \text{ if } n > n_{-1} \rightarrow V_n(n_h, z_l) = -\tau_d, \\ &= -\tau_d \text{ if } n < n_{-1} \rightarrow V_n(n_l, z_h) = \tau_c. \end{aligned}$$

BUT here is the problem: I cannot apply this if $n = n_{-1}$ since the adjustment cost function has a kink at this point. Hence I cannot say what $V_n(n_h, z_h)$ is; Indeed, I cannot even say that the value function is differentiable at that point. From what will follow it is τ_c, but that's not obvious from our expressions.

This is why C&F do it the way they do, which is to give the law of motion for the expected present value of the MPL minus the wage, a variable they call v_h (or v_l) in the good (or bad) state:

$$\begin{aligned} v_h &= \alpha z_h n_h^{\alpha-1} - W + \beta \left((1-p) v_h + p v_l \right), \\ v_l &= \alpha z_l n_l^{\alpha-1} - W + \beta \left((1-p) v_l + p v_h \right), \end{aligned}$$

and they note that $v_h = \tau_c$ and $v_l = -\tau_d$, because when the state is high jobs are created to the point where the marginal value equals the cost. (That's a bit informal, but for the full details you may want to look at the general case of section 2 which has an appendix with proofs.)

With the last two equations you can solve for n_h, n_l and find:

$$\begin{aligned} \alpha z_h n_h^{\alpha-1} &= w + \tau_c \left(1 - \beta (1-p) \right) + \beta p \tau_d, \\ \alpha z_l n_l^{\alpha-1} &= w - \tau_d \left(1 - \beta (1-p) \right) - \beta p \tau_c. \end{aligned}$$

The intuition here is that jobs are created up to the point where the MPL equals wage plus the cost of creating the job, plus the expected cost of destroying later on when productivity reverses. Jobs are retained up to the point where MPL equals the wage minus the destruction cost you avoid paying now, minus the future creation costs you avoid paying later on. Because of the kink in the adjustment cost, there is an asymmetry in the policy function.

In this economy, there are at each point in time $\frac{1}{2}$ firms in good state and $\frac{1}{2}$ firms in bad state.[5] Hence each period there are $\frac{p}{2}$ firms which shift from z_l to z_h and $\frac{p}{2}$ firms which shift from z_h to z_l. Of course aggregate employment is constant. Gross job creation is $\frac{p}{2}(n_h - n_l)$. For simplicity I follow C&F and use a log-measure: $JC = \frac{p}{2}\ln\left(\frac{n_h}{n_l}\right)$. (In the paper they use precisely the measures used by Davis & Haltiwanger.)

Now that we have solved for the model with only idiosyncratic shocks, we can consider the impact of a jump in the wage W by 1% at time $t = 0$. Suppose this jump is permanent and unexpected. Then firms policies at time 0 will shift from the old policies to the new ones: adjust employment to n_h' (resp. n_l') when z shifts from z_l to z_h (resp. vice-versa).

To compute the aggregate response following this wage shock, note that firms are initially not in the stationary distribution, so that there will be some aggregate impact. The convergence to the new steady-state distribution will occur as firms change technology state and adjust to the new employment targets n''s.

Clearly,

$$(\alpha - 1)\frac{\Delta n_h}{n_h} = \frac{w}{w + \tau_c(1 - \beta(1-p)) + \beta p \tau_d}\frac{\Delta w}{w},$$

$$(\alpha - 1)\frac{\Delta n_l}{n_l} = \frac{w}{w - \tau_d(1 - \beta(1-p)) - \beta p \tau_c}\frac{\Delta w}{w}.$$

Hence in response to a 1% increase in the wage, both n_h and n_l fall but n_l falls by more than n_h since the wage is a bigger portion of the total cost (because adjustment costs are deducted at the retaining margin, and added at the hiring margin):

$$\frac{w}{w + \tau_c(1 - \beta(1-p)) + \beta p \tau_d} < 1 < \frac{w}{w - \tau_d(1 - \beta(1-p)) - \beta p \tau_c}.$$

Now a small trick: if there's a small fixed cost of changing employment on top of the linear adjustment cost, then firms which were in state z_h and had n_h employees will not find it optimal to adjust to n_h' as long as their z does not change - they will wait for an idiosyncratic shock change to adjust. As a result, aggregate job creation at the time of the shift $t = 0$ will be (in logs)

$$JC_{t=0} = \frac{p}{2}\ln\left(\frac{n_h'}{n_l}\right),$$

[5] This is the only invariant (i.e. stationary) distribution, i.e. the only solution to $\mu = \mu P$ where P is the transition matrix of the Markov process.

and job destruction will be

$$JD_{t=0} = -\frac{p}{2}\ln\left(\frac{n_l'}{n_h}\right).$$

We can now infer the change in JC and JD in response to the shock. Before they were equal. Now since W rises, $\Delta n_h < 0$, and $\Delta n_l < 0$ so job destruction rises, whereas job creation falls. This pattern of "cyclicality" is right. To assess the relative importance of JC and JD, note that

$$\begin{aligned}
\Delta JC &= \frac{p}{2}\Delta\ln(n_h) = -\frac{1}{1-\alpha}\frac{p}{2}\frac{w}{w+\tau_c(1-\beta(1-p))+\beta p\tau_d}\Delta\log W, \\
\Delta JD &= -\frac{p}{2}\Delta\ln(n_l) = \frac{1}{1-\alpha}\frac{p}{2}\frac{w}{w-\tau_d(1-\beta(1-p))-\beta p\tau_c}\Delta\log W,
\end{aligned}$$

$$\rightarrow \frac{\Delta JD}{\Delta JC} = -\frac{w+\tau_c(1-\beta(1-p))+\beta p\tau_d}{w-\tau_d(1-\beta(1-p))-\beta p\tau_c},$$

which is greater than 1 in absolute value. Hence JD is more volatile than JC because n_h reacts less than n_l to an increase in W as explained above.

An exercise is to compute (by hand) the response for JD and JC for $t = 1, 2, ...$ and not only $t = 0$ following this increase in W. (Hint: compute the cross-sectional distribution at each date before that.)

Beyond the example: the Bellman equation that they study quantitatively is

$$\begin{aligned}
V(n_{-1}, z, W) &= \max_{n\geq 0}\left\{\begin{array}{c} zn^\alpha - Wn - \tau_c(n-n_{-1})1_{n\geq n_{-1}} \\ -\tau_d(n_{-1}-n)1_{n<n_{-1}} + \beta E_{\substack{z'/z \\ W'/W}} V(n, z', W') \end{array}\right\}, \\
\ln z' &= \ln z + \mu + \sigma\varepsilon', \\
& \quad \varepsilon' \text{ iid } N(0,1) \\
& \quad W'/W \text{ Markov with transition } \Pi.
\end{aligned}$$

Note a useful "trick" to reduce the size of the state space is to use homogeneity. More precisely, let $x = n_{-1}/z^{1/(1-\alpha)}$, and let $y = n/z^{1/(1-\alpha)}$, then we can rewrite the Bellman equation as

$$g(x,z,W)=\max_{y\geq 0}\left\{\begin{array}{c} z.z^{\alpha/(1-\alpha)}y^{\alpha}-Wyz^{1/(1-\alpha)} \\ -\tau_c z^{1/(1-\alpha)}\left(y-x\right)1_{y\geq x}-\tau_d z^{1/(1-\alpha)}\left(x-y\right)1_{y<x} \\ +\beta E_{\substack{z'/z \\ W'/W}}\, g\left(y\left(\frac{z'}{z}\right)^{-\frac{1}{1-\alpha}},z',W'\right)) \end{array}\right\}.$$

Now we guess and verify that $g(x,z,W)=z^{1/(1-\alpha)}v(x,W)$. To do his, plug the guess on the RHS of the Bellman equation and check that you the expression will be of the same form. In our case:

$$\begin{aligned} & \max_{y\geq 0}\left\{\begin{array}{c} z.z^{\alpha/(1-\alpha)}y^{\alpha}-Wyz^{1/(1-\alpha)} \\ -\tau_c z^{1/(1-\alpha)}\left(y-x\right)1_{y\geq x}-\tau_d z^{1/(1-\alpha)}\left(x-y\right)1_{y<x} \\ +\beta E_{\substack{z'/z \\ W'/W}}\, g\left(y\left(\frac{z'}{z}\right)^{-\frac{1}{1-\alpha}},z',W'\right) \end{array}\right\} \\ = \; & z^{1/(1-\alpha)}\max_{y\geq 0}\left\{\begin{array}{c} y^{\alpha}-Wy-\tau_c\left(y-x\right)1_{y\geq x}-\tau_d\left(x-y\right)1_{y<x} \\ +\beta E_{\substack{z'/z \\ W'/W}}\left[\left(\frac{z'}{z}\right)^{\frac{1}{1-\alpha}}v\left(y\left(\frac{z'}{z}\right)^{-\frac{1}{1-\alpha}},W'\right)\right] \end{array}\right\} \\ = \; & z^{1/(1-\alpha)}\max_{y\geq 0}\left\{\begin{array}{c} y^{\alpha}-Wy-\tau_c\left(y-x\right)1_{y\geq x}-\tau_d\left(x-y\right)1_{y<x} \\ +\beta E_{\substack{\varepsilon' \\ W'/W}}\left[e^{[\mu+\sigma\varepsilon']\frac{1}{1-\alpha}}v(ye^{-[\mu+\sigma\varepsilon']\frac{1}{1-\alpha}},W')\right] \end{array}\right\}, \end{aligned}$$

hence we can simplify the equation to obtain

$$v(x,W)=\max_{y\geq 0}\left\{\begin{array}{c} y^{\alpha}-Wy-\tau_c\left(y-x\right)1_{y\geq x}-\tau_d\left(x-y\right)1_{y<x} \\ +\beta E_{\substack{\varepsilon' \\ W'/W}}\left[e^{[\mu+\sigma\varepsilon']\frac{1}{1-\alpha}}v(ye^{-[\mu+\sigma\varepsilon']\frac{1}{1-\alpha}},W')\right] \end{array}\right\}.$$

In the end we only have 2 state variables, so this is relatively easy to solve. A few remarks beyond this "trick":

(a) exogenous prices make this a tractable problem, and it's not clear what would be gained by going to an equilibrium wage determination;

(b) by making the idiosyncratic shock permanent rather than transitory, they go against their mechanism since it will work better with transitory shocks (since you know you will have to reverse your job creation or destruction);

(c) an interesting issue that they can examine is the effect of the trend μ, which is emphasized by Foote's paper in QJE 1998. (Foote has some suggestive empirical work but somehow he doesn't write down precisely a model, or show simulations to compare with the empirical work.) The idea is the following: when the idiosyncratic shock has a negative trend, units will be shrinking on average, so many units will have "too many" employees on average, i.e. the cross-sectional distribution of employment will have a lot of weight to the right. This implies that a negative aggregate shock will lead many of these units to adjust employment down, so job destruction would be more volatile than job destruction. This argument is however lose because optimal behavior means that, given the known trend, units will change their behavior (i.e. optimal points of adjustment), and so it is not clear exactly how the cross-sectional distribution would change. In this C&F paper the effect is there, but it is quantitatively small.

[Make all this paragraph more precise!]

(d) I find the argument of this paper natural and appealing. One relatively weak point is that adjustment costs need to be expressed in units of output not time, for the argument to work. Moreover it will be interesting to think back about this paper in light of the "new" findings on the cyclicality of hires and separations by Shimer and Hall.

Part III

Lecture 4

Outline: (0) End of Campbell-Fisher, (1) the Mortensen-Pissarides model, (2) Data on worker flows and the cyclicality of separations and hires, (3) Shimer 2005 AER and the effect of productivity shocks in the M-P model; the role of rigid wages.

16. The Mortensen-Pissarides model in continuous time

A popular model of unemployment. Possible sources of further reading: Pissarides' book: "Equilibrium Unemployment" (MIT Press); see chapter 1 for the basic model in discrete time. Also covered in Sargent-Ljungqvist. Here I cover the deterministic version, w/o aggregate shocks.

Workers and firms are risk neutral and maximize the expected discounted value of wages and profits respectively, using the discount rate r.

If you are not familiar with the continuous time model, the easiest way is to think of it as the limit of a discrete time model when the period becomes small (i.e. the period is between time t and time $t+dt$ with dt a small number). During each period dt, some workers have a job with a firm, which creates an output y, and for which they receive a wage w; they lose this job with probability s, upon which they join the unemployed pool. Firms each period create a certain number of "vacancies" i.e. job openings. Creating a vacancy for one period costs c. If there are V vacancies and U unemployed workers, then in a given period there will be $m(U,V)$ matches i.e. job created. m is a matching function: it stands for the process of search and matching, which is not modeled. It is typically assumed that m is increasing in both arguments and exhibits constant returns to scale, and I will maintain these assumptions. From the point of view of a firm then, a vacancy has a probability $m(U,V)/V$ to be filled. From the point of view of an unemployed worker on the other hand, the probability of finding a job is $m(U,V)/U$. I design by z the unemployment benefit that the worker receives will unemployed.

Let $\theta = V/U$ the vacancy-unemployment ratio (sometimes called "labor market tightness"). Remark that the probability of filling a vacancy is $m(U,V)/V = m(1/\theta,1) \stackrel{def}{=} q(\theta)$ with $q' < 0$. Moreover the job-finding rate is $m(U,V)/U = m(1,\theta) = f(\theta) = \theta q(\theta)$ which is increasing in θ: the probability of filling a vacancy is higher when the labor market is not tight and the probability of finding a job is higher when the labor market is tight.

I can now write the Bellman equation for the value of being unemployed U and the value of being employed E, where value means the expected discounted stream of income.

$$U = zdt + \frac{1}{1+rdt}\left(f(\theta)dt \times E + (1 - f(\theta)dt)\,U\right)$$

Intuition: value = income today (flow per unit of time, times the length of the period), plus expected value tomorrow: discount, then add the two possible events: w/prob $f(\theta)dt$ I will find a job, then I get the value E, or I stay unemployed, and get the value U.

Rearranging, suppressing second order terms (i.e. $dt^2 << dt$), and dividing by dt yields:

$$rU = z + f(\theta)\,(E - U)\,. \tag{16.1}$$

Note that the interpretation of this equation is also as the return on the "capital value" U is a dividend plus an expected capital gain (probability times possible gain).

Similarly for an employed worker I get:

$$rE = w + s\left(U - E\right).$$

Next we look at a firm value. The value of a firm with one worker is F and the value of an open vacancy is J, with

$$\begin{aligned} rJ &= -c + q(\theta)\left(F - J\right), && (16.2) \\ rF &= y - w + s\left(J - F\right). && (16.3) \end{aligned}$$

To complete the model we need to make assumptions regarding (1) the creation of new jobs and (2) the determination of the wage. Re (1), we assume free entry: $F = 0$. Re (2), the usual assumption in the litterature is Nash Bargaining. The motivation for this is that once a match has been formed, there is a bilateral monopoly situation where the worker and the firm have to negotiate and outside options are worse than an agreement because both parties need to search again. The Nash solution is to set w to maximize the product of surpluses,

$$\left(F - J\right)^{\beta}\left(E - U\right)^{1-\beta},$$

where β is a bargaining parameter.[6] The FOC for this w.r.t w is

$$\frac{\partial F}{\partial w}\beta\left(F - J\right)^{\beta-1}\left(E - U\right)^{1-\beta} + (1-\beta)\frac{\partial E}{\partial w}\left(F - J\right)^{\beta}\left(E - U\right)^{-\beta} = 0. \qquad (16.4)$$

Note that implicitly I assume that the two parties negotiating do not take into account the general equilibrium effect on U and J: the wage that they bargain over concerns only their current match.

First note that

$$\begin{aligned} \frac{\partial E}{\partial w} &= \frac{1}{r+s}, \\ \frac{\partial F}{\partial w} &= \frac{-1}{r+s}. \end{aligned}$$

(This reflects that an increase in the wage increases the PDV of being employed, discounted at rate $r+s$, and simultaneously reduces the PDV of a filled vacancy.)

[6] This functional form can be derived for an arbitrary bargaining function that satisfies (i) Pareto optimality, (ii) independence of irrelevant alternatives, (iii) invariance by affine transforms.

The equation (16.4) can be simplified to yield

$$\begin{aligned}\beta(E-U) &= (1-\beta)(F-J), \\ &= (1-\beta)F. \end{aligned} \tag{16.5}$$

The dynamics for the number of unemployed people, assuming a population of measure 1, is, if U now denotes the number of unemployed and E the number of employees:

$$\begin{aligned} dU_t &= U_{t+dt} - U_t = sE_t dt - f(\theta_t)U_t dt \\ \frac{dU}{dt} &= sE_t - f(\theta_t)U_t \\ &= s(1-U_t) - f(\theta_t)U_t \\ &= s - (s + f(\theta_t))U_t. \end{aligned}$$

In particular, unemployment will converge to the steady-state value U^* of this ODE, given by

$$U^* = u^* = \frac{s}{s + f(\theta)}. \tag{16.6}$$

This (implicit) relation between unemployment and vacancies is called the Beveridge curve. It is decreasing. Diagram $v, u+$ Beveridge curve : given θ find u.

Hence equilibrium unemployment is driven by the exogenous separation rate s and the tightness θ : a higher θ means less unemployment.

To solve for the steady-state, we have a system of 6 equations in 6 unknowns: w, E, U, F, J, θ.

$$F = 0$$

$$\beta(E-U) = (1-\beta)F$$

$$F = \frac{c}{q(\theta)}$$

$$\begin{aligned} rU &= z + f(\theta)(E-U) \\ rE &= w + s(U-E) \\ rJ &= -c + q(\theta)(F-J) \\ rF &= y - w + s(J-F). \end{aligned}$$

The last equation reads

$$F = \frac{y - w}{r + s},$$

which implies

$$\frac{c}{q(\theta)} = \frac{y - w}{r + s}. \tag{16.7}$$

This last equation is the "Job creation curve" (JC): it equates the cost of creating a filled vacancy (LHS) with the benefit of a filled vacancy. The cost is the vacancy cost times the expected time it takes to find a match. The benefit is the PDV of output minus the wage, discounted using the separation rate. Along JC, wages and market tightness are negatively related since a high wage means less profit from a filled vacancy and thus less vacancies posted from firms, so a lower θ.

Looking at the other side of the market:

$$\begin{aligned} J &= \frac{\beta}{1-\beta}(E - U) \\ \frac{c}{q(\theta)} &= \frac{\beta}{1-\beta}\frac{w - z}{r + s + f(\theta)} \\ w - z &= \frac{c\,(r + s + f(\theta))}{q(\theta)}\frac{1-\beta}{\beta} \end{aligned} \tag{16.8}$$

Note the RHS is increasing in θ. Here a higher market tightness leads to a higher wage: this is because a higher tightness pushes up the opportunity of unemployed which allows the workers to get a bigger wage from the bargaining process. The strength of this effect depends on the bargaining weight β.

Clearly, I can solve now for w and θ the system (16.7) and (16.8), which will yield a unique solution. With this in hand, I can perform simple comparative statics. For instance, an increase in the "unemployment benefit" z will increase w for given θ according to (16.8), so the outcome will be a higher w (because workers have better "outside threats" or bargaining points) and a lower θ (because profitability falls and so does job creation). With a lower θ, equation (16.6) implies a higher unemployment rate.

What's interesting about this model?

- Intuitive realism: the model captures the large ongoing flows of workers separating from their jobs, and of simultaneous new hires. It is a theory of "equilibrium unemployment" that gives you a reasonable reason (matching frictions) why there is on average 5% of unemployed people in the US.

- Tractability: it is easy to extend it and to do comparative statics. For instance one can incorporate endogenous destruction, idiosyncratic productivity shocks, and real-life regulations. (Search models tend to be a bit more complicated.)

- Finally (and more controversially), the MP model allows you to think of "involuntary unemployment", a concept many people like. As is apparent from our Bellman equations, workers don't make a choice, and the only economics are (1) the free entry condition and (2) the bargaining process. Unemployed in this model are "unlucky guys" who didn't drew a good match from the random matching process. On the other hand, search models and of course labor supply model (Lucas Rapping 1969 JPE and RBC model, esp. Hansen-Rogerson 1985-1988 JME) lead to the question, why do unemployed people turn down jobs? Why don't they search more? Why don't they lower their wage or job quality prospects?

- Note that this MP model does not make a distinction between worker flows and job flows.

17. Facts on Worker Flows and the Business Cycle Behavior of Unemployment, Vacancies, and the Job-Finding Rate

Part of the contribution of Shimer (and related work by Hall) is to highlight the business cycle behavior of the US labor market.[7] These aggregate facts can be summarized as:

Fact #1: the unemployment rate u is volatile, with a standard deviation of about 20%.

Fact #2: the number of vacancies v is also volatile, with a standard deviation of sligthly less than 20%.

Fact #3: u and v are strongly negatively correlated, so $\theta = v/u$ is highly volatile, with a standard deviation of about 35%.

Fact #4: the job-finding probability is volatile and highly positively correlated with θ. (This suggests a matching function yielding $f(\theta) = m(1, \theta)$ fits well - but

[7] Since these findings regarding the business cycle frequency movements in labor market variables are relatively strong, I suspect that they would hold for other countries. However this remains to establish.

as noted in class, we don't see the curvature of this matching function on Shimer's plots.)

Fact #5: the job-separation probability is not very volatile and not very cyclical.

Underlying data: u is measured by the BLS in a survey to find people "out of work, ready to start work, and actively looking for work". v is measured by JOLTS since December 2000, before by the help-wanted ads index. [See Shimer's slides.]

The job finding rate is measured from the equation for unemployment:

$$u_{t+1} = u_t - f_t u_t + u_{t+1}^{st}, \tag{17.1}$$

where u_{t+1}^{st} is short-term ($\leq$ 4 weeks) unemployment i.e. new entrants into unemployment. This yields

$$f_t = 1 - \frac{u_{t+1}^{st} - u_{t+1}}{u_t}.$$

The fact that this job-finding rate is highly countercyclical relative to u is directly due to u^{st} being more cyclical than u.

The separation rate is then inferred from

$$u_{t+1}^{st} = s_t e_t \left(1 - \frac{f_t}{2}\right),$$

given f_t, e_t, u_{t+1}^{st}. This equation reflects that a portion s_t of the stock of employees separates each month, but $f_t/2$ finds a job before the end of the month when unemployment is counted. Shimer in "reassessing the ins and outs of unemployment" looks at the issue of "what happens within the month" in more detail by going to continuous time. This is motivated by the huge number of very short-term jobs that people take - "the typical job lasts one day". The results thus obtained seem to amplify the facts above.

The key intuition is that when the job-finding rate is high, many separated workers find jobs again within the month. Hence the rate of separation that is naturally measured as $s_t e_t$ is biased: it will be higher when the job finding rate is low.

An issue that would seem to be important is the heterogeneity in job-finding rate and the job-to-job transitions. For the highly employable workers, the finding rate is very high, so I would guess that the job finding rate of the short-term unemployed is much higher than the finding rate of the average unemployed. This doesn't seem to show up in Shimer's computations.

Reconciliation between Davis-Haltiwanger and Shimer-Hall?

(1) Difference in sample: D-H cover manufacturing in the 70s to early 80s whereas S-H look at the whole economy and a longer sample. In particular since the mid 80s the pattern seems to be closer to the S-H view than to the D-H view.

(2) Difference in concepts: job flows at establishments vs. worker flows. [This is not much mentioned. First remember that for a given plant, Change in Employment = Hires − Separations = Job creation if Hires > Separations, and Hires − Separations = −Job destruction if Hires < Separations. One *can* imagine that the cyclical pattern of total hires and total separation is not informative re the pattern of JC and JD: suppose that in an expansion, hiring and separation increase at establishments losing jobs, whereas hiring increase and separation decrease at establishments creating jobs - then, the separation rate would be acyclical, hiring would be procyclical, JC would not vary much and JD would be volatile. This story doesn't seem very appealing though.)

(3) Difference in measurement: Shimer uses data on short-term unemployment to construct f_t and s_t, whereas D-H count the jobs across firms. (Previous work with flows of workers looked at the levels of job finding and separations, not the rates; examined the flows with the raw data, which is problematic because many transitions from U to E or E to U are measurement errors; and?).

Further remarks

One aspect of job separations that varies markedly over the business cycle is whether separations are layoffs or quits. The later are much more likely in upturns than in downturns. (Of course, because of the Coase theorem, the distinction between a layoff and a quit is not very easy to make for an economist.) See Davis, X and Haltiwanger 2005 JEP for more.

18. Shimer 2005 AER: the business cycle failure of Mortensen-Pissarides. Some possible fixes.

Shimer 2005 AER key point: given the small volatility of productivity, the MP model for the standard calibration dramatically underpredicts the volatility of vacancies and unemployment: there is not enough "amplification" of shocks. (Part of the contribution of this paper is to document the facts, which I did above.)

Shimer uses a variation of the M-P model above with aggregate shocks to productivity, but the key result is actually easy to see from a comparative statics exercise: what is the impact of an un expected, once-and-for-all increase in y? I can compute this using the deterministic model described above.

First, combine the two equations 16.7 and 16.8 to find θ as a function of exogenous variables:

$$\begin{aligned}
\frac{c}{q(\theta)} &= \frac{y-w}{r+s} \to y = w + \frac{(r+s)c}{q(\theta)} \\
w - z &= \frac{c\left(r+s+f(\theta)\right)}{q(\theta)} \frac{1-\beta}{\beta} \to \\
y - z &= \frac{(r+s)c}{q(\theta)} + \frac{c\left(r+s+f(\theta)\right)}{q(\theta)} \frac{1-\beta}{\beta} \\
y - z &= \frac{c}{q(\theta)} \left\{ r+s+\left(r+s+f(\theta)\right) \frac{1-\beta}{\beta} \right\} \stackrel{def}{=} \phi\left(\theta\right).
\end{aligned}$$

Next, differentiate to find the impact of a one-percent increase in y :

$$\frac{y}{y-z} d \log y = \frac{\theta \phi'(\theta)}{\phi(\theta)} d \log \theta.$$

$$\begin{aligned}
\frac{\theta \phi'(\theta)}{\phi(\theta)} &= -\frac{\theta q'}{q} + \frac{\theta f'(\theta) \frac{1-\beta}{\beta}}{r+s+\left(r+s+f(\theta)\right) \frac{1-\beta}{\beta}} \\
&= -\frac{\theta q'}{q} + \frac{\theta f'(\theta)}{f(\theta)} \frac{f(\theta)}{\left(r+s\right) \frac{\beta}{1-\beta} + \left(r+s+f(\theta)\right)} \\
&= -\left(\frac{\theta f'(\theta)}{f(\theta)} - 1 \right) + \frac{\theta f'(\theta)}{f(\theta)} \frac{(1-\beta) f(\theta)}{r+s+(1-\beta) f(\theta)} \\
&= 1 + \frac{\theta f'(\theta)}{f(\theta)} \left(\frac{(1-\beta) f(\theta)}{r+s+(1-\beta) f(\theta)} - 1 \right) \\
&= 1 + \frac{\theta f'(\theta)}{f(\theta)} \left(\frac{r+s}{r+s+(1-\beta) f(\theta)} \right)
\end{aligned}$$

Hence the elasticity is:

$$\begin{aligned}
\frac{d \log \theta}{d \log y} &= \frac{y}{y-z} \times \frac{1}{1 + \eta(\theta) \frac{r+s}{r+s+(1-\beta) f(\theta)}} \\
&= \frac{y}{y-z} \times \frac{r+s+\left(1-\beta\right) f(\theta)}{\left(r+s\right)\left(1+\eta(\theta)\right) + \left(1-\beta\right) f(\theta)},
\end{aligned}$$

where $\eta(\theta) = \frac{\theta f'(\theta)}{f(\theta)}$.

The first term reflects that a 1% increase in y holding z fixed will increase the net surplus from working by more than 1%. The second term reflects that job creation will increase (according to 16.7) if wages do not increase too much.

Plugging in some "reasonable" parameter values, we observe that the elasticity is not very big, around 1.5, so that θ and y will have roughly the same volatility. The parameters used are $y = 1$ (normalization), $z = 0.4$, $f(\theta) = 0.4$, $s = 0.03$, $r = 0.003$, $\eta(\theta) = 1 - \beta = 1 - 0.72 = 0.28$. The elasticity is the product of

$$\frac{r + s + (1 - \beta) f(\theta)}{(r + s)(1 + \eta(\theta)) + (1 - \beta) f(\theta)} \simeq 1,$$

and

$$\frac{y}{y - z} \simeq \frac{1}{0.6} \simeq 1.3,$$

which gives a product less than two.

Note that the reason for the result in the first term is simply that $f(\theta) >> r, s$.

The parameter choice that is most controversial is z and β. The later is not measured but is set equal to $\eta(\theta)$. This is known as Hosios' condition, it implies that the model's equilibrium is constrained efficient. (The two market failures - search and bargaining - balance each other in this case.)

Several reactions to this failure of the "reasonably calibrated" M-P model to account for the business cycle facts:

(1) Some have questioned the calibration. In particular, Manovskii et al. show that setting a very high z (unemployment benefit/home production) and a low β yields more amplification. Roughly, a high z implies that a percent increase in y makes it much more attractive to work in the market rather than at home. (This is the term $y/(y - z)$ in our formula.) A low β implies furthermore that the wage increase will not capture the increase in productivity. This solution amounts to saying that small differences in market productivity induces large changes in the relative value of working. This remains controversial among economists.

(2) A second possibility is to question the validity of looking at unemployment, and look at employment instead. We know that a standard RBC model does reasonably well at capturing the volatility of total hours worked (= hours per employee times the number of employees). If you believe that most unemployment is not related to search or matching frictions, but is a labor supply choice, this may be reasonable to do.

(3) The most common reaction is to blame the "flexibility" of wages in this model. Suppose that the wage is constant, set at some exogenous level $\overline{w}$. Then we can get rid of Nash bargaining since we don't need it to determine the wage, and θ is determined by JC:

$$\frac{c}{q(\theta)} = \frac{y - \overline{w}}{r + s}.$$

This will make θ much more responsive to a change in productivity, because y and $\overline{w}$ are similar:

$$\frac{\Delta \log \theta}{\Delta \log y} = \frac{1}{1 - \eta(\theta)} \frac{y}{y - \overline{w}} >> 1.$$

This nicely captures the "stylized facts" that wages are smooth (not very volatile) and profits are strongly procyclical. This result with an exogenous wage is discussed in Shimer 2003 JEEA, Hall 2005 AER, and it has motivated much recent interest in models of endogenous wage rigidity. One common story is that firms insure (imperfectly) workers against productivity fluctuations. (My general understanding of this endogenous wage rigidiy literature is that while it is easy to alter the pattern of wage, it is much harder to obtain significant deviations from the efficient allocation.)

It is important to note that the wage that matters is the wage of the new matches - the new hires. The wage of the people previously hired has no real impact on the allocation. (This, indeed, is why it is easy to think of a fixed wage in this setup, as Hall argued.) Understanding firms' compensation policies may be useful to resolve this puzzle.

Shocks to the separation rate do not appear very interesting because (1) the evidence suggests they are not very important, and (2) they tend to generate a negative correlation between u and v as firms create more vacancies to compensate for shorter jobs.

Part IV

Lecture 5

Outline of the lecture: (1) quick review of some recent work on firm-level volatility, (2) discussion of PS2, (3) basic search model and comparative statics, (4) equilibrium search (Lucas-Prescott).

19. Time-varying volatility

In the past few years there has been some interesting work on the evolution of firm-level volatility. I believe this is an active area for research, since (1) not too many formal models have been used to think about it and (2) the empirical work remains to be improved, and other data sources can be used.

Motivation

We've seen in this class that idiosyncratic shocks matter a lot more for firm-level dynamics than aggregate shock. As a result, it seems interesting to assess how the amount of "risk" that firms face has changed over time. (Note that "risk" is to be taken with caution, because it may be that what appears to us as unpredictable "risk" given our data, may be predictable to the manager/owner of the firm.)[8]

Facts

(1) there has been an increase in the volatility of firm-level sales (or employment) since the mid-1970s.

This is the key finding of studies by Comin and Philippon (NBER macro), who summarize several papers written recently.

(2) Campbell et al. (JF 2001) document that the idiosyncratic volatility of stocks has risen. (But some argue that it has started falling back since 2001.) Hence the explanatory power of the CAPM has fallen.

(3) Finally, Eisfeldt and Rampini (2006 JME) suggest that the cross-sectional distribution of productivity may widen in recessions. (The evidence for this seems weaker.) They view this as a sign of "good times to reallocate capital across firms"; since there is not much reallocation in recessions, it would appear that the costs of reallocating must be strongly countercyclical.

While these studies are intruiguing, one limitation is that they use Compustat instead of establishment-level data. I am not aware of any work using

[8] The parallel with similar work on consumption or income risk for households is perhaps useful. The mainstream conclusions from that research are essentially, at least for the US:

(1) there has been an increase in individual-level risk, since 1970, mostly in the permanent shocks, not the transitory ones. This has led to an increase in ex-post inequality. (See the labor literature, e.g. Moffitt.)

(2) income risk is also countercyclical: there seems to be more idiosyncratic risk in recessions (Storesletten, Telmer and Yaron JPE 2003.)

(3) there is some debate I think about shocks are very persistent or not too much. (Heaton and Lucas JPE 1996 found different results from Storesletten et al. who assume permanent shocks.)

establishment-level data that examines volatility, or even how business cycle patterns differ across establishments. Philippon and Comin look at industries and find that the correlation between industries has declined, not the volatility of each industry.

This pattern fits with the intuition that times are "more turbulent" today (because of deregulation?, more links with financial markets?, more trade?, technology?). However some pieces of the evidence don't seem to fit this. For instance, the average turnover has rather declined, so jobs appear to last longer. (See Davis, Haltiwanger and X JEP 2005 who also report some facts from the Census data which is at odds with the idea of increased volatility.)

[To check and elaborate]

20. Search: the basic model

Many assumptions to start (but we'll relax a lot of them):

No savings: current consumption = current income.

Each period draws a job offer from a c.d.f. F. This offer is iid over time.

Decide to accept or reject the job offer. (No recall.)

Expected discounted separable utility $E\sum_{t\geq 0}\beta^t u(c_t)$

Let $V(w)$ = value of a job at wage w = expected discounted utility, given has a job at wage w now.

Since job lasts forever:

$$V(w) = u(w) + \beta V(w) \rightarrow V(w) = \frac{u(w)}{1-\beta}$$

Value of being unemployed today: get benefit, and tomorrow draw a job offer, decide to accept or reject it:

$$\begin{aligned} U &= u(z) + \beta E \max(U, V(w)) \\ &= u(z) + \beta \int_0^\infty \max(U, V(w))\, dF(w). \end{aligned}$$

Let w^* be such that $U = V(w^*) = \frac{u(w)}{1-\beta}$. Then $\max(U, V(w)) = V(w)$ iff $w \geq w^*$ i.e. the searcher will take the job offer if he draws a wage above w^*, otherwise he will keep searching.

Thus:

$$\begin{aligned}
U &= u(z)+\beta\int_0^{w^*} U dF(w)+\beta\int_{w^*}^{\infty} V(w)dF(w) \\
U(1-\beta) &= u(z)+\beta\int_{w^*}^{\infty}(V(w)-U)\,dF(w) \\
u(w^*) &= u(z)+\frac{\beta}{1-\beta}\int_{w^*}^{\infty}(u(w)-u(w^*))\,dF(w) \\
(1-\beta)(u(w^*)-u(z)) &= \beta\int_{w^*}^{\infty}(u(w)-u(w^*))\,dF(w),
\end{aligned}$$

an equation which determines w^*, the reservation wage.

Comparative statics immediate with this equation:

- An increase in β will increase w^*.
- A FOSD shift in F will increase w^*.
- An increase in z will increase w^*.

In some cases it can be useful to rewrite the RHS of the previous formula using integration by parts:

$$\begin{aligned}
\beta\int_{w^*}^{\infty}(u(w)-u(w^*))\,dF(w) &= \lim_{A\to\infty}\left\{\beta\left[(u(w)-u(w^*))\,F(w)\right]_{w=w^*}^{w=A}-\beta\int_{w^*}^{A}u'(w)F(w)dw\right\} \\
&= \lim_{A\to\infty}\left\{\beta\,(u(A)-u(w^*))\,F(A)-\beta\int_{w^*}^{A}u'(w)F(w)dw\right\} \\
&= \lim_{A\to\infty}\left\{\beta\int_{w^*}^{A}u'(w)\,(1-F(w))\,dw\right\} \\
&= \beta\int_{w^*}^{\infty}u'(w)\,(1-F(w))\,dw.
\end{aligned}$$

The next comparative statics is a bit more interesting: assume $u(c)=c$ i.e. risk neutrality (or complete markets, which gives the same). Consider the effect of a shift of F in the SOSD sense. This will increase $\int_{w^*}^{\infty}F(x)dx$ by definition, and so this will decrease the RHS and increase w^* : because there is more "option value" of waiting for a better draw, people will have a higher reservation wage. Clearly if we don't have risk-neutrality this result will not be true in general since risk aversion will play a role.

Hazard rate $\stackrel{def}{=}$ probability of finding a job $p=1-F(w^*)$ constant over time.

Average unemployment duration is:

$$\sum_{t\geq 1}(1-p)^{t-1}pt=\frac{1}{p}.$$

[Proof: $\frac{d}{dx}\sum_{k\geq 1}x^k = \sum_{k\geq 1}kx^{k-1}$. But $\sum_{k\geq 1}x^k = \frac{x}{1-x}$ so $\frac{d}{dx}\sum_{k\geq 1}x^k$ is also $\frac{d}{dx}\frac{x}{1-x} = \frac{1}{(1-x)^2}$.]

Consider now the effect of unemployment benefits on output. Let b be the "home production" of an unemployed worker and let tr be a transfer ("unemployment insurance") from those that have a job to those that don't have one. What is the optimal tr? Question left for your study!

The "lake model" with exogenous job destruction→ predictions for unemployment duration and job-finding probability,

So overall this search model delivers intuitive results, and it is easy to extend the model to relax some assumptions:

(1) relax the assumption that jobs last forever: let λ be the probability of separation, then

$$V(w)=u(w)+\beta\left(\lambda U+(1-\lambda)V(w)\right).$$

Exercise: find the equation determining w^* and derive the comparative statics.

(2) relax the assumption that workers don't search for jobs. Assume a worker has a probability $\delta<1$ of getting a job offer. Then his Bellman equation is now

$$V(w)=u(w)+\beta\delta\int_0^\infty \max\left(V(u),V(w)\right)dF(u)+\beta(1-\delta)V(w).$$

Clearly with on-the-job search, searchers will take jobs sooner since they can keep searching thereafter.

[Can write down the model in continuous time as well. See Shimer-Rogerson-Wright who have a nice survey in the JEL (see Shimer's webpage).]

Question however: where does the wage distribution F come from. Diamond's paradox: if employers know that workers have this wage reservation rule, then it is clearly optimal for them to post wage offer only at w^*. Hence F becomes degenerate at w^*. But then everybody finds with the first draw and there is no "interesting" search problem.

Endogeneizing the wage offer distribution is thus an important step. Sargent and Ljungqvist cover some models by Jovanovic and Derek Neal. We will discuss another model due to Lucas and Prescott.

Application of the simple search model: Sargent and Ljungqvist (JPE 1998) show how the effect of an increase in "turbulence" on unemployment depends on institutions (unemployment benefits). "Turbulence" means that workers lose more of their human capital when they lose their job. In a turbulent economy, unemployment benefits become more attractive because the wages one can get after losing one's job are low. In this model the wage distribution is exogenous, and skills are upgraded or lost stochastically. The interesting decision is the reservation wage, as a function of how long you've been unemployed (which determines your benefits) and your current skills.

Other applications of search: money (Kiyotaki-Wright), marriage, etc.

21. Search: Equilibrium (Lucas-Prescott 1974 JET)

This is a theory of the natural rate of unemployment.

(Reading: if you prefer there is a presentation in Sargent and Ljungqvist.)

Setup

There is a continuum of competitive labor markets ("islands") with production function $sF(n)$. Islands are subjected to persistent idiosyncratic shocks s which are iid across islands. (There is no aggregate shock.) Agents choose optimally to leave an island when the island's productivity (and thus the wage) is too low. When they leave the island, they lose today's return, but they will arrive next period (after much paddling!) on another island. L&P assume that search is directed i.e. workers which change islands will be allocated in an efficient manner, i.e. to maximize the return to working. (Concretely, they "know" all the other opportunities - the other islands' shocks - and they allocate themselves so as to be indifferent across all the islands.)

Again here workers are risk-neutral. (And we can make them risk averse if there is no savings / insurance possible; it seems a good topic to look at the impact of the extent of insurance on the search behavior. See Alvarez and Veracierto?)

(1) Individual decision, given that searching has a return λ

Let $v(s, y; \lambda)$ be the expected discounted utility of a worker on an island with productivity x, where there are y workers at the beginning of the period, given that the expected discounted utility of going searching for a job on another island is λ (which we take as given now, and will determine in equilibrium).

Let w be the wage, given by the MPL if there are n workers and productivity is $x : w(s, n) \overset{def}{=} sF'(n)$, where n is the number of people working during this year (i.e., initial population minus the ones who leave the island).

There are three exclusive possibilities:

(a) Some workers leave the island (and some may come too):

$$v(s, y; \lambda) = \lambda.$$

Note that since some workers leave the island, it must be that it is not worth staying, or $w(s, y) + \beta Ev(s', y; \lambda) \leq \lambda$.

(b) All workers stay on the island, some workers come next period. The workers that do come must earn λ, since they are allocated across islands to have the same return everywhere. Thus the value for the workers on the island today is

$$v(s, y; \lambda) = w(s, y) + \lambda.$$

[Note λ includes discounting.]

(c) All workers stay on the island, nobody comes next period:

$$v(s, y; \lambda) = w(s, y) + \beta Ev(s', y; \lambda).$$

Note that since nobody comes, $\beta Ev(s', y; \lambda) \leq \lambda$.

Combining the three possibilities yields the Bellman equation:

$$v(s, y; \lambda) = \max \{\lambda, w(s, y) + \min(\lambda, \beta Ev(s', y; \lambda))\}$$

It is easy to check that the operator defined by the RHS is a contraction (Blackwell sufficient conditions work). Hence there is a unique solution to this Bellman equation. Moreover the usual monotonicity argument implies that v is increasing in s and λ and decreasing in y. From this monotonicity, I deduce that

- some agents leave, if $s \leq s^l(y)$, then next period's beginning of period population is given by $n^*(s, y; \lambda)$, where

$$w(s, n^*(s, y, \lambda)) + \beta \mathbb{E} v(s', n^*(s, y, \lambda); \lambda) = \lambda,$$

and note that n^* is independent of y.

- all workers stay, some come next period, with next period's beginning of period population given by

$$\beta \mathbb{E} v(s', n^*(s, y, \lambda); \lambda) = \lambda.$$

if $s \geq s^u(y)$. Here too n^* independent of y.

- all agents stay, nobody comes, then $n^*(s, y, \lambda) = y$. This occurs if $s \in [s^l(y), s^u(y)]$.

(2) Determination of the equilibrium λ

Plot $n^*(s, y, \lambda)$ against y.

Then $X = (s, n)$ is a Markov process. Can prove there is a unique invariant distribution under suitable technical assumptions, for given λ. Call it μ_λ.

Total employment is $N_\lambda = \int n(s, y; \lambda) d\mu_\lambda(s, y)$.

Assume a fixed supply of workers N per island. To find λ, equate the supply and the "demand" N_λ :

$$N_\lambda = N.$$

Under technical conditions, can show that N_λ is decreasing and continuous in λ so there is a unique equilibrium.

Remark 1: directed vs. undirected search → Alvarez and Veracierto (NBER macro annual) study a model where agents are allocated in equal numbers across islands (they randomly paddle from one island to another), they do not know which islands are "good" or "bad".

Remark 2: wage determination→ in Lucas-Prescott it is competitive, as opposed to the ex-post bilateral monopoly of matching.

Remark 3: Jovanovic, Derek Neal.

Part V

Lecture 6

This last lecture will introduce real option models. The idea is that firms have opportunities to invest, and can not only decide to invest or not invest, but also to delay investment. These models have attracted interest since the late 80s; the book by Dixit and Pindyck (Investment under Uncertainty) is a good summary. This will also be the occasion to have a look at some continuous time techniques. The key reading is Dixit 1989 JPE.

22. Simple introductory example

A firm can pay a cost c to do some investment, which will yield a profit $\widehat{\pi}$ today ($t = 0$) and starting tomorrow will yield forever a flow π each period. π however is unknown and distributed according to some c.d.f. F. Hence all uncertainty will be resolved tomorrow.

The net PDV of the project, if it is undertaken today, and assuming we discount the future at rate β, is

$$\begin{aligned} PDV^{today} &= -c + \widehat{\pi} + E\sum_{t\geq 0}\beta^t \pi \\ &= -c + \widehat{\pi} + \frac{\beta}{1-\beta}\int_0^{\infty} \pi dF(\pi). \end{aligned}$$

The standard decision rule is "invest if this net present value is positive, otherwise don't invest". However Dixit and Pindyck point out that you can also wait until tomorrow and then decide to invest only if π is high enough, i.e.

$$\begin{aligned} PDV^{wait} &= \beta E \max\left(\sum_{t\geq 0}\beta^t \pi - c, 0\right) \\ &= \beta \int_{c(1-\beta)}^{+\infty} \left(\frac{\pi}{1-\beta} - c\right) dF(\pi). \end{aligned}$$

Clearly there will be cases when waiting is better than investing today. To simplify further, let's assume that π can take two values, either 0 (w. probability $1-p$) or $\overline{\pi}$ (w. probability p); and assume that if $\pi = \overline{\pi}$, it is worth doing the investment.

In this case, it is better to wait if and only if

$$\begin{aligned} -c + \widehat{\pi} + \frac{\beta}{1-\beta}\int_0^{\infty} \pi dF(\pi) &\leq \beta \int_{c(1-\beta)}^{+\infty} \left(\frac{\pi}{1-\beta} - c\right) dF(\pi) \\ -c + \widehat{\pi} + \frac{\beta}{1-\beta} p\overline{\pi} &\leq \beta p \left(\frac{\overline{\pi}}{1-\beta} - c\right) \\ \widehat{\pi} &\leq c\left(1 - \beta p\right), \end{aligned}$$

which is more likely if the profits $\widehat{\pi}$ foregone by waiting are small relative to the gain of not doing the investment if the investment opportunity turns out to be bad (i.e. the $c(1-\beta p)$ which is saved; not the β reflects the fact that the investment is done only tomorrow, which also saves money).

This class of model suggests an important role for the arrival of new information: when there is more uncertainty, firms will typically want to wait more, and when the uncertainty gets resolved, decisions will be made. Something that does sound strange is the idea that firms have opportunities for investment that they can wait to activate - this seems to preclude perfect competition.

23. Continuous time stochastic processes (Brownian motion) and Bellman equation in continuous time

[This is really a list of recipes. Much better references: Nancy Stokey's lecture notes [web], books by Sheldon Ross "Stochastic Processes" [Wiley], Karlin and Taylor, Brzezniak and Zastawniak (Springer for undegraduates, esp. recommendended), and appendix of Dixit and Pindyck.]

- Definition: a Wiener process (or Brownian motion) is a process $Z(t)$ satisfying:

(i) [Independent increments] for any $0 \leq t_1 < t_2 < ... < t_N$, $Z(t_i) - Z(t_{i-1})$ is independent of $Z(t_{i-1}) - Z(t_{i-2})$

(ii) [Normal increments] for any $0 \leq t$, $Z(t) - Z(0)$ is normal, mean 0 and variance t

- Properties: (a) it is continuous; (b) it is a Markov process since its conditional distribution at time $t+s$, for all $s \geq 0$, depends only on its value at time t, $Z(t)$.

- Approximation: during a small interval of time t, $\Delta Z = \varepsilon_t \sqrt{\Delta t}$ with ε_t iid $N(0,1)$.

- In particular $E\Delta Z = 0$, $V\Delta Z = \sigma^2 \Delta t$. This implies that it is continous (but not differentiable since...).

- Notation: $dX = \sigma dZ$ is a brownian motion multiplied by a constant number σ.

- Brownian motion with drift: $dX = \mu dt + \sigma dZ$.

- General Ito process is $dX = \mu(X,t)dt + \sigma(X,t)dZ$ where μ andσ are functions of both the current state and time.

- Binomial approximation: one way to approximate a Brownian motion in the interval $[0,t]$ is to use a simple random walk: fix n a large integer, and let $\Delta t = t/n$ be a small time interval, and let $x(0) = 0$. Next we'll say that

$$
\begin{aligned}
x(n+1) &= x(n)+h \text{ w.probability } p \\
&= x(n)-h \text{ w. probability } q
\end{aligned}
$$

We now choose p, q to satisfy

$$
\begin{aligned}
E\Delta x &= \mu\Delta t, \\
V\Delta x &= \sigma^2\Delta t.
\end{aligned}
$$

i.e.

$$
\begin{aligned}
(p-q)\,h &= \mu\Delta t \\
h^2 &= \sigma^2\Delta t \rightarrow h=\sigma\sqrt{\Delta t}
\end{aligned}
$$

or:

$$
\begin{aligned}
2p-1 &= \frac{\mu}{\sigma}\sqrt{\Delta t} \\
p &= \frac{1}{2}\left(1+\frac{\mu}{\sigma}\sqrt{\Delta t}\right).
\end{aligned}
$$

With this choice of p and h, we have an approximation to a Brownian motion (μ, σ^2).

- More general Ito process:

$$
\begin{aligned}
dx &= \mu(x,t)dt+\sigma(x,t)dZ(t) \\
E\,(dx) &= \mu(x,t)dt \\
V\,(dx) &= \sigma^2(x,t)dt
\end{aligned}
$$

  μ is the (instantaneous) drift, σ the volatility. All these processes are continuous and Markov.

- Geometric brownian motion (analog to an iid growth rate):

$$
dx=\mu x dt+\sigma x dZ(t).
$$

  It corresponds to an iid growth rate. Using Ito's lemma, it is easy to see that $\ln x$ is a brownian motion with drift $\mu-\sigma^2/2$ and volatility σ; see below; hence $\ln x(t)$ is conditional on $x(0)$, a normal with mean $(\mu-\sigma^2/2)\,t+\ln x(0)$ and variance $\sigma^2 t$. In particular, $E(x(t)\mid x(0))=e^{\mu t}$ (log-normal formula).

- Ohrnstein-Uhlenbeck process (analog to an AR(1)):

$$dx(t) = \theta\left(x(t) - \overline{x}\right) + \sigma dZ(t).$$

Can show that $E(x(t) \mid x(0)) = e^{-\theta t}(x(0) - \overline{x})$, and can derive more formulas using the function $\phi(w) = E\left(e^{-wX}\right)$; see e.g. Dixit and Pindyck.

- Square-root process (stays always > 0): $dx(t) = \theta\left(x(t) - \overline{x}\right) + \sigma\sqrt{x(t)}dZ(t)$.

- Ito's Lemma: assume that $dx = \mu(x,t)dt + \sigma(x,t)dZ(t)$, and let $y = F(x,t)$. Then:

$$dy = F_x\left(x,t\right)\mu(x,t)dt + F_t\left(x,t\right)dt + \frac{F_{xx}(x,t)}{2}\sigma^2(x,t)dt + F_x\left(x,t\right)dZ(t).$$

A generalization of the usual formula. Note the F_{xx} term which comes from the fact that $dZ(t)^2 = dt$. [And $dZ(t)dt = 0$, $dt^2 = 0$.] Using Ito's lemma to find the relation b/w a Brownian motion and a Geometric Brownian motion: start with a brownian motion,

$$dx = \mu dt + \sigma dZ,$$

and let $y = e^x$. Apply Ito's lemma:

$$\begin{aligned} dy &= e^x \mu dt + \frac{e^x}{2}\sigma^2 dt + e^x \sigma dZ(t), \\ dy &= \left(\mu + \frac{\sigma^2}{2}\right) y dt + \sigma y dZ(t). \end{aligned}$$

- Bellman equation in continuous time. Suppose you wish to maximize $E\int_0^\infty e^{-rt}\pi(x(t),u(t))dt$ where u is a control, x is a state evolving as:

$$dx(t) = \mu(x(t),u(t))dt + \sigma(x(t),u(t))dZ(t).$$

Let $V(x_0)$ be the value of this problem. Then we have the Hamilton-Jacobi-Bellman (HJB) equation:

$$rV(x) = \sup_u \left\{\pi(x,u) + V'(x)\mu(x,u) + \frac{1}{2}\sigma^2(x,u)V''(x)\right\}$$

[Intuitive derivation:

$$V(x) = \sup_u \left\{ \pi(x,u) + \frac{1}{1+rdt} EV(x+dx) \right\}$$

Using Ito's lemma:

$$V(x+dx) = V(x) + V'(x)\mu(x,u)dt + V''(x)\frac{\sigma^2(x,u)}{2}dt + V'(x)\sigma(x,u)dZ(t).$$

And using $E(dt) = dt, E(dZ(t)) = 0$ yields the result above.]

- Cross-sectional distributions and Kolmogorov forward/backward equations: to be added...

24. Dixit's model of entry and exit under uncertainty (JPE 1990)

I highly recommend this article: it is very clear and relatively simple. Industry model with exogenous price following a geometric brownian motion

$$\frac{dP}{P} = \mu dt + \sigma dZ.$$

Firms can enter at a cost k into this market. If enter can produce one unit (flow per unit of time) at cost w. Can choose to exit at a cost l (which could be negative; we need to impose that $k \geq -l$). Clearly firms will decide to enter the industry at price P_H (to be found) and to leave at price P_L.

Read the discussion of the "Marshallian" (Vinerian?!) theory - in this case $P_H = w + rk$ and $P_L = w - rl$. The usual issue of recoverable costs. This can generate some inaction, but only with "surprises" (irrational expectations).

Let's now solve the decision problem of a firm to find P_H and P_L. Let $V_0(P)$ be the value of a firm inactive (i.e. not producing today) given today's price is P. The Bellman equation satisfied by V is, following our derivation above:

$$rV_0(P) = \mu P V_0'(P) + \frac{\sigma^2}{2} P^2 V_0''(P).$$

If the firm is active, the value function $V_1(P)$ satisfies:

$$rV_1(P) = P - w + \mu P V_1'(P) + \frac{\sigma^2}{2} P^2 V_1''(P).$$

The boundary conditions are given by the optimal switching from active to inactive when the price reaches the upper threshold (or lower one):

$$\begin{aligned} V_1(P_H) - k &= V_0(P_H), \\ V_1(P_L) &= V_0(P_L) - l. \end{aligned}$$

Moreover, the thresholds are chosen optimally. This implies the following "smooth pasting" condition that a marginal change in the threshold brings no gain:

$$\begin{aligned} V_1'(P_H) &= V_0'(P_H), \\ V_1'(P_L) &= V_0'(P_L). \end{aligned}$$

To solve the system of ODES, we need to add a general solution of the homogeneous equation to a particular solution. The particular solutions are easy to find:

$$\begin{aligned} \widetilde{V}_1(P) &= \frac{P}{r-\mu} - \frac{w}{r}, \\ \widetilde{V}_0(P) &= 0. \end{aligned}$$

The general solutions can be found as functions P^{ξ} where ξ satisfies necessarily

$$r = \mu\xi + \frac{\sigma^2}{2}\xi(\xi - 1).$$

It is easy to see that one root of this equation is negative, call it $-\alpha$, and one root is greater than 1, call it β. The general solutions are thus

$$V_i(P) = A_i P^{-\alpha} + B_i P^{\beta} + \widetilde{V}_i(P),$$

and we need only to find P_H, P_L, A_0, A_1, B_0, B_1 that satisfy the boundary conditions. Clearly we can "kill" some terms since as $P \to 0$, $V_0(P) \to 0$ since there will likely not be any entry, or very far into the future. Hence $A_0 = 0$. Similarly as $P \to \infty$, $V_1(P) \simeq \widetilde{V}_1(P)$ since the firm will likely never exit. (The precise interpretation of this limit condition is that the option value of closing is worthless.) Hence $B_1 = 0$. At the end we're thus left with a system of 4 nonlinear equations in 4 unknowns.

Key theoretical results: if there is no uncertainty, the entry and exit prices are as usual $w + rk$ and $w - rl$. If there is some uncertainty, $P_H > w + rk$ and $P_L < w - rl$: firms become less responsive to prices because the value of waiting for more information is important. This generates "hysteresis" i.e. the quantity supplied depends not only on today's price but also on the past prices. (These models attracted interest in the late 80s in the face of big changes in exchange rates w/o corresponding adjustment of trade flows.)

If there are no sunk costs of course P_H and $P_L \to w$ so uncertainty does not matter.

Finally an important result, that shows up also in the numerical results, is that $dP_H/dk \to \infty$ as $k \to 0$ so a small sunk cost has large effect.

Extension: different processes for P lead to very different behavior. In the exchange rate (XR) application, if XR are thought to mean revert in the long term, this means that the response to an increase in the XR today will not lead to much supply reaction because firms think the XR will revert. Can also extent for more complex firms problems (variable output, etc.) and for a risk-adjustment in the discount rate.

Limitation: exogenous price. The supply reaction is shown to exhibit hysteresis, given this price process, but of course we don't know if this price process would be generated in equilibrium by firms' actions→ this motivates the presentation of the Caballero-Pindyck model.

Follow-up paper: Dixit 1991 studies in a related model investment with price ceilings (e.g. rent controls). This makes firms' investment even more sluggish.

Exercise: try and write this in discrete time! (You might not have so many nice formulas, but it should be very easy to put on the computer.)

25. Making the price endogenous: Caballero and Pindyck IER 1996 (Sketch)

I simplify a bit the Caballero-Pindyck model by suppressing idiosyncratic shocks and aggregate productivity shocks. The only uncertainty is on the demand curve. The aggregate demand curve is $Q_t = M_t P_t^{-\eta}$ for $\eta > 0$ and M_t follows a geometric brownian motion:

$$\frac{dM_t}{M_t} = \mu_m dt + \sigma_m dZ_t.$$

Each firm produces one unit, yielding a total supply $Q_t = N_t$. The total number of firms N_t evolves as

$$dN_t = dE_t - \delta N_t dt,$$

with $dE_t \geq 0$: firms can enter but can not choose to exit, they must wait for a "death" which arrives at rate δ. This irreversibility will lead to some option value, since firms will not enter exactly when they can break even - they will wait for the price to rise further. The free entry condition is:

$$F \geq E_t \int_0^\infty e^{-(r+\delta)(s-t)} P_{t+s} ds,$$

with equality if $dE_t \geq 0$. F is an entry cost, independent of the number of entrants.

Let $P(M,N)$ be the price that prevails given that there are N firms and that the demand shock is M :

$$Q = MP^{-\eta} = N \rightarrow P(M,N) = \left(\frac{N}{M}\right)^{\frac{1}{\eta}}.$$

Let $x = \frac{N}{M}$. Ito's lemma can be adapted for a function of two-variables:

$$z = F(x,y) \rightarrow dz = F_x dx + F_y dy + \frac{1}{2} F_{xx} dx^2 + \frac{1}{2} F_{yy} dy^2 + F_{xy} dx dy.$$

Hence the following process for x :

$$\begin{aligned} dx &= \frac{dN}{M} - \frac{NdM}{M^2} + \frac{1}{2}\frac{2NdM^2}{M^3} - \frac{dNdM}{M^2} \\ &= \frac{dE - \delta N dt}{M} - x\left(\mu_m dt + \sigma dZ_m(t)\right) + x\sigma_m^2 dt \\ &= \frac{dE}{M} - \left(\mu_m + \delta - \sigma_m^2\right) x dt - x\sigma dZ_m(t), \end{aligned}$$

where $dE/M = 0$ if there is no entry, and dE/M adjusts to keep $E_t \int_0^\infty e^{-(r+\delta)(s-t)} P_{t+s} ds = F$ otherwise (enough entry to make the price decrease to this value).

I conjecture the following equilibrium: x is the only state variable, and the present discounted value $V(x) = E_t \int_0^\infty e^{-(r+\delta)(s-t)} P_{t+s} ds$ will be either F, if there is entry, in which case x will be constant, equal to some threshold x^*, with supply growing to meet the demand, or $< F$, in which case x will move stochastically in $[0, x^*]$. Thus we have for $x < x^*$

$$\begin{aligned} (r+\delta)V(x) &= P(x) + V'(x)\mu(x) + \frac{1}{2}V''(x)\sigma^2(x), \\ P(x) &= x^{1/\eta}, \\ \mu(x) &= -\left(\mu_m + \delta - \sigma_m^2\right) x, \\ \sigma(x) &= -x\sigma. \end{aligned}$$

Moreover since x^* is the point at which entry occurs, we have $V(x^*) = F$, and by optimality of the threshold of entry $V'(x^*) = 0$.

This is now an ODE that I can solve.

Particular solution: $V(x) = Kx^{1/\eta}$ with

$$(r+\delta)K = 1 - \left(\mu_m + \delta - \sigma_m^2\right)\frac{K}{\eta} + \frac{1}{2}\sigma_m^2\frac{K}{\eta}\left(\frac{1}{\eta} - 1\right)$$

$$\to K = \left(r + \frac{\mu_m}{\eta} + \frac{1+\eta}{\eta}\delta - \frac{1}{2}\frac{\eta+1}{\eta^2}\sigma_m^2\right)^{-1}.$$

General solution:

$$V(x) = A_0 x^{\lambda_0} + A_1 x^{\lambda_1},$$

where the λ's solve

$$\begin{aligned}(r+\delta) &= -\lambda\left(\mu_m + \delta - \sigma_m^2\right) + \frac{1}{2}\sigma_m^2\lambda(\lambda - 1)\\ &= \lambda\left(-\mu_m - \delta + \frac{\sigma_m^2}{2}\right) + \frac{1}{2}\sigma_m^2\lambda^2,\end{aligned}$$

i.e.

$$\lambda_{0,1} = \frac{\mu_m + \delta - \frac{\sigma_m^2}{2} \pm \sqrt{\left(\mu_m + \delta - \frac{\sigma_m^2}{2}\right)^2 + 2(r+\delta)\sigma_m^2}}{\sigma_m^2}.$$

Note that we can rule out the negative root since it would imply an infinite value for low x. Hence

$$V(x) = A_0 x^{\lambda_0} + Kx^{\frac{1}{\eta}},$$

with $\lambda_0 > 0$, and we now need only to find x^* and A_0 such that

$$\begin{aligned}V'(x^*) &= 0,\\ V(x^*) &= F,\end{aligned}$$

i.e.:

$$\begin{aligned}A_0\lambda_0 x^{*\lambda_0 - 1} + \frac{K}{\eta}x^{*\frac{1}{\eta}-1} &= 0 \to A_0 x^{*\lambda_0} = -\frac{K}{\lambda_0\eta}x^{*\frac{1}{\eta}} = 0\\ A_0 x^{*\lambda_0} + Kx^{*\frac{1}{\eta}} &= F \to Kx^{*\frac{1}{\eta}}\left(1 - \frac{1}{\lambda_0\eta}\right) = F\end{aligned}$$

$$Kx^{*\frac{1}{\eta}}\left(1-\frac{1}{\lambda_0\eta}\right)=F$$

$$x^{*\frac{1}{\eta}}=\frac{F}{K}\frac{\lambda_0\eta}{\lambda_0\eta-1}$$

Plot x and the PDV $V(x)$: both are asymmetric, falling when demand falls, but getting back up to x^* (resp. F) when demand rises. Given the path for $x = N/M$ and given the realization of M can compute what happens to $Y = N$ and thus to the price.

Comparative statics appear to be a bit complicated, but Caballero and Pindyck do derive some, so this should be doable using this expression and the formula for λ_0.

But I think this model is most interesting because it shows how an industry equilibrium with endogenous prices can be constructed, while preserving the microeconomic intuition: firms will wait until there is a lot of "excess demand" to enter (i.e. they wait until the price rises quite a bit above the level necessary to cover their cost in a static context), because they rationaly expect that when demand falls, the price will fall and they will make losses, being unable to exit the industry easily.

C&P do a little bit of empirical work to study the impact of firm- or industry-level uncertainty on investment behavior.

26. Dixit's approximation to the menu cost model (RES 1991)

Akerlof and Yellen and Mankiw simultaneously (1985) pointed out that around an optimum, the cost of inaction is small because first-order effects are zero:

$$x^*\in\arg\max_x \pi(x)\rightarrow\pi'(x^*)=0,$$

and

$$\pi(x)\simeq\pi(x^*)+\pi'(x^*)(x-x^*)+\pi''(x^*)\frac{(x-x^*)^2}{2}=\pi(x^*)+\pi''(x^*)\frac{(x-x^*)^2}{2},$$

i.e. first order changes in x have only second order changes in π. The point is that "small" costs of inaction, be they physical or mental, will generate inaction, since at some point the costs will outweigh the benefit of adjusting.

Stylized version of the menu cost model (justified by a second-order approximation). A state variable oscillates according to

$$dx = \sigma dz.$$

At a cost g, you can reset x to any desired value. Your objective is to minimize the PDV of adjustment costs plus $kx^2 e^{-\rho t}$, i.e. you wish to keep the x close to zero.

Clearly the optimal policy is to reset x to 0 whenever you reach some threshold x^* (or $-x^*$). To determine x^*, we will need to use the value function v defined as

$$\begin{aligned}
v(x_0) &= \min_{T_i} \left\{ \int_0^{\infty} e^{-\rho t} k x_t^2 dt + \sum_{i \geq 0} e^{-\rho T_i} g \right\}, \\
s.t. &: \\
dx &= \sigma dz, \\
x(T_i^+) &= 0.
\end{aligned}$$

The Hamilton-Jacobi-Bellman (HJB) equation is

$$\rho v(x) = kx^2 + v''(x) \frac{\sigma^2}{2}, \tag{26.1}$$

for all $x \in [-x^*, x^*]$.

Moreover, for $x = x^*$ we adjust x to zero at cost g, hence we have the following condition (called "value matching"):

$$v(x^*) = v(0) + g.$$

Finally, a last condition is given by the following "smooth pasting" condition which ensures that x^* is chosen optimally:

$$v'(x^*) = 0.$$

Solution: clearly v is symetric around zero. First, let $\alpha = \sqrt{2\rho/\sigma^2}$, then the solution of the ODE (26.1) is

$$v(x) = A \left[\exp(\alpha x) + \exp(-\alpha x) \right] + \frac{kx^2}{\rho} + \frac{k\sigma^2}{\rho^2},$$

[I have added a particular solution to the general solution of the 2nd order ODE, and imposed the symetry.] Next we need only to find A and x^*, i.e. solve the following system of equations in A, x^* :

$$\begin{aligned} A\left[\exp\left(\alpha x^*\right)+\exp\left(-\alpha x^*\right)\right]+\frac{k x^{*2}}{\rho} &= A\times 2+g, \\ \alpha A\left[\exp\left(\alpha x^*\right)-\exp\left(-\alpha x^*\right)\right]+\frac{2 k x^*}{\rho} &= v'(x^*)=0. \end{aligned}$$

We can do a Taylor approximation to these equations, putting in enough terms, to find an approximation for x^* and A. You actually need 4 orders, and you get

$$x^*=\left(\frac{6\sigma^2 g}{k}\right)^{1/4}.$$

The key point: now a fourth-order cost g will generate first-order inaction. This is because the Akerlof-Yellen-Mankiw effect is magnified by the option value: now you may not want to adjust x because you think it will probably mean-revert on its own. (This of course depends heavily on the fact that the shocks tend to mean revert.)