TAKEHOME FINAL EXAM

(1) (60 minutes) Consider an economy in which there are two states indexed by $i \in \{1,2\}$. Each agent is endowed with one unit each of two kinds of real assets, indexed by $j \in \{1,2\}$. Asset j pays y_j in state j and nothing in the other state. The agents in the economy all have U functions that are linear in consumption. (Whether these U's are "utility" or not under multiplier preferences is arguable.) They have multiplier preferences, so their objective function is to maximize

$$-\theta \log\left(E\left[e^{-c/\theta}\right]\right), \tag{1}$$

where c is a random variable taking on the value c_i in state i. Everyone agrees that the probability of state 1 is π, and that of state 2 is $1-\pi$. The price of asset 1 in units of asset 2 is q, and individuals see themselves as able to trade the two assets in a competitive market at that price. There is no storage, so an individual's consumption in state i is just the total yield on the asset of type i that he holds.

(a) Assume $\theta = 1$, $\pi = .25$. Plot two representative indifference curves in c_1, c_2 space for these individuals.
Most people drew indifference cures that were correct, but few made them "representative". I was looking for an indication that you had realized the somewhat unusual properties of these preferences. In particular, if V is the level of the objective function corresponding to an indifference curve $c_1 = f(c2; V)$, then for $V > \log(4) = 1.39$ the indifference curves do not cut the axes, but instead asymptote to vertical and horizontal lines that are away from the axes. When $V < \log(4/3) = .29$, the indifference curves cut both axes, and for intermediate values of V they cut the c_1 axis but not the c_2 axis. The three plots below show the indifference curves when c_1, c_2 ranges over the rectangle with corners $(0,0)$ and $(\bar{C},\bar{C})$, where $\bar{C} \in \{1, 10, 100\}$ with $\bar{C} = 1$, all the curves cut the $c_1 = 0$ axes. While some do not cut the c_2 axis it is hard to see that on the graph. At the other extreme, with $\bar{C} = 100$, the indifference curves are nearly rectangular, with corners near the $c_1 = c_2$ line. This means that they are willing to pay extremely high prices for insurance when their consumption is high, but not so much when their consumption is low in both states.

(b) Assume $y_1 = 2$, $y_2 = 1$. Find the equilibrium value of q, and compare it to what that price would be if individuals simply maximized expected U (of course with this same linear U function).

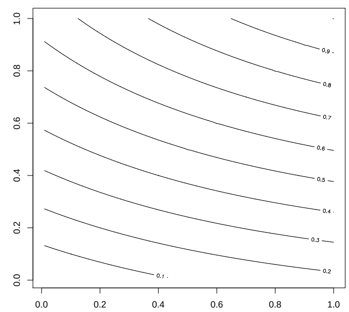

FIGURE 1. Indifference curves with c small

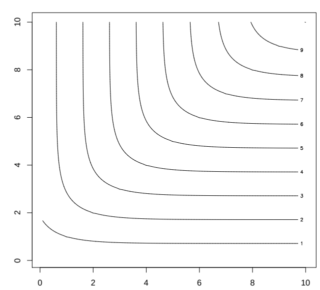

FIGURE 2. Indifference curves with medium levels of c

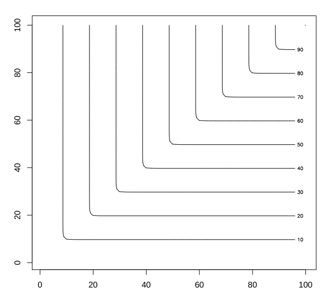

FIGURE 3. Indifference curves with large values of c

The individual's budget constraint is

$$qc_1 + c_2 = q + 1 .$$

Maximizing the multiplier preferences objective function subject to this constraint gives the FOC's

$$\frac{.25e^{-c_1}}{.25e^{-c_1} + .75e^{-c2}} = q\lambda$$
$$\frac{.75e^{-c_2}}{.25e^{-c_1} + .75e^{-c2}} = \lambda$$

and therefore

$$q = \frac{1}{3}e^{c_2 - c_1} = .1226\,.$$

Under expected utility with linear utility, the FOC's are simply $.25 = q\lambda$, $.75 = \lambda$ and thus $q = 1/3$. Thus as expected, the value of the security that delivers in the better state is reduced relative to what it would be with expected utility.

(c) Find a utility function $U(c)$ such that if agents were maximizing the expected value of this utility function, with the same probability π as above, the equilibrium relative price of the assets would be the same as in part 1b.
If $W(c)$ is this new utility function, we simply require that at $c_1 = 2$, $c_2 = 1$, $W'(2)/W'(1) = q = 1/(3e)$. One approach, which at least some people took, is to make W CRRA, i.e. $W(c) = c^{1-\gamma}/(1-\gamma)$. Then q at this point in c_1, c_2 space emerges as $2^{-\gamma}$, from which we can solve to find that $\gamma = 3.03$. However another approach, which at least one person took, delivers an answer to the next part as a side effect: Use CARA utility, i.e. $W(c) = 1 - e^{-c}$. It is easy to see that this delivers the same q under expected utility as does multiplier preferences with linear utility, for every (c_1, c_2) pair and indeed also for every probability π.

(d) If, after observing the behavior of these agents in the environment described above, you could take a few agents of this same type and put them in just one different environment, in which there were still two states, but the payoffs y_i and/or the probability π were different, could you determine whether the agents were maximizing expected U with risk aversion or were instead making choices with linear utility and on multiplier preferences as in parts 1a and 1b? What changes in y or π would you make, and what would you look for to tell the difference? Or why would observing the equilibrium in the new environment not be enough to distinguish multiplier preferences from risk aversion?
If you didn't notice that CARA utility exactly mimics multiplier preferences here, you might have held the utility function you used to answer the previous part fixed, and asked how could you distinguish expected utility maximization with that utility from the multiplier preferences derived from linear

utility. But once you see that CARA utility works for every (c_1, c_2) pair and every π, you realize that expected utility with this CARA utility implies behavior that is indistinguishable from multiplier preferences. In this light, we see that multiplier preferences in this static problem are a way to map an initial utility function $u(c)$ into a new, more risk-averse utility function $-e^{-u(x)/\theta}$,to which ordinary expected utility maximization can be applied.

(2) (20 minutes) Suppose

$$c_t = (1-\rho)\sum_{s=0}^{\infty}\rho^s y_{t-s} + \varepsilon_t\,, \tag{2}$$

where ε_t is i.i.d. $N(0, \sigma_\varepsilon^2)$ and independent of y at all leads and lags. Also suppose

$$y_t = (1-\alpha)\bar{y} + \alpha y_{t-1} + \nu_t\,, \tag{3}$$

with ν_t i.i.d. $N(0, \sigma_\nu^2)$.

If at time t all values of c_s and y_s for $s \le t$ are known, what is the mutual information between c_{t+1} and y_{t+1}? To what limits does it converge as $\rho \to 0$ and as $\rho \to 1$?

The conditional variance of c_{t+1} given knowledge of data from t and earlier is $(1-\rho)^2\sigma_\nu^2 + \sigma_\varepsilon^2$. After observation of y_{t+1} the variance is σ_ε^2. The mutual information is therefore

$$\tfrac{1}{2}\log_2\left(1 + \frac{(1-\rho)^2\sigma_\nu^2}{\sigma_\varepsilon^2}\right)$$

This converges to $\log_2(1+\sigma_\nu^2/\sigma_\varepsilon^2)$ as $\rho \to 0$, and to zero as $\rho \to 1$. This illustrates how delayed, smoothed reactions can imply arbitrarily low information flows while still preserving long run relations. Here the long run response of c to y is unchanged at 1.0 as ρ varies.

(3) (30 minutes)Suppose that we are trying to track an exogenously given stochastic process y_t with a decision variable c_t. Each period we must move c_t up by one unit or down by one unit — i.e. $c_t = c_{t-1} \pm 1$. (We cannot leave c unchanged.) All values of c_s, y_s for $s < t$ are known when we choose c_t.

(a) Show that the mutual information between y_t and c_t can be no larger than one bit per time unit.

Since c_t takes on only two values, the entropy of its distribution cannot exceed one bit, which is its entropy when both values are equally likely, and cannot be less than zero, which is its entropy when one of the two values has probability one. The mutual information is the expected reduction in entropy between the unconditional distribution for c and its conditional distribution given y. If the unconditional distribution has one bit entropy, while knowledge of y allows us to determine c's value exactly, we would achieve the one bit per time unit mutual information. If the unconditional entropy of

c is less than one bit, or if the conditional entropy is greater than zero, the mutual information is less than one bit.

(b) Suppose $y_t \sim N(0,1)$ i.i.d. across t. If we are minimizing $E[(y_t - c_t)^2]$, what is the optimal way to set c_t given the $c_t = c_{t-1} \pm 1$ constraint? What is expected loss under the optimal policy? What is the actual rate of information flow?

A good answer to this is beyond what could reasonably have been asked on an in-class exam, because I hadn't thought through the whole problem. However, many people made an assumption that trivialized the problem (assuming c_t and y_t independent, which makes information flow zero in all parts of the problem). Since y_t is i.i.d. in these examples, the problem becomes trivial if y_t is independent of c_t. Some people assumed this, and the wording of the problem left it somewhat ambiguous. In the more interesting case where we allow dependence between y_t and c_t, the optimal policy is clearly to move c up when $y_t > c_{t-1}$ and move it down otherwise. This does not always move c_t closer to y_t, but it always picks the $c_t \pm 1$ value that is closest to y_t. By observing c_t, we discover whether y_t is greater than or less than c_{t-1}. The mutual information between them is therefore

$$\Phi(c_{t-1}) \log(\Phi(c_{t-1})) + (1 - \Phi(c_{t-1})) \log(1 - \Phi(c_{t-1})) ,$$

where Φ is the standard normal cdf.

To find the average rate of information flow between the two, we need to know the marginal distribution of c. A natural assumption is that the distribution of c is concentrated on the integers, and over 99% of the probability will be concentrated on integers less than or equal to 3 in absolute value. (To get to 4, we would first have to be at 2, get a probability-less-than-.03 draw of $y > 2$, then draw a probability-less-than-.002 draw of $y > 3$.) The transition probability matrix among the 7 values from -3 to 3 (treating 3 as returning to 2 with probability one, as an approximation) is

-3	-2	-1	0	1	2	3	
-3	0.000	1.000	0.000	0.000	0.000	0.000	0.000
-2	0.023	0.000	0.977	0.000	0.000	0.000	0.000
-1	0.000	0.159	0.000	0.841	0.000	0.000	0.000
0	0.000	0.000	0.500	0.000	0.500	0.000	0.000
1	0.000	0.000	0.000	0.841	0.000	0.159	0.000
2	0.000	0.000	0.000	0.000	0.977	0.000	0.023
3	0.000	0.000	0.000	0.000	0.000	1.000	0.000

Finding the left eigenvector of this matrix corresponding to the unit eigenvalue value gives us the ergodic distribution of the states, which is

	-3	-2	-1	0	1	2	3
1	0.001	0.040	0.249	0.419	0.249	0.040	0.001

The vector of mutual informations at c_{t-1} values from -3 to 3 is

	-3	-2	-1	0	1	2	3
1	0.015	0.157	0.631	1.000	0.631	0.157	0.015

The crossproduct of this list of mutual informations with the ergodic distribution probabilities gives us the average rate of information flow: .746 bits per time period.

$$\sum_{c=-3}^{3} p(c)\left((1-\Phi(c))E[(y-c+1)^2|y>c] + \Phi(c)E[(y-c-1)^2 \mid y<c]\right).$$

The expected losses can be calculated analytically as .670 per period, if I haven't made any algebra slips.

(c) Repeat your answer to the previous question if the constraint is simply $c_t = \pm 1$. That is, instead of the change in c being plus or minus one, c itself must be plus or minus one.
This is easier. Obviously one wants $c = 1$ when $y > 0$, $c = -1$ when $y < 0$. The marginal distribution of c then puts equal probability on 1 and -1, and observing y tells us exactly the value of c. Thus there is one bit per time unit of mutual information. The expected losses are The expected losses are $E[(y-1)^2 \mid y > 0] = .404$.

(d) What is the optimal policy for setting c if the constraint is simply that mutual information between c_t and y_t cannot exceed one bit per time period, with no restrictions on the values c can take on. What are expected losses under that policy?
We know that with quadratic loss the optimal form of uncertainty is Gaussian, and in our case it will be optimal to make $E[y_t \mid c_t] = c_t$. Therefore we will have an unconditional variance of y_t of 1 and a variance of $y_t \mid c_t$ of $\sigma^2 < 1$. To make the information flow be one bit per time period, we must make $\log_2(1) - \log_2(\sigma) = 1$, i.e. $\sigma = .5$. Expected losses are then .25.

(4) (45 minutes) Consider a model in which representative agents solve

$$\max_{C,M} E[\sum_{t=0}^{\infty} \beta^t \log C_t] \quad \text{subject to} \tag{4}$$

$$C_t(1+\gamma v_t) + \frac{M_t}{P_t} = \bar{Y} + \frac{M_{t-1}}{P_t} \tag{5}$$

$$v_t = \frac{P_t C_t}{M_t} \tag{6}$$

The government budget constraint is

$$\frac{M_t - M_{t-1}}{P_t} = \bar{G}\,. \tag{7}$$

Monetary policy sets $M_t = \theta M_{t-1}$ in every period.

(a) Show that, subject to certain restrictions on the parameters, this model has a competitive equilibrium with a uniquely determined price path.

The FOC's are

$$\partial C: \qquad \frac{1}{C_t} = \lambda_t(1+2\gamma v_t)$$

$$\partial M: \qquad \frac{\lambda_t}{P_t}(1-\gamma v_t^2) = \beta\frac{\lambda_{t+1}}{P_{t+1}}$$

The social resource constraint is

$$C_t(1+\gamma v_t) + \bar{G} = \bar{Y}\,.$$

From the FOC's and the policy rule we can derive, letting $z_t = 1/(v_t(1+2\gamma v_t))$,

$$(1-\gamma v_t^2)z_t = \frac{\beta}{\theta}z_{t+1}\,.$$

As long as $\beta < \theta$, there is a steady state solution to this equation in which

$$v_t \equiv \bar{v} = \sqrt{\frac{\theta-\beta}{\gamma\theta}}$$

Since v and z are monotonically related, with $z \to \infty$ as $v \to 0$ and vice versa, this solution is unique. Otherwise, because the difference equation in z is unstable any deviation of v_t from $\bar{v}$ leads to $v \to \infty$ or $v \to 0$. Deviations of v to infinity are impossible because the $1-\gamma v_t^2$ term in the z equation would become negative, so there is no solution. Deviations of v to zero are impossible because they would imply $m_t = M_t/P_t$ going to infinity, which would violate transversality. (Real balances, which are useful only to reduce transactions costs, become arbitrarily large; transactions costs become arbitrarily small; eventually it must be clear that spending real balances is feasible and improves utility.)

(b) Show that there is a one-to-one relation between $\bar{G}$ and θ. What is the level of $\bar{G}$ when θ is set to maximize $\bar{G}$?

The government budget constraint, together with the policy rule, tells us that

$$m_t(1-\theta^{-1}) = \bar{G}\,.$$

In equilibrium, $C(1+\gamma\bar{v}) + \bar{G} = \bar{Y}$, so

$$\bar{C} = \frac{\bar{Y}-\bar{G}}{1+\gamma\bar{v}}$$

Therefore

$$m_t \equiv \frac{\bar{C}}{\bar{v}} \quad \text{and} \quad (1-\theta^{-1})\frac{\bar{Y}-\bar{G}}{\bar{v}(1+\gamma\bar{v})} = \bar{G}$$

After some algebra, this lets us arrive at

$$\bar{G} = \frac{(1-\theta^{-1})\bar{Y}}{\bar{v}(1+\gamma\bar{v}) + 1 - \theta^{-1}}$$

This is monotone increasing in θ, but because $\bar{v}$ depends on θ, proving monotonicity is more messy algebra than I realized. It is not hard to see that it converges to a number less than $\bar{Y}$ as $\theta \to \infty$ and that it is always less than $\bar{Y}$.

(c) Apply Bassetto's approach to this model. Show that there is a problem with infeasibility of the policy at off-equilibrium situations. What kind of an addendum to the policy (if any) for these off-equilibrium situations could support the competitive equilibrium?

If at some date the public refused to sell goods to the government at the equilibrium price level, the government could not both make real expenditures $\bar{G}$ and finance them entirely by increasing money by the factor θ. Specifying a market game in which government behavior in this off-equilibrium state is feasible is not quite so easy as in the pure FTPL model. This model does have an equilibrium albeit a horrible one, with valueless money and $barG$ forced to zero. Velocity is infinite, transactions costs take up the entire endowment, $C = 0$, and utility is $-\infty$. But once in this equilibrium, no policy to manipulate the stock of fiat money can get the economy out of it. So to support the equilibrium the policy authority has to have the power to tax directly and the willingness to use this power if necessary to maintain the value of money. It can stand ready to trade money for goods at some higher-than-equilibrium price level $\bar{p}\theta^t$ at any date. (This introduces a "trading post" that is usually inactive.) If the price level moves above $\bar{p}\theta^t$, private agents who held money will have undergone a large capital loss on their real money balances, but they will have no reason to project that these losses will continue. If the economy returns to the equilibrium path the next period, anyone who holds real balances will have a large capital gain, and even if the above-normal price level repeats, the capital loss will only be at the equilibrium rate $(1-\theta^{-1})$, so the incentives to hold money to reduce transactions costs will be as strong as on the equilibrium path.

This illustrates the point that in both "fiscal theory" and "monetary" equilibria, valued money depends on a commitment to back money or nominal debt with taxes off the equilibrium path.

(5) (30 minutes) Cochrane argues that monetary policy behavior cannot be identified by looking at historical data. Sargent-Williams-Zha, Sims-Zha, and Primiceri all claim to estimate monetary policy behavior, indeed in some cases elaborate dynamic evolution of monetary policy behavior, from historical data. Is Cochrane saying that what these papers claim to do is impossible, so that their claims to do so cannot be taken seriously? If so, is he right? [There is no correct answer to this question. You are graded on your argument.]

Several people seemed to think Cochrane's arguments apply only to New Keynesian models. In fact they apply to any dynamic, stochastic competitive equilibrium model in which nominal debt or fiat money appears. When such a model is driven by an exogenously evolving exogenous disturbance vector ε_t with state vector s_t, Every variable i in the model solution can be written as a function $f_i(s_t)$. In particular, if the model contains both an interest rate r_t and a money stock M_t, it will be true on the solution path that $M_t = f_M(s_t)$ and $r_t = f_r(s_t)$. If the former is the monetary policy rule, then the model generally has a unique equilibrium price path if fiscal policy responds positively to the level of debt, because a commitment to making money stock unresponsive to other nominal variables is a form of active monetary policy. But if $r_t = f_r(s_t)$ is a policy commitment, this is a form of passive monetary policy, and equilibrium price will be unique only if fiscal policy is active. Either situation could be the truth, and we cannot tell which is which by looking at the time paths of variables realized in an equilibrium. Uniqueness of the price level is determined by *beliefs* of the public about *off equilibrium path* behavior of the policy authorities. If inflation did start increasing,above the equilibrium path rate, would money stock continue to follow its stationary equilibrium path, or would instead r continue to follow its stationary equilibrium path? There is no way to tell. In the equilibrium both follow stationary paths that are functions of the exogenous state, but either of their two $_i$'s, or neither, might represent the actual policy commitment. That is Cochrane's point.

Broadly speaking, his point is that on an equilibrium path there will be many equations containing r, M, and/or fiscal variables that hold exactly. There is no way in general of determining which of these are the policy rule or rules. So his point is a very broad one about identification and it certainly applies to the SZ, SWZ, and Primiceri papers. Each of those papers assumes restrictions on the behavior of monetary authorities and the private sector that, if true, allow identification of monetary policy behavior. So in that sense Cochrane's critique does not apply. But SZ achieve identification by exclusion restrictions on contemporaneous coefficients in the monetary policy rule and in private sector behavior. Cochrane might argue that these restrictions are doubtful, which they are. SZ do verify that their identified monetary policy equation, when modified, has effects on the behavior of the economy like what might be expected from a monetary policy shift. SWZ assume that monetary policy directly controls inflation, without any intermediate role for interest rates. This is clearly counterfactual for the period they study, and they do not check that shifts in what they identify as monetary policy behavior produce effects on the behavior of the economy that would be expected from a monetary policy shift. Primiceri assumes policy-makers directly control a real variable that impacts unemployment directly. Certainly no policy authority actually thought it had such control in this period. Primiceri also does not verify that changing his estimated policy behavior has expected effects on the behavior of the economy. So in all these papers there are implicit identification problems of the type Cochrane points to. This is a reason for not accepting likelihood of the model as direct indicator of whether the story the model tells about the evolution of policy is correct.