FINAL EXAM

There are three questions that will be equally weighted in grading. Since you may find some questions take longer to answer than others, and partial credit will be given for incomplete answers, be sure not to spend disproportionate time on any one question before you have worked on other parts of the exam.

(1) Consider Barro's simple model of optimal distorting taxation. Revenues are τ_t, and deadweight losses are a quadratic function $\phi\tau_t^2$ of revenues. The government therefore solves

$$\min_{\tau_t} E \sum_{t=0}^{\infty} \beta^t \tfrac{1}{2}\phi\tau_t^2 \qquad \text{subject to} \tag{1}$$

$$B_t = \rho B_{t-1} + g_t - \tau_t \,. \tag{2}$$

and a requirement that $E[B_t\rho^{-t}] \xrightarrow[t\to\infty]{} 0$.

(a) Show that if $\rho\beta = 1$, this setup (in which g_t is taken as exogenously given) implies that both B_t and τ_t are martingales.

This part duplicates the lecture and Barro's article, except that I failed to give you a needed assumption: For the martingale property to hold for B_t, it must be that the expected value of discounted (at the rate β) future g is always the same. The FOC's of the problem are

$$\phi\tau_t = \lambda_t$$
$$\lambda_t = \beta\rho E_t\lambda_{t+1} \,.$$

These deliver the conclusion that the solution must make τ_t a martingale. Transversality for the government (since it dislikes both large positive and large negative debt) or the given constraint on the rate of growth of B implies that we can solve the budget constraint forward, and using the martingale property of τ_t we get

$$B_t = E_t \sum_{s=1}^{\infty} \rho^{-s}(\tau_t - g_{t+s})$$

If we add the assumption that $E_t g_{t+s}$ is a constant, say g^* for all $s > 1$, then

$$B_t = \frac{\tau_t - g^*}{\rho - 1} \,,$$

which, being the sum of a constant and a martingale, is a martingale. The same conclusion would hold if we assumed g_t was itself a martingale.

(b) For this and the remaining parts, suppose g_t is i.i.d., with $g_t = 0$ with probability .95, $g_t = 1$ with probability .05, each period. Describe and sketch a graph of the time paths of τ and B under these assumptions, being sure to distinguish periods during which g_t remains at 0 from periods where it is positive.

Using our calculations above, we find that $\tau_t = (\rho - 1)B_t + E[g_t]$, and plugging that into the GBC gives us $B_t = B_{t-1} + g_t - .05$. Thus the debt decreases by .05 during each period in which $g_t = 0$, and τ_t, being a linear function of the debt, decreases along with it. When $g_t = 1$, debt increases by .95, and τ_t increases along with it, by the amount $.05(\rho - 1)$. The sketch threfore should look sawtooth, with occasional sharp jumps upward followed by periods of steady, slow decline, both for debt and taxes.

(c) Now suppose that τ and B must both remain non-negative. Show that this implies that optimal τ and B are not martingales. Again characterize the time paths of τ and B.

Now we need to consider Lagrange multipliers on the $\tau_t > 0$ and $B_t > 0$ constraints. The FOC's become

$$\phi\tau_t = \lambda_t + \mu_t$$
$$\lambda_t + \nu_t = E_t\lambda_{t+1}\,,$$

with μ the multiplier on the $\tau \geq 0$ constraint and ν the multiplier on the $B \geq 0$ constraint. The constraints can be thought of as inequalities ($\geq$ inequalities, with an objective function we are minimizing), so the Lagrange multipliers must all be non-negative. The $\tau_t > 0$ constraint therefore can never bind, because it can only bind when $\mu_t > 0$, yet it is also equal to $\phi^{-1}(\lambda_t + \mu_t)$, which is always positive when $\mu_t > 0$. So $\phi\tau_t = \lambda_t$. Thus we can conclude that $\tau_t = E_t\tau_{t+1}$ except when the $B_t \geq 0$ constraint binds ($\nu_t > 0$). At such times $\tau_t < E_t\tau_{t+1}$. Thus τ_t is not a a martingale, though it follows a martingale process during periods when $B_t > 0$. We could also have argued this directly: If τ were a martingale, it would be a bounded martingale (bounded below by zero) and thus would have to converge with probability one. But if it converges to some constant value, then B_t will be driven to $\pm\infty$ with probability one, at the rate ρ^t, violating the given constraint on the growth rate (and also transversality for the government).

The time paths will have $\tau_t > (\rho - 1)B_t$, so that B_t shrinks when $g_t = 0$, and both τ_t and B_t jumping upward when $g_t = 1$. The difference from the previous part is that when $g_t = 0$ and debt is very small, The debt will be set to zero. Thereafter, both τ_t and B_t will remain at zero until $g_t = 1$, at which point $B_t > 0$ and $\tau_t > 0$ and the process restarts. The condition for the debt being retired is that the τ_t value required to retire the debt be less than the expected value of taxes next period, given that next period's debt will be zero.

(d) Now suppose that instead of ordinary risk-free debt, the government issues only contingent debt. Everyone understands that it will be totally repudiated, i.e. become worthless, at any date t where $g_t > 0$. We assume investors are risk-neutral, so that the expected gross rate of return on this contingent debt is the same $\rho = \beta^{-1}$ as that on risk-free debt. Assuming τ and B are not required to be always non-negative, is it optimal for τ and B to be martingales in this case?

(e) Which form of debt, risk-free or contingently repudiated, delivers the lowest expected dead-weight loss? (Assume no non-negativity constraint for this part.)

This is probably most naturally answered by treating the rate of return on government debt as a random variable R_t, equal to $\bar{R}$ when $g_t = 0$ and 0 when $g_t > 0$. In order for the expected return on this debt to match $\beta^{-1} = \rho$, we must have $.95\bar{R} = \rho$. The FOC's now are

$$\lambda_t = \tau_t$$
$$\lambda_t = \beta E_t[R_{t+1}\lambda_{t+1}]\,.$$

This implies

$$\tau_t = .95\beta\bar{R}E_t[\tau_{t+1} \mid g_{t+1} = 0] + 0\,.$$

This implies that over a span of dates during which $g_t = 0$, the conditional distribution of τ_t given the $g_t \equiv 0$ condition is a martingale. It does not imply that $E_t\tau_t = E_t\tau_{t+1}$ unless $E_t[\tau_{t+1} \mid g_{t+1} = 0] = E_t[\tau_{t+1} \mid g_{t+1} = 1]$. Since g is the only source of randomness here, it seems natural to suppose that this would require that τ_t be constant, say $\bar{\tau}$. In this case we can solve the GBC forward in the usual way to obtain

$$B_t = \frac{\bar{\tau} - .05}{\rho - 1}\,.$$

When $g_t = 1$ and therefore last period's debt has been repudiated, $\tau_t + B_t = g_t = 1$. Putting this together with the equation above, lets us solve to find

$$\bar{\tau} = \frac{\bar{R} - 1}{\bar{R}}\bar{g}\,.$$

This value of $\bar{\tau}$ implies that when $g_t = 0$, revenue exactly matches the interest on the debt, $(\bar{R} - 1)B$. Thus debt does in fact stay constant. The government maintains a conventionally balanced budget (though a positive primary surplus) until the date when the debt is repudiated.

The losses under this policy are just $\frac{1}{2}\bar{\tau}^2/(1-\beta)$. With non-contingent debt the optimal policy is $\tau_t = (\rho - 1)B_t + .05$. With this policy B_t is a martingale with increments $g_t - .05$. The variance of $g_t - .05$ is $.05 \cdot .95^2 + .95 \cdot .05^2 = .0475$. Thus $\text{Var}_t(B_{t+s}) = .0475s$ and the expected losses under this with zero initial debt is

$$\tfrac{1}{2}\left(\frac{.05^2}{1-\beta} + .0475\beta^{-1}\right)\,,$$

where we have used the fact that $\sum_0^\infty sa^s = a/(1-a)^2$.

The losses cannot be unambiguously ranked for arbitrary β. When $\beta = .95$, losses are about twice as big with contingent debt, while with $\beta = .9$ they are about 30% bigger with risk-free debt. The contingent debt keeps τ constant, which avoids the cost from future growing variances, but it makes the initial level of taxes higher — much higher if the ratio of $1 - \beta$ to $((1/(1-p) - \beta$ is small, where p is the probablility of non-zero g_t. For β close to one, the latter effect dominates, whereas for smaller β the former effect dominates.

No one got the correct answer to this part, which is perhaps not surprising. Considerable credit was given for just getting the first-order condition right (i.e., in a form that recognized that R_t and g_t could in principle be correlated, so the martingale conclusion is at least not obvious).

(2) Consider a continuous time search model of unemployment, similar to what we discussed in class, with one slightly more realistic variation: not everyone is eligible for unemployment benefits. We will suppose that upon losing a job, a worker is randomly determined to be eligible for unemployment benefits, or not, and remains in that status throughout the unemployment spell. So now we have the following parameters:

$F()$:	cdf of wage offers
w:	an actual wage offer
$W(w)$:	present value of wages, after acceptance of wage offer w
b:	dollar value per unit time of benefits in unemployed-with-benefits state
U:	present value of wages when unemployed, with benefits, with no offer
X:	present value of wages when unemployed, without benefits, with no offer
α:	rate of arrival of job offers when unemployed (same whether with or without benefits)
λ:	rate of arrival of job terminations when employed
r:	rate of time discount
δ:	probability, conditional on his becoming unemployed, that a worker is eligible for benefits.

As in the simpler in-class model, agents are assumed to maximize the expected discounted present value of their earnings from all sources, wages and benefits. In the in-class model we solved for $W(w)$ and U by calculating utility conditional on a known termination date for a job or the unemployed state, then averaging across those termination dates using appropriate probability weights.

(a) Find equations that can be solved for X, U, and $W(w)$ as functions of the given parameters.

$$W(w,T) = \int_0^T e^{-rt} w\,dt + e^{-rT}(\delta U + (1-\delta)X)$$

$$\therefore W(w) = \int_0^\infty \lambda e^{-\lambda T} W(w,T)\,dT = \frac{w + \lambda(\delta U + (1-\delta)X)}{\lambda + r} \qquad (\dagger)$$

$$U(T) = \int_0^T e^{-rt} b\,dt + e^{-rT} \int \max(W(w),U)dF(w)$$

$$\therefore U = \int_0^\infty \alpha e^{-\alpha t} U(T)\,dT = \frac{b + \alpha \int \max(W(w),U)dF(w)}{\alpha + r}$$

$$X = \frac{\alpha \int \max(W(w),X)dF(w)}{\alpha + r}\,.$$

Note that since $X > 0$, it not generally true (contrary to what many exams asserted) that an unemployed person receiving no benefits accepts every job offer. It is clear that $U > X$ and therefore that the reservation wage of those getting benefits will be higher. The reason we can be sure that $U > X$ is that if we had $X > U$, then those unemployed with benefits could adopt the reservation wage of those without benefits, obtaining the same expected return from job offer arrivals,

while also getting b while they were waiting, which would imply that $U > X$ after all.
Though the derivation above follows what was done in class, and the hint in the problem statement, it was perfectly ok, if you did it correctly, to arrive at similar equations by "dt" approximations.

(b) In what sense, if any, does the reservation wage property hold in this model?
Since U and X depend on w's distribution, but not on any particular value of w, in this setup as in the standard one $W(w)$ is increasing in w, as can be seen in (†). Therefore there will be two reservation wages, one for those unemployed with benefits and one for those unemployed without benefits.

(3) An economist presents simulated time paths for K, C, I, and L (but not A) from a solution to the following growth model:

$$\max_{C,K,L} E\left[\sum_{t=0}^{\infty}\beta^t \frac{(C_t^{\alpha}(1-L_t)^{1-\alpha})^{1-\gamma}}{1-\gamma}\right] \qquad \text{subject to}$$

$$(C_t^{\phi} + \nu I_t^{\phi})^{1/\phi} = A_t(K_{t-1}^{\theta} + \mu L_t^{\theta})^{1/\theta}$$

$$I_t = K_t - (1-\delta)K_{t-1}$$

To keep the model well-behaved, we need $\gamma > 0$, $\alpha \in (0,1)$, $\phi > 1$, $\theta < 1$, $\mu > 1$, $\nu > 1$, $\delta \in (0,1)$. We assume $\log(A_t/A_{t-1})$ is i.i.d. $N(0,\sigma^2)$. The economist presents numerical values for all the parameters along with the simulated data.

You note that there is a lot of complicated non-linearity in this model and suspect the economist might not have solved it accurately. How could you use the simulated data to check the accuracy of the solution? Be specific about what you would calculate and how you would assess the strength of evidence against an accurate solution.

Quite a few people took this question as an invitation to talk about how to *solve* a model, suggesting that they would check accuracy by solving the model themselves and seeing if their solution, in simulations, delivered a higher value of the objective function than the simulated data presented by the economist. This is an extremely inefficient approach, and would deliver ambiguous conclusions. We emphasized in lectures and class notes that the sharpest test of accuracy is a check on whether the expectational first-order conditions hold. Of course it should also be true that the deterministic FOC's, the identities, and the assumed normal distribution of $\Delta \log A_t$ should hold, but these are very likely to hold exactly, even if the solution algorithm used is incorrect or has not converged. Their accuracy does not depend on whether the solution for decision rules in the model is correct. (Unless the economist has used backsolving, in which case the expectational FOC's would hold exactly and the $\Delta \log A_t$ distribution would be the right thing to check.)

A good answer would have presented FOC's, shown how to eliminate Lagrange multipliers from them, and discussed how to check whether the Euler equation expectational errors did in fact have zero conditional expectation. If done with the simulated data, this would involve estimating regressions of the expectational errors on simulated data dated at and before the time at which the expectations were supposedly being formed. A reasonable criterion for whether the solution is accurate then is whether, when the regressions are estimated with the simulated data, the R^2's of the regressions are highly statistically significant. If so, the agents populating the model would have been able to see that the expectational errors could be improved upon. Quite a few exams described something close this procedure. Some credit was given for recognizing that the check was going to have something to do with the Euler equations, and deriving them, even if what followed that ran off track.