# EC334 Fourth Seminar
# Principal-Agent Problem

Junxi Liu

# Cheap Talk Model

- Player 1 (or the Agent) has private information and both players′ payoffs depend on player 1′s private information. Play 2 (or the Principle) makes the decision.
- Player 1′s action is a message that has no direct effect on payoffs. (Cheap)
- General way of proceeding
  - Nature selects a type which is the private information that Player 1 has.
  - Player 1 chooses some message to send to Player 2
  - Player 2 observes the message while not knowing the true type of nature
  - Payoffs for both players are realized

# A Motivating Example — Finite Case

- Player 1 lives in Coventry area so she knows the true traffic conditions around Warwick Uni, which is one value of 1 (bad traffic), 3 (medium traffic), and 5 (good traffic). *(The type of Nature $\theta$)*

- Player 2 wants to join Warwick and is choosing where to live. He would prefer to live in one of the five locations: 1, 2, 3, 4, 5, where 1 is the closest to the uni but also most boring, and 5 is the furthest but also most lovely place.

- Player 1 is a good friend of Player 2 so she wants Player 2 to live closer to her, and Player 1 is living in 5 currently.

- Their respective payoffs are:
  - $v_1(a_2, \theta) = 5 - (\theta + 1.1 - a_2)^2$ → If *$\theta$=5*, player 1 would want player 2 to choose 5; if *$\theta$ = 3*, player 1 would prefer player 2 to choose 5 over 3:
  - $v_2(a_2, \theta) = 5 - (\theta - a_2)^2$

# A Motivating Example — Finite Case

**Claim 18.2** *There exists a* ***babbling equilibrium*** *in which player 1's message reveals no information and player 2 chooses an action to maximize his expected utility given his prior belief.*

***Proof*** To construct the babbling (perfect Bayesian) equilibrium let player 1's strategy be to send a message $a_1 \in \{1, 3, 5\}$ with equal probability of $\frac{1}{3}$ each regardless of $\theta$. This means that the message is completely uninformative: player 2 knows that regardless of the message, $\Pr\{\theta\} = \frac{1}{3}$ for all $\theta \in \{1, 3, 5\}$. This implies that player 2 maximizes his expected payoff,

$$\max_{a_2 \in \{1,2,3,4,5\}} Ev_2(a_2, \theta) = 5 - \tfrac{1}{3}(-(1 - a_2)^2) + \tfrac{1}{3}(-(3 - a_2)^2) + \tfrac{1}{3}(-(5 - a_2)^2),$$

which is maximized when $a_2 = 3$.[6] Because each of player 2's information sets is reached with positive probability, player 2's beliefs are well defined by Bayes' rule everywhere, and player 1 cannot change these beliefs by changing his strategy.[7] Hence player 1 is indifferent between each of the three messages and is therefore playing a best response. ■

**Construct and Verify**

# A Motivating Example — Finite Case

- Notice that the reason truth-telling is not an equilibrium is because
  - Were it to be an equilibrium, player 1 could always achieve his most preferred outcome by misleading player 2 into choosing a "higher" action.
  - The reason for this is precisely because player 1 is biased toward higher actions compared to player 2.
- If we can construct an equilibrium if we put an upper bound on how high a number player 2 will choose, and player 1 will partially tell the truth
  - We know from player 1's bias that he would always pre- fer that player 2 choose a somewhat higher action than player 2 would like, except for when player 2 wishes to choose the highest possible action
  - At the same time, for a low enough state, player 1 may prefer player 2's optimal action over much higher actions.

# A Motivating Example — Finite Case

**Claim 18.3** *There is a perfect Bayesian equilibrium in which player 1 partially reports the true state of the world. In particular player 1 truthfully reveals state $\theta = 1$ but pools information in states $\theta = 3$ and $\theta = 5$.*

***Proof*** To construct the equilibrium let player 1's strategy be to send a message $a_1 = 1$ when $\theta = 1$ and $a_1 \in \{3, 5\}$ with equal probability of $\frac{1}{2}$ each if $\theta \in \{3, 5\}$. This implies that player 2's posterior belief is $\Pr\{\theta = 1|a_1 = 1\} = 1$ while $\Pr\{\theta = 3|a_1\} = \Pr\{\theta = 5|a_1\} = 0.5$ for $a_1 \in \{3, 5\}$. As a consequence, if $a_1 = 1$ then player 2 will choose $a_2 = 1$, while if $a_1 \in \{3, 5\}$ then player 2 maximizes his expected payoff,

$$\max_{a_2 \in \{1,2,3,4,5\}} Ev_2(a_2, \theta) = 5 - \tfrac{1}{2}(-(3 - a_2)^2) + \tfrac{1}{2}(-(5 - a_2)^2),$$

which is maximized when $a_2 = 4$. To see that player 1's strategy is a best response it is easy to confirm that if $\theta \in \{3, 5\}$ then player 1 prefers $a_2 = 4$ over $a_2 = 1$, implying that he is indifferent between selecting $a_1 = 3$ and $a_1 = 5$ and strictly prefers either over $a_1 = 1$. If $\theta = 1$ then player 1 prefers $a_2 = 1$ over $a_2 = 4$ so he strictly prefers $a_1 = 1$ over his other possible actions. ■

# A Motivating Example — Continuous Case

The game is basically the same as the one described in the previous section, with the exception of $\theta \in \Theta = [0, 1]$ and the assumption that the state of the world $\theta$ is uniformly distributed on $[0, 1]$. Let player 2's action set include all real numbers, $a_2 \in \mathbb{R}$. Player 1's action set can be any arbitrary set of messages $A_1$, but it will be convenient to let $A_1 = [0, 1]$ so that the message space conforms with the state space $\Theta$. Player 2's payoff is $v_2(a_2, \theta) = -(a_2 - \theta)^2$, while player 1's payoff is $v_1(a_2, \theta) = -(a_2 + b - \theta)^2$, which implies that for any given value of $\theta \in [0, 1]$, player 2's optimal choice is $a_2 = \theta$, while player 1's is $a_2 = \theta + b$.

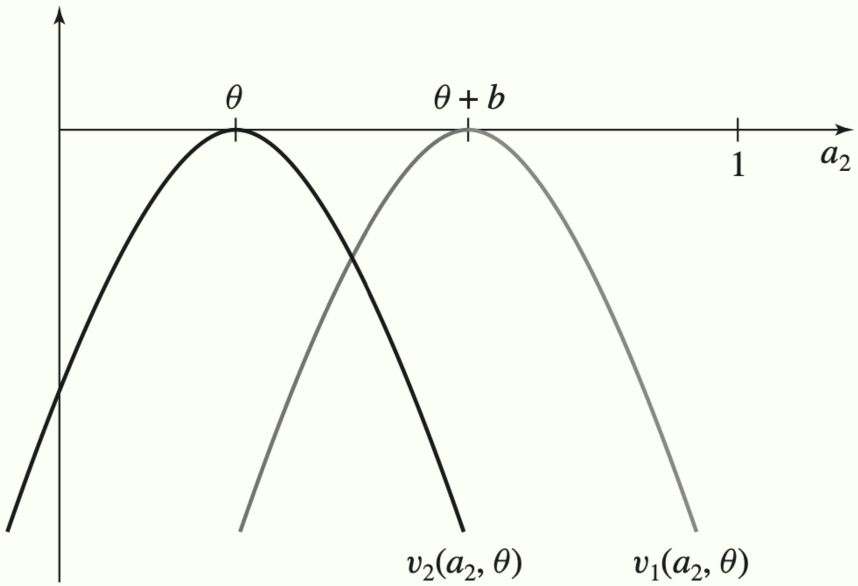


**FIGURE 18.2** Payoffs in the cheap-talk game.

# A Motivating Example — Continuous Case

**Claim 18.4** *There exists a babbling perfect Bayesian equilibrium in which player 1's message reveals no information and player 2 chooses an action to maximize his expected utility given his prior belief.*

***Proof*** We construct the equilibrium in a similar way to the finite case. Let player 1's strategy be to send a message $a_1 = a_1^B \in [0, 1]$ regardless of $\theta$. This means that the message is completely uninformative and player 2 believes that $\theta$ is distributed uniformly on [0, 1]. This implies that, conditional on receiving the message $a_1^B$, player 2 maximizes his expected payoff,

$$\max_{a_2 \in \mathbb{R}} E v_2(a_2, \theta) = \int_0^1 -(\theta - a_2)^2 d\theta = -\frac{1}{3} + a_2 - a_2^2,$$

which is maximized when $a_2 = \frac{1}{2}$. Let player 2's off-equilibrium-path beliefs be $\Pr\{\theta = \frac{1}{2} | a_1 \neq a_1^B\} = 1$ so that his off-the-equilibrium-path best response to any other message is $a_2 = \frac{1}{2}$ as well. It is easy to see that player 1 is indifferent between any of his messages and hence choosing $a_1 = a_1^B$ is a best response. ■

# Cheap Talk: Any Other Way to Improve?

- Two extreme cases: here is no truthful equilibrium and there is always a babbling equilibrium
  - What about something in between, i.e. not completely true but also truthful in terms of **some** information?
- Consider a two-message model: given the continuous type, play 1 can only report one of two message: low or high
- Recall the discrete example:
  - If the true type of the world is somewhere near 5, I would have the incentive to make it looks a bit higher to maximize my payoff
  - However, if the true type of the world is extremely low, I would not actually gain too much payoff by inducing player 2 to choose too high a number — I don't want to be too bad a friend!
- Let $a_2(a_1') < a_2(a_1'')$

# Cheap Talk: Two Message Model

**Claim 18.5** *In a two-message equilibrium player 1 must use a threshold strategy as follows: if* $0 \le \theta \le \theta^*$ *he chooses* $a_1'$*, whereas if* $\theta^* \le \theta \le 1$ *he chooses* $a_1''$.

***Proof*** For any $\theta$ player 1's payoffs from $a_1'$ and $a_1''$ are as follows:

$$v_1(a_2(a_1'), \theta) = -(a_2(a_1') + b - \theta)^2$$

$$v_1(a_2(a_1''), \theta) = -(a_2(a_1'') + b - \theta)^2,$$

which implies that the extra gain from choosing $a_1''$ over $a_1'$ is equal to

$$\Delta v_1(\theta) = -(a_2(a_1'') + b - \theta)^2 + (a_2(a_1') + b - \theta)^2.$$

The derivative of $\Delta v_1(\theta)$ is equal to $2(a_2(a_1'') - a_2(a_1')) > 0$ because $a_2(a_1') < a_2(a_1'')$. This implies that if type $\theta$ prefers to send message $a_1''$ over $a_1'$ then every type $\theta' > \theta$ will also prefer $a_1''$. Similarly if type $\theta$ prefers to send message $a_1'$ over $a_1''$ then so will every type $\theta' < \theta$. This in turn implies that if two messages are sent in equilibrium then there must be some threshold type $\theta^*$ as defined in claim 18.5. It follows that when $\theta = \theta^*$ player 1 must be indifferent between sending the two messages. ■

# Cheap Talk: Two Message Model

Now that we know what restrictions apply to player 1's strategy in a two-message equilibrium, we can continue to characterize the strategy of player 2 in a two-message perfect Bayesian equilibrium as follows:

**Claim 18.6** *In any two-message perfect Bayesian equilibrium in which player 1 is using a threshold $\theta^*$ strategy as described in claim 18.5, player 2's equilibrium best response is $a_2(a_1') = \frac{\theta^*}{2}$ and $a_2(a_1'') = \frac{1-\theta^*}{2}$.*

***Proof*** This follows from player 2's posterior belief and from him playing a best response. In equilibrium player 2's posterior following a message $a_1'$ is that $\theta$ is uniformly distributed on the interval $[0, \theta^*]$, and his posterior following a message $a_1''$ is that $\theta$ is uniformly distributed on the interval $[\theta^*, 1]$. Player 2 plays a best response if and only if he sets $a_2(a_1) = E[\theta | a_1]$, which proves the result. ■

# Cheap Talk: Two Message Model

Using claims 18.5 and 18.6 we can now characterize the two-message perfect Bayesian equilibrium as follows:

**Claim 18.7** *A two-message perfect Bayesian equilibrium exists if and only if* $b < \frac{1}{4}$.

***Proof*** From claim 18.5 we know that when $\theta = \theta^*$ player 1 must be indifferent between his two messages so that

$$v_1(a_2(a_1'), \theta^*) = v_1(a_2(a_1''), \theta^*),$$

which from claim 18.6 and from the fact that $\frac{\theta^*}{2} < \theta^* < \frac{1-\theta^*}{2}$ is equivalent to

$$\theta^* + b - \frac{\theta^*}{2} = -\left(\theta^* + b - \frac{1-\theta^*}{2}\right). \tag{18.1}$$

The solution to (18.1) is $\theta^* = \frac{1}{4} - b$, which can result in a positive value of $\theta^*$ only if $b < \frac{1}{4}$. To complete the specification of off-the-equilibrium-path beliefs, let player 2's beliefs be $\Pr\{\theta = \frac{\theta^*}{2} | a_1 \notin \{a_1', a_1''\}\} = 1$, so that he chooses $a_2 = \frac{\theta^*}{2}$, which causes player 1 to be indifferent between sending the message $a_1'$ and any other message $a_1 \notin \{a_1', a_1''\}$, implying that his threshold strategy is a best response. ■

# What About More Messages?

**Three Messages:** Consider the three-message equilibrium described in Section 18.2. Find the threshold values $\theta'$ and $\theta''$ and show that for this to be an equilibrium it must be that $b < \frac{1}{12}$.

To construct a three-message equilibrium we first divide the interval $[0, 1]$ into three segments with a message for each: $a_1'$ for $[0, \theta']$, $a_1''$ for $[\theta', \theta'']$, and $a_1'''$ for $[\theta'', 1]$. Player 2's best response must be $a_2(a_1') = \frac{\theta'}{2}$, $a_2(a_1'') = \frac{\theta''+\theta'}{2}$, and $a_2(a_1''') = \frac{1+\theta''}{2}$. Finally player 1 must be indifferent between $a_1'$ and $a_1''$ when $\theta = \theta'$, and indifferent between $a_1''$ and $a_1'''$ when $\theta = \theta''$. This last condition yields two equations with two unknowns that will determine the equilibrium thresholds $\theta'$ and $\theta''$.

You are left to solve this equilibrium in exercise 18.3 and show that it exists if and only if $b < \frac{1}{12}$. This should not be too surprising because the bias is what hampers player 1's ability to transmit more credible information in equilibrium. As the bias drops, we can construct finer and finer partitions of the interval $[0, 1]$ so that there is more meaningful communication between the sender and the receiver. In particular define $b_2 = \frac{1}{4}$ and $b_3 = \frac{1}{12}$. We have shown that if $b > b_2$ we have only a babbling equilibrium, if $b < b_2$ then we can construct a two-message equilibrium, and if $b < b_3$ then we can construct a three-message equilibrium.

# Quick Recap on Cheap Talk Models

- First, no matter how small the bias is, as long as $b > 0$ there is no fully truthful equilibrium, which is what we demonstrated in claim 18.1.
- There exists a series $b_2 > b_3 > b_4 > \ldots > b_M$ such that if $b < b_M$, then an equilibrium with M partitions exists.
- In any equilibrium there must be some loss of information that depends on the magnitude of the bias b.
- If $b < b_M$ then there are M different perfect Bayesian equilibria, starting with the babbling equilibrium up until the "most informative" equilibrium with M partitions.

# Principle-Agent Problem, in General

- Principal proposes contract and pay money
- Agent possesses information and takes action
- Outcomes for both depend on action and type
- Two usual constraints:
  - Incentive-compatibility (i.e. no incentive to deviate)
  - Participation (i.e. taking action should be better than the outside option)
- Most importantly, don't freak out: these problems essentially just consists of several parts:
  - Calculate expected payoff, either using summation (discrete) or integral (continuous)
  - Utility maximization problem solving
  - Check incentive and participation compatibility to make sure it's an equilibrium

# Example from 2021 Exam

A government procurement officer is trying to decide how many doses of a new coronavirus vaccine to order. This decision will depend on the effectiveness of the vaccine, which will be determined by clinical trials conducted by a scientific advisor. You may assume that the effectiveness of the vaccine is given by a random variable $\varepsilon$, uniformly distributed on the interval $[E_0, E_0 + 1]$. The scientific advisor believes the utility of ordering a quantity $Q$ is $U_A(Q|\varepsilon) = 1 + \varepsilon Q - Q^2$; if perfectly-informed about effectiveness, the procurement officer would value $Q$ at $U_G(Q|\varepsilon) = 1 + (\varepsilon + \beta)Q - Q^2$ where $\beta$ is a non-negative constant. After the trials, the government officer asks the scientific advisor to report on the vaccine's effectiveness and purchases the quantity that maximises his expected utility. The scientific advisor is not paid for his efforts.

a. How would you set up this problem? Can the government advisor be sure of purchasing the optimal quantity (according to his preferences)? If so, how? If not, why not? How does your answer depend on the size of $\beta$? (**15 marks**)

b. Suppose that the minimum effectiveness is $E_0 = 25\%$ and that $\beta = 5\%$. Find the 'babbing equilibrium' for this situation – how much will the government order and what expected utilities will the two parties get? (**7 marks**)

c. Now construct a two-part equilibrium – depending on the advice they receive, the government will place either a small order $Q^S$ or a large order $Q^L$. At what reported level of effectiveness will the government switch its order size, and what are the values of $Q^S$ and $Q^L$? (**10 marks**)

d. How would you find the most efficient equilibrium (you do not have to compute it explicitly, but should say how it could be identified)? (**10 marks**) a

This is a cheap talk problem; should note that first-best can be achieved only if $\beta = 0$. They should note that there is always an equilibrium in which the government ignores the advisor and purchases the a priori optimal amount $Q_0(E_0, E_0 + 1)$, which they compute in the next part. The optimal strategy is to partition the range of effectiveness into intervals $[x, y]$ and associate to each interval the order that maximises expected utility $Q^*(x, y) = \underset{Q}{\operatorname{argmax}} \int_x^y U_G(Q|\varepsilon)d\varepsilon$. The more intervals, the more efficient is the outcome, but the number (and thus the efficiency) are bounded above by a decreasing function of $\beta$. Finally, they should note that for any two adjacent intervals $[x, y]$ and $[y, z]$, the scientific adviser would be indifferent between the purchase levels for both intervals if she was convinced that the true effectiveness was exactly $y$ – in other words $U_A(Q^*(x, y)|y) = U_A(Q^*(y, z)|y)$.

b

In this case, there is only one purchase level regardless of report. If the government believes that the true state is uniformly distributed on $[a, b]$, it's expected utility for purchasing $Q$ is

$$EU_G(Q) = 1 + \left( \frac{\int_a^b \varepsilon d\varepsilon}{b - a} + \beta \right) Q - Q^2 = 1 + \left( \frac{a + b + 2\beta}{2} \right) Q - Q^2$$

Optimal $Q$ is $\frac{a+b+2\beta}{4}$. In this case, $Q = \frac{.25+1.25+.1}{4} = 0.4$, $U_G = 1.16$, $U_A = 1.14$.

c

Denote the critical report level by $\varepsilon^*$. The two order sizes are

$$Q^S = \frac{.25 + \varepsilon^* + 2 * .05}{4} = \frac{.35 + \varepsilon^*}{4}$$

$$Q^L = \frac{1.25 + \varepsilon^* + 2 * .05}{4} = \frac{1.35 + \varepsilon^*}{4}$$

$\varepsilon^*$ is defined by the condition that the advisor should be indifferent between $Q^S$ and $Q^L$ when the true state is $\varepsilon^*$. Solving $U_A\left(\frac{.35+\varepsilon^*}{4}\middle|\varepsilon^*\right) = U_A\left(\frac{1.35+\varepsilon^*}{4}\middle|\varepsilon^*\right)$ for $\varepsilon^*$ gives $\varepsilon =$ (in general, for any level of $\beta$), $\varepsilon^* = 0.75 + 2\beta$; in this case,

$$\varepsilon^* = 0.85$$

$$Q^S = 0.3$$

$$Q^L = 0.55$$

d

The most efficient equilibrium is the one with the greatest number of intervals, so they should look for the largest n s.t. there exists a sequence $0.25 = \varepsilon^1, \dots, \varepsilon^n = 1.25$ (or .26 for the 1% case) where

$$Q^i = \frac{\varepsilon^i + \varepsilon^{i+1} + 0.1}{4}$$

And for each $i = 1, \dots, n-1$

$$1 + \varepsilon^{i+1}\left[\frac{\varepsilon^i + \varepsilon^{i+1} + 0.1}{4}\right] - \left[\frac{\varepsilon^i + \varepsilon^{i+1} + 0.1}{4}\right]^2 = 1 + \varepsilon^{i+1}\left[\frac{\varepsilon^{i+1} + \varepsilon^{i+2} + 0.1}{4}\right] - \left[\frac{\varepsilon^{i+1} + \varepsilon^{i+2} + 0.1}{4}\right]^2$$

# Example from 2019 Exam

**3.** A government ministry is planning to invest in a facility to carry freight between the UK and the EU in the event of a no-deal Brexit. The appropriate level of investment $I$ depends on the 'state of the freight market' – a random variable $\omega$ uniformly distributed on the interval [2,3]. The gross value of the deal is given by $R(\omega) = 10\omega I - 2I^2$. A unit of investment costs $P = 2$; investment costs are paid by the government.

To explain answers, say that $\omega$ is uniformly distributed on $[\alpha, \beta]$ and $R(\omega) = A\omega I - BI^2$

a. If the government chooses $I(\omega)$ to maximise $V(\omega) = R(\omega) - PI(\omega)$ after learning $\omega$, what are the optimal $I(\omega)$ and *expected* values of $R$ and $V$? [**15 marks**]

$$I(\omega) = \frac{A\omega - P}{2B} = \frac{10\omega - 2}{4}$$

$$E(Revenue) = \frac{A^2(\beta^3 - \alpha^3)}{12 * B} - \frac{P^2(\beta - \alpha)}{4B} = \frac{236}{3} = 78.6667$$

$$E(V) = E(Revenue) - \frac{PA(\beta^2 - \alpha^2) - 2P^2(\beta - \alpha)}{4B} = 67.1667$$

b. If the government must choose $I$ before knowing what $\omega$ will be, what level of investment will they choose and what are the expected values of $R$ and $V$? [**25 marks**]

$$I = \frac{A\frac{\alpha + \beta}{2} - P}{2B} = 5.75$$

$$E(Revenue) = \frac{AI(\beta^2 - \alpha^2)}{2} - BI^2(\beta - \alpha) = 77.625$$

$$E(V) = E(Revenue) - PI = 66.125$$

Suppose instead the government contracts with a ferry company to make the investment. You may assume that the ferry company does not know $\omega$ before signing the contract, but will learn it before recommending an investment level. *The government pays for the investment* and offers the ferry company a share $\theta$ of gross value [the ferry company gets $U = \theta R$ and the government gets $V(\omega) = (1-\theta)R(\omega) - PI(\omega)$]. The ferry company has an outside option worth 1.

c. If the government is committed to following the ferry company's recommendation, what share $\theta$ will the government offer, what level of investment $I(\omega)$ will the ferry company recommend and what are the expected values of $R$, $V$ and $U$? [**30 marks**]

$$\theta = \frac{12U_0B}{A^2(\beta^3 - \alpha^3)} = 0.012632$$

$$I(\omega) = \frac{A\omega}{2B}$$

$$E(Revenue) = \frac{A^2(\beta^3 - \alpha^3)}{12 * B} = 79.16667$$

$$E(V) = (1-\theta)E(Revenue) - \frac{PA(\beta^2 - \alpha^2)}{4B} = 65.66667$$

$$U = 1$$

Finally, suppose the ferry company observes $\omega$ and then makes a recommendation as to the right level of investment, which the government is free to reject or modify.

d. How would your answer differ from that found in part c and how would it be affected by a change in the price $P$ of investment? [**30 marks**]

This is a cheap-talk problem; the government would set a finite number if investment levels that the ferry company could choose between. P establishes the level of 'bias' – the higher it is, the less the government wants to invest, the smaller the maximum number of investment levels and the less efficient the outcome.

# Example from 2020 Exam

**3.** A The UK government has embarked on an ambitious plan to build a new high-speed rail network. The work is controlled by a dedicated firm set up to acquire land, commission construction and procure rolling stock. The firm has superior information regarding the actual costs and capabilities of the project.

First, suppose that the firm submits a report to the government detailing the current cost and revenue projections (simplified to $reported\ NPV = V$), along with a check equal to $V$. The true value may be higher or lower; if it is higher, the firm keeps the surplus. When the contract was signed it was believed that $V$ was uniformly distributed on an interval $[V_0, V_1]$. The government can either accept $V$ or call in an auditor. This will incur an audit cost of $C^{audit}$ (paid by the government). Both the firm and the government are risk neutral. The firm has an outside offer worth $W_0$; if it does not expect at least this much, it will not sign the contract.

a. If the government must either accept or reject the report, what will its optimal strategy look like (i.e. which reports would it threaten to reject to maximise the expected (net of audit cost) payment by the firm)? Justify your answer. [**10 marks**]

To optimal strategy will be to set a 'trigger' level V, reject any report below this level and retain all NPV in the case of an audit. The reason why this is optimal is standard (see slides); the trigger level is set to balance the losses at higher values of V against the increased cost of audit. The government retains all income in the case of an audit because it could replace any contract that allowed the firm residual profit in the case of audit with an equivalent one that retained all income but audited less frequently. The government's problem is to choose $V$ to maximise

$$U_{gov}(V) = \int_{V_0}^{V} (w - C)\, dw + V \int_{V}^{V_1} dw \;\; s.t.$$

$$U_{firm}(V) = \int_{V}^{V_1} (w - V) dw \geq W_0$$

b. If $V_0 = 10, V_1 = 100, C^{audit} = 30, and\ W_0 = 0$, what are the expected audit cost, expected government receipts (gross) and expected profit for the firm? [**12 marks**]

The interior optimum (where the firm's participation constraint does not bind) is $V = V_1 - C$. The requested quantities are:

$$EC^{audit} = (V_1 - V_0 - C)C = 1800$$

$$EU_{gov} = (V - V_0)V_0 + \frac{(V - V_0)^2}{2} + (V_1 - V)V = 2700$$

$$EU_{firm} = \frac{(V_1 - V)^2}{2} = 450$$

c. What would the trigger value and associated expected net payoff to the government be if $W_0 = 1250$? [**8 marks**]

In this case the government would have to set V at $V = V_1 - \sqrt{2W_0} = 50$ and expected payoff drops to 2500.

d. Would the government prefer a contract that called the auditors in with a probability that decreased with the firm's reported value? (justify your answer) [**5 marks**]

Yes; ideally it would audit with probability 1 at $V = V_0$ and drop to 0 at $V = V_1$; this would economise on audit cost and give the government a greater share of profits above the old trigger value. The audit cost savings mean that both the firm and the government are better off.

Now suppose instead that the firm assesses the state of demand and cost for the project, and reports to the government the amount $C^B$ that the project should cost to build. The government will respond to this report by providing a budget $B$ for the project. The 'real' cost is given by a random variable $C$ that is uniformly distributed on an interval $[C_0, C_1]$. The value to society and to the firm depend on the true cost and the budget as follows

$$V^{soc} = V^* - (B - C)^2$$
$$V^{firm} = W_0 - (B - C - \beta)^2$$

Where $\beta$ represents the 'inaccuracy' of the government's post-project audits; the more the firm asks for compared to the actual cost, the greater the risk of a demand for repayment and a penalty.

e. Under the optimal contract, how would the government respond to the firm's different possible budget requests? How would the gap between maximum social value (if the government knew $C$ exactly in advance) and the amount achieved by the optimal contract to change if $\beta$ were to decrease (post-project audits became more accurate)? (You need not provide a mathematical solution, but should clearly justify your answer) [**15 marks**]

This si a standard cheap talk problem (lightly disguised). The discussion should note that the optimum involves a finite set of budget levels the government would approve; each is the midpoint of an interval of possible reports, set so that a firm with a true value at the boundary of the interval is just indifferent between the next highest and next lowest 'acceptable' values. For any value of $\beta$ there are a set of equilibria with one interval (babbling equilibrium = least efficient) up to a maximum number n*. The lower the value of $\beta$ (the more accurate the post-project audits), the greater this maximum number and the more efficient the 'best' feasible equilibrium.