September 17

0.0.7.4 Punishment

Our definition of a trigger strategy has a player switch to a Nash equilibrium strategy in the event that punishment is triggered in the game. Assuming that the discount factors are sufficiently large so that the trigger strategies form a Nash equilibrium, the assumption that the players switch to a Nash equilibrium if punishment is triggered insures that the equilibrium is subgame perfect, i.e., the punishment is credible. We obtain a lower bound on the discount factor for Nash equilibrium if we dispense with the requirement of subgame perfection. Let's think instead about the worst punishment that one player can impose on the other because this will serve as the most effective deterrent.

Player i's **minmax value** m_i is the lowest payoff in the stage game that player $-i$ can impose on him through his choice of a strategy s_{-i} given that player i can choose his own strategy to maximize his own payoff:

$$m_i = \min_{s_{-i}} \max_{s_i} \pi_i\left(s_i, s_{-i}\right).$$

Here, $\pi_i\left(s_i, s_{-i}\right)$ is the payoff to player i in the stage game given the strategy profile $\left(s_i, s_{-i}\right)$. The "max" represents the capability of player i to choose his own strategy s_i to maximize his payoff given the other player's strategy s_{-i}, and the "min" represents player $-i$'s choice of s_{-i} minimize this "best response" payoff for player i. Let $\widetilde{s}_{-i}$ denote the strategy of player $-i$ at which the minmax value is obtained. The strategy $\widetilde{s}_{-i}$ is the worst choice of a strategy by player $-i$ from the perspective of player i. Existence of $\widetilde{s}_{-i}$ is not a problem in finite games. Campbell (p. 61) calls this value player i's **security level**.

Notice that player i cannot receive less than his minmax value in any Nash equilibrium of the stage game. The value m_i is the lowest payoff that player i receives when he chooses his strategy in his own best interest.

We can of course define a minmax value m_{-i} for player $-i$ and the strategy $\widetilde{s}_i$. For a desired outcome $\left(s_1^*, s_2^*\right)$ in the stage game with corresponding payoffs $\left(\pi_1^*, \pi_2^*\right)$, define the trigger strategy of player i as follows:

Player i: Choose s_i^* to start the game and as long as $\left(s_1^*, s_2^*\right)$ is played by the players; if $\left(s_1^*, s_2^*\right)$ is ever not played, then switch to $\widetilde{s}_i$ for ever more.

The choice of $\widetilde{s}_i$ in every future stage is the worst thing that player i can do to player $-i$ and it therefore is the most effective deterrent. Similar remarks hold for player $-i$'s choice of $\widetilde{s}_{-i}$. We obtain a different and looser lower bound on the discount factors in this case. Player i compares his payoff if he follows his trigger strategy to his best possible deviation in stage k (assuming that $d_i > \pi_i^*$):

$$\begin{aligned}(1-\delta_i)\sum_{t=0}^{\infty}\delta_i^t\cdot\pi_i^* &\geq (1-\delta_i)\left[\sum_{t=0}^{k-1}\delta_i^t\cdot\pi_i^* + \delta_i^k\cdot d_i + \sum_{t=k+1}^{\infty}\delta_i^t\cdot m_i\right] \Leftrightarrow \\ \delta_i &\geq \frac{d_i-\pi_i^*}{d_i-m_i}\end{aligned}$$

We substitute his maxmin payoff for his Nash equilibrium payoff π_i. Notice that

$$\frac{d_i-\pi_i^*}{d_i-m_i} \leq \frac{d_i-\pi_i^*}{d_i-\pi_i}$$

because $m_i \leq \pi_i$ (a player's minmax value is less than or equal to his payoff in any Nash equilibrium).

These new trigger strategies do not necessarily form a subgame perfect Nash equilibrium when they do form a Nash equilibrium. Consider a history in which (s_1^*, s_2^*) does not occur in some stage. In the subgame defined by that history, the strategies specify that the players play $(\widetilde{s}_i, \widetilde{s}_{-i})$ in each and every stage. This need not define a Nash equilibrium in the subgame (in particular, there is no reason that $\widetilde{s}_i$ must be a best response to $\widetilde{s}_{-i}$ in the stage game).

Through the use of a more severe punishment by each player in their trigger strategies, we have obtained a lower bound on the discount factor δ_i that is sufficient to insure that (π_1^*, π_2^*) is played in each and every stage of a Nash equilibrium of the supergame. We have sacrificed subgame perfection of the equilibrium, however, in that punishment may not be credible.

Example 21 *Let's reconsider the following stage game:*

1\2	*L*	*C*	*R*
T	*1,-1*	*2,1*	*1,0*
M	***3,4***	*0,1*	*-3,2*
B	*4,-5*	*-1,3*	*1,1*

We'd like to implement $(3,4)$ as the outcome in each stage. As above, we don't need to worry about player 2 deviating from his trigger strategy. Let's focus on player 1's minmax value m_1 and the strategy $\widetilde{s}_2$ that player 2 should use if he really wants to hurt player 1:

$$m_1 = 1, \widetilde{s}_2 = R.$$

Before we had the following bound on player 1's discount factor:

$$\delta_1 \geq \frac{d_1 - \pi_1^*}{d_1 - \pi_1} = \frac{4-3}{4-2} = \frac{1}{2}.$$

If player 2 punishes with R instead of C, we have the bound

$$\delta_1 \geq \frac{d_1 - \pi_1^*}{d_1 - m_1} = \frac{4-3}{4-1} = \frac{1}{3}.$$

Don't worry about any more of section 7 of Campbell's book than I have discussed here!

We'll jump forward next to sections 6 and 7 of Chapter 2 in Campbell's book. The first 5 sections of Chapter 2 are mainly a math review. I presume you don't need this. If mathematical topics come up that you haven't seen, then just let me know and we'll go over them.

0.0.8 Chapter 2, Section 6: Decision Making Under Uncertainty

Section 6 discusses how we make decisions when we are faced with uncertain or random events and section 7 concerns insurance against random events. A basic issue in section 6 is, "Why do people buy insurance when the insurance company prices the policies so that they on average make a profit (and on average, the policy holder loses)? For simplicity, most of Campbell's book is restricted to the case of binary outcome, i.e., "high or low", "good or bad", or "succeed or fail".

The basic model is as follows. A person receives wealth x with probability π and wealth y with probability $1-\pi$.

Examples:

1. a person has wealth y and faces the possibility of a loss $y-x$ due to robbery, fire, etc. with probability π.
2. A person has wealth z, and buys an asset (or makes an investment) that results in a loss of size $z-x$ with probability π or a gain $y-z$ with probability $1-\pi$.

   asset =investment or security

   expected monetary value of the asset (EMV) Campbell, page 112

Example 22 *This is inspired by Example 6.1 in Campbell, p. 113. A person has wealth \$100. She may be robbed of \$60 with probability 0.3. Here expected wealth is therefore*

$$0.3 \cdot 40 + 0.7 \cdot 100 = 82.$$

How much should she be willing to pay for an insurance policy that completely compensates her for a loss if it occurs? Knowing the probability of a robbery, at what price could an insurance profitably sell such a policy to the person?

Let's forgo the first question for a moment. The insurance company has to pay out a \$60 claim with probability 0.3. The expected payout of the insurance company is therefore $0.3 \cdot 60 = 18$. *It would have a positive expected profit at any price that exceeds \$18 for the policy. At the price of \$18, the company has an expected profit of* 0. *As a general rule, firms avoid lines of businesses in which they don't expect to make any money. We have also ignored the costs of the insurance company (rent, paying the salesman, the management, etc.). It needs to make money on its policies, not simply break even on them in an expected value sense.*

Would the woman be willing to pay more than \$18 dollars for a policy? If she buys the policy at \$18, then her wealth will surely be $100 - 18 = 82$. *She would buy the insurance at \$18+ε only if she preferred the certain outcome of 82-ε to the expected (but uncertain) outcome of 82. We want to model this idea that a certain outcome may be preferred to a larger but uncertain expected outcome.*

Definition. A person is **risk averse** if, for all m, he prefers the monetary outcome of m to any asset that provides him with an expected monetary outcome of **m** but with positive probability gives him an outcome strictly less than **m**.

Another characteristic of risk averse behavior: Consider the following two assets:

asset 1: results in wealth x with probability $1/2$ and y with probability $1/2$, where $x < y$;

asset 2: for some $\delta > 0$, results in wealth $x - \delta$ with probability $1/2$ and in wealth $y + \delta$ with probability $1/2$

The two assets have the same expected value (the $1/2$ plays a role here – it isn't true for arbitrary probabilities). The greater "spread" in the payoffs in asset 2 makes a risk averse person choose asset 1 over asset 2.

The risk averse person chooses asset 1 over asset 2 because there is "greater uncertainty" in asset 2.

We want to model a person's decision making over uncertain events so that we can model risk aversion. We do this by assuming that a person has a utility function of money or monetary outcomes, $u(m)$. It is obvious that $u(\cdot)$ is increasing in m.

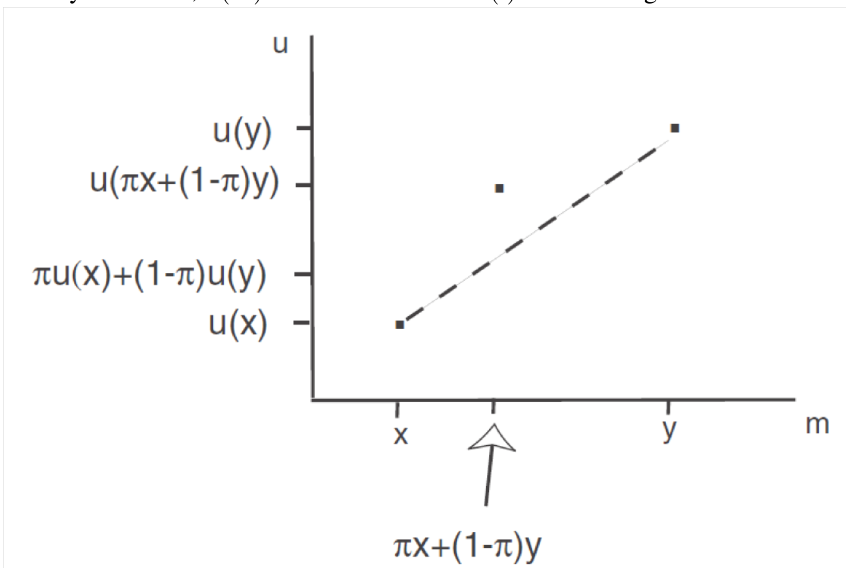

If the person makes his choices according to expected utility, then we need

$$u(\pi x + (1-\pi)\,y) > \pi u\,(x) + (1-\pi)\,u\,(y)$$

to indicate his choice of the certain payoff of $\pi x + (1-\pi)\,y$ over the uncertain payoff of $\pi u\,(x) + (1-\pi)\,u\,(y)$. This is depicted above, and suggests the following shape for the graph of u:

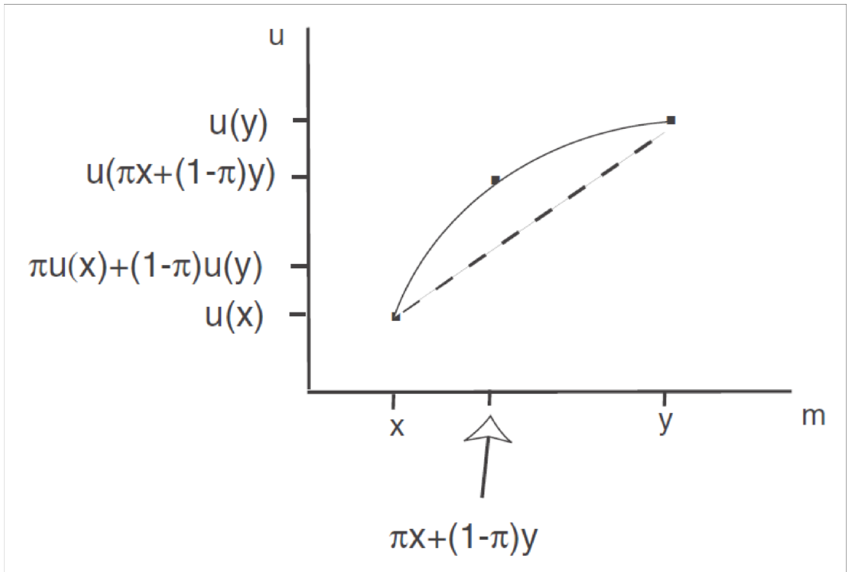

The expected utilities from assets that pay off with either x dollars or y dollars are represented by the dotted line. The horizontal coordinate of any point on this line is the expected monetary value of the asset. The vertical coordinate is the expected utility. We want the utility of the expected monetary value to exceed the expected utility so that the individual prefers the "sure thing" over the randomized return. The graph of the utility of a risk averse person therefore has the property that any line segment between any two points on the graph lies below the graph, i.e., the graph is concave down. We therefore model risk aversion by assuming that the utility of money is stricly concave down.

We can also see that the individual weighs asset 1 over asset 2, where:

asset 1: results in wealth x with probability $1/2$ and y with probability $1/2$, where $x < y$;

asset 2: for some $\delta > 0$, results in wealth $x - \delta$ with probability $1/2$ and in wealth $y + \delta$ with probability $1/2$.

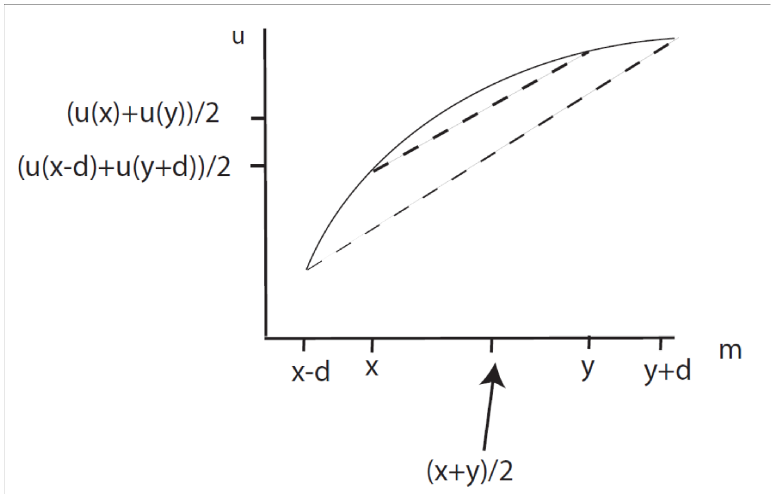

Risk neutral behavior is modeled by a linear utility of money, i.e., $u(m) = km$ for some $k > 0$. It's typical to take $k = 1$. The risk neutral person compares assets purely based upon their expected monetary values.

Comment. Modeling risk aversion using a concave down utility function of money is the work of Oscar Morgenstern and John Von Neumann. Von Neumann was one of the greatest mathematical minds of the 20th century. His accomplishments include starting the field of game theory with the economist Oscar Morgenstern, important contributions to probability theory and analysis, work on the Manhattan Project, and contributions to the foundations of computing and artificial intelligence.

Problem 1, p. 122. The individual's utility-of-wealth function is $U(w) = \sqrt{w}$ and current wealth is \$10,000. Is this individual risk averse? What is the maximum premium that this individual would pay to avoid a loss of \$1900 that occurs with probability $1/2$? Why is this maximum premium not equal to half of the loss?

The utility function is strictly concave down and therefore models risk averse behavior. Let p denote the maximum premium. The value of p equates the expected utility from having insurance to the expected utility from not having insurance; at a higher premium, he would be worse off from having insurance, and at a lower premium, he would be better off to have insurance. Therefore, we can solve this as

$$\begin{aligned} \frac{\sqrt{8100}}{2} + \frac{\sqrt{10000}}{2} &= \sqrt{10,000 - p} \\ \frac{90}{2} + \frac{100}{2} &= \sqrt{10,000 - p} \\ 95 &= \sqrt{10,000 - p} \\ 9025 &= 10,000 - p \\ p &= 975. \end{aligned}$$

This is strictly more than the expected loss of $1900/2 = 950$, which is a reflection of the individual's risk aversion.

Problem 2, 122. An individual has a utility-of-wealth function $U(w) = \ln(w + 1)$

and a current wealth of \$20. Is this individual risk averse? How much of this wealth will the person use to purchase an asset that yields zero with probability $1/2$ and with probability $1/2$ returns \$4 for every dollar invested?

Yes, he's risk averse. We can check by noting that

$$U'(w) = \frac{1}{1+w}$$

and therefore

$$U''(w) = -\frac{1}{(1+w)^2} < 0.$$

The function $U(w)$ is therefore strictly concave down. There are two common and tractable risk averse utility functions – $\ln$ and $\surd$. Campbell will clearly rely on variations of these two examples quite a bit.

Let C be the amount that the individual invests in the asset. His expected utility from acquiring the asset is

$$\begin{aligned} & \frac{1}{2}\ln(1+20-C) + \frac{1}{2}\ln(1+20-C+4C) \\ = & \frac{1}{2}\ln(21-C) + \frac{1}{2}\ln(21+3C). \end{aligned}$$

He should choose $C \in [0, 20]$ to maximize this expression. Setting its derivative with respect to C equal to 0 produces

$$\begin{aligned} 0 &= \frac{1}{2}\cdot\frac{1}{21-C}(-1) + \frac{1}{2}\cdot\frac{1}{21+3C}(3) \\ 0 &= -\frac{1}{21-C} + \frac{3}{21+3C}. \end{aligned}$$

Multiplying through by $(21-C)(21+3C)$ produces

$$\begin{aligned} 0 &= -(21+3C) + 3(21-C) \\ 0 &= 42 - 6C \Rightarrow C = 7. \end{aligned}$$

The second derivative of his expected utility is

$$-\frac{1}{(21-C)^2} - \frac{9}{(21+3C)^2}$$

which is negative. The solution $C = 7$ therefore maximizes the individual's expected utility.

Notice that the expected return on each dollar invested is \$2, for an expected profit of \$1 for every dollar invested. Risk aversion is the reason that the individual does not invest all of his money in the asset (as a risk neutral person would do).

Exercises: p. 122-124, problems 6, 7, 8 and 9
p. 135, problems 4 and 5