Optimal Contracting with a Single Privately-Informed Party

Dr. Margaret Meyer

2021

Optimal contracting with a single privately-informed party

This class of hidden-information model is applicable to study design of tax policy, regulatory policy, price discrimination schemes (with respect to quality or quantity), and incentive schemes

Here, we'll study a principal-agent model with hidden information, where the principal P is an uninformed buyer of labor and the agent A is a privately-informed seller of labor.

- A's private information: his cost of producing output
- Good A's have lower marginal and total costs of producing output, and P wants both to pay these agents less and to have them produce more.
- How should P design an incentive scheme?
- P can induce self-selection of different A types by offering a menu of different (output, payment) combinations.
- Applicable to i) performance-related pay for employees or ii) incentive contract for independent contractors

Principal-agent model with hidden information

- Agent's preferences: $U(t, x, s) = t - C(x, s)$, where x = output, t = payment, and s = agent's privately-observed type, representing a characteristic of A's preferences or technology
  - Assume: $C_x > 0$; $C_{xx} \geq 0$; $C(0, s) = 0 \ \forall s$; $C_{xs} > 0$
  - $C_{xs} > 0$ implies that slope of indifference curve is increasing in s: $\frac{dt}{dx}\big|_{U=\bar{U}} = C_x$, $\uparrow$ in s
  - utility from not signing contract (reservation utility) is 0
- Principal's preferences: $V(t, x) = R(x) - t$
  - Assume: $R_x > 0$, $R_{xx} \leq 0$
  - Note that P's payoff does not depend on s directly: A's private info. is about cost of producing output but not about value of output to P.

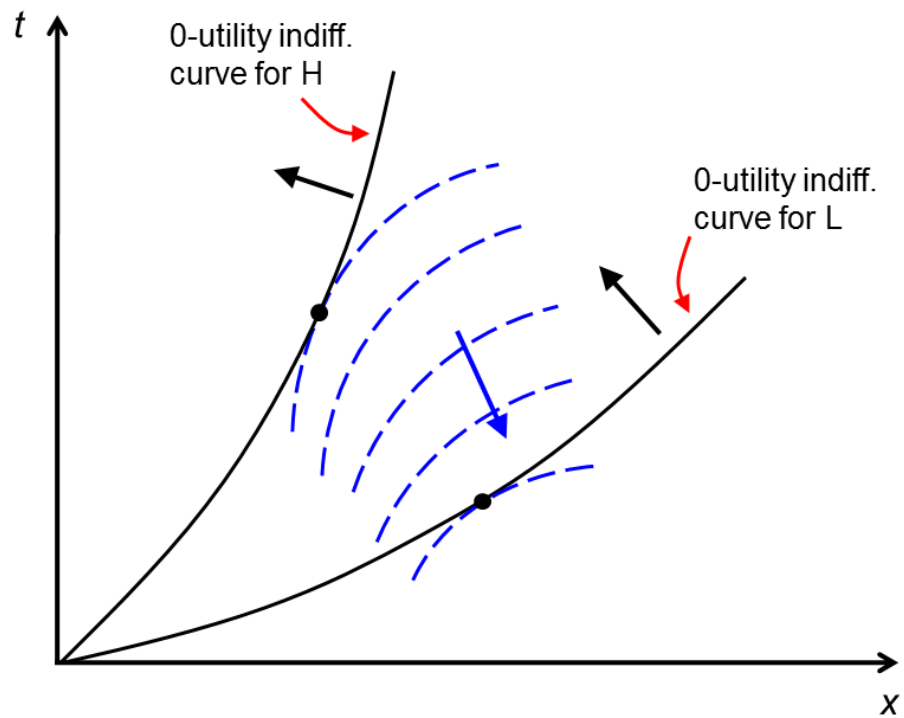

Principal-agent model with hidden information

Timing of the contracting game:

1. A observes s privately; P believes s is drawn from probability distribution F
2. P offers A a contract - terms are binding on both parties if accepted
3. A accepts or rejects contract
4. A chooses output; P makes payment to A; only A observes total costs incurred

P's objective in designing the contract is to maximize expected profit. In the absence of hidden information, P's optimal contract would both

1. achieve the allocatively efficient output level x from each type s of A: $R'(x) = C_x(x, s)$.
2. give each type of A utility equal to his reservation utility, i.e. avoid leaving rents to A.

When A is privately informed before signing the contract, P faces a tradeoff between efficient production and rent extraction.

Delegation approach to the optimal contracting problem

- P offers A a contract described by a function $t = T(x)$
- Given s and contract $t = T(x)$, A chooses x to maximize $T(x) - C(x, s) \Longrightarrow$ A's strategy is denoted $x = X(s)$.
- *Example*: Suppose s takes on only 2 possible values, L, H, with $L < H$. Then function $T(x)$ needs to be defined only at 2 possible values of x: $T(x_1) = t_1$, $T(x_2) = t_2$.

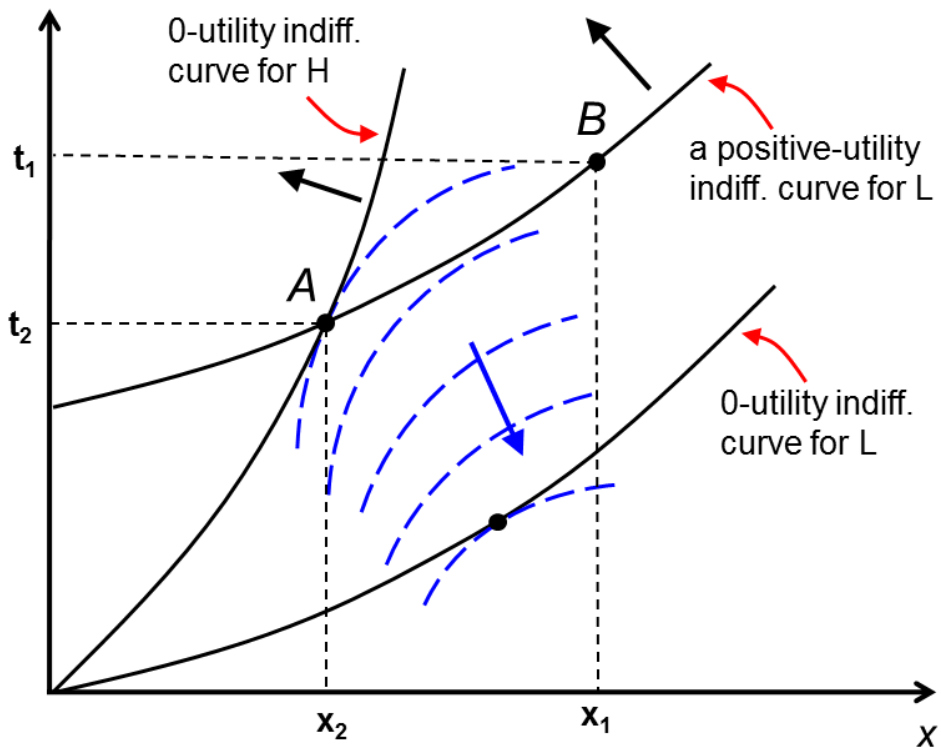

- H-type strictly prefers A to B; L-type indifferent between A and B

Revelation approach to the optimal contracting problem

- The "revelation approach" is an alternative approach. It is more useful analytically, while the delegation approach is more intuitive.
- Revelation approach: P offers A a direct revelation mechanism (DRM).
- A DRM commits P to rules specifying x and t as functions of the report, $\hat{s}$, that A makes about his type s: $\{x(\hat{s}), t(\hat{s})\}$.
- Given s and the DRM, A chooses $\hat{s}$ to maximize $t(\hat{s}) - C(x(\hat{s}), s) \Longrightarrow$ A's strategy is denoted $\hat{s} = S(s)$.
- **Definition:** A *truthful DRM* is a DRM for which truthful reporting (i.e. $S(s) = s$) is an optimal strategy for all types s of A.
- In the two-type example above, contract (A, B) can be implemented as a truthful DRM:
  - $x(\hat{s} = H) = x_2, \;\; t(\hat{s} = H) = t_2$
  - $x(\hat{s} = L) = x_1, \;\; t(\hat{s} = L) = t_1$

Revelation Principle (1)

Revelation Principle: Given any contract $t = T(x)$ or any DRM $\{x(\hat{s}), t(\hat{s})\}$, there exists a truthful DRM which, for every type of agent, results in the same (output, payment) pair being chosen as under the original contract or DRM.

Implication: In designing optimal contracts/mechanisms, we can restrict attention to truthful DRM's; if an outcome is not achievable by a truthful DRM, it is not achievable at all.

Proof.

- Given an arbitrary contract $t = T(x)$ that induces A to choose $x = X(s)$, define the outcome function $(x^*(s), t^*(s)) \equiv (X(s), T(X(s)))$.
- Given an arbitrary DRM $\{x(\hat{s}), t(\hat{s})\}$ that induces A to choose $\hat{s} = S(s)$, define the outcome function $(x^*(s), t^*(s)) \equiv (x(S(s)), t(S(s)))$.
- Key observation: The DRM $\{x^*(\hat{s}), t^*(\hat{s})\}$ is the truthful DRM we seek.

Revelation Principle (2)

Proof (cont.)

Need to show

- $\{x^*(\hat{s}), t^*(\hat{s})\}$ is a truthful DRM:
  - If $x = X(s)$ was optimal for A given $t = T(x)$, then optimal feasible (x, t) pair for type s was $(X(s), T(X(s))) = (x^*(s), t^*(s)))$ by definition. This implies that, given DRM $\{x^*(\hat{s}), t^*(\hat{s})\}$, reporting $\hat{s} = s$ is optimal for A.
  - If $\hat{s} = S(s)$ was optimal for A given $\{x(\hat{s}), t(\hat{s})\}$, then optimal feasible (x, t) pair for type s was $(x(S(s)), t(S(s))) = (x^*(s), t^*(s))$ by definition. This implies that, given DRM $\{x^*(\hat{s}), t^*(\hat{s})\}$, reporting $\hat{s} = s$ is optimal for A.
- Outcome under $\{x^*(\hat{s}), t^*(\hat{s})\}$ is same as under original contract/ DRM:
  - When $\hat{s} = s$ under DRM $\{x^*(\hat{s}), t^*(\hat{s})\}$, outcome is $(x^*(s), t^*(s))$, which is same as under original contract/ DRM.

□

Revelation Principle (3)

Revelation Principle: Given any contract $t = T(x)$ or any DRM $\{x(\hat{s}), t(\hat{s})\}$, there exists a truthful DRM which, for every type of agent, results in the same (output, payment) pair being chosen as under the original contract or DRM.

Interpretation: Suppose the original DRM is not a truthful DRM, i.e. for some s, A strategically manipulates his report to P and chooses $\hat{s} \neq s$. By committing to the DRM $\{x^*(\hat{s}), t^*(\hat{s})\}$, P induces A to report truthfully because P is in effect committing to do A's manipulation for him.

- The assumption that P can **commit** to contract/DRM is crucial.

The Revelation Principle holds in any optimal contracting model with hidden information where the uninformed party moves first and can commit to a contract/mechanism.

- The Revelation Principle does *not* imply that the uninformed party can costlessly extract the privately-informed party's information.
- It is merely a *tool* for analyzing the optimal mechanism under hidden information.

Revelation Principle: intuitive example

Regulator (P) chooses a policy/mechanism to regulate a monopolist (A)

- Monopolist's cost function is $C(q) = K + cq$
  - Constant marginal cost c is monopolist's private info.
  - K is a commonly known fixed cost
- Market demand is commonly known, hence so is monopoly price *function* $p = p^m(c)$
- Policy/mechanism yields a price p and lump-sum subsidy $T \ \forall c$.

Example of a (naive) policy (delegation approach): $T(p) = K$

- Monopolist chooses a price p; subsidy determined by $T(p) = K$.
- What price will a monopolist with marginal cost c choose?
- $p = p^m(c)$, the monopoly price given privately-observed c

Example of a (naive) DRM: $p(\hat{c}) = \hat{c}$ and $T(\hat{c}) = K$

- Monopolist chooses a report $\hat{c}$ about c; (p, T) determined by DRM.
- What $\hat{c}$ will a monopolist with marginal cost c choose?
- $\hat{c} = p^m(c)$, since this yields the monopoly price given c, and report has no effect on T

This policy and DRM yield same outcome: $p = p^m(c)$ and $T = K \ \forall c$.

Revelation Principle: intuitive example

Example of a (naive) policy (delegation approach): $T(p) = K$

Example of a (naive) DRM: $p(\hat{c}) = \hat{c}$ and $T(\hat{c}) = K$

This policy and DRM yield same outcome: $p = p^m(c)$ and $T = K \ \forall c$.

The following truthful DRM yields same outcome as the naive policy and the naive DRM:

$$p(\hat{c}) = p^m(\hat{c}) \quad \text{and} \quad T(\hat{c}) = K$$

- Since report $\hat{c}$ has no effect on T, monopolist chooses report, given c, to achieve monopoly price $p^m(c)$; hence reporting $\hat{c} = c$ is optimal
- This DRM is truthful because it commits the regulator to distort price, given report, in the same way monopolist would distort price under original naive policy or DRM.

Using the Revelation Prin. to solve for the optimal DRM in the principal-agent model

- Suppose there are only two types of A: $s \in \{H, L\}$, with $H > L$.
- Let $Pr(s = H) = q$ and $Pr(s = L) = 1 - q$.
- Given the revelation principle, P need only optimize over truthful DRM's.
  - With only two types of A, a truthful DRM is a pair $\{(x_H, t_H), (x_L, t_L)\}$ that induces H to choose $\{(x_H, t_H)\}$ and L to choose $\{(x_L, t_L)\}$
- Hence P solves

$$\begin{aligned} \max_{(x_H,t_H),(x_L,t_L)} \quad & q(R(x_H) - t_H) + (1-q)(R(x_L) - t_L) \\ \text{s.t.} \quad & t_H - C(x_H, H) \geq 0 && (IR_H) \\ & t_L - C(x_L, L) \geq 0 && (IR_L) \\ & t_H - C(x_H, H) \geq t_L - C(x_L, H) && (IC_H) \\ & t_L - C(x_L, L) \geq t_H - C(x_H, L) && (IC_L) \end{aligned}$$

- IR_s denotes the "individual rationality" (i.e. participation) constraint for type s, and IC_s the "incentive compatibility" (i.e. truthtelling) constraint for type s.

Benchmark case: optimal contract under full information

If P could observe s before contracting, P would choose (x_s, t_s) to solve

$$\max_{(x_s,t_s)} R(x_s) - t_s \quad \text{s.t.} \quad t_s - C(x_s, s) \geq 0 \quad (IR_s)$$

Optimal contract under full information, $\{(x_H^*, t_H^*), (x_L^*, t_L^*)\}$, satisfies, for each type s, $R'(x_s^*) = C_x(x_s^*, s)$ (output is allocatively efficient) and $t_s^* = C(x_s^*, s)$ (A earns no rent).

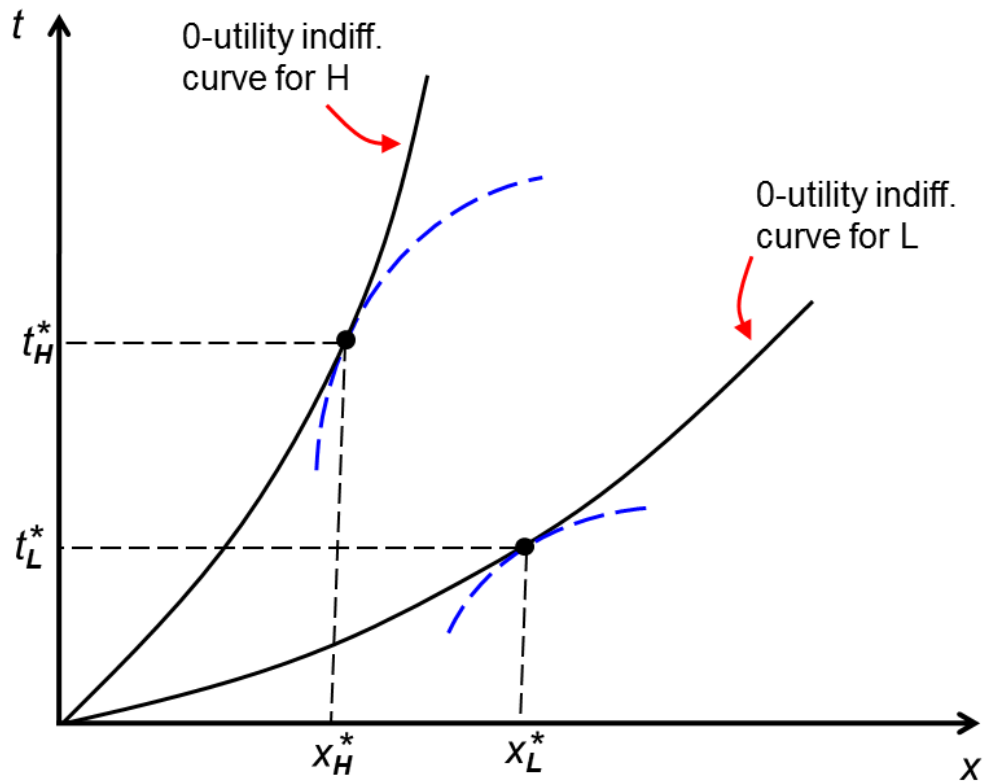

Is optimal outcome under full info achievable under asymmetric info?

Now suppose P **cannot** observe s before contracting. The Revelation Principle tells us that if the outcome under full info is achievable under asymmetric info, then it is achievable via a truthful DRM.

But the DRM $\{(x_H^*, t_H^*), (x_L^*, t_L^*)\}$ is **not** truthful, because type L strictly prefers (x_H^*, t_H^*) to (x_L^*, t_L^*), i.e. IC_L is violated.

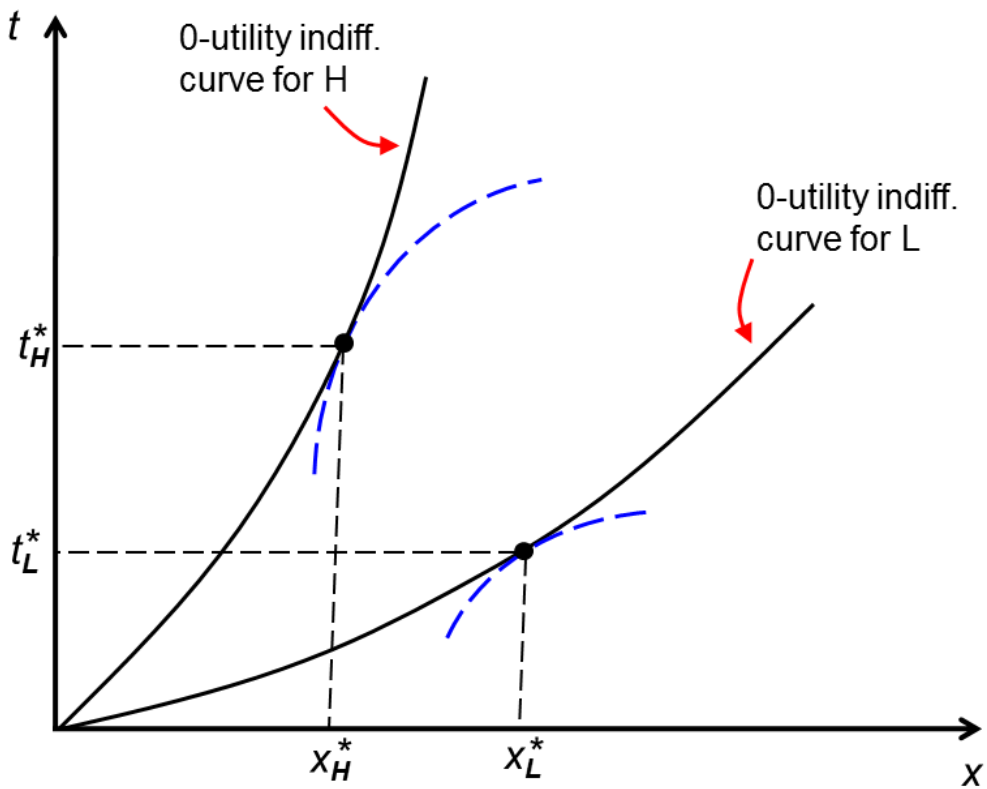

Using the Revelation Prin. to solve for the optimal DRM

- Given the revelation principle, P solves

$$\begin{aligned} \max_{(x_H,t_H),(x_L,t_L)} \quad & q(R(x_H)-t_H)+(1-q)(R(x_L)-t_L) \\ \text{s.t.} \quad & t_H - C(x_H,H) \geq 0 && (IR_H) \\ & t_L - C(x_L,L) \geq 0 && (IR_L) \\ & t_H - C(x_H,H) \geq t_L - C(x_L,H) && (IC_H) \\ & t_L - C(x_L,L) \geq t_H - C(x_H,L) && (IC_L) \end{aligned}$$

- Denote solution to this problem by $(\hat{x}_H,\hat{t}_H),(\hat{x}_L,\hat{t}_L)$.
- *Claim:* This problem can be simplified as follows:
  1. We can ignore IR_L if the other 3 constraints are satisfied. (*Proof:* $t_L - C(x_L,L) \geq t_H - C(x_H,L) \geq t_H - C(x_H,H) \geq 0$.)
  2. IR_H will be binding at the optimum.
  3. IC_L will be binding at the optimum, but IC_H will be slack.
  4. $\hat{x}_L = x_L^*$: L-type's output will be allocatively efficient.
- Last 3 claims can be proved by contradiction.

Special case of the model

It is helpful, initially, to make some simplifying assumptions:

- $C(x, s) = sx$
- $R(x) = px$ for $x \in [0, \bar{x}]$, $R(x) = p\bar{x}$ for $x > \bar{x}$
- $p > H > L$
- $Pr(s = H) = q$ and $Pr(s = L) = 1 - q$

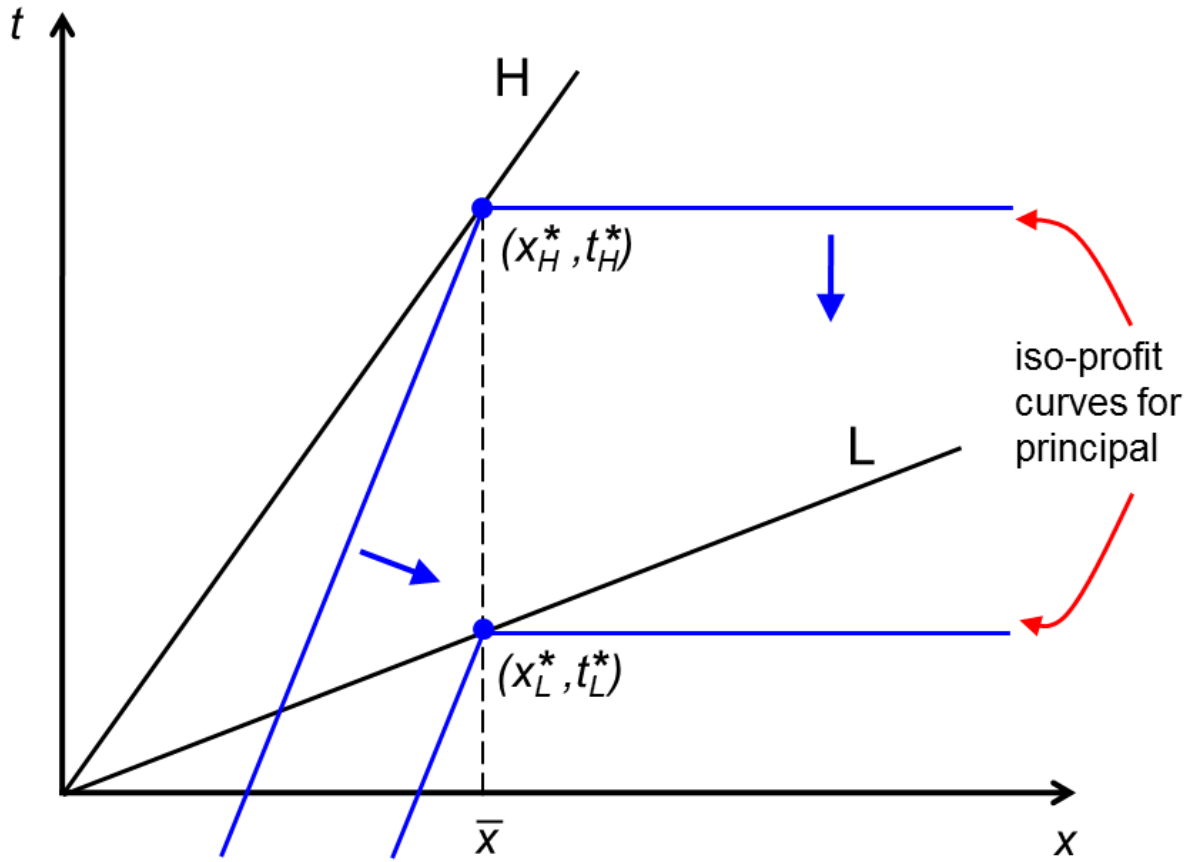

Optimal contract under asymmetric information

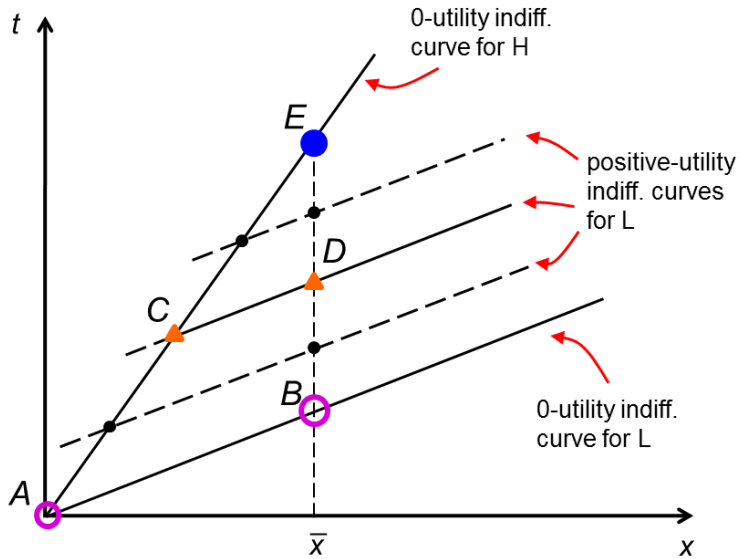

$$\max_{(x_H,t_H),(x_L,t_L)} q(R(x_H)-t_H)+(1-q)(R(x_L)-t_L)$$

s.t. $t_H - Hx_H = 0 \quad (IR_H)$ and $t_L - Lx_L = t_H - Lx_H \quad (IC_L)$

- IR_H binding $\Longrightarrow$ $(\hat{x}_H, \hat{t}_H)$ lies along segment AE
- $\hat{x}_L = x_L^* \Longrightarrow (\hat{x}_L, \hat{t}_L)$ lies along segment BE
- IC_L binding $\Longrightarrow$ $(\hat{x}_H, \hat{t}_H)$ and $(\hat{x}_L, \hat{t}_L)$ lie on same indifference curve of L
- So optimal contract is in family with "representative member" (C, D).
  - Given contract (C, D), H-type chooses C and L-type chooses D.
  - H-type earns no rent but produces inefficiently low output; L-type earns positive rent but produces efficient output.
  - For H-type , x is *lower* but t/x is *higher* than for L-type .

Optimal contract under asymmetric information

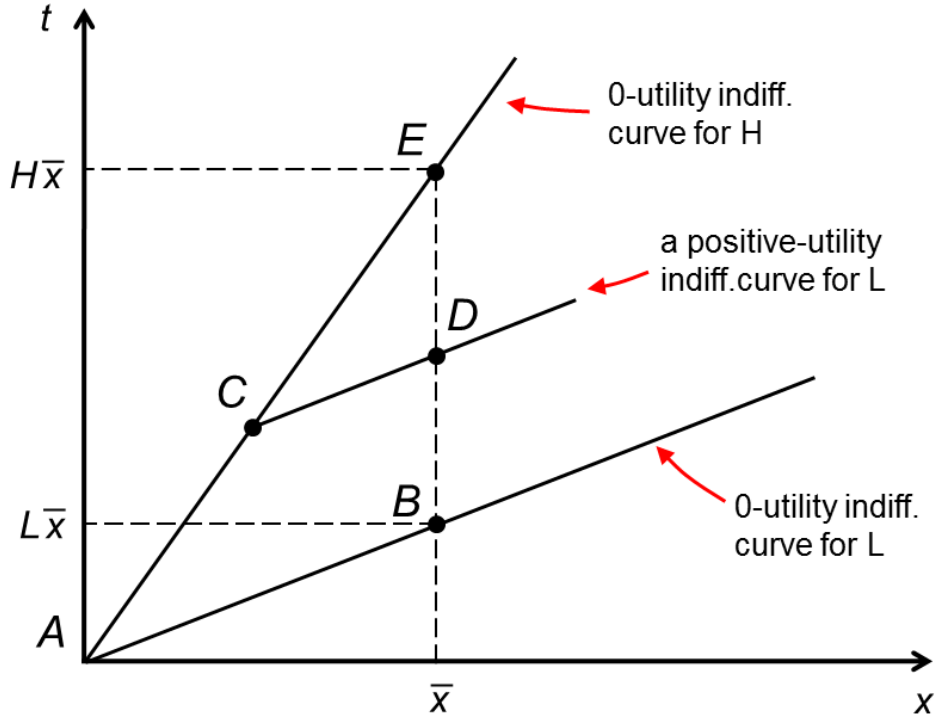

Within the family of contracts with representative member (C, D):

- At one extreme is (A, B), the most extreme *separating* contract.
- At other extreme is (E, E), which is a *pooling* contract.

Given linearity assns. on P's revenue and A's cost function, P's exp. profit from (C, D) is convex combination of that from (A, B) and (E, E). Hence P's optimal contract is a "corner solution".

- If $E\pi(A, B) > E\pi(E, E)$, optimal contract is (A, B).
- If $E\pi(A, B) < E\pi(E, E)$, optimal contract is (E, E).

Optimal contract under asymmetric information

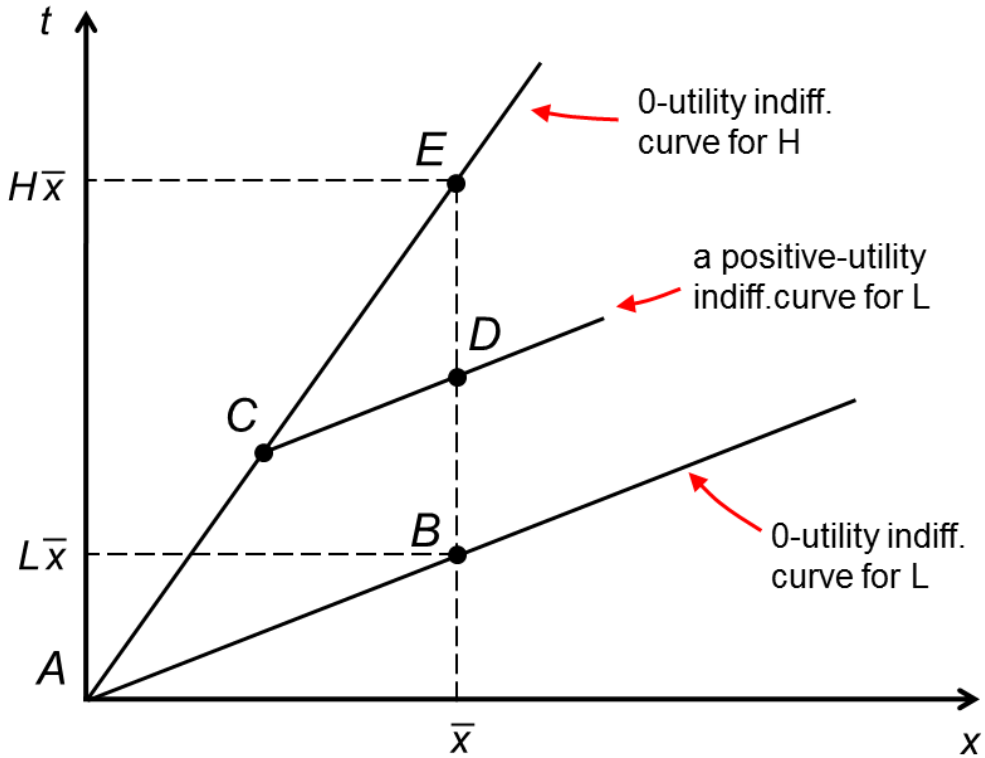

$E\pi(A,B) = (1-q)(p-L)\bar{x}$ and $E\pi(E,E) = (p-H)\bar{x}$.

Hence optimal contract is (A,B) (extreme separating contract) if $1-q > \frac{p-H}{p-L}$ and (E,E) (pooling contract) if $1-q < \frac{p-H}{p-L}$.

- As L-type more likely ($q \downarrow$), P more likely to prefer saving rents and risking 0 output (from H-type), hence more likely to prefer (A,B).
- As $p \uparrow$, opportunity cost to P of getting 0 output from H-type rises, so P more likely to prefer (E,E).

Model with more general cost and revenue functions

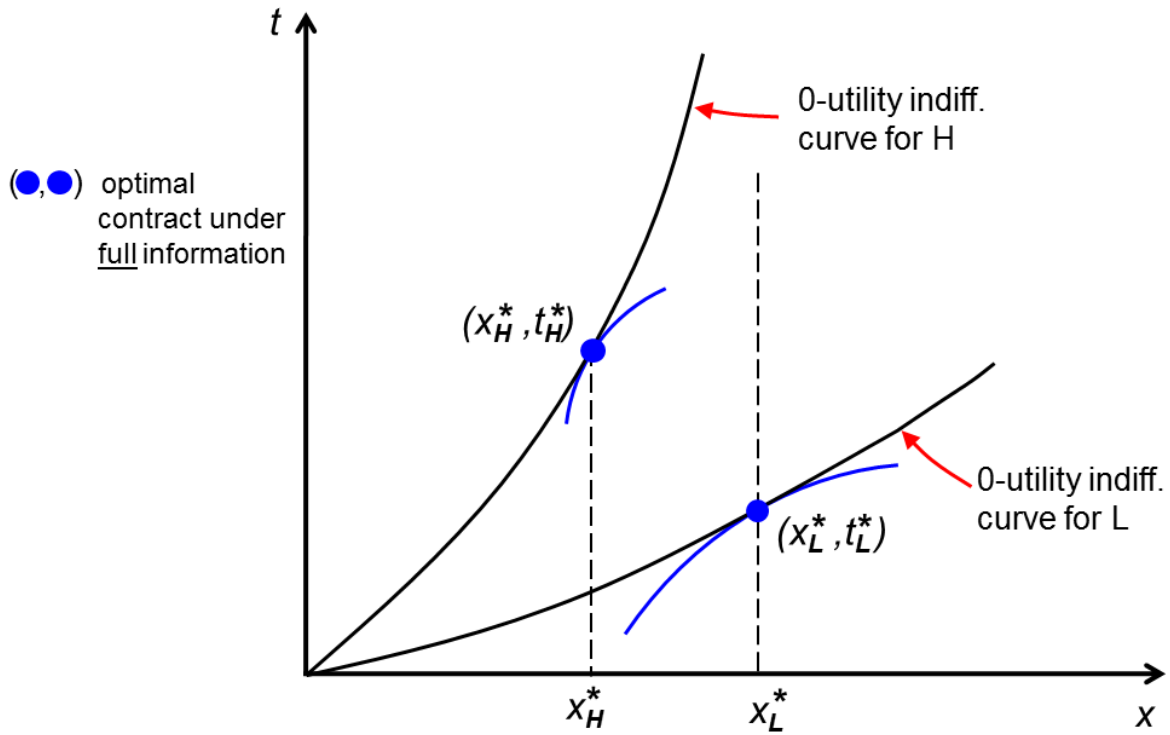

- Full-info optimal contract (output efficient for both types; neither type earns rents) violates IC_L.
- As in the linear model, the only contraints which bite are IC_L and IR_H.
- Optimal contract: H's output is inefficiently low, but L's is efficient.
- Optimal contract: L earns rent, but H does not.
- In general model, P's optimal contract is an *interior* solution to tradeoff between goals of i) allocative efficiency and ii) minimizing rents to A.

Model with more general cost and revenue functions

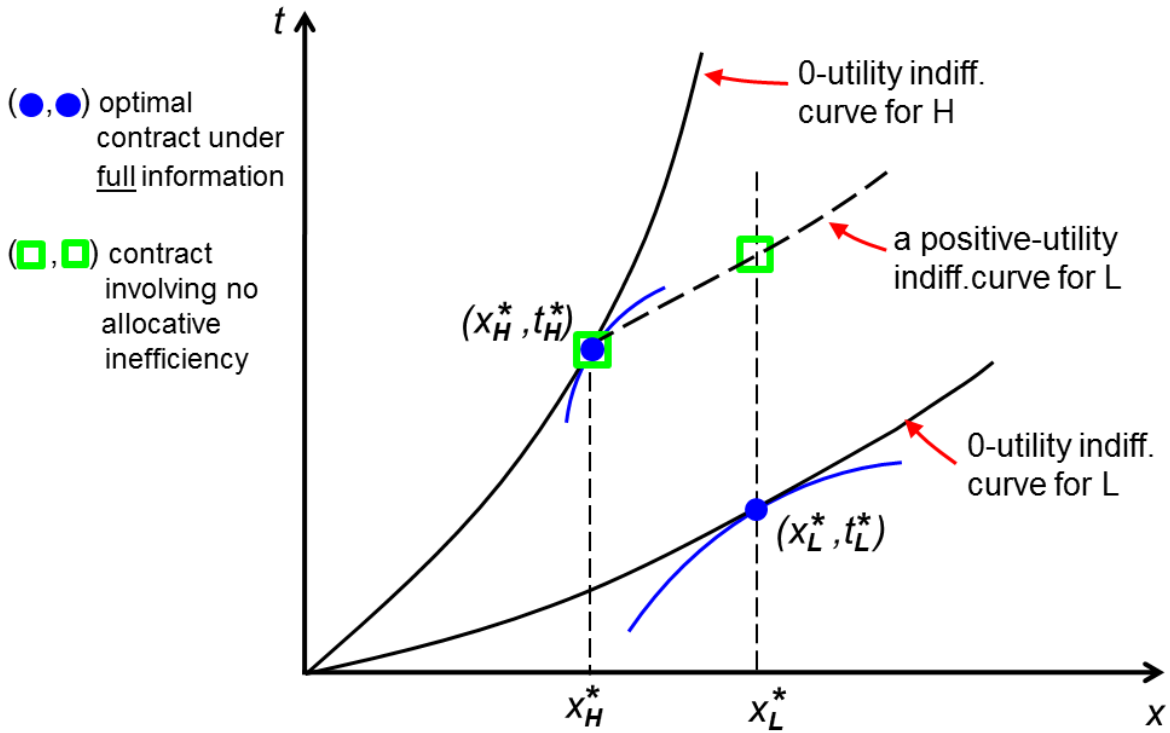

- Full-info optimal contract (output efficient for both types; neither type earns rents) violates IC_L.
- As in the special case, the only contraints which bite are IC_L and IR_H.
- Optimal contract: H's output is inefficiently low, but L's is efficient.
- Optimal contract: L earns rent, but H does not.
- In general model, P's optimal contract is an *interior* solution to tradeoff between goals of i) allocative efficiency and ii) minimizing rents to A.

Model with more general cost and revenue functions

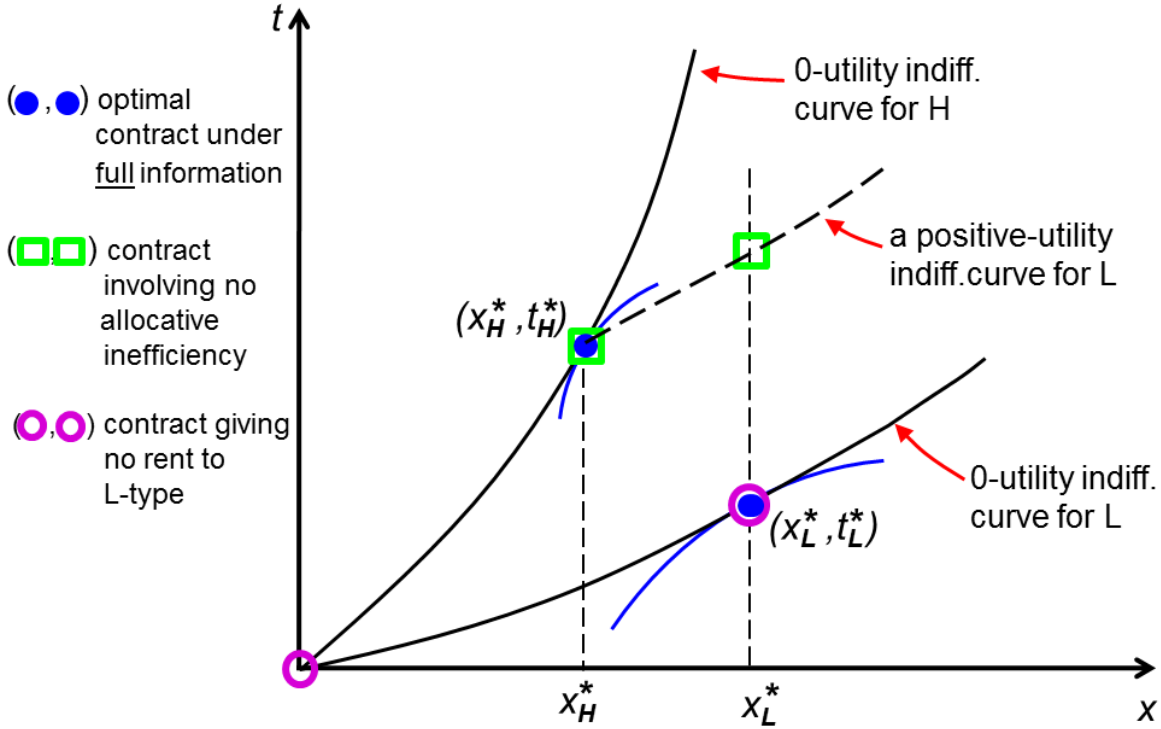

- Full-info optimal contract (output efficient for both types; neither type earns rents) violates IC_L.
- As in the special case, the only contraints which bite are IC_L and IR_H.
- Optimal contract: H's output is inefficiently low, but L's is efficient.
- Optimal contract: L earns rent, but H does not.
- In general model, P's optimal contract is an *interior* solution to tradeoff between goals of i) allocative efficiency and ii) minimizing rents to A.

Model with more general cost and revenue functions

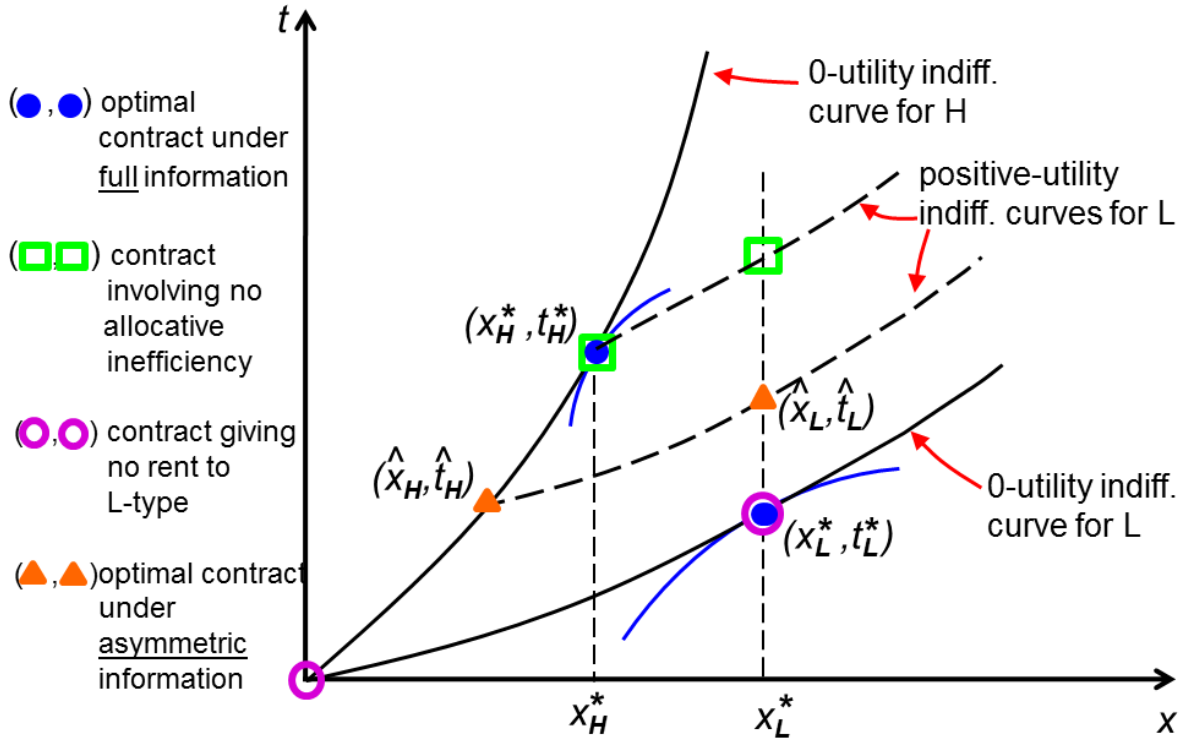

- Full-info optimal contract (output efficient for both types; neither type earns rents) violates IC_L.
- As in the special case, the only contraints which bite are IC_L and IR_H.
- Optimal contract: H's output is inefficiently low, but L's is efficient.
- Optimal contract: L earns rent, but H does not.
- In general model, P's optimal contract is an *interior* solution to tradeoff between goals of i) allocative efficiency and ii) minimizing rents to A.

What would a social planner do in the P-A model?

P's objective is to maximize expected profit. In the absence of hidden information, P's optimal contract would both

1. achieve the allocatively efficient output level x_s^* from each type s of A: $R'(x_s^*) = C_x(x_s^*, s)$.
2. give each type of A utility equal to his reservation utility, i.e. avoid leaving rents to A.

When A is privately informed before signing the contract, P faces a tradeoff between efficient production and rent extraction.

What would a social planner do in this environment?

- Suppose the planner cares about allocative efficiency but does not have explicit distributional objectives.
- So planner cares about payments from P to A only because of their effect on participation and self-selection decisions by A types.

What would a social planner do?

Planner would solve

$$\max_{(x_H,t_H),(x_L,t_L)} q(R(x_H) - C(x_H, H)) + (1-q)(R(x_L) - C(x_L, L))$$

$$\begin{aligned} \text{s.t.} \quad & t_H - C(x_H, H) \geq 0 && (IR_H) \\ & t_L - C(x_L, L) \geq 0 && (IR_L) \\ & t_H - C(x_H, H) \geq t_L - C(x_L, H) && (IC_H) \\ & t_L - C(x_L, L) \geq t_H - C(x_H, L) && (IC_L) \end{aligned}$$

- Planner wants to achieve the allocatively efficient output level x_s^* from each type s of A: $R'(x_s^*) = C_x(x_s^*, s)$.
- Planner will use transfers to satisfy the (IR) and (IC) constraints.

What would a social planner do?

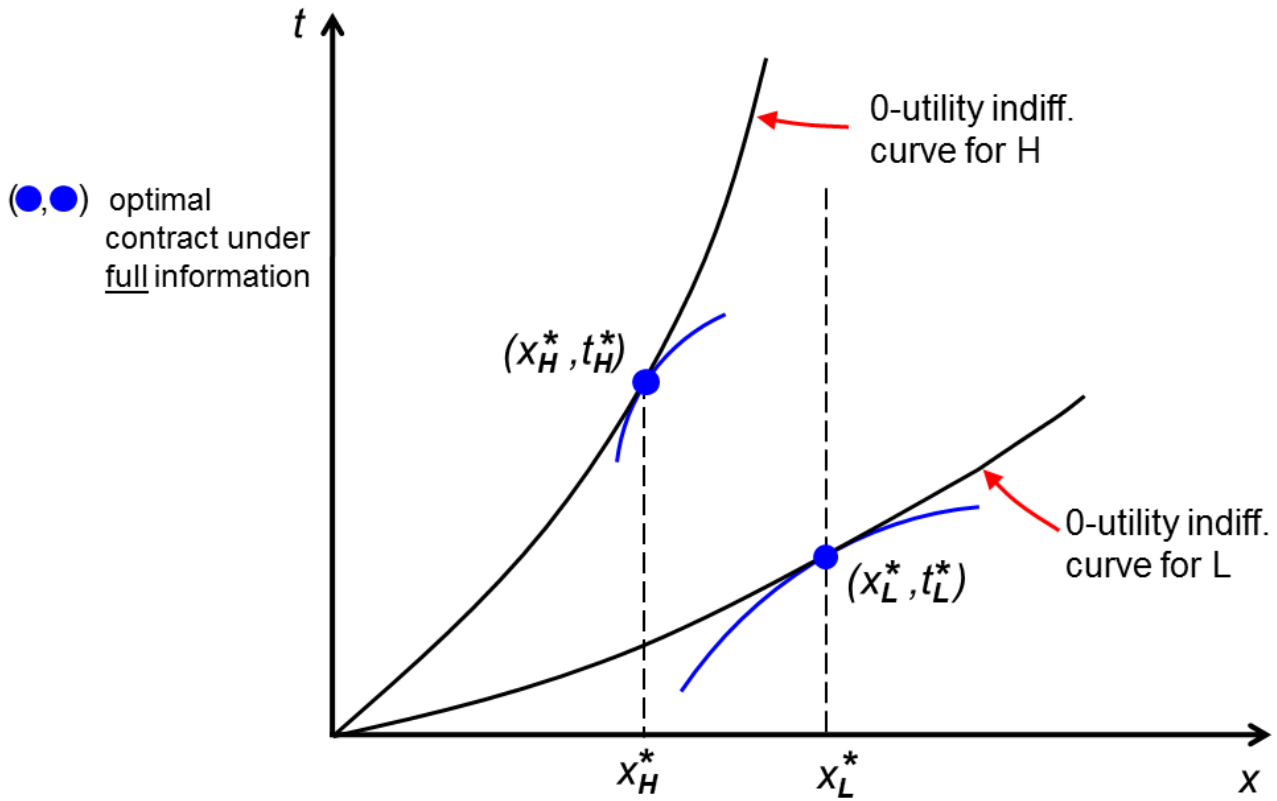

What would a social planner do?

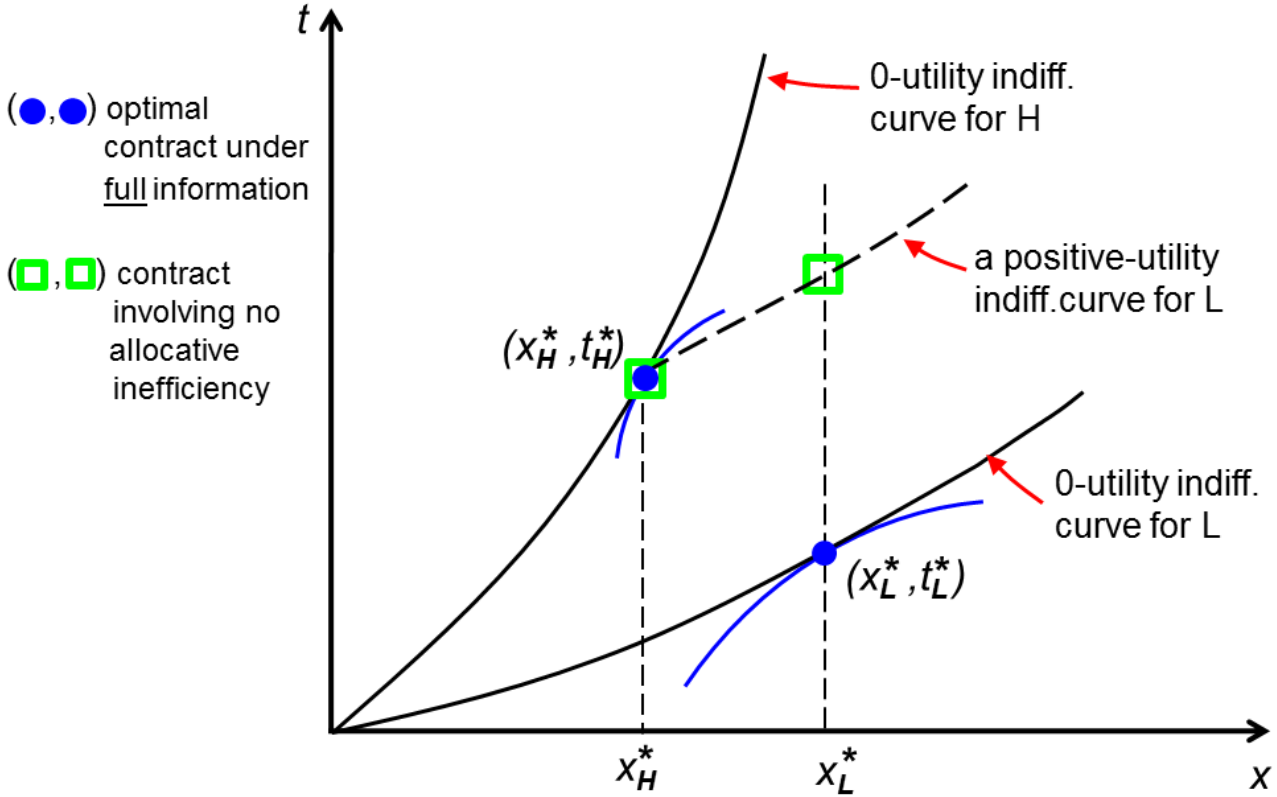

- The green squares are an optimal contract for the planner: H-type produces x_H^* and L-type produces x_L^*.
- The planner's optimum is not unique here: keeping x_L at x_L^*, t_L could be raised somewhat while still satisfying all constraints.

Regulation of a privately-informed monopolist

Same qualitative features as in principal-agent model emerge in optimal contract/mechanism in other settings with hidden information.

Regulation of a privately informed monopolist:

- monopolist's (M's) cost function $C(q) = cq + K$, where marginal cost c is privately known by M and fixed cost K is commonly known
- market demand function $Q(p)$, commonly known
- regulator's (R's) objective function $= CS + \alpha\pi$, where CS is consumer surplus, π is M's profit, and $\alpha \in [0, 1)$. $\alpha < 1 \implies$ R wants to minimize transfers from consumers to M, e.g. because taxes are distortionary.
- M will refuse to operate if profit under regulation negative, so R faces an IR constraint for M: $\pi \geq 0$
- regulatory instruments: price per unit p and lump-sum subsidy T

Full-info benchmark: If R knows c as well as K, then for all types of M (all c), R sets $p = c$ and $T = K$.

- **no tradeoff** between R's allocative and distributional objectives
- allocative efficiency achieved since price $=$ marginal cost
- distributional objective achieved, since costly transfers from consumers to M are minimized—no type of M earns rents

Regulation of a privately-informed monopolist

Hidden information: M privately observes c, then R commits to a regulatory policy, then M chooses whether or not to operate.

- Regulatory policy (mechanism) consists of a menu of (price, subsidy) pairs, $\{p(\hat{c}), T(\hat{c})\}$, and M's of different type (c) self-select by choosing the most profitable pair (formally, by choosing what $\hat{c}$ to report).
- The revelation principle implies that the optimal policy can be described as a truthful DRM.

Suppose there are only two types of M: $c \in \{c_L, c_H\}$, with $c_L < c_H$.

- Full-info optimal policy $\{(c_H, K), (c_L, K)\}$ is not achievable: low-cost M would report high cost (i.e. choose (c_H, K)).
- Optimal regulatory policy represents **tradeoff** between R's allocative and distributional objectives:
  - allocative: $p_H > c_H$ (output inefficiently low for high-cost M); $p_L = c_L$ (output efficient for low-cost M)
  - distributional: low-cost M earns rent; high-cost M does not

Regulation of a privately-informed monopolist

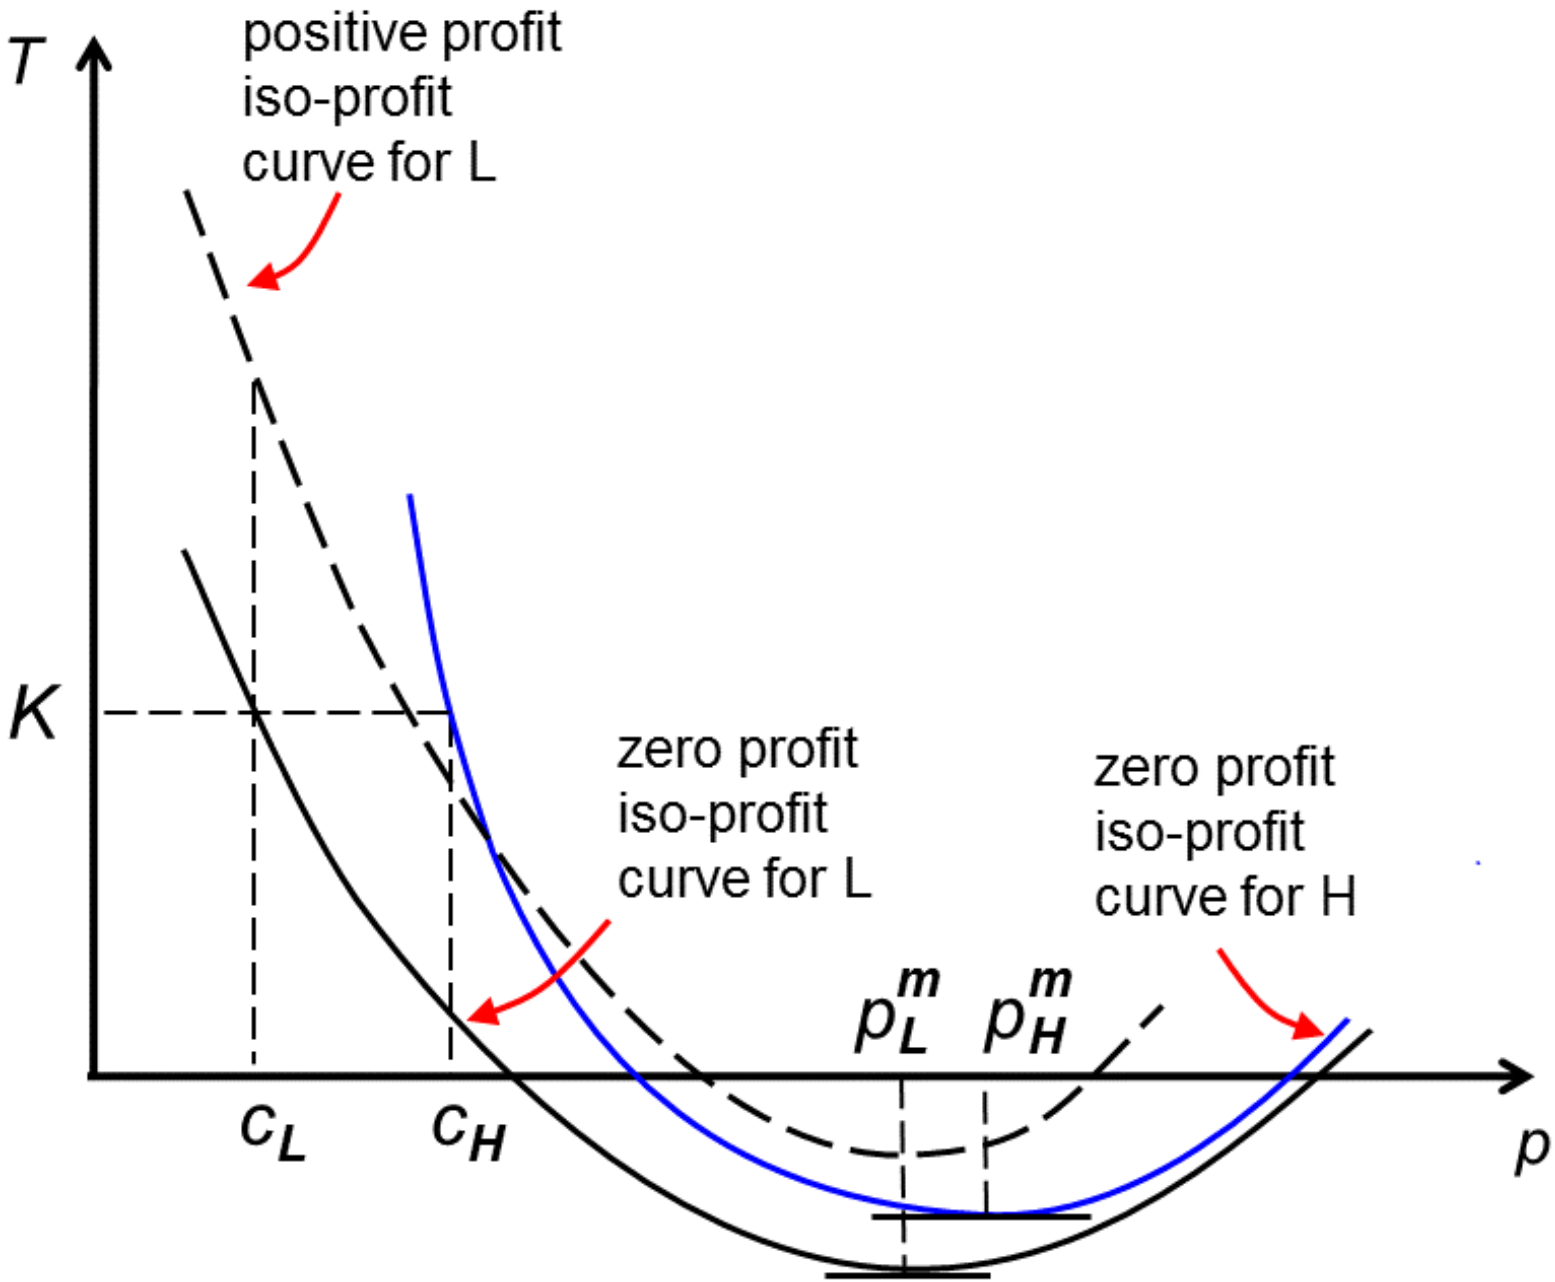

Regulation of a privately-informed monopolist

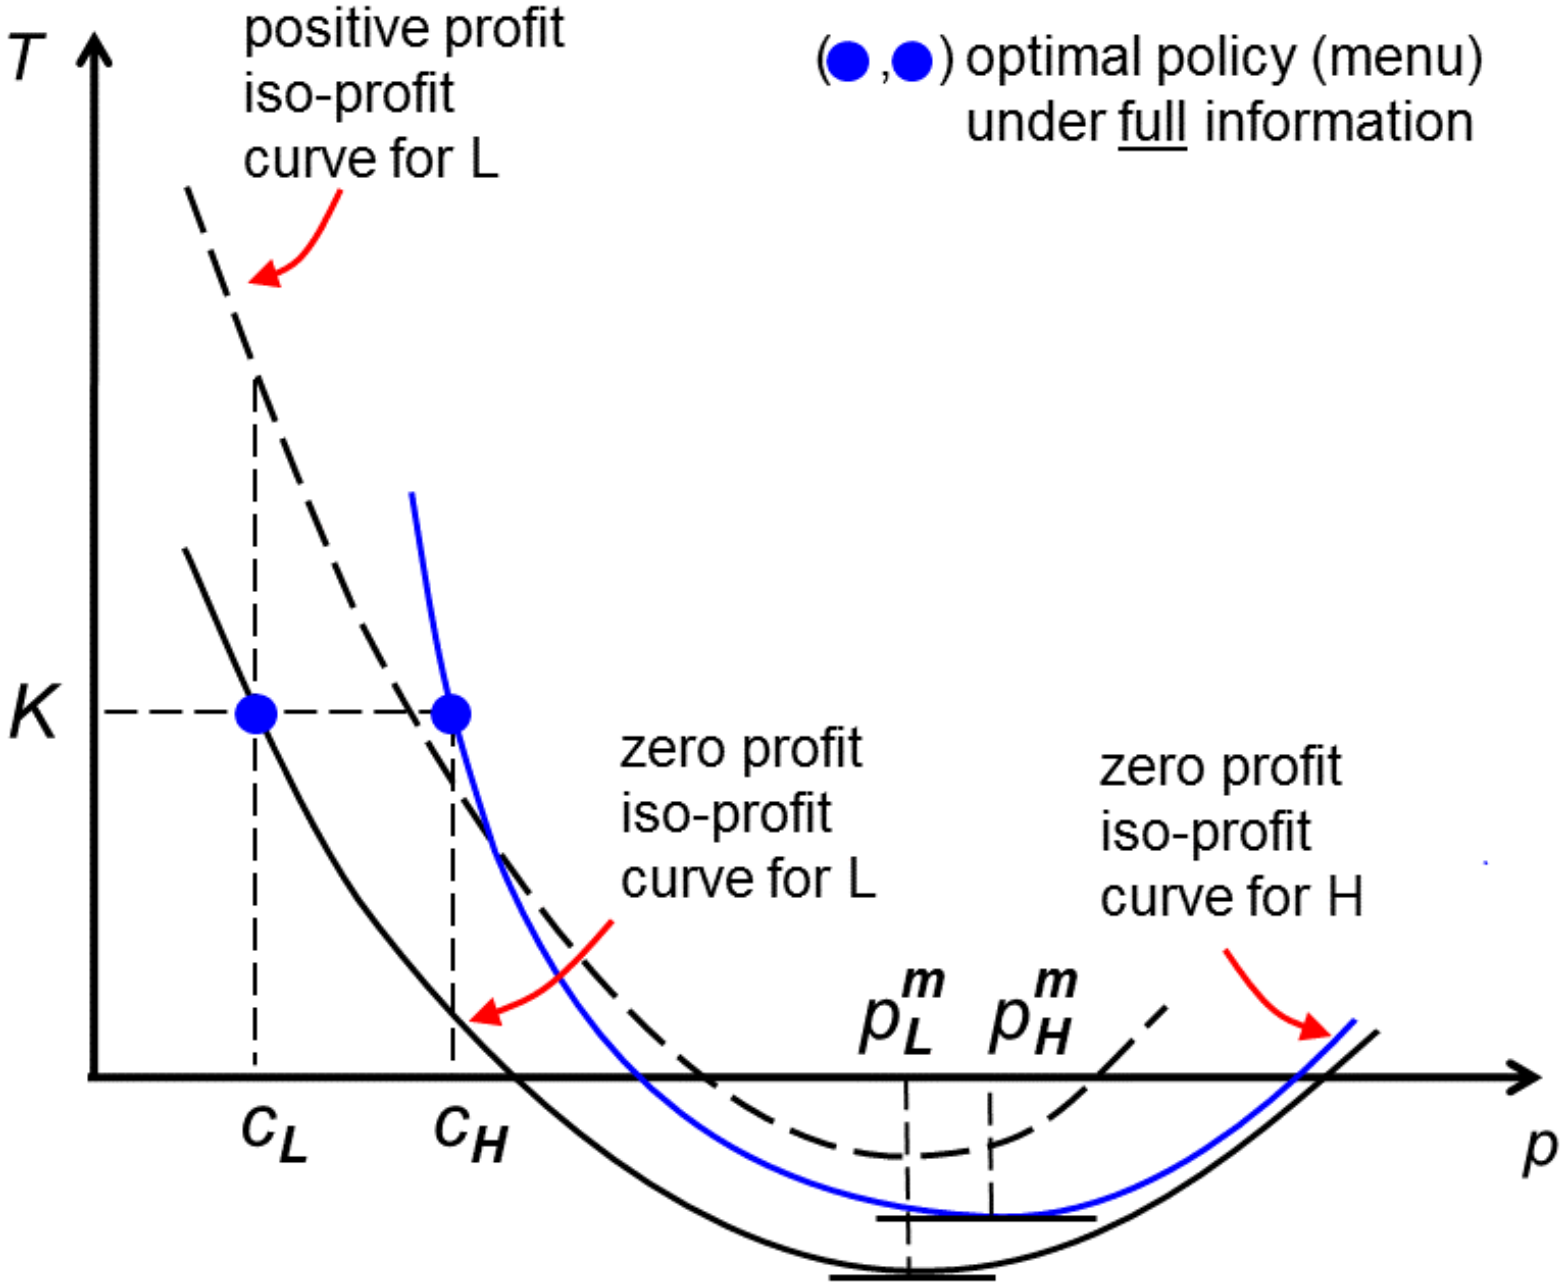

Regulation of a privately-informed monopolist

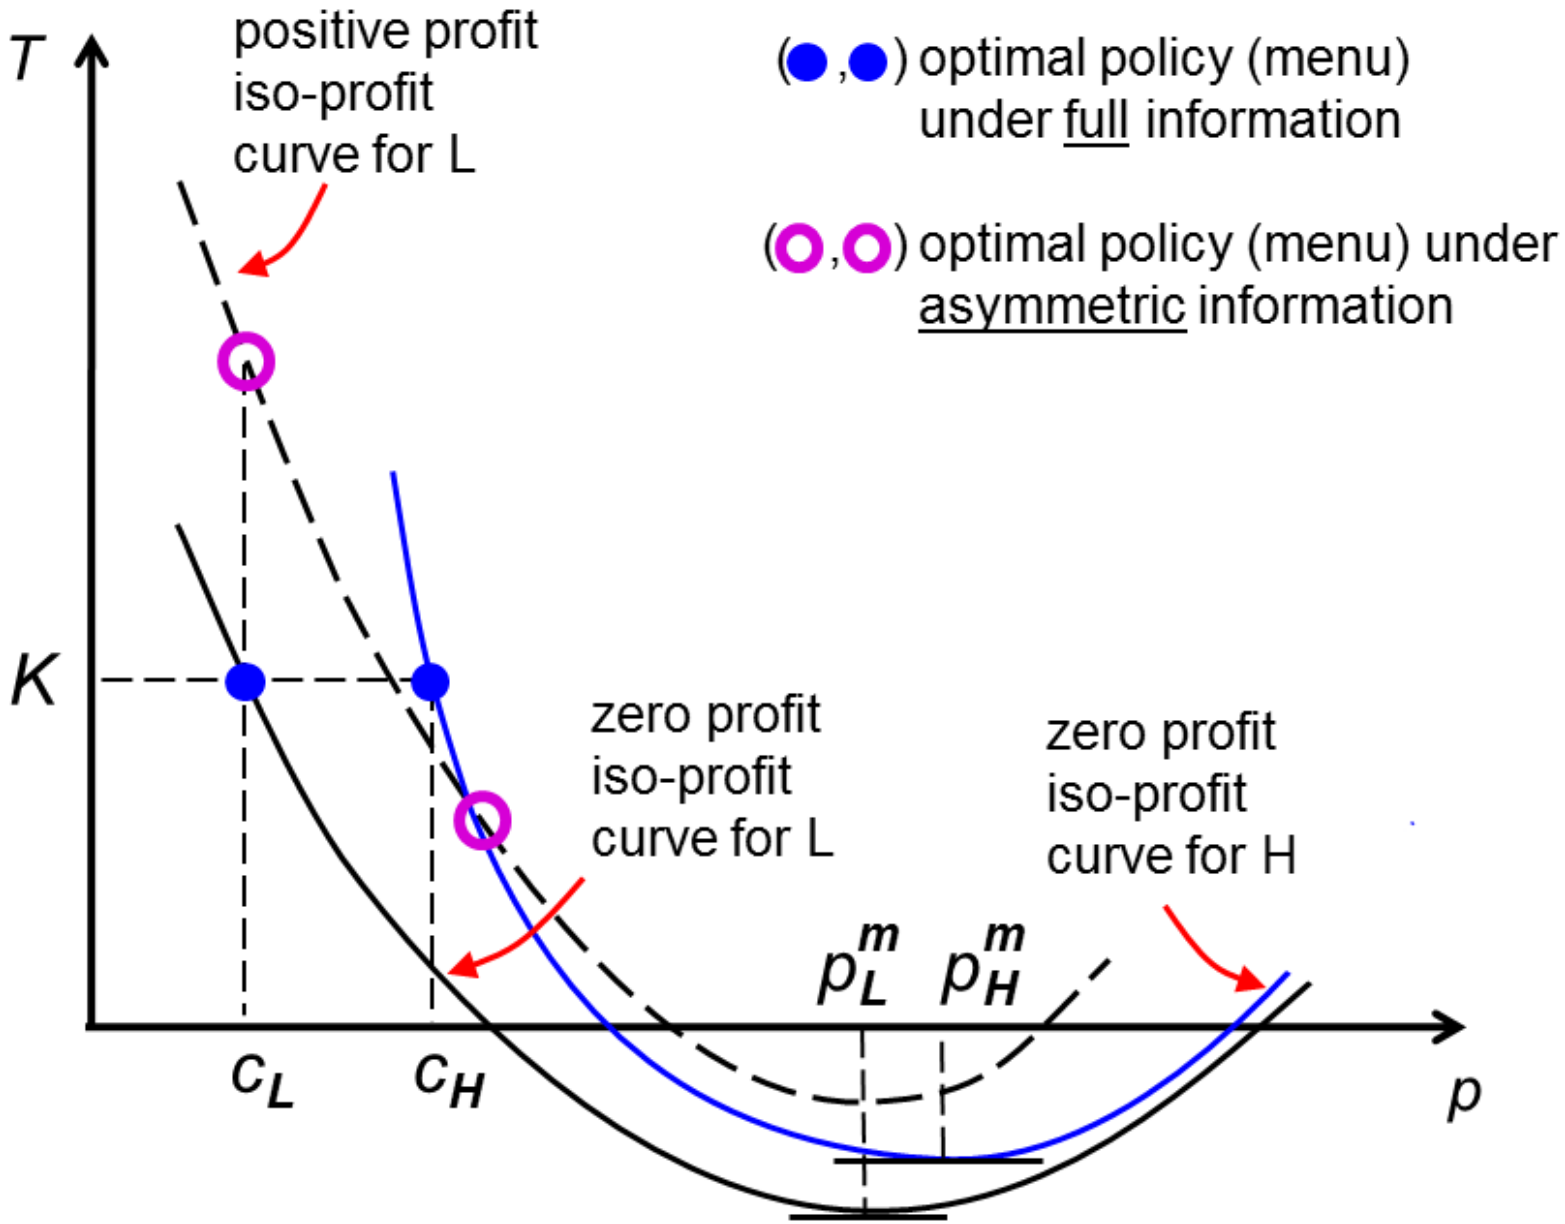

The timing of acquisition of hidden info is crucial

Does hidden info always result in losses relative to full info? No!

Example: Suppose timing of contracting game btw. P and A is changed:

1. P offers contract to risk-neutral A **before** A learns s
2. A accepts or rejects
3. s is drawn from F and privately observed by A
4. A chooses output and receives payment

Now, optimal mechanism for P solves:

$$\max_{(x_H,t_H),(x_L,t_L)} q(R(x_H)-t_H)+(1-q)(R(x_L)-t_L)$$

$$\begin{aligned} \text{s.t.} \quad & q(t_H - C(x_H,H)) + (1-q)(t_L - C(x_L,L)) \geq 0 \quad (\textbf{ex ante IR}) \\ & t_H - C(x_H,H) \geq t_L - C(x_L,H) \qquad (IC_H) \\ & t_L - C(x_L,L) \geq t_H - C(x_H,L) \qquad (IC_L) \end{aligned}$$

Now A does not have any informational advantage over P *at the time he signs the contract*. Hence, the participation constraint is *weaker*.

The timing of acquisition of hidden info is crucial

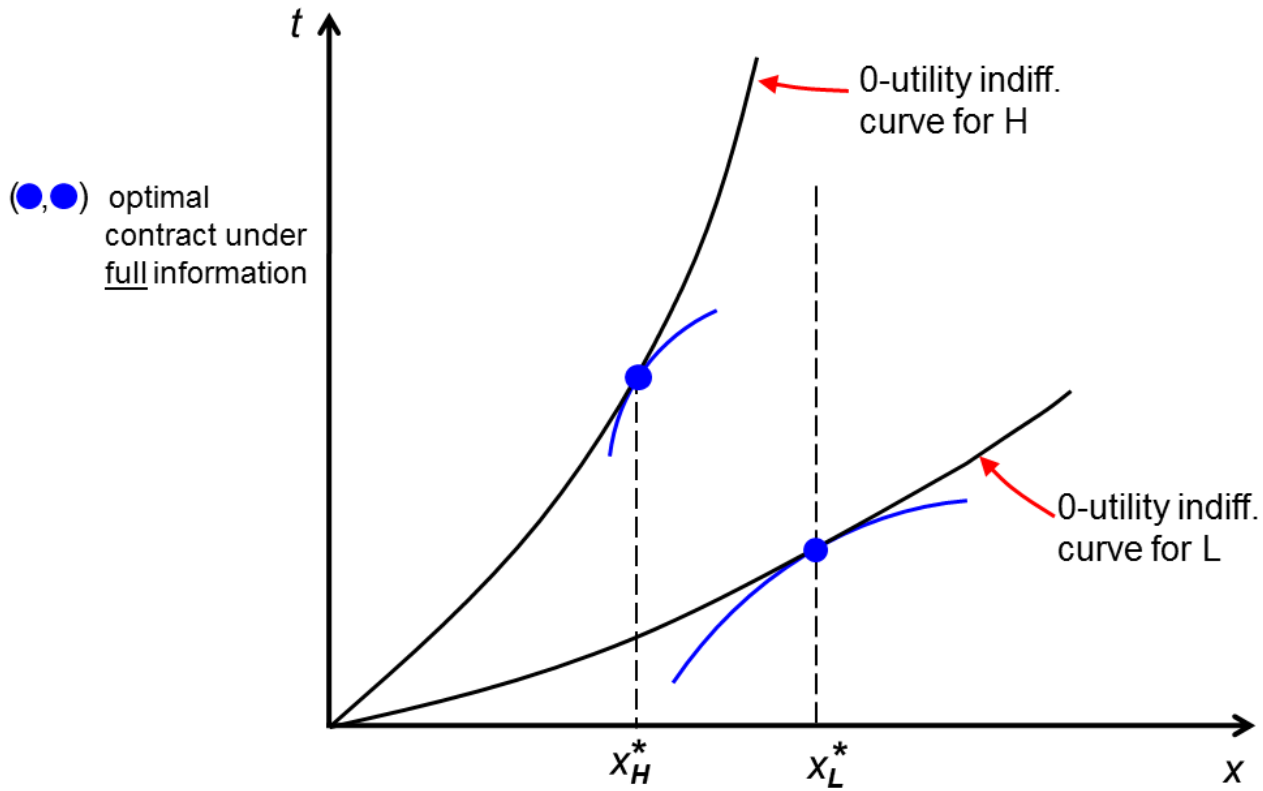

The timing of acquisition of hidden info is crucial

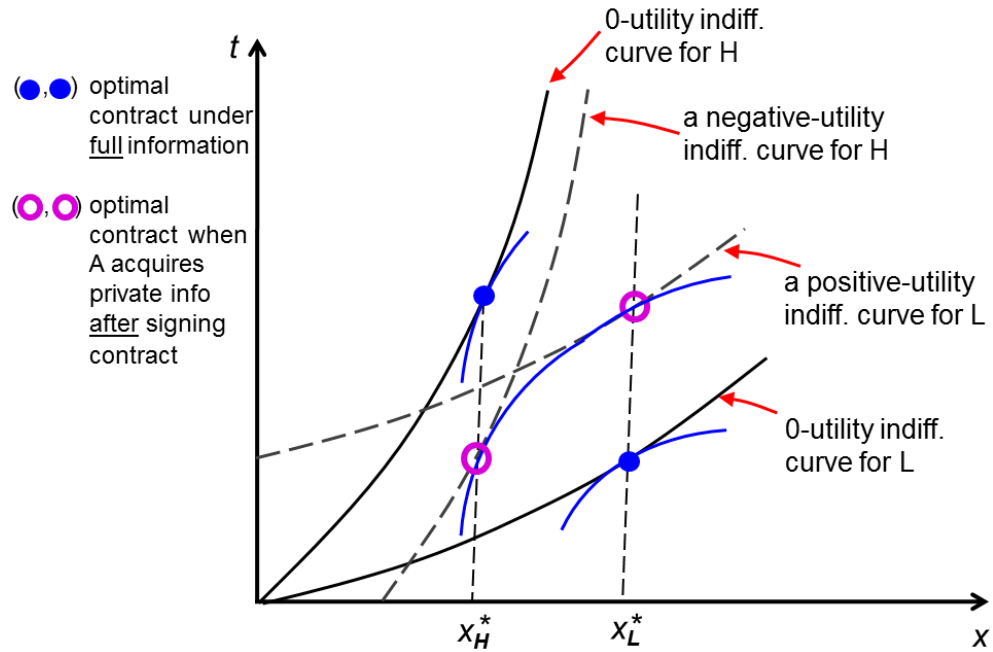

- Consider a contract (purple circles) in which P's profit is the same whether A turns out to be H or L and in which A's ex ante IR constraint is binding. (NB: ex post, H's payoff is negative, while L's is positive.)
- Allocatively efficient output for both types satisfies both IC constraints.
- A's ex ante expected utility ($= 0$) is the same as his ex ante expected utility in full-information benchmark.
- Hence, P's ex ante expected payoff is **the same** as in full-information benchmark: When A acquires hidden information **after** signing the contract, P suffers no loss, and no inefficiency results.