EconS 503 - Microeconomic Theory II
Homework #7 - Answer key

1. **Cheap talk when the expert receives imprecise signals.** Consider the following cheap talk model between an expert (E), such a special interest group, and a decision maker (DM), such as a politician. For simplicity, assume that the state of the world is discrete, either $\theta = 1$ or $\theta = 0$ with prior probability $p \in (0, 1)$ and $1 - p$, respectively. The expert privately observes an informative but noisy signal s, which also takes two discrete values $s \in \{0, 1\}$. The precision of the signal is given by the conditional probability
$$\text{prob}(s = k|\theta = k) = q,$$
where $k = \{0, 1\}$, and $q > \frac{1}{2}$. In words, the probability that the signal s coincides with the true state of the world θ is q (precise signal), while the probability of an imprecise signal where $s \neq \theta$ is $1 - q$. The time structure of the game is as follows:

   1) Nature chooses θ according to the prior p.
   2) Expert observes signal s and reports a message $m \in \{0, 1\}$
   3) Decision maker observes m and responds with $x \in \{0, 1\}$
   4) θ is observed and payoffs are realized

   The payoff function for the decision maker is
$$u(x, \theta) = \left(\theta - \frac{1}{2}\right) x$$
   while that of the expert is
$$v(m, \theta) = \begin{cases} 1, & \theta = m \\ 0, & \theta \neq m \end{cases}$$
   which, in words, indicates that the expert's payoff is 1 when the message she sends coincides with the true realization of the state of the world, but becomes zero otherwise. Importantly, her payoff is unaffected by the signal, which she only uses to infer the actual realization of parameter θ. Intuitively, $v(m, \theta)$ is often understood as a "reputation function" since it provides the expert with a payoff of 1 only when his message was an accurate representation of the true state of the world (which in this model he does not precisely observe).

   (a) Is there a Perfect Bayesian equilibrium (PBE) in which the expert reports his signal truthfully?

   - Is there a Perfect Bayesian equilibrium (PBE) in which the expert reports his signal truthfully?

- *Updated beliefs.* In a strategy profile where the expert sends a message that coincides with the signal she receives (that is, sending message $m = 1$ after receiving signal $s = 1$, but sending message $m = 0$ after receiving signal $s = 0$)[1], the decision maker and the expert sustain the same beliefs about θ since $m = s$. Specifically, after receiving a signal of $s = 1$ (a message of $m = 1$), both expert and decision maker use Bayes' rule to update their beliefs yielding
$$\mu_1 = \frac{pq}{pq + (1-p)(1-q)}$$
while after receiving a signal of $s = 0$ (a message of $m = 0$), both expert and receiver updated their beliefs as follows
$$\mu_0 = \frac{p(1-q)}{p(1-q) + (1-p)q}$$

- *Decision maker's response.* Given the above beliefs, after receiving a message $m = 1$ from the expert, the decision maker responds with $x = 1$ if
$$\mu_1 \left(1 - \frac{1}{2}\right) 1 + (1-\mu_1)\left(0 - \frac{1}{2}\right) 1 \geq \mu_1 \left(1 - \frac{1}{2}\right) 0 + (1-\mu_1)\left(0 - \frac{1}{2}\right) 0$$
or
$$\mu_1 \frac{1}{2} + (1-\mu_1)\left(-\frac{1}{2}\right) \geq 0$$
or, simplifying, $\mu_1 \geq \frac{1}{2}$. From the above expression of posterior belief μ_1, this condition holds if
$$\frac{pq}{pq + (1-p)(1-q)} \geq \frac{1}{2}$$
or, after rearranging, $p \geq 1 - q$, which holds by assumption. That is, the decision maker responds with $x = 1$ after receiving message $m = 1$ for all admissible parameter values. Similarly, after receiving a message $m = 0$, the decision maker responds with $x = 1$ if
$$\mu_0 \left(1 - \frac{1}{2}\right) 1 + (1-\mu_0)\left(0 - \frac{1}{2}\right) 1 \geq \mu_0 \left(1 - \frac{1}{2}\right) 0 + (1-\mu_0)\left(0 - \frac{1}{2}\right) 0$$
or, after simplying, $\mu_0 \geq \frac{1}{2}$. From the above expression of posterior belief μ_0, this condition holds if
$$\frac{p(1-q)}{p(1-q) + (1-p)q} \geq \frac{1}{2}$$
or $p \geq q$. In words, the decision maker responds with $x = 1$ after observing message $m = 0$ when the probability of $\theta = 1$, p, is higher than the probability of the expert receiving precise signals, q. Otherwise (when $p < q$), the decision maker responds with $x = 0$ after observing message $m = 0$. Therefore, when $p < q$ we can say that the decision maker responds with an action $x(m) = m$ to every message $m \in \{0, 1\}$ he receives from the sender.

[1] This truthful reporting of signals can be described more compactly by saying that the expert's strategy is a message $m(s) = s$ for every signal $s \in \{0, 1\}$.

- *Expert's messages - After receiving signal $s = 1$.* If the expert reports her signal truthfully (sending message $m = 1$), her expected payoff is

$$\mu_1 v(1,1) + (1 - \mu_1)v(1,0) = \mu_1$$

  Intuitively, the above expression says that the expert sends a message $m = 1$ but does not know if the state of the world is $\theta = 1$, which yields a payoff of 1 since $\theta = m$; or if the state of the world is $\theta = 0$, which yields a payoff of zero for her since $\theta \neq m$. If, instead, she misreports her signal (sending message $m = 0$), her expected payoff becomes

$$\mu_1 v(0,1) + (1 - \mu_1)v(0,0) = 1 - \mu_1$$

  Therefore, the expert truthfully reports her signal if $\mu_1 \geq 1 - \mu_1$, or $\mu_1 \geq \frac{1}{2}$. Using the expression of posterior belief μ_1, we obtain that

$$\frac{pq}{pq + (1-p)(1-q)} \geq \frac{1}{2}$$

  which collapses to $p \geq 1 - q$. In words, after receiving a signal of $s = 1$, the expert truthfully conveys her signal if the probability of receiving such a signal is higher than the probability of an imprecise signal, $1 - q$.

- *Expert's messages - After receiving signal $s = 0$.* If the expert reports his signal truthfully (that is, sending message $m = 0$), her expected payoff is

$$\mu_0 v(0,1) + (1 - \mu_0)v(0,0) = \mu_0 0 + (1 - \mu_0)1 = 1 - \mu_0$$

  Intuitively, the above expression says that the expert sends a message $m = 0$ but does not know if the state of the world is $\theta = 1$, which yields a payoff of zero for her since $\theta \neq m$; or if the state of the world is $\theta = 0$, which yields a payoff of 1 since $\theta = m$. If, instead, the expert sends message $m = 1$ (lying about her message), her expected payoff becomes

$$\mu_0 v(1,1) + (1 - \mu_0)v(1,0) = \mu_0 v(1,1) + (1 - \mu_0)v(1,0) = \mu_0$$

  Therefore, the expert truthfully reports her signal if $1 - \mu_0 \geq \mu_0$, or $\frac{1}{2} \geq \mu_0$. Examining the expression of posterior belief μ_0, we find that

$$\frac{p(1-q)}{p(1-q) + (1-p)q} \leq \frac{1}{2}$$

  simplifies to $p \leq q$. In words, after receiving a signal of $s = 0$, the expert truthfully conveys her signal if the probability of an accurate signal, q, is higher than the probability of receiving a signal of $s = 1$. Combining the above conditions $p \geq 1 - q$ and $p \leq q$, we obtain $1 - q \leq p \leq q$.

- *Summary*:
  - When $p \geq q$, a PBE where the expert truthfully reports her signal can be sustained if $1 - q \leq p \leq q$ (from the expert) and $p \geq q$ (from the decision maker), which are only compatible when $p = q$. In words, the

prior probability of the state of the world being $\theta = 1$, p, must coincide with the probability with which the expert receiving precise signals, q. While the expert truthfully reports her signals to the decision maker, the decision maker does not follow the expert's advise when observing a message of $m = 0$.

– When $p < q$, a PBE where the expert truthfully reportsher signal can be sustained if $1 - q \leq p \leq q$ (from the expert) and $p < q$ (from the decision maker), where the expert sends a message $m(s) = s$ for every signal $s \in \{0, 1\}$ she received, while the decision maker responds with an action $x(m) = m$ for every message $m \in \{0, 1\}$ he receives. In this PBE, the expert truthfully reports her signals to the decision maker, and the decision maker follows the expert's advise after every message.

2. **Policy announcements as signals.** Consider Downs' (1957) model of voting with a continuum of voters with policy ideals in the interval $[0, 1]$, distributed according to cummulative distribution function $F(x)$ with positive and continuous density in $[0, 1]$. The median voter $x = m$ satisfies $F(m) = \frac{1}{2}$, and is either low (L) or high (H), where $L < H$, with equal probabilities. The time structure of the game is the following:

1) Political candidate 1 privately observes the position of the median voter (that is, $m = L$ or $m = H$), and announces a policy position p_1.
2) Candidate 2 observes p_1, and updates its beliefs about the position of the median voter. Candidate 2 then responds announcing his own policy p_2.
3) After observing policies p_1 and p_2, voters vote for the candidate who is closest to their ideal policy. In case of a tie, you can assume that candidates evenly share votes.

Candidates only care about winning the election and assign a payoff of 1 to winning, $\frac{1}{2}$ to a tie, and 0 to losing.

(a) Find at least one separating Perfect Bayesian Equilibria (PBEs).

- In a separating strategy profile, candidate 1's policies are different after observing a low and a high median voter, that is, $p_1(L) \neq p_1(H)$.
- *Candidate 2's beliefs.* Candidate 2's beliefs must be concentrated, that is, after observing policy $p_1(L)$, he beliefs that $\mu\left(L|p_1(L)\right) = 1$; while after observing policy $p_1(H)$ he beliefs that $\mu\left(H|p_1(H)\right) = 1$. All other policies $p_1 \neq p_1(L) \neq p_1(H)$ are regarded as off-the-equilibrium messages, and Bayes' rule yields an undefined belief $\frac{0}{0}$, entailing that we need to leave off-the-equilibrium beliefs unrestricted, that is, $\mu\left(H|p_1\right) \in [0, 1]$.
- *Candidate 2's response.* Given the above updated beliefs, candidate 2 responds with policy $p_2 = L$ after $p_1(L)$, and with policy $p_2 = H$ after $p_1(H)$ since that maximizes his chances of winning the election (if player 1 announces $p_1(L) \neq L$ and $p_1(H) \neq H$ respectively) or his chances of tying the election (if player 1 announces $p_1(L) = L$ and $p_1(H) = H$ respectively). After any off-the-equilibrium policy announcement from candidate 1, $p_1 \neq p_1(L) \neq p_1(H)$,

candidate 2 exhibits off-the-equilibrium belief $\mu \equiv \mu(H|p_1) \in [0,1]$, and thus responds with the *expected* median policy

$$\mu H + (1-\mu)L$$

where the first term represents the probability that the median voter is high type, while the second term indicates the probability he is low type.

- *Candidate 1's messages.* Candidate 1, when choosing his policy announcement p_1 in the first stage, anticipates candidate 2's optimal responses; as discussed in the previous bullet point. When he observes that the median voter is $m = L$, he chooses a policy $p_1(L) = L$, because otherwise candidate 2 could win for sure by responding with $p_2 = L$ when player 1 chooses any other policy $p_1(L) \neq L$. Similarly, after observing that the median voter is $m = H$, candidate 1 chooses a policy $p_1(H) = H$; as otherwise candidate 2 could win for sure by responding with $p_2 = H$ after $p_1(H) \neq H$.
- Therefore, in the separating equilibrium, after observing that the median voter is $m = k$, where $k = \{H, L\}$, candidate 1 responds with a policy that coincides with the median voter's ideal, $p_1(k) = k$, and candidate 2 responds with the same policy $p_2(p_1(k)) = k$. As a result, there is a tie in the election, and each player receives an expected payoff of $\frac{1}{2}$.

(b) Find at least one pooling Perfect Bayesian Equilibrium (PBE).

- In a pooling strategy profile, candidate 1 announces the same policy after observing a low and a high median voter, that is, $p_1(L) = p_1(H) = \overline{p}_1$.
- *Candidate 2's beliefs.* Candidate 2 posterior beliefs after observing the pooling policy $\overline{p}_1$ cannot be updated with Bayes' rule, so they coincide with his priors, $\frac{1}{2}$, that is, $\mu(L|\overline{p}_1) = \mu(H|\overline{p}_1) = \frac{1}{2}$. (Recall that both types of median voters are equally likely.) Like in part (a) of the exercise, all policies different than the pooling policy $\overline{p}_1$, $p_1 \neq \overline{p}_1$, are regarded as off-the-equilibrium messages, and off-the-equilibrium beliefs are left unrestricted, that is, $\mu(H|p_1) \in [0,1]$.
- *Candidate 2's response.* Given the above updated beliefs, after extremely low policies from candidate 1, $\overline{p}_1 < L$, candidate 2 responds with policy $p_2 = L$, since by doing so he win the election for sure. Intuitively, even if the median voter was $m = L$, candidate 1's policy is so radicalized that responding with $p_2 = L$ is enough to win the election, both if the median voter is L and H. Likewise, if candidate 1 announces extremely high policies, $\overline{p}_1 > H$, candidate 2 responds with $p_2 = H$, since that also lets candidate 2 to win the election with certainty. A similar argument as above explains this optimal response from candidate 2. Finally, when candidate 1 announces intermediate policies, $H \geq \overline{p}_1 \geq L$, candidate 2 can optimally respond by mimicking candidate 1's announcement, $\overline{p}_1$. Doing so leads to a tie, with an expected payoff of $\frac{1}{2}$; while choosing a different response in that interval, $p_2 \in [L, H]$, leads to candidate 1 winning the election with probability 50% (which yields a payoff of zero to candidate 2) or a tie (with expected payoff $\frac{1}{2}$), but never wins the election. Therefore, player 2 prefers to mimic candidate 1's announcement, $\overline{p}_1$.
- *Candidate 1's messages.* When candidate 1 observes that the median voter is $m = L$, he chooses a pooling policy $\overline{p}_1$, where $\overline{p}_1 \in [L, H]$. This leads

candidate 2 to respond mimicking that announcement, $p_2 = \overline{p}_1$, ultimately producing a tie in the election and an expected payoff of $\frac{1}{2}$. If, instead, candidate 1 deviates to another policy $\overline{p}_1 < L$, player 2 responds with $p_2 = L$ and wins the election for sure. Similarly, if candidate 1 deviates to $\overline{p}_1 > H$, player 2 responds with $p_2 = H$ and wins the election for sure. A similar argument applies when candidate 1 observes that the median voter is $m = H$, where he chooses a pooling policy $\overline{p}_1 \in [L, H]$, which is also mimicked by candidate 2, $p_2 = \overline{p}_1$, leading to a tie as well.

- Therefore, in the pooling equilibrium, after observing that the median voter is $m = k$, where $k = \{H, L\}$, candidate 1 responds with a pooling policy $\overline{p}_1 \in [L, H]$, and candidate 2 responds with the same policy $p_2 = \overline{p}_1$. As a result, there is a tie in the election, and each player receives an expected payoff of $\frac{1}{2}$.

3. **Collusion with probability of being caught - Harrington (2014).**[2] Consider an industry with N firms. For generality, we do not assume whether they compete in quantities or prices yet, nor the inverse demand function or costs they face. Consider that firms are symmetric and in the Nash equilibrium of the unrepeated game, every firm earns profits π^N, so we label the present value of the noncollusive stream as

$$V^N \equiv \pi^N + \delta\pi^N + ... = \frac{1}{1-\delta}\pi^N.$$

When firms collude, each of them earns profit π^C, where $\pi^C > \pi^N$. When a firm unilaterally deviates from the collusive outcome, it earns a deviating profit of π^D, where $\pi^D > \pi^C$ in that period. Consider a standard Grim-Trigger strategy (GTS) where every firm chooses to collude in period $t = 1$, and continues to do so in subsequent periods $t > 1$ if all firms colluded in previous periods. If one firm did not cooperate in previous periods, however, all firms revert to the Nash equilibrium of the unrepeated game, earning π^N thereafter (permanent punishment scheme). For simplicity, assume that all firms exhibit the same discount factor $\delta \in (0, 1)$.

   (a) Find the minimal discount factor δ that sustains this GTS as a subgame perfect equilibrium of the game.

   - After a history of cooperation, every firm i keeps cooperating as long as

$$\underbrace{\frac{1}{1-\delta}\pi^C}_{\text{Cooperation}} \geq \underbrace{\pi^D}_{\text{Deviation}} + \underbrace{\frac{\delta}{1-\delta}\pi^N}_{\text{Permanent punishment}}$$

   which, after solving for discount factor δ, yields

$$\delta \geq \widehat{\delta} \equiv \frac{\pi^D - \pi^C}{\pi^D - \pi^N}$$

   The above number is a positive number since the profit from deviating, π^D, satisfies $\pi^D > \pi^C$ and $\pi^D > \pi^N$ by definition.

[2] Harrington, Joseph E. Jr. (2014) "Penalties and the Deterrence of Unlawful Collusion," *Economic Letters*, 124, pp. 33-36.

- *Comparative statics of cutoff $\widehat{\delta}$*:
  - Cutoff $\widehat{\delta}$ is increasing in π^N. Intuitively, as the profits from reverting to the Nash equilibrium of the unrepeated game π^N increase, the punishment from deviation become less severe, ultimately making the deviation more attractive.
  - In contrast, cutoff $\widehat{\delta}$ is decreasing in π^C. In words, this indicates that, when the profits from cooperation π^C increase, deviation becomes less attractive and can be sustained under larger values of discount factor δ.
  - Finally, cutoff $\widehat{\delta}$ is increasing in the deviating profit π^D since $\frac{\partial\widehat{\delta}}{\partial\pi^D} = \frac{\pi^C-\pi^N}{(\pi^D-\pi^N)^2} > 0$, since profits from cooperation π^C satisfy $\pi^C > \pi^N$. Intuitively, when deviation becomes more attractive, the GTS can only be sustained for more restrictive conditions on discount factor δ.

(b) For the rest of the exercise, let us assume that the cartel faces a exogeneous probability p of being discovered, prosecuted, and convicted, by a regulatory agency such as the Federal Trade Commission. If caught and convicted in period t, a firm must pay a fine F^t, where $F^t = \beta F^{t-1} + f$. Parameter $1-\beta$ can be understood as the depreciation rate, which we assume to satisfy $\beta \in (0,1)$ to guarantee that the penalty is bounded. In addition, assume that $F^0 = 0$, so that $F^1 = f$, $F^2 = \beta F^1 + f$, and similarly for subsequent periods. Find the collusive value $V^C(F)$ given an accumulated penalty F. [*Hint*: Solve for $V^C(F)$ recursively.]

- The collusive value $V^C(F)$ is defined recursively, for period $t=1$ and $t=2$, as follows

$$V^C(F^1) = \pi^C + \underbrace{p\left[\frac{\delta}{1-\delta}\pi^N - F^2\right]}_{\text{Detected}} + \underbrace{(1-p)\delta V^C(F^2)}_{\text{Not detected}}$$

  where penalties are $F^1 = \beta F^0 + f = f$ and $F^2 = \beta F^1 + f = \beta F + f$. More generally for any two periods t and $t+1$, the collusive value is defined as

$$V^C(F) = \pi^C + \underbrace{p\left[\frac{\delta}{1-\delta}\pi^N - (\beta F + f)\right]}_{\text{Detected}} + \underbrace{(1-p)\delta V^C(\beta F + f)}_{\text{Not detected}}$$

  In words, the collusive value includes the collusive profit, π^C, and:
  - If the cartel is detected, which happens with probability p, every firm must pay a penalty $\beta F + f$ after being detected, and a generate stream of Nash equilibrium profits π^N thereafter since we assumed that after being detected firms can never form a cartel in the future.
  - If the cartel is undetected, which occurs with probability $1-p$, the stream of collusive payoffs starts in the next period, giving rise to continuation payoff $\delta V^C(\beta F + f)$.

- Since the collusive value $V^C(F)$ shows up in both the left- and right-hand side of the above expression, we can solve for $V^C(F)$ to obtain

$$
\begin{aligned}
V^C(F) &= \pi^C + p\left[\frac{\delta}{1-\delta}\pi^N - (\beta F + f)\right] + (1-p)\delta V^C(\beta F + f) \\
&= \pi^C + p\left[\frac{\delta}{1-\delta}\pi^N - (\beta F + f)\right] + (1-p)\delta\pi^C + p(1-p)\delta\frac{\delta}{1-\delta}\pi^N \\
&- p(1-p)\delta\left(\beta\left(\beta F + f\right) + f\right) + (1-p)^2\delta^2 V^C(\beta\left(\beta F + f\right) + f) \\
&= \left(\pi^C + p\pi^N\frac{\delta}{1-\delta}\right)\left[1 + \delta\left(1-p\right)\right] - pF\left[\beta + (1-p)\delta\beta^2\right] \\
&- pf\left[1 + \delta\left(1-p\right) + \delta\left(1-p\right)\beta\right] + (1-p)^2\delta^2 V^C(\beta\left(\beta F + f\right) + f) \\
&= \left(\pi^C + p\pi^N\frac{\delta}{1-\delta}\right)\sum_{t=0}^{T}\left[\delta\left(1-p\right)\right]^t - p\beta F\sum_{t=0}^{T}\left[\left(1-p\right)\beta\delta\right]^t \\
&- pf\sum_{t=0}^{T}\delta^t\left(1-p\right)^t\left(\sum_{s=0}^{t}\beta^s\right) + (1-p)^{T+1}\delta^{T+1}V^C\left(\beta^{T+1}F + f\sum_{t=0}^{T}\beta^t\right)
\end{aligned}
$$

From the transversality condition that

$$
\lim_{T\to\infty}(1-p)^{T+1}\delta^{T+1}V^C\left(\beta^{T+1}F + f\sum_{t=0}^{T}\beta^t\right) = 0
$$

which means the present discounted value of collusion at infinity is zero, we obtain

$$
\begin{aligned}
V^C(F) &= \left(\pi^C + p\pi^N\frac{\delta}{1-\delta}\right)\frac{1-\left[\delta\left(1-p\right)\right]^{T+1}}{1-\delta\left(1-p\right)} - p\beta F\frac{1-\left[\left(1-p\right)\beta\delta\right]^{T+1}}{1-\beta\delta\left(1-p\right)} \\
&- pf\sum_{t=0}^{T}\delta^t\left(1-p\right)^t\frac{1-\beta^{t+1}}{1-\beta} \\
&= \left(\pi^C + p\pi^N\frac{\delta}{1-\delta}\right)\frac{1-\left[\delta\left(1-p\right)\right]^{T+1}}{1-\delta\left(1-p\right)} - p\beta F\frac{1-\left[\left(1-p\right)\beta\delta\right]^{T+1}}{1-\beta\delta\left(1-p\right)} \\
&- \frac{pf}{1-\beta}\frac{1-\left[\delta\left(1-p\right)\right]^{T+1}}{1-\delta\left(1-p\right)} + \frac{p\beta f}{1-\beta}\frac{1-\left[\beta\delta\left(1-p\right)\right]^{T+1}}{1-\beta\delta\left(1-p\right)}
\end{aligned}
$$

For an infinite horizon game, we take $T \to \infty$, such that

$$
\begin{aligned}
V^C(F) &= \left(\pi^C + p\pi^N \frac{\delta}{1-\delta}\right) \frac{1}{1-\delta(1-p)} - \frac{p\beta F}{1-\beta\delta(1-p)} \\
&\quad - \frac{pf}{1-\beta}\frac{1}{1-\delta(1-p)} + \frac{p\beta f}{1-\beta}\frac{1}{1-\beta\delta(1-p)} \\
&= \frac{\pi^C + p\pi^N \frac{\delta}{1-\delta}}{1-\delta(1-p)} - \\
&\quad \frac{p\beta F(1-\beta)[1-\delta(1-p)] + pf[1-\beta\delta(1-p)] - p\beta f[1-\delta(1-p)]}{(1-\beta)[1-\delta(1-p)][1-\beta\delta(1-p)]} \\
&= \underbrace{\frac{\pi^C + p\pi^N \frac{\delta}{1-\delta}}{1-\delta(1-p)}}_{\substack{\text{Expected present value} \\ \text{of profits from the product market}}} - \underbrace{\frac{p(\beta[1-\delta(1-p)]F + f)}{[1-\delta(1-p)][1-\beta\delta(1-p)]}}_{\text{Expected discounted penalty}}
\end{aligned}
$$

Intuitively, the first term represents the expected present value of profits from the product market, which includes the possibility of earning collusive profits π^C for the periods that the cartel is undetected, or Nash equilibrium profits π^N for all periods after the cartel is detected. The second term indicates the expected discounted penalty once the firm is discovered, prosecuted, and convicted.

(c) Write down the condition (inequality) expressing that every firm has incentives to collude, obtaining $V^C(F)$ rather than deviating. For simplicity, you can assume that if the cartel is convicted during the deviation period, it has no chances of being caught during the permanent punishment phase.

- Every firm cooperates if and only if

$$
\underbrace{V^C(F)}_{\text{Collusion}} \geq \underbrace{\pi^D - p(\beta F + f)}_{\text{Expected profit from deviation}} + \underbrace{\frac{\delta}{1-\delta}\pi^N}_{\text{Permanent punishment}}
$$

Intuitively, when the firm deviates from the collusive price the cartel can still be detected by regulatory authorities, which occurs with probability p, and thus charged with a penalty $\beta F + f$.

(d) The steady-state penalty is $F = \frac{f}{1-\beta}$, which is found by solving $F = \beta F + f$. Evaluate the collusive value $V^C(F)$ at this penalty, and insert your result in the condition you found in part (c) of the exercise. Rearrange and interpret.

- Evaluating the collusive value $V^C(F)$ at the steady-state penalty $F = \frac{f}{1-\beta}$, we obtain

$$
V^C\left(\frac{f}{1-\beta}\right) = \frac{\pi^C + p\pi^N\left(\frac{\delta}{1-\delta}\right) - \frac{pf}{1-\beta}}{1-\delta(1-p)}
$$

- Substituting the above results into part (c) yields

$$
\frac{\pi^C + \frac{p\delta\pi^N}{1-\delta} - \frac{pf}{1-\beta}}{1-\delta(1-p)} \geq \pi^D - \frac{pf}{1-\beta} + \frac{\delta\pi^N}{1-\delta}
$$

Rearranging, we find

$$\pi^C + \delta(1-p)\frac{pf}{1-\beta} + \delta(1-p)\pi^D \geq \pi^D + \frac{\delta\pi^N}{1-\delta}\left[1-\delta(1-p)-p\right]$$

and solving for discount factor δ, we obtain

$$\delta \geq \widehat{\delta}(p) \equiv \frac{\pi^D - \pi^C}{(1-p)\left(\frac{pf}{1-\beta} + \pi^D - \pi^N\right)}$$

Note that cutoff $\widehat{\delta}(p)$ collapses to $\frac{\pi^D-\pi^C}{\pi^D-\pi^N}$ when the probability of cartel detection is zero ($p=0$) as in part (a) of the exercise. In contrast, when detection is perfect, $p=1$, cutoff $\widehat{\delta}(p)$ approaches infinity, thus indicating that condition $\delta \geq \widehat{\delta}(p)$ cannot hold for any admissible discount factor $\delta \in [0,1]$.

- Next, we differentiate the cutoff $\widehat{\delta}(p)$ with respect to p,

$$\frac{\partial\widehat{\delta}(p)}{\partial p} = \frac{\left(\pi^D-\pi^C\right)\left[\pi^D-\pi^N-(1-2p)\frac{f}{1-\beta}\right]}{\left[(1-p)\left(\frac{pf}{1-\beta}+\pi^D-\pi^N\right)\right]^2} \geq 0$$

Intuitively, as detection becomes more likely, the minimal discount factor sustaining collusion $\widehat{\delta}(p)$ increases, since firms are less attracted to collude.

- Lastly, we differentiate the cutoff $\widehat{\delta}(p)$ with respect to f,

$$\frac{\partial\widehat{\delta}(p)}{\partial f} = -\frac{p\left(\pi^D-\pi^C\right)}{(1-\beta)(1-p)\left(\frac{pf}{1-\beta}+\pi^D-\pi^N\right)^2} \leq 0$$

Intuitively, as the penalty becomes more severe, the minimal discount factor sustaining collusion $\widehat{\delta}(p)$ decreases, since firms have to face a larger penalty in expectation and have less to gain from collusion.

(e) *Bertrand competition.* Assume that firms compete a la Bertrand, selling homogeneous products with inverse demand function $p(Q) = 1 - Q$ where Q denotes aggregate output. All firms face a symmetric marginal cost $c > 0$. In this setting, every firm obtains zero profits in the Nash equilibrium of the unrepeated game, entailing $\pi^N = 0$. If a firm unilaterally deviates from the collusive price (charging a price infinitely close, but below, the collusive price), it captures all industry sales, earning a profit $\pi^D = N\pi^C$ during the deviating period. Evaluate your results from part (d) of the exercise in this context. Then discuss whether collusion becomes easier to sustain when the penalty f increases; and when the number of firms N increases.

- We first need to find the profits that firms obtain from cooperating by setting a collusive price. Since the inverse demand function is $p(Q) = 1 - Q$, the direct demand function becomes $Q = 1 - p$, implying that the joint-profit maximization problem in this setting is

$$\max_{p\geq 0}\ pQ - cQ = p(1-p) - c(1-p)$$

Differentiating with respect to price p yields $1-2p+c=0$, and solving for p we obtain a collusive price of $p^C = \frac{1+c}{2}$. Therefore, collusive profits for the industry are

$$N\pi^C = p^C(1-p^C) - c(1-p^C) = \left(1 - \frac{1+c}{2}\right)\left(\frac{1+c}{2} - c\right) = \frac{(1-c)^2}{4}$$

which implies that every firm's collusive profits, π^C, is $\frac{(1-c)^2}{4N}$.

- Evaluating our above results from part (d) at profits $\pi^N = 0$, $\pi^C = \frac{(1-c)^2}{4N}$, and $\pi^D = N\pi^C = \frac{(1-c)^2}{4}$, we find that every firm cooperates after a history of cooperation if and only if

$$\delta \geq \overline{\delta} = \frac{\frac{(1-c)^2}{4} - \frac{(1-c)^2}{4N}}{(1-p)\left[\frac{pf}{1-\beta} + \frac{(1-c)^2}{4}\right]}$$

$$= \frac{1}{1-p} \cdot \frac{\frac{(1-c)^2}{4}\left(1-\frac{1}{N}\right)}{\frac{4pf+(1-\beta)(1-c)^2}{4(1-\beta)}}$$

$$= \frac{N-1}{N(1-p)} \cdot \frac{(1-\beta)(1-c)^2}{4pf + (1-\beta)(1-c)^2}$$

- *Comparative statics of the cutoff* $\overline{\delta}$:
  - Differentiating the cutoff $\overline{\delta}$ with respect to N, we obtain

  $$\frac{\partial \overline{\delta}}{\partial N} = \frac{1}{N^2(1-p)} \cdot \frac{(1-\beta)(1-c)^2}{4pf + (1-\beta)(1-c)^2} \geq 0$$

  so that as the number of firm increases, a higher discount rate is needed to sustain collusion because the firm can capture a larger profit from deviation, π^D, relative to their collusion profit, π^C.
  - Differentiating the cutoff $\overline{\delta}$ with respect to f, we obtain

  $$\frac{\partial \overline{\delta}}{\partial f} = -\frac{N-1}{N(1-p)} \frac{4p(1-\beta)(1-c)^2}{\left[4pf + (1-\beta)(1-c)^2\right]^2} \leq 0$$

  Intuitively, as the penalty becomes more severe, the minimal discount factor sustaining collusion decreases, since firms have to face a larger penalty in expectation and have less to gain from collusion.

(f) *Cournot competition.* Assume now that firms compete a la Cournot, selling homogeneous products with inverse demand function $p(Q) = 1-Q$ where Q denotes aggregate output. All firms face a symmetric marginal cost $c > 0$. In this setting, every firm obtains profit π^N in the Nash equilibrium of the unrepeated game. If a firm unilaterally deviates from the collusive output, its deviating profit becomes π^D by capturing profit of the whole industry. Evaluate your results from part (d) of the exercise in this context. Then discuss whether collusion becomes easier to sustain when the penalty f increases; and when the number of firms N increases.

- *Cooperation.* We first need to find the profits that firms obtain from cooperating by setting a collusive level of output. Since the inverse demand function is $p(Q) = 1 - Q$, the joint profit maximization problem in this setting is

$$\max_{Q \geq 0} pQ - cQ = (1 - Q - c)\, Q$$

Differentiating with respect to price Q yields

$$Q^C = \frac{1-c}{2}$$

entailing that every individual firm produces $q^C = \frac{1-c}{2N}$ units of output. The collusive price is then

$$p^C = 1 - Q^C = 1 - \frac{1-c}{2} = \frac{1+c}{2}$$

Therefore, collusive profits for the industry are

$$\begin{aligned}\left(p^C - c\right) Q^C &= \left(\frac{1+c}{2} - c\right) \frac{1-c}{2} \\ &= \frac{(1-c)^2}{4}\end{aligned}$$

which implies that every firm's collusive profit is $\pi^C = \frac{(1-c)^2}{4N}$.

- *Unilateral deviation.* When all firms choose the collusive output $\frac{1-c}{2N}$, firm i's optimal deviation is found by inserting the collusive output profile of its rivals into firm i's best response function, $q_i(Q_{-i}) = \frac{1-c}{2} - \frac{1}{2}Q_{-i}$, as follows

$$\begin{aligned}q_i^{Dev} \equiv q_i(Q_{-i}^C) &= \frac{1-c}{2} - \frac{1}{2}\underbrace{(N-1)\,\frac{1-c}{2N}}_{Q_{-i}^C} \\ &= \frac{N+1}{4N}\,(1-c)\end{aligned}$$

which yields a market-clearing price of

$$\begin{aligned}p^{Dev} &= 1 - Q_{-i}^C - q_i^{Dev} \\ &= 1 - (N-1)\,\frac{1-c}{2N} - \frac{N+1}{4N}\,(1-c) \\ &= \frac{4N - (3N-1)\,(1-c)}{4N}\end{aligned}$$

such that the deviating firm generates a profit of

$$\begin{aligned}\pi^D &= \left(p^{Dev} - c\right) q_i^{Dev} \\ &= \left(\frac{4N - (3N-1)\,(1-c)}{4N} - c\right) \frac{N+1}{4N}\,(1-c) \\ &= \frac{(N+1)^2\,(1-c)^2}{16N^2}\end{aligned}$$

- Evaluating our above results from part (d) at profits $\pi^N = \left(\frac{1-c}{N+1}\right)^2$, $\pi^C = \frac{(1-c)^2}{4N}$, and $\pi^D = \frac{(N+1)^2(1-c)^2}{16N^2}$, we find that every firm cooperates after a history of cooperation if and only if

$$\begin{aligned}
\delta \geq \overline{\delta} &= \frac{\frac{(N+1)^2(1-c)^2}{16N^2} - \frac{(1-c)^2}{4N}}{(1-p)\left[\frac{pf}{1-\beta} + \frac{(N+1)^2(1-c)^2}{16N^2} - \left(\frac{1-c}{N+1}\right)^2\right]} \\
&= \frac{\frac{(1-c)^2}{16N^2}\left[N^2+2N+1-4N\right]}{\frac{1-p}{16N^2(N+1)^2}\left[\frac{16pfN^2(N+1)^2}{1-\beta} + \left((N+1)^4 - 16N^2\right)(1-c)^2\right]} \\
&= \frac{(1-c)^2(1-\beta)}{1-p} \cdot \frac{(N+1)^2(N-1)^2}{16pfN^2(N+1)^2 + (N-1)^2(N^2+6N+1)(1-c)^2(1-\beta)}
\end{aligned}$$

- *Comparative statics of the cutoff* $\overline{\delta}$:
  - Differentiating the cutoff $\overline{\delta}$ with respect to N, we obtain

$$\begin{aligned}
\frac{\partial\overline{\delta}}{\partial N} &= \frac{(1-c)^2(1-\beta)}{1-p} \cdot \frac{4N(N+1)(N-1)}{16pfN^2(N+1)^2 + (N-1)^2(N^2+6N+1)(1-c)^2(1-\beta)} \\
&\quad - \frac{(1-c)^2(1-\beta)}{1-p} \cdot \frac{32pfN(N+1)(N+2) + 4(N-1)(N^2+4N-1)(1-c)^2(1 -}{\left[16pfN^2(N+1)^2 + (N-1)^2(N^2+6N+1)(1-c)^2(1-\beta)\right.} \\
&\quad = \frac{4(1-c)^2(1-\beta)}{1-p} \cdot \\
&\qquad \frac{(N+1)(N-1)\left[4pf(N+1)^2(2N^3-N^2+1) + (N-1)^4(1-c)^2(1-\beta)\right]}{\left[16pfN^2(N+1)^2 + (N-1)^2(N^2+6N+1)\right]^2} \geq 0
\end{aligned}$$

  so that as the number of firm increases, a higher discount rate is needed to sustain collusion because the firm can capture a larger profit from deviation, π^D, relative to their collusion profit, π^C.
  - Differentiating the cutoff $\overline{\delta}$ with respect to f, we obtain

$$\begin{aligned}
\frac{\partial\overline{\delta}}{\partial f} &= -\frac{(1-c)^2(1-\beta)}{1-p} \cdot \\
&\qquad \frac{16pN^2(N+1)^2}{\left[16pfN^2(N+1)^2 + (N-1)^2(N^2+6N+1)(1-c)^2(1-\beta)\right]^2} \leq 0
\end{aligned}$$

  Intuitively, as the penalty becomes more severe, the minimal discount factor sustaining collusion decreases, since firms have to face a larger penalty in expectation and have less to gain from collusion.