FINAL EXAM ANSWERS

(3) (45 points) Consider an overlapping generations model with constant population and a linear investment technology with gross rate of return θ. The constraints for an agent in generation t are

$$C_1(t) + S_t + B_t + \tau_t = \bar{Y} \tag{1}$$

$$C_2(t+1) = \theta S_t + R_t B_t + \omega_{t+1} \tag{2}$$

$$S_t \geq 0\,, \qquad B_t \geq 0\,, \tag{3}$$

where $C_1(t)$ and $C_2(t+1)$ are, respectively, consumption in the first and second periods of life by an agent born at t, S_t is real savings, B_t is government bonds acquired in the first period of life by generation t, R_t is the interest rate on government bonds issued at t, τ_t is a per capita lump sum tax on the young at t, and ω_{t+1} is the per capita lump sum transfer to the old at $t+1$. The utility function of each generation is $\log(C_1(t)C_2(t+1))$. The government budget constraint is

$$B_t + \tau_t = R_{t-1}B_{t-1} + \omega_t\,. \tag{4}$$

(a) (12 points) Consider an equilibrium in which there is a "pay-as-you-go" government retirement plan, so that $B_t \equiv 0$ and $\tau_t \equiv \bar{\tau} \equiv \omega_t$. Show that if $\theta > 1$, in a steady state equilibrium where C_1 and C_2 are each constant across generations, both steady state S and utility are decreasing functions of $\bar{\tau}$, at least over some range of values for $\bar{\tau}$. Show in particular that equilibrium without pay-as-you-go retirement, i.e. with $\bar{\tau} = 0$, delivers higher utility than any equilibrium with $\bar{\tau} > 0$.

The FOC's are:

$$\partial C_1: \qquad \frac{1}{C_1(t)} = \lambda_1(t) \tag{A3.1}$$

$$\partial C_2: \qquad \frac{1}{C_2(t+1)} = \lambda_2(t+1) \tag{A3.2}$$

$$\partial S: \qquad \lambda_1(t) = \theta E_t \lambda_2(t+1) \qquad (\text{if } S_t > 0) \tag{A3.3}$$

$$\partial B: \qquad \lambda_1(t) = R_t E_t \lambda_2(t+1) \qquad (\text{if } B_t > 0)\,, \tag{A3.4}$$

which can be reduced, in this model with no uncertainty, to

$$\frac{C_2(t+1)}{C_1(t)} = \theta \qquad (\text{if } S_t > 0) \tag{A3.5}$$

$$\frac{C_2(t+1)}{C_1(t)} = R_t \qquad (\text{if } B_t > 0)\,. \tag{A3.6}$$

Using the FOC's, the private budget constraints, and the assumptions $\tau_t = \omega_t \equiv \bar{\tau}$, $B_t \equiv 0$, *we can solve to obtain*

$$C_1(t) = \tfrac{1}{2}\left(\bar{Y} - \left(1 - \frac{1}{\theta}\right)\bar{\tau}\right)\,. \tag{A3.7}$$

This is the usual result with log utility that first period consumption is half the present value of total endowment. Though the transfer from the government when old is equal to the tax revenue extracted when young, because the transfer comes later and $\theta > 1$, it has lower present value. Hence the budget constraint of the private agent, in terms of $C_1(t)$ and $C_2(t+1)$, shifts inward as $\bar{\tau}$ increases. Since the relationship is linear, steady state welfare is clearly higher with $\bar{\tau} = 0$ than with any positive value of $\bar{\tau}$.

(b) (5 points) Show that this conclusion does not hold if $\theta < 1$, and explain why.

The coefficient on $\bar{\tau}$ in (A3.7) switches sign when θ drops below 1, so in this case increasing τ shifts the budget constraint out and increases welfare. Increasing τ reduces the need for retirement savings. When $\theta > 0$, the total resources in any period in steady state are higher, the higher is steady state savings. But when $\theta < 0$, with constant population, we are in a situation of dynamic inefficiency. Because saved resources dissipate rather than grow, total resources per period in steady state are higher when there is no steady state savings.

(c) (7 points) Show that if at some particular date policy unexpectedly changes from pay-as-you-go, with $\bar{\tau} > 0$, to autarky, with $\bar{\tau} = 0$ and still $B_t \equiv 0$, not all generations benefit. Who gains and who loses?

The short answer: Those old at the date of the policy switch lose, because they do not get their expected retirement benefits, and everyone thereafter goes to the autarky solution, which is better (assuming $\theta > 1$). If $\theta < 1$ everyone loses: The current old lose their expected retirement benefits, and everyone thereafter goes to the autarky solution, which in this case is worse.

(d) (7 points) How would your answer to (3c) change if the switch in policy were anticipated one period in advance?

If they know the change is coming when they are young, those who will be old at the time of the switch will recognize at $T-1$ (if T is the switch date) that their budget set is smaller and will therefore save more than they otherwise would have. This will reduce the impact on their utility of the switch in policy. This will have no effect on those old at $T-1$, because they will still be consuming their transfer payments plus their storage, neither of which changes.

(e) (14 points) Consider instead an unanticipated switch at $t = 0$ from pay-as-you-go with $\theta > 1$ and $S_t > 0$ to no retirement benefits, with the transition debt-financed. That is, a switch from $\omega_t \equiv \bar{\tau} \equiv \tau_t$ for $t < 0$ to $\omega_t = 0$ for $t > 0$, with $\omega_0 = \bar{\tau}$, $\tau_t \equiv \tau^*$ for $t \geq 0$, and $\omega_0 = B_0 + \tau^*$. We assume that τ^* is set so that the government budget constraint (4) delivers constant $B_t \equiv B_0$ for all $t > 0$. Who gains and loses under this policy switch?

The really quick answer to this question, which either nobody saw or nobody had the courage to rely on, is this: So long as $S > 0$ and agents are faced with θ as the rate of return on savings, which matches the actual rate at which resources can be moved between periods, the equilibrium is efficient. That means no generation can be made better off without making another worse off. Since in this part's scenario the old at the switch date have the same welfare as before, and since all generations after that have the same welfare (facing the same pattern of taxes and the same rate of return on saving), everyone must have the same welfare.

A more constructive answer: Since the old at $t = 0$ get exactly their anticipated transfer payment, their welfare is unaffected by the switch. The discounted present value of endowments for the young at t is $\bar{Y} - \tau^*$. Assuming that we have $S_t > 0$ after the switch, we will have $R_t \equiv \theta$. The constancy of the debt then implies

$$B_0 = \theta B_0 - \tau^* \,, \tag{A3.8}$$

which together with the problem's assumption that $B_0 + \tau^* = \bar{\tau}$ implies

$$\tau^* = \left(1 - \frac{1}{\theta}\right)\bar{\tau} \,. \tag{A3.9}$$

But since the return on savings is θ before and after the switch, this just says that the position of the budget set is exactly the same *before and after the switch. Before, the agent perceives a tax when young and a transfer when old, with present value $(1 - 1/\theta)\bar{\tau}$. After the switch, the agent sees only a much smaller tax when young, but this smaller tax has a current value equal to the present value of the old tax-transfer pair. Thus the consumption choices will be the same before and after the switch. The debt-financed policy change has no effect on consumption, welfare, or real savings at any date.*

It is perhaps worth noting, though it's not part of the answer to this question, that if taxes were distorting instead of lump-sum, there would be an efficiency-increasing effect of the switch, because tax rates on the young would be lower.

(4) (45 points) Suppose an economy has a Phillips curve of the form

$$u_t = \bar{u} - \alpha(\pi_t - \hat{\pi}_t) + \varepsilon_t$$

and a policy authority that maximizes

$$-\frac{1}{2} E\left[\sum_{t=0}^{\infty} \beta^t (u_t^2 + \theta \pi_t^2)\right] .$$

Suppose that $\hat{\pi}_t$, the public's subjectively expected inflation, is formed by the simple adaptive rule $\hat{\pi}_t = \pi_{t-1}$.

(a) (10 points) Find first order conditions for an optimum in the government's planning problem, assuming the government chooses at each date t both π_t and u_t knowing the history of all variables up to time t, is subject to the Phillips curve as a constraint, and rationally takes account of the effects of its current actions on the public's expectations of future inflation.

The Euler equations are:

$$\partial\pi: \qquad \theta\pi_t = \alpha\lambda_t - \beta\alpha E_t \lambda_{t+1} \tag{A4.1}$$

$$\partial u: \qquad u_t = \lambda_t \tag{A4.2}$$

which give us, after eliminating λ, the two-equation system, in matrix notation,

$$\begin{bmatrix} \alpha & 1 \\ 0 & \beta \end{bmatrix} \begin{bmatrix} \pi_t \\ u_t \end{bmatrix} = \begin{bmatrix} \alpha & 0 \\ -\theta & \alpha \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \varepsilon_t + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \eta_t \,. \tag{A4.3}$$

The TVC is

$$\limsup_{t\to\infty} \beta^t E[(\alpha u_t - \theta\pi_t) d\pi_t] \le 0 \,. \tag{A4.4}$$

The conditions in the notes for the application of the standard TVC are not met here, because the constraints are equality constraints and because the choice variables are not constrained to be always positive. If we separate the equality constraint into two inequalities with separate, always positive multipliers, we still end up with (A4.4). *But the "$d\pi$" terms require careful consideration, because they are not naturally bounded and could have any sign.*

(b) (10 points) *Here and in all the parts of this question below, assume $\theta = \alpha = 1$, $\beta = .9$.* Show that the Euler equations and constraint have a unique non-explosive solution. [This will involve finding the roots of a quadratic equation. One-decimal-place accuracy is enough.] Show the solution for π_t and u_t as functions of lagged π and u and current ε.

The matrix approach to solving this starts by finding the eigenvectors of

$$\begin{bmatrix} \alpha & 1 \\ 0 & \beta \end{bmatrix}^{-1} \begin{bmatrix} \alpha & 0 \\ -\theta & \alpha \end{bmatrix} = \frac{1}{.9} \begin{bmatrix} 1.9 & -1 \\ -1 & 1 \end{bmatrix} = A\,, \tag{A4.5}$$

where the two matrices on the left are from (A4.3). *This boils down to finding the roots of the quadratic equation $1 - 2.9\lambda + .9\lambda^2$, which are 2.83 and $.393$. To suppress the unstable root of 2.83, we must have*

$$\begin{bmatrix} 1 & x \end{bmatrix} \begin{bmatrix} \pi_t \\ u_t \end{bmatrix} = constant\,, \tag{A4.6}$$

where $[1\ x]$ is the left eigenvector of A corresponding to the unstable root. We can find x (among other ways) by solving the linear equation $1.9 - x = .9 \cdot 2.83$, to get $x = -.647$. This stability condition then together with the original Phillips Curve constraint becomes

$$\pi_t - .647\bar{u} + .647\pi_t - .647\pi_{t-1} - .647\varepsilon_t = constant$$
$$\therefore \quad \pi_t = constant + .39\pi_{t-1} + .39\varepsilon_t\,. \tag{A4.7}$$

The Phillips curve itself tells us that in a stable solution (where the unconditional expectation $E\pi_t$ is constant), $Eu_t = \bar{u}$. Then the Euler equations tell us that $E\pi_t = (\alpha/\theta)(1-\beta)\bar{u}$, and thus that the constant term in (A4.7) *is $.1\bar{u} \cdot .61 = .061\bar{u}$. The solution of u_t can then be read off by substituting the right-hand side of* (A4.7) *into the Phillips Curve.*

This part could also be answered by substituting the Phillips curve into the first order condition to obtain a single second-order difference equation in π_t, which has the same stable and unstable root as that we solved above.

(c) (5 points) Explain why the explosive solutions to the system are not optima. [A direct argument that they don't produce good values of the objective is probably easier than relying on formal transversality here.]

On an unstable path at least one of u_t and π_t has to grow exponentially at 2.83^t, eventually. Because the objective function is discounted only at $.9^t$, and because it is a sum of squares, it is easily seen to be minus infinity along any such path. Since we know that a stationary solution delivering bounded losses is possible, it is clear that the explosive solutions can't be optimal. An argument based on transversality is possible. Because we have used only sufficiency of transversality, we need here to show that the TVC for the stable solution to the Euler equations is satisfied, which *guarantees that it is the optimum. The stationary solution makes the coefficient on* $d\pi_t$ *in the TVC positive, at least most of the time. A positive, exponentially growing* π_t *is certainly technically feasible according to the problem statement. So we have to go further and check whether the Euler equations and constraint could be satisfied on such an upwardly explosive path for* π_t*. From the Phillips curve, we see that exponentially upward explosion in* π *implies exponentially downward explosion, at the same rate, in u. But the Euler equation implies that upward explosion in* π *has to be associated with upward explosion in u. Thus the two cannot be satisfied simultaneously on an upwardly explosive path for u, and the TVC is indeed satisfied for the stable solution.*

(d) (7 points) Suppose instead that policy makers change every period, and that each new policy maker at t takes $\hat{\pi}_t$ as given, ignoring the effect of her π_t choice on future $\hat{\pi}$'s. What is the resulting behavior of π and u as a function of lagged π and u and current ε?

This just yields the familiar no-commitment solution. The $E_t\lambda_{t+1}$ *term disappears from* (A4.1) *and leads to the solution* $u_t = \frac{1}{2}(\bar{u} + u_{t-1} + \varepsilon_t)$, $\pi_t = u_t$.

(e) (7 points) Suppose now that the policy authority (mistakenly) believes the public to be rational, so that she solves for optimal policy assuming $\hat{\pi}_t = E_{t-1}\pi_t$, here again taking full account of current actions on future expectations. What is the resulting behavior of π and u as a function of lagged π and u and current ε?

This is the familiar commitment policy rule, though its consequences here are different because expectations are not in fact rational. The solution follows exactly the Lecture 10 notes to arrive at (in periods other than the first)

$$\pi_t = \frac{\alpha\tilde{\varepsilon}_t}{\theta + \alpha^2}$$
$$u_t = \bar{u} + \frac{\theta\tilde{\varepsilon}_t}{\theta + \alpha^2},$$

where $\tilde{\varepsilon}_t$ *is the shock to the natural rate* as perceived by this policy maker with the false model. *Since the policy-maker assumes, based on this solution, that* $\hat{\pi}_t \equiv 0$*, we get*

$$\tilde{\varepsilon}_t = u_t - \bar{u} + \alpha\pi_t$$
$$\therefore \quad \pi_t = \frac{\alpha}{\theta}(u_t - \bar{u}). \tag{A4.8}$$

Combining this last policy rule for π *in terms of observable variables with the true Phillips curve, we get*

$$\pi_t = \frac{\alpha^2}{\theta + \alpha^2}\pi_{t-1} + \frac{\alpha}{\theta + \alpha^2}\varepsilon_t. \tag{A4.9}$$

The corresponding solution for u_t *follows by substitution into* (A4.8).

(f) (6 points) For each of the three policies you have computed, calculate the period loss function $u_t^2 + \theta\pi_t^2$ in the deterministic steady state (where $\varepsilon \equiv 0$ and both π and u have settled down to constant values). Show that these steady state losses are smallest for the last policy, that mistakenly assumes a rational public. Does this show that a false assumption that the public is rational can lead to better policy? If so, discuss why this is so; if not, explain why not.

When there are no shocks, in steady state $\pi_t = \pi_{t-1}$ *and* $\varepsilon_t = 0$, *so* $u_t = \bar{u}$ *and losses are just* $\bar{u}^2 + \bar{\pi}^2$, *where* $\bar{\pi}$ *is the steady state value of* π. *The last, "mistaken commitment" policy just sets* $\bar{\pi} = 0$ *and results in losses in steady state of* $\bar{u}^2$. *The no-commitment policy results in* $\bar{\pi} = \alpha\bar{u}/\theta$, *while the optimal policy results in* $\bar{\pi} = (1-\beta)\alpha\bar{u}/\theta$. *With* $\beta \in (0,1)$, *the optimal policy clearly delivers steady state losses in between the other two policies, with the mistaken commitment policy giving the lowest steady state losses. The optimal policy delivers a better discounted loss function, however. The optimal policy does not move all the way to zero inflation, because the future gains from being at zero inflation are small when inflation is small, and do not outweigh the current costs, in terms of higher unemployment, of lowering inflation. If the "commitment" solution were actually implemented in this economy at some particular date, the losses would be very high, because the commitment solution involves high initial inflation. This looks attractive in the commitment solution because that (false) model assumes that the initial inflation has no consequences for future expectations. Since in this economy high initial inflation leads to high expected inflation next period, the initial gains from the first-period surprise inflation will be almost entirely offset by a loss the next period, when the decline in inflation to zero (or near zero if* $\varepsilon_t \neq 0$*) creates a surprise (to the public) deflation.*

(5) (45 points) **Short answers.**

(a) (11 points) Explain why, in a wide range of dynamic models, if the rate of tax on capital is constant, the derivative of steady state utility with respect to the capital tax rate is less than zero, even evaluated at the point where the tax rate is zero. Also explain why this is not true of, say, a constant tax on labor. Does this imply that when a given level of expenditure must be financed, only distorting taxes are available, but all tax rates must be held constant, the optimal capital tax rate will be zero? Why or why not?

The effect of a labor tax is to lower consumption and increase leisure in steady state. In the neighborhood of a zero tax, the tradeoff is compensating, so that the increased leisure offsets the decreased consumption. There is no effect on the equilibrium capital/output or capital/labor ratios, because the equilibrium condition that the marginal product of capital equal the inverse of the discount factor is unaffected. The effect of a capital tax is to lower equilibrium steady state K/L *and increase the marginal product of capital. This lowers steady state consumption even when labor supply is inelastic, as we showed in class, and by more than is compensated for by increased leisure when labor supply is elastic. On the other hand, in the neighborhood of a zero capital tax, the effect on discounted utility of going to a small, perpetual positive capital tax is zero to first order. It will decrease future utility, but by encouraging current consumption, will create a matching current utility gain. Thus a constant* K *tax, like a constant* L *tax, has zero-to-first-order effects on utility, and one would expect that with taxes constrained to be constant, it will be optimal to have both non-zero.*

(b) (11 points) For most advanced economies it has been true for some years that direct measures of arbitrage opportunities across government bond markets in different countries imply that flows of funds between countries' bond markets are essentially costless. Assuming that countries had complete and perfect internal asset markets and product markets, would the fact that funds flow freely across international bond markets imply that capital markets should produce optimal international risk sharing? Why or why not?

As we saw in classroom discussion and an exercise, international capital flows via the bond market only generally do not suffice to produce perfect risk sharing. They allow for consumption smoothing, but not for sharing of country-level risk.

(c) (11 points) Neither Barro's original analysis of optimal debt policy nor the analysis in our classroom extension of it along the lines of the "Fiscal Consequences for Mexico. . . " model takes any account of the implications of sticky prices. If prices were sticky, along the lines of a Dixit-Stiglitz/Calvo model of stickiness, Barro's conclusion that taxes should be a random walk would still be a good approximation, while the conclusion of our classroom model that taxes should be constant would break down. Do you agree or disagree? Why or why not?

Barro's model is (except in a brief section at the end of his paper) entirely in real terms. Its conclusions therefore hold even if prices are constant, so long as its assumption of a quadratic deadweight loss from taxation is correct. A really good answer (no actual answer was this good) might have pointed out that Barro's argument depends on the marginal deadweight loss being a function of the tax rate alone. Cyclical fluctuations in the tax base or in other sources of distortion in the economy (of which sticky prices are one) could undermine the simple version of Barro's conclusion. But this qualification aside, Barro's argument for random walk tax rates does not depend on whether stickiness is present.

The argument for constant or near-constant tax rates, on the other hand, depends on having a way to collect lump-sum, non-distorting taxes as responses to fiscal shocks. The "Mexico" model shows that this can be done with a very simple monetary policy if prices are flexible. But if they are not flexible, relying on surprise deflation and inflation to adjust the size of the real debt would have real costs.

(d) (11 points) Modern Keynesian sticky price models are usually constructed to deliver the conclusion that expansionary monetary policy delivers a temporary increase in output and employment. In the model of "neutral stickiness" we discussed in class, prices are sticky, but expansionary monetary policy has no real effects. That model assumes that workers are identical and own the representative firm and that risk sharing is perfect. If we relaxed these restrictive assumptions, would the model deliver the usual Keynesian conclusions? Why or why not?

The neutral stickiness model depends on its labor contracts being treated like bonds. Inflation and deflation surprises change contract values ex post, but don't affect the real costs to firms or returns to workers of increased labor effort at the margin. That is determined in the market for current labor contracts. This basic result does not depend on perfect risk-sharing or on workers owning the firm. Dropping those assumptions would give inflation and deflation real effects, but they would be distributional, and there is no particular argument that these real effects would make inflation expansionary. The expansionary effects in standard stickiness models arise from the fact that firms with low prices expand output to meet demand — the sticky prices reflect actual marginal costs somewhere in the economy.

(e) (1 point) You don't have to work for this point; it's free.