Solutions to Problem Set 5
Political Science 152/352

Question 1

(a)

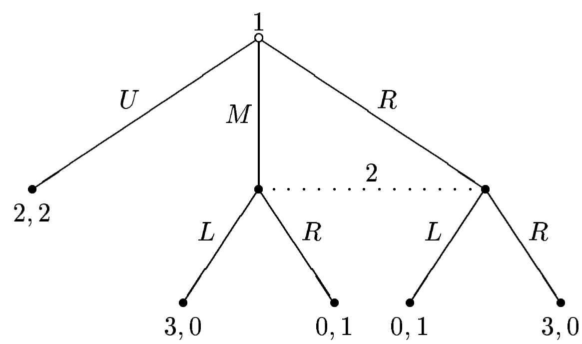

Pure Strategy Nash Equilibria

This game has no pure strategy equilibria. When Player 2 chooses L, Player 1's best reply is to choose M, but this cannot be Nash since then Player 2 has an incetive to instead choose R. Similarly, when Player 2 chooses R, Player 1's best reply is to choose D, but then Player 2 has an incentive to choose L. This is easiest to see when we look at the normal form for this game:

		Player 2 L	Player 2 R
Player 1	U	2,2	2,2
	M	3,0	0,1
	D	0,1	3,0

Mixed Strategy Nash Equilibria

The Nash equilibria of this game are of the form $\sigma_\lambda = (U; \lambda L, 1 - \lambda R)$ for all $\lambda \in [1/3, 2/3]$. Player 1 chooses the pure strategy U. Given Player 2's mixed strategy, Player 1's best reply is U. Since $EU(U) = 2, EU(M) = 3\lambda, EU(D) = 3 - 3\lambda$, when $\lambda \in [1/3, 2/3], EU(U) \geq EU(M)$ and $EU(U) \geq EU(D)$. When Player 1 chooses U, Player 2 is indifferent between her two strategies since she never gets the move, and so has no incentive to deviate from these mixed strategies.

To see why there are no other mixed strategy equilibria, we need to consider mixed strategies for Player 1. First, suppose Player 1 mixes only over M and D, so he puts zero probability on U. For this to happen, Player 2 must play a mixed strategy such that $EU(M) = EU(D)$. This implies $\lambda = 1/2$, but then $EU(M) = EU(D) = 3/2 < EU(U) = 2$, so Player 1 deviates to the pure strategy U. This also implies that 1 will never mix over all of his pure strategies. Second, suppose Player 1 mixes only over U and M, then λ must be $2/3$. This cannot be part of an equilibrium, however, since Player 2 has an incentive to deviate to the pure strategy R. Similarly, if 1 mixes only over U and D, 2 must mix so that $\lambda = 1/3$ and has no incentive to do so since her best reply is L. This exhausts all other types of mixed strategy equilibria.

Subgame Perfect Nash Equilibria

This game has only 1 subgame, so all of the Nash equilibria are also subgame perfect.

Perfect Bayesian Equilibria

Since we know that perfect Bayesian equilibria are a subset of subgame perfect Nash equilibria, we need only consider the strategy profiles σ_λ for $\lambda \in [1/3, 2/3]$. Note, however, that Player 2's information set is off the equilibrium path, so Bayes Rule does not apply. In order for Player 2 to use a mixed strategy, his posterior beliefs must be that $\mu = Pr(M|M \text{ or } D) = 1/2$. Thus, the set of perfect Bayesian equilibria of the game are the set of assessments $(\sigma_\lambda, \mu = 1/2)$ such that $\lambda \in [1/3, 2/3]$.

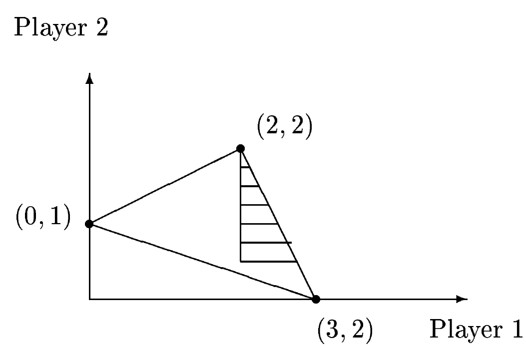

(b) There are three possible pairs of payoffs, so the set of feasible payoffs is the convex hull around the points $(2, 2), (3, 0), (0, 1)$. The minmax values are $v_1 = 2$ and $v_2 = 1/2$. Note that when Player 2 uses any mixed strategy with $\lambda \in [1/3, 2/3]$, Player 1 can choose U and get a payoff of 2 and when $\lambda \in [0, 1/3]$, Player 1 can get a payoff greater than 2 by choosing R, and when $\lambda \in [2/3, 1]$, Player 2 can get a payoff greater than 2 by playing L. When Player 1 uses a mixed strategy $1/2M, 1/2D$, he can hold Player 2 to an expected payoff of $1/2$. Thus, the average per period payoffs are shown in the triangle with horizontal lines.

Question 2

(a) One way to choose payoffs to represent risk averse preferences is to first let $U(\$10) = 10$ and $U(-\$10) = -10$. The lottery is a $1/2$ chance of getting \$10 and a $1/2$ chance of losing \$10, so the expected utility of the lottery is 0 (and the expected value of the lottery is \$0). Risk averse preferences imply that the utility from getting the expected value for sure will be greater than the expected utility from the lottery itself, so let $U(\$0) = 1$. More generally, any payoff for not betting that is greater than 0 will work. Here, both players are equally risk averse (and the level of risk aversion is small). The extensive form is then:

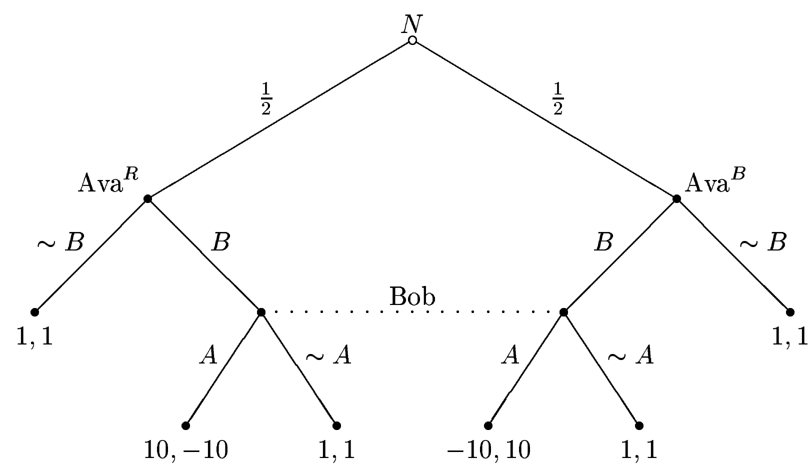

(b) A bet occurs when Ava offers it and Bob accepts. There are three possible strategies for Ava (2 separating, 1 pooling):

1. Ava offers B if the card is either black or red – Bayes Rule implies that the posterior beliefs equal the prior beliefs. Let $\mu = Pr(\text{Ava}^R|B)$, so in this case, $\mu = 1/2$. Since Bob's choice is between the lottery or getting nothing for sure, risk aversion implies that his best reply is to not accept.

2. Ava offers B if and only if the card is black – The posterior belief is that $\mu = 0$ and Bob accepts, but Ava^B clearly has an incentive to choose $\sim B$ because she wouldn't want to take a losing bet for sure.

3. Ava offers B if and only if the card is red – Bob's posterior belief is $\mu = 1$, in which case he knows that Ava has offered him a bet he is sure to lose, so his best reply is $\sim A$.

Thus, there are no PBEs where a bet is offered and accepted. The result is quite intuitive. If Bob knows that Ava always offers a bet, he won't accept because he is risk averse. If Ava offers the bet conditional on the card she sees, she only offers it when she knows she will win, and since Bob knows this, he won't accept a bet he knows he will lose.

(c) There is a separating PBE where Ava^R chooses B, Ava^B chooses $\sim B$, Bob plays $\sim A$, and $\mu = 1$. Given that Bob will choose $\sim A$ at his information set, Ava is indifferent between B and $\sim B$ regardless of the card she sees, so she has no incentive to deviate from her strategies. In part (b)3, we saw that in this case, $\mu = 1$ and Bob's best reply is $\sim A$. Part (b)2 shows why there is no other separating equilibrium.

Both types of pooling PBEs are possible. If Ava offers B regardless of the card she has, then $\mu = 1/2$, and Bob chooses $\sim A$. Given Bob's strategy, Ava is again indifferent between B and $\sim B$ regardless of the card, so she has no incentive to deviate. If Ava never offers the bet (always chooses $\sim B$), then Bob's information set is never reached. For these particular preferences, Bob's strategy is optimal for any posterior $\mu \leq 11/20$, and these beliefs are consistent because Bayes Rule does not apply.

Question 3

(a)

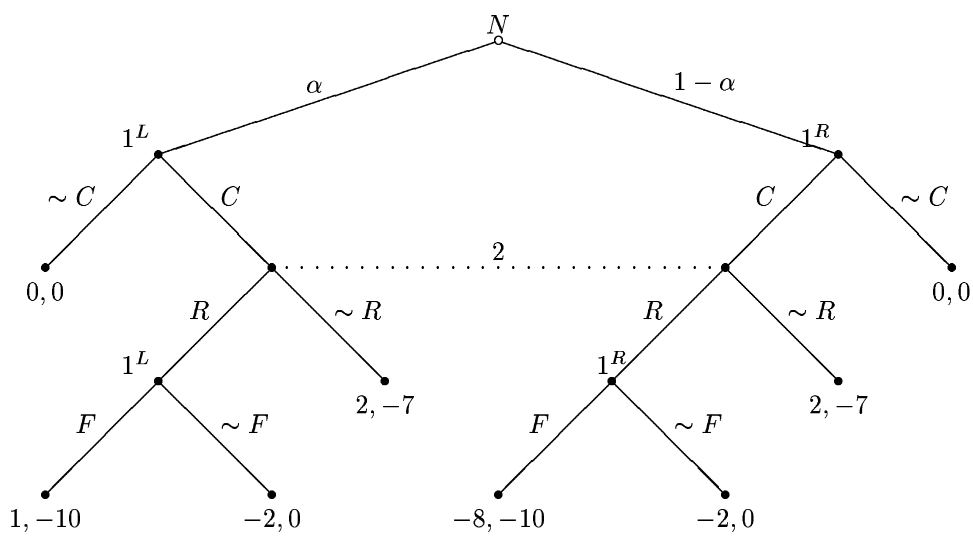

(b) We can simplify the game by noting that at the end of the game, if reached, 1^L will choose F, while 1^R will choose $\sim F$. Next, we note that regardless of State 2's action, 1^R prefers C. Thus, any pooling PBE must have both types of State 1 choosing C. In order for 1^R to also prefer C, State 2 must choose $\sim R$ in equilibrium. Finally, since it is a pooling equilibrium where State 2 will be given the move on the equilibrium path, the posterior belief $\mu = Pr(1^L|C) = \alpha$, and State 2 prefers $\sim R$ when $EU(R) \leq EU(\sim R) \Rightarrow -10\mu \leq -7 \Rightarrow \mu \geq 7/10$. Therefore, there is a pooling equilibrium when $\alpha \geq \alpha^* = 7/10$. A complete characterization of the equilibrium is therefore the assessment (σ, μ), where

$$\sigma = \begin{cases} 1^L: & C, F \\ 1^R: & C, \sim F \\ 2: & \sim R \end{cases}$$

$$\mu = \alpha$$

provided that $\alpha \geq 7/10$. The intuition is that the type 1^R bluffs and mimics 1^L, challenging State 2 even though it knows it won't carry through with its threat because State 2 has prior beliefs that it is more likely to be facing 1^L (which it is sure will carry through with its threat), leading State 2 to prefer not to resist.

(c) Since we know that 1^L always wants to play C, then in a semi-separating equilibrium, it must be that 1^R that mixes. In order for 1^R to mix, State 2 must play a mixed strategy that induces $EU_{1^R}(C) = EU_{1^R}(\sim C)$. Let $r = Pr(R)$, then $-2r + 2(1-r) = 0 \Rightarrow r = 1/2$. In turn, we need to figure out the posterior beliefs and mixed strategy of State 1^R that induces State 2 to mix. From part (c), we know that State 2 is indifferent between R and $\sim R$ when $\mu = 7/10$. Let q be the probability that 1^R plays C. Using Bayes' Rule,

$$\begin{aligned} \mu &= \frac{Pr(C|1^L)Pr(1^L)}{Pr(C|1^L)Pr(1^L) + Pr(C|1^R)Pr(1^R)} \\ \Rightarrow \frac{7}{10} &= \frac{\alpha}{\alpha + q(1-\alpha)} \\ \Rightarrow q &= \frac{3\alpha}{7(1-\alpha)} \end{aligned}$$

A complete characterization of the semi-separating PBE when $\alpha < 7/10$ is therefore given by:

$$\sigma = \begin{cases} 1^L: & C, F \\ 1^R: & \frac{3\alpha}{7(1-\alpha)}C, 1 - \frac{3\alpha}{7(1-\alpha)} \sim C, \sim F \\ 2: & \frac{1}{2}R, \frac{1}{2} \sim R \end{cases}$$

$$\mu = \frac{7}{10}$$

Intuitively, 1^R knows that State 2 is less sure a priori that it is facing 1^L and therefore cannot afford to always bluff. Notice that by mixing, State 2 exposes itself to the possibility of going to war with 1^L while also allowing the possibility that it may yield to type 1^R state.