DECISION THEORY NOTES

1. The dynamic inconsistency example

That example can be made simpler — almost trivial — without bringing in a nefarious second agent.

Suppose our decision maker can at a time $t = 0$ buy or sell one unit of a contingent claim to one dollar in any of the states $\{\omega_1, \ldots, \omega_4\}$ for 25 cents each. (This is an option to buy or sell, or neither buy nor sell, exactly one unit, not a price at which arbitrary amounts can be bought or sold.) She maximizes expected return and puts probability .25 on each of the four states. She also knows that next period, no matter what information arrives, she will again have the same opportunity to buy or sell one unit of these contingent claims. Next period information will arrive that reveals which of $\{\omega_1, \omega_2\}$ or $\{\omega_3, \omega_4\}$ contains the true ω_i. Also, it is known that next period the same opportunity to buy or sell contingent claims will arrive, and that the prices will then be either $(.51, .5, 0, 0)$, in the case where it has been revealed that the state is in $\{\omega_1, \omega_2\}$, or $(0, 0, .51, .5)$ otherwise. If she makes plans according to her probability distribution over states and executes those plans, she will sell the state 1 claim at time 1 if the state has been shown to be in $\{\omega_1, \omega_2\}$, and in the other case will sell the state 3 claim. At time 0 she may neither buy nor sell, or she may buy or sell any of the claims, since they all cost the same as their expected yield. If, for example, she neither buys nor sells at time 0, her pattern of payoffs across the four states is $(-.49, .51, -.49, .51)$. That is, she gets the .51 return from selling the contingent claim at time 1 in all states, but has to pay out 1.0 in states 1 and 3. Her overall expected return is .01.

But now suppose that at time 1, instead of updating her probabilities by simply conditioning on the information that arrives at this date, she "re-twists" the probabilities, now giving the states the probabilities $(.55, .45, 0, 0)$ when the information restricts the state to $\{\omega_1, \omega_2\}$ and giving them the probabilities $(0, 0, .55, .45)$ in the other case. Now in the first case, for example, she sees the contingent claim on the first state, priced at .51, as a good buy, and that on the second state, priced at .5, as a profitable sale. Accordingly she buys the first and sells the latter, and does the analogous thing with the other information set. The sale and buy have a net cost of .01, which applies in all states, so her pattern of returns across states is $(.99, -1.01, .99, -1.01)$, with an expected return of -.01 based on her original probabilities, but an expected return as of time 1, based on her new probabilities and the information at time 1, of .09. But this pattern of behavior has resulted in an inadmissible decision rule. She could have, at time 0, bought and sold contingent claims to deliver the return pattern $(1, -1, 1, -1)$ for a zero net cost, and this pattern of returns is strictly better, in every state, than what she has ended up with. Another way to look at this is that, using her new conditional probabilities at time 1, she can look back at time 0 and kick herself — why didn't she buy the $(1, -1, 1, -1)$ pattern of returns when she had the chance? It is this type of "remorse", in which updated conditional probabilities make a decision-maker look back at earlier behavior and wish it could have been different, that is ruled out by the Epstein-Schneider rectangularity condition.

2. The Weitzman example

There is an asset whose yield at time t is y_t. If we assume that this is the only asset, and that consumption of a representative agent at time t is just e^{y_t}, then a standard Euler equation can be thought of as determining the time-t price of the asset, P_t, from

$$P_t U'(C_t) = \beta E_t \left[e^{y_{t+1}} U'(C_{t+1}) \right] .$$

Assuming CRRA utility, so that $U'(C_t) = C_t^{-\gamma}$, and using the assumption $C_t \equiv e^{y_t}$, we can reduce this to

$$P_t e^{-\gamma y_t} = \beta E_t \left[e^{(1-\gamma) y_{t+1}} \right] = \beta E_t \left[\exp((1-\gamma)\hat{\mu} + \tfrac{1}{2}(1-\gamma)^2 \sigma^2) \right] , \qquad (*)$$

where the right-most expression has applied the law of iterated expectations to replace $e^{y_{t+1}}$ by its conditional expectation given μ and σ^2, assuming y_{t+1} conditional on information at t is distributed as $N(\mu, \sigma^2)$.

The likelihood function for T i.i.d. draws of y_t from a $N(\mu, \sigma^2)$ distribution is

$$(2\pi)^{-T/2} \sigma^{-T} \exp\left(-\frac{T\hat{\sigma}^2}{2\sigma^2} - \frac{T(\mu - \hat{\mu})^2}{2\sigma^2} \right) ,$$

and the posterior distribution for μ and σ^2 has this form, so long as the prior is "conjugate". (One way to define a conjugate prior is to say it is a distribution such that the posterior is the same shape as the likelihood, for a somewhat larger sample.) In this expression $\hat{\mu}$ and $\hat{\sigma}^2$ are the sample mean and the sample variance of $\{y_1, \ldots, y_T\}$. The expectation of the right hand side of the Euler equation $(*)$ is the expression above, normalized to integrate to one in μ and σ^2 jointly (i.e. the posterior density), times the expression in $(*)$, integrated over μ and σ^2.

We can integrate the expression with respect to μ analytically, since the product of the right-hand side of $(*)$ with the posterior density has as exponent a quadratic form in μ, and is therefore proportional to a Normal pdf. Collecting terms in μ and integrating with respect to μ, we arrive at an expression proportional to

$$\int_0^\infty \sigma^{-T+1} \exp\left(\tfrac{1}{2} \left((1-\gamma)^2 \sigma^2 - T \frac{\hat{\sigma}^2}{\sigma^2} \right) \right) .$$

Note that the exponent consists of one piece that converges to zero as $\sigma^2 \to \infty$, and another that goes to ∞, unless $\gamma = 1$. The σ^{-T+1} component of the integrand goes to zero as $t \to \infty$, but at a slower than exponential rate. So the integral does not converge. The asset has infinite value. When $\gamma > 1$ this occurs because the possibility of large negative returns is so dire (the utility function goes to $-\infty$ as $C \to 0$) that putting aside savings to use in that state is extremely important. When $\gamma > 1$ instead the infinite value comes from the high probability of extremely high returns, which have high marginal utility because of the low risk aversion.

It is this problem with combining inference about yield distributions with CRRA utility, first pointed out by John Geweke, that leads Weitzman to suggest that asset pricing anomalies can be explained by uncertainty about the tails of yield distributions.