Feeback Design In Contests

Jeffrey Ely

November 14, 2019

In A Nutshell

Information design with multiple competing long-lived players whose actions affect an evolving payoff-relevant state.

Effort Maximizing Contests

Setup

- Two players, choosing effort over time $t \in [0, T]$.
  - Effort is binary.
  - Stopping game.
- $F(t)$ is the probability that player i has a success on or before time t.
  - F is strictly increasing and continuously differentiable
  - F has a declining hazard rate
  - $F(0) = 0$.
- Successes are independent across players.
- Effort cost is linear in effort duration $c \cdot t$ where $c > 0$.

Contest Structure

- There is a prize with common total value 1.
- The principal observes successes but not effort.
- The players observe their own effort, do not observe successes.
- Player earns

$$q - c \cdot t$$

  if q is the probability he is awarded the prize.

Change of Variables

- The total cost of success probability p is

$$c \cdot F^{-1}(p)$$

- Marginal cost

$$\mathrm{MC}(p) = c \cdot \frac{d}{dp} F^{-1}(p).$$

We will assume that MC is strictly increasing.

Two Prize Structures

1. Winner Take All: the first to succeed earns the prize exclusively.
2. Egalitarian: The prize is equally shared by all who succeed.

Feedback Design

The principal can provide arbitrary feedback to each player about past successes.

- Private feedback
- Public feedback
- Feedback about feedback
- Random feedback

We will consider the problem of maximizing effort.

Winner Take All

- If player 2 works until time t_2^*, she succeeds with probability $p_2^* = F(t_2^*)$
- If player 1 succeeds with probability $p_1 = F(t_1)$, then he wins the prize with the following total probability.

$$q = \begin{cases} \frac{1}{2}\left[p_1 + (1 - p_1)\, p_1\right] & \text{if } t_1 \leq t_2^* \\ \frac{1}{2}\left[p_2^* + (1 - p_2^*)\, p_2^*\right] + (1 - p_2^*)\left[p_1 - p_2^*\right] & \text{if } t_1 > t_2^*. \end{cases}$$

Marginal Benefit

And thus the marginal benefit of increased success probability is

$$\mathsf{MB}^{\mathrm{WTA}}(p_1) = \begin{cases} 1 - p_1 & \text{if } p_1 \leq p_2 \\ 1 - p_2^* & \text{if } p_1 > p_2. \end{cases}$$

Egalitarian Contest

Player 1 wins the prize with probability

$$q = p_1 (1 - p_2^*) + \frac{1}{2} p_1 p_2^*$$

yielding marginal benefit

$$\mathrm{MB}^{\mathrm{EGA}}(p_1) = 1 - p_2^*/2,$$

No Feedback Benchmark

If no feedback is given, then symmetric equilibrium in either contest is characterized by a successprobability $p_1 = p_2^* = p$ satisfying

$$\mathrm{MB}(p) = \mathrm{MC}(p)$$

Equilibrium Illustration

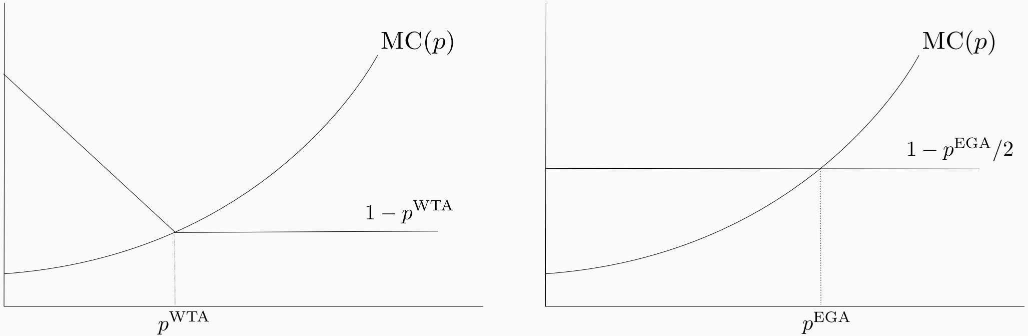

(a) Winner Take All

(b) Egalitarian

Egalitarian Induces More Effort

Because this is impossible:

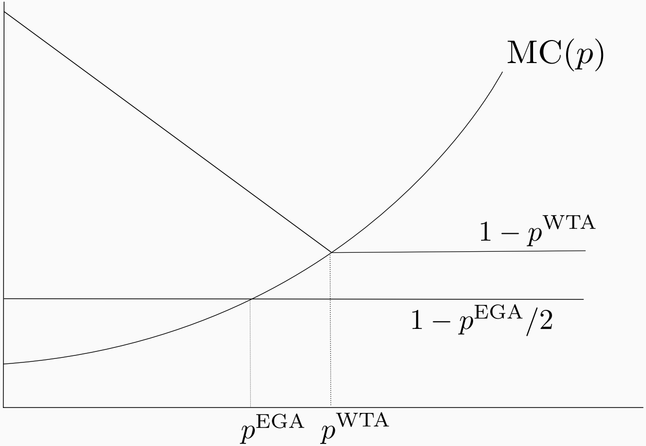

Feedback Design: Egalitarian Contest

We will discretize time when necessary.

Minmax Payoffs

Suppose $p_2 = \bar{p} = F(T)$.

Minmax Payoffs

If player 1 chooses p_1 her payoff is

$$q = p_1 (1 - \bar{p}) + \frac{1}{2} p_1 \bar{p}$$

yielding *minmax* marginal benefit

$$\overline{\text{MB}}(p_1) = 1 - \bar{p}/2.$$

Minmax Payoff

Let $\overline{v}^{\mathrm{EGA}}$ be player 1's payoff from a best-response (characterized by $\overline{\mathrm{MB}} = \mathrm{MC}$).

Proposition

Regardless of the disclosure policy, neither player's equilibrium payoff in the egalitarian contest can be less than $\overline{v}^{\mathrm{EGA}}$.

Sufficient Condition for Optimal Feedback

Proposition

Suppose there exists a feedback policy and corresponding equilibrium strategies such that

1. *The prize is earned with the maximum probability* $\overline{q} = 1 - (1 - \bar{p})^2$
2. *Each player earns the minmax payoff* $\overline{v}^{\mathrm{EGA}}$.

Then the feedback policy is optimal.

Welfare Accounting

Player i's utility is

$$u_i = q_i - c \cdot \mathbf{E} t_i$$

So,

$$\begin{aligned}\sum_i u_i &= \sum_i q_i - c \cdot \mathbf{E} \sum_i t_i \\ &= \mathrm{TW} - c \cdot \mathbf{E} \sum_i t_i\end{aligned}$$

Re-arranging,

$$c \cdot \mathbf{E} \sum_i t_i = \mathrm{TW} - \sum_i u_i$$

Discrete Time

Let f^t be the conditional probability of success in period t given no success so far.

One-Shot Deviations

Suppose at time t, player 1 believes he has succeeded with probability p^t. Then

$$\left(1 - p^t\right) f^t \left[\frac{\bar{p}}{2} + (1 - \bar{p})\right]$$

is the gain from spending effort for one more period.

One-Shot Deviations

Define π^t to satisfy

$$\left(1-\pi^t\right) f^t \left[\frac{\bar{p}}{2} + (1-\bar{p})\right] = c$$

- In the continuous time limit this is just MB = MC.
- In the case of a constant hazard rate π^t is constant.

Leading Them On

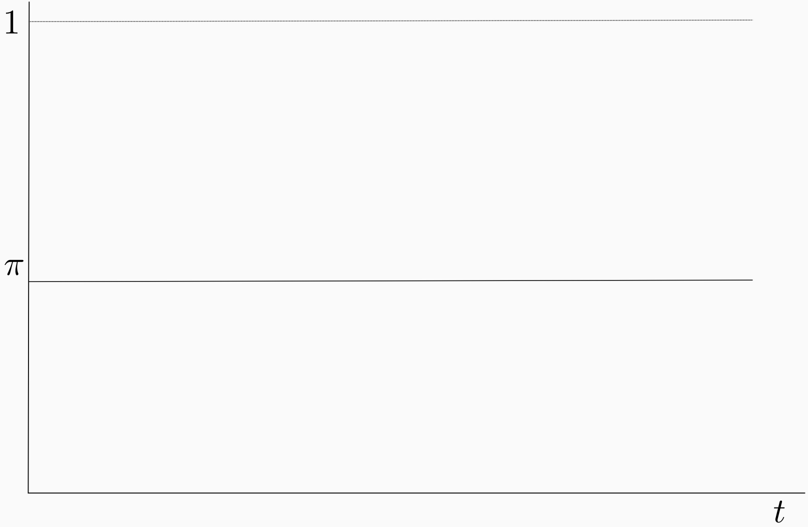

Leading Them On

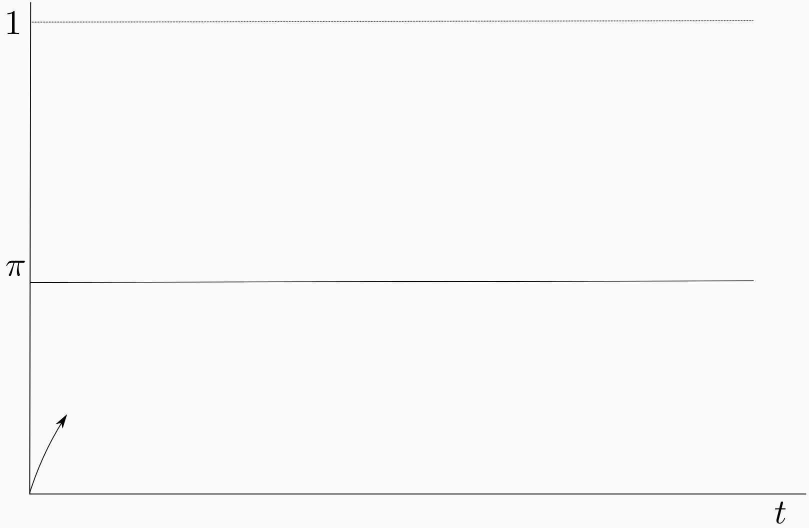

Leading Them On

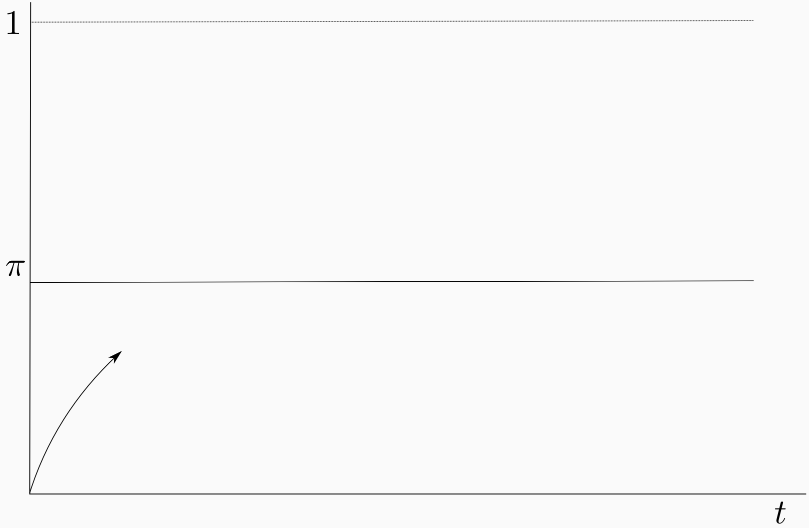

Leading Them On

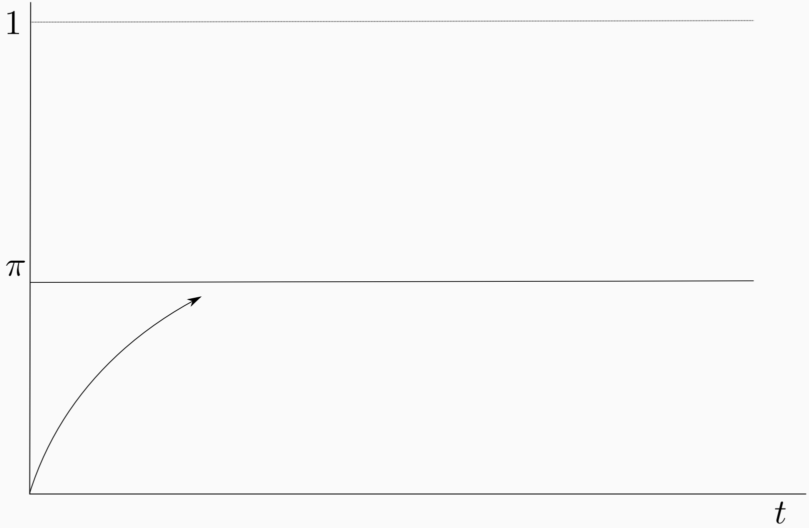

Leading Them On

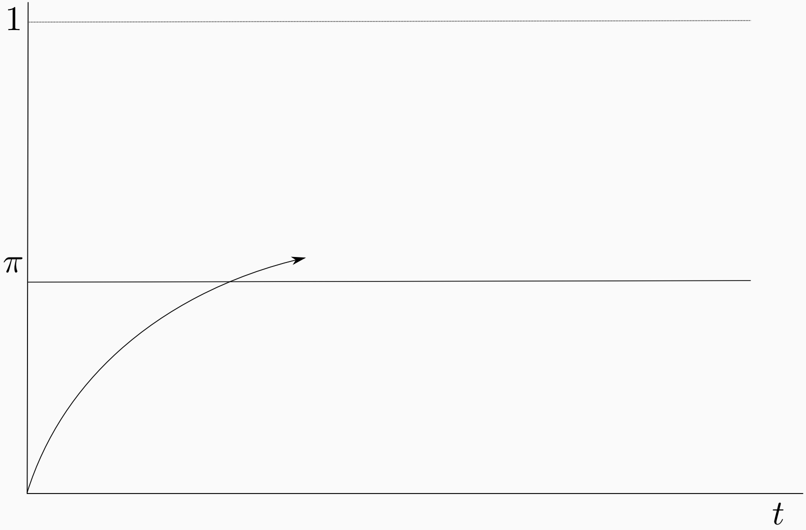

Leading Them On

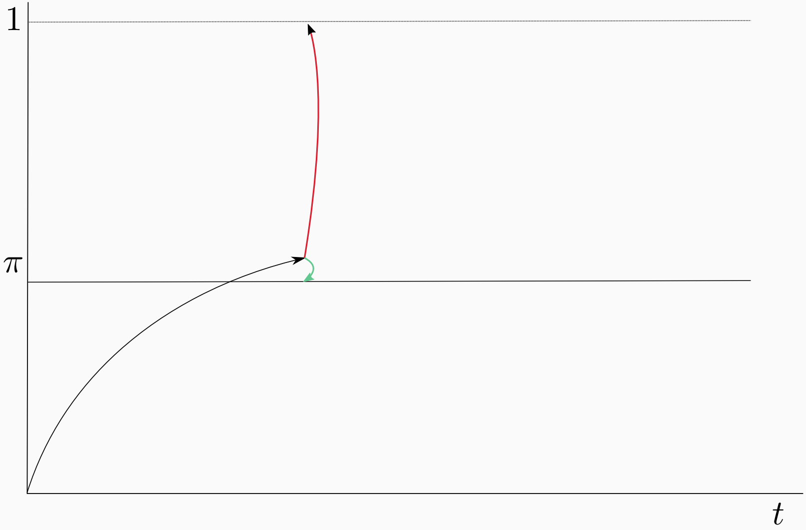

Leading Them On

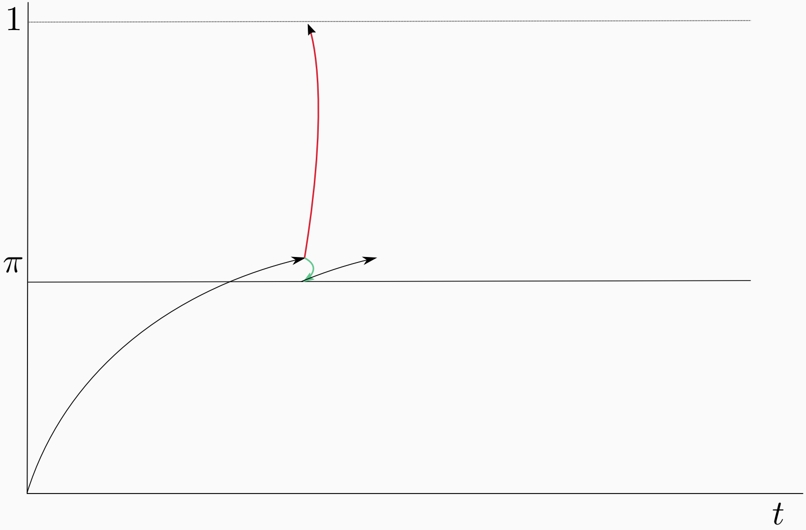

Leading Them On

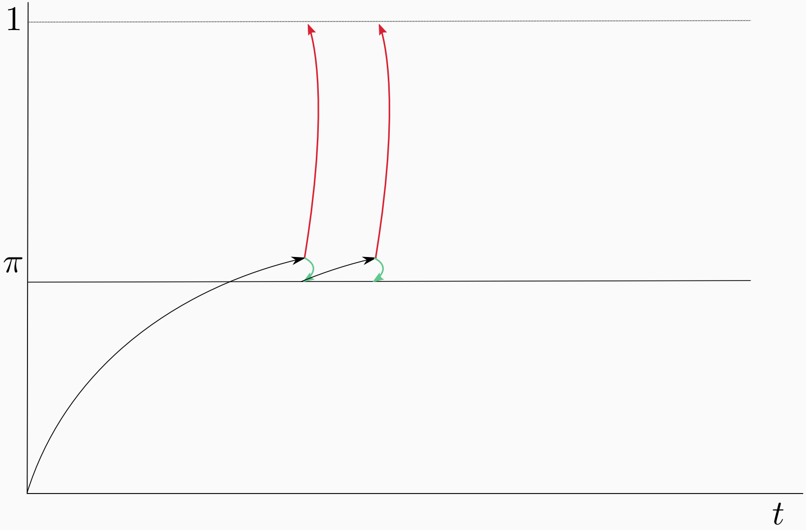

Leading Them On

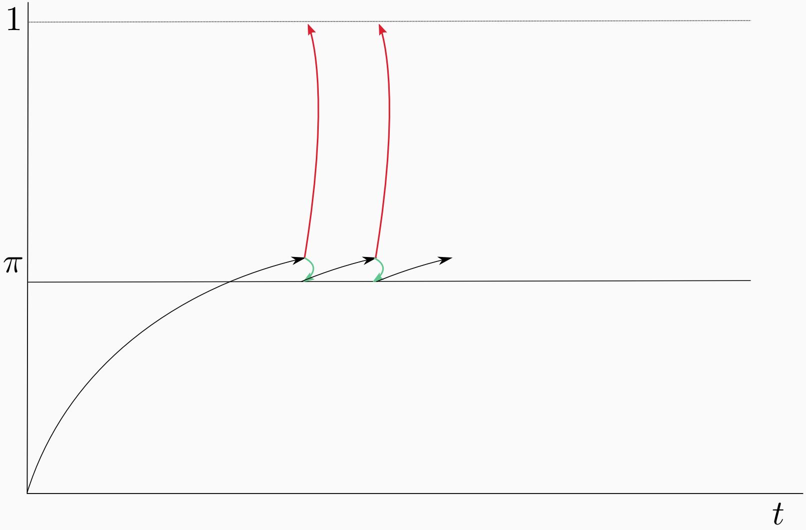

Leading Them On

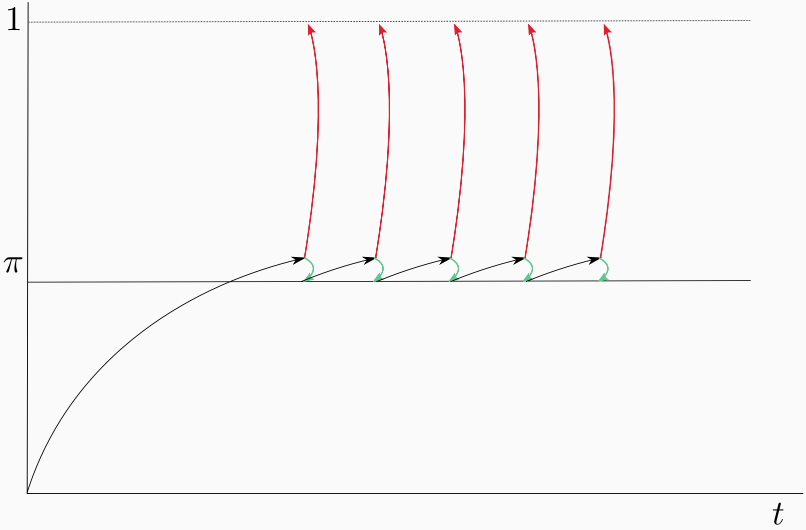

One-Shot Deviation Principle

By the One-Shot Deviation Principle, player 1's best response is to spend effort until the principal "tells him to quit."

Joint Mechanism

- Note that in so doing Player 1 succeeds with the maximum probability $\bar{p}$.
- Thus, while using this mechanism for Player 1, we can repeat the analysis for Player 2.
- The principal's messages will be conditionally independent across players.

Optimality

In the continuous-time limit, the players are held to their minmax values.

Leading Them On

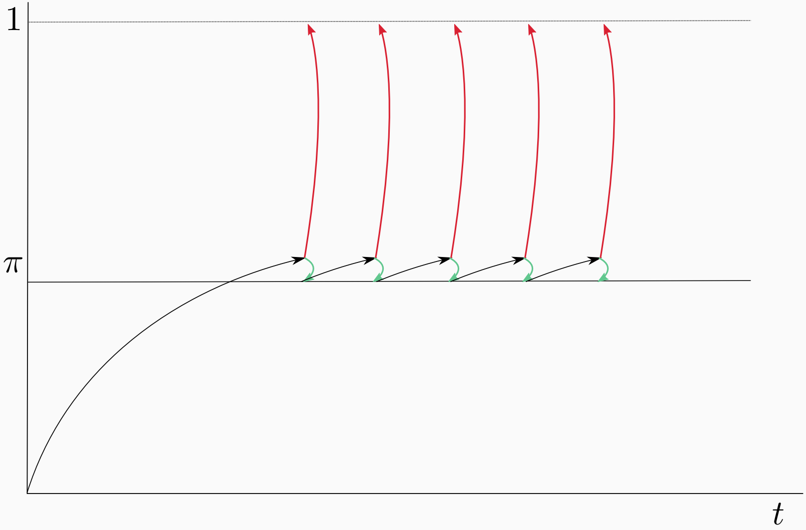

Leading Them On

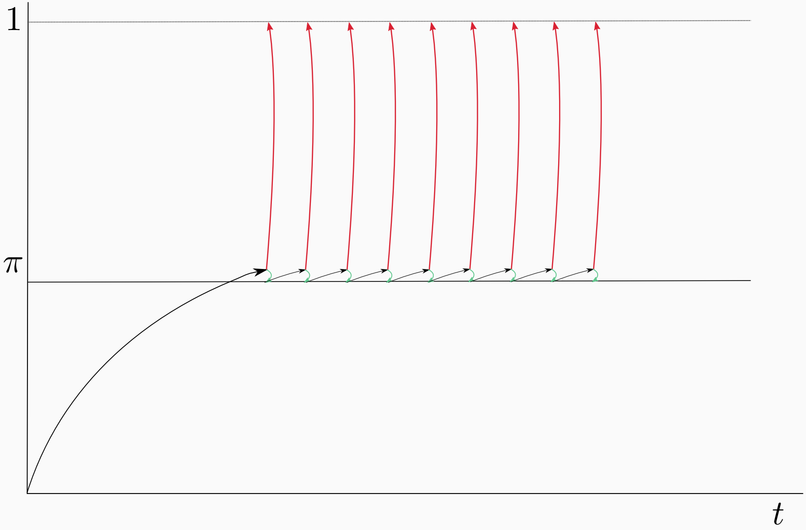

Leading Them On

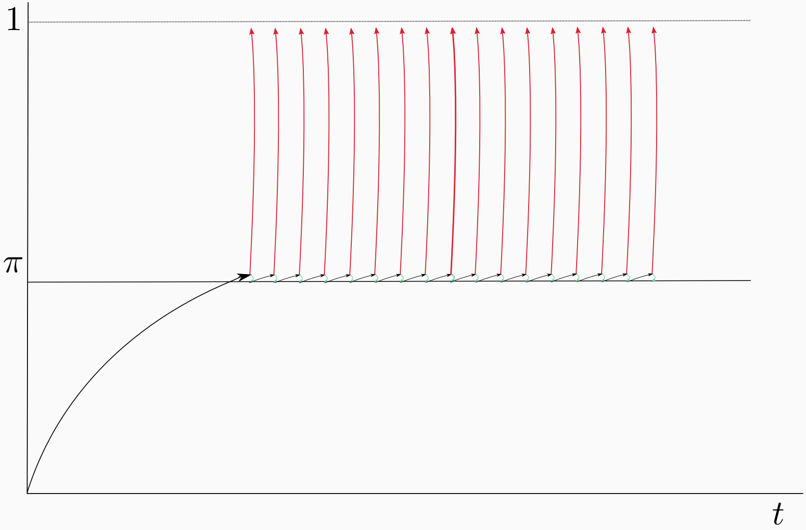

Leading Them On

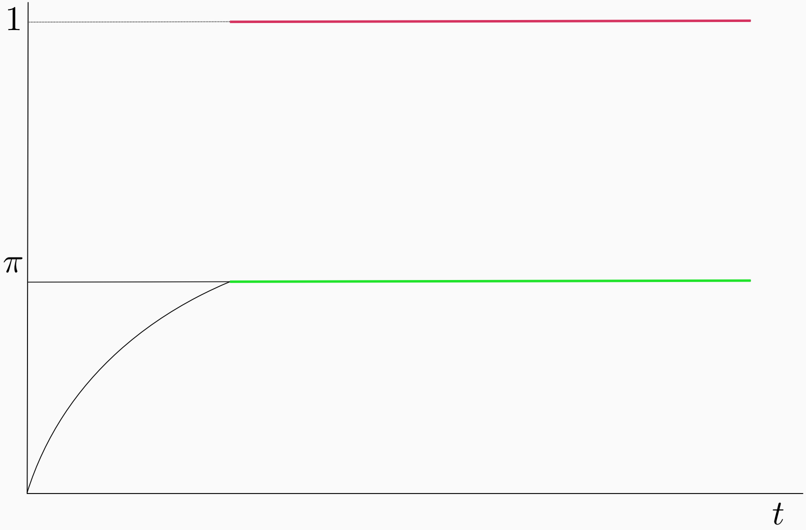

Leading Them On

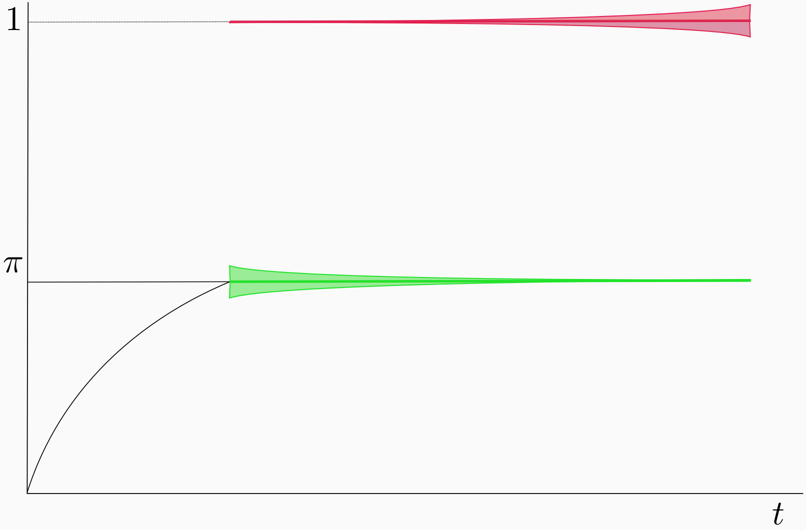

Leading Them On

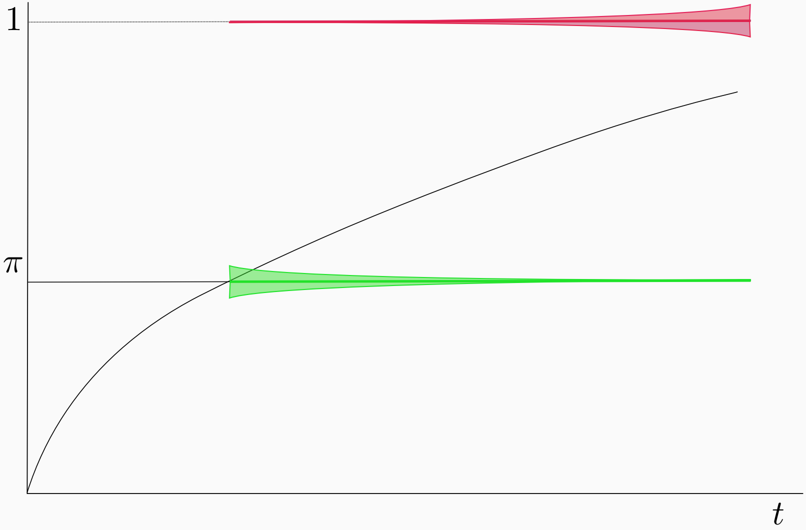

We Can Do The Same Thing for Winner-Take-All

In WTA the minmax is a little different. Let's assume that player 2 never quits before player 1 does.

We Can Do The Same Thing for Winner-Take-All

When player 1 chooses to succeed with probability p_1 her payoff is

$$q = \frac{1}{2}\left[p_1 + (1 - p_1)p_1\right]$$

so that the marginal benefit is

$$\overline{\text{MB}}(p_1) = 1 - p_1$$

WTA: One-Shot Deviations

Suppose at time t, player 1 believes *at least one player* has succeeded with probability p^t. Then

$$\left(1-p^t\right) f^t \left[\frac{f^t}{2} + (1-f^t)\right]$$

is the gain from spending effort for one more period. Define π^t to satisfy

$$\left(1-\pi^t\right) f^t \left[\frac{f^t}{2} + (1-f^t)\right] = c$$

WTA: Leading Them On

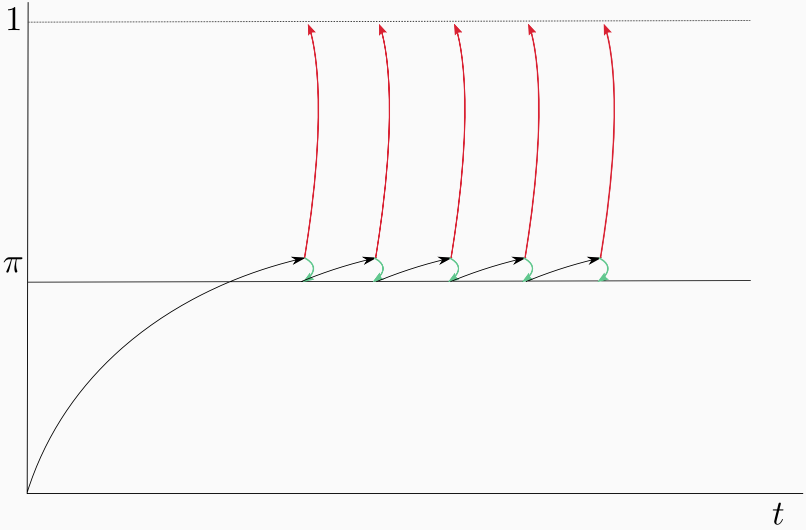

WTA: Leading Them On

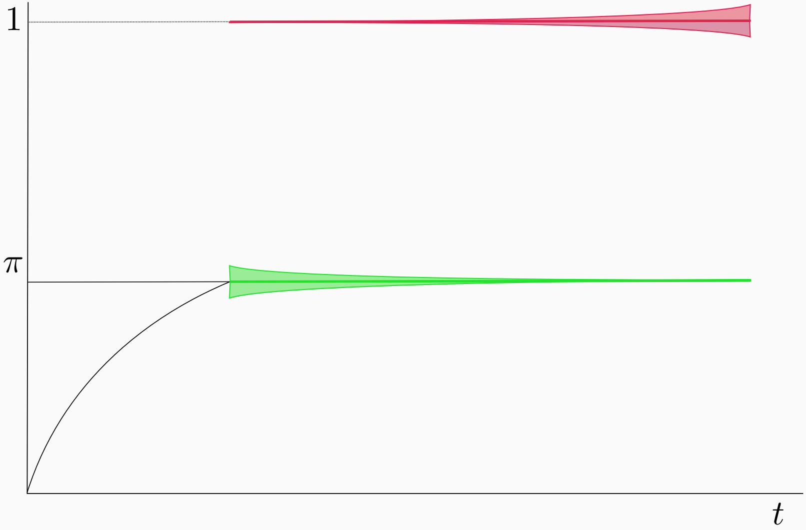

WTA: Leading Them On

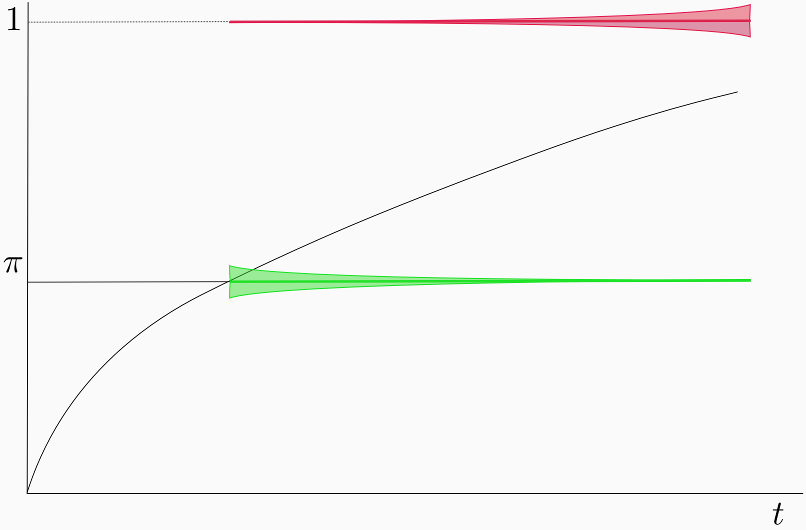

Ranking The Two Contests

- Each awards the prize with the maximum feasible probability.
- Each holds the players to their minmax values.
- Thus, whichever yields a lower minmax value is better for the principal

Ranking The Two Contests

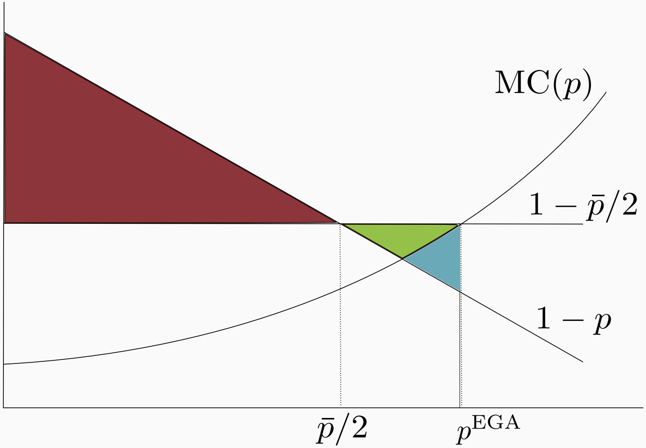

Figure 1: Comparison Under Optimal Disclosure

What's Not Part of The Optimal Feedback Policy

1. Feedback about feedback.
2. Feedback about the opponent's success.

Intuition

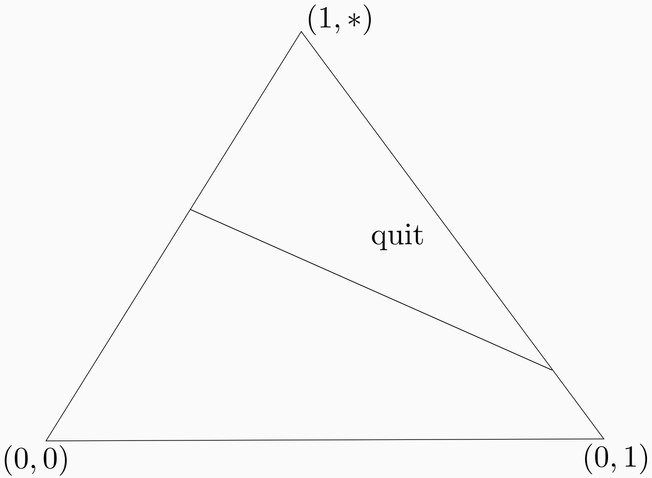

The threshold for willingness to work one more period

Intuition

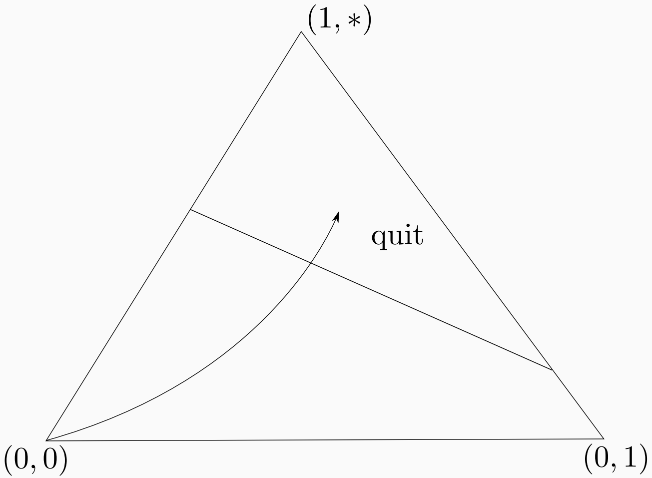

At some point the threshold is crossed.

Intuition

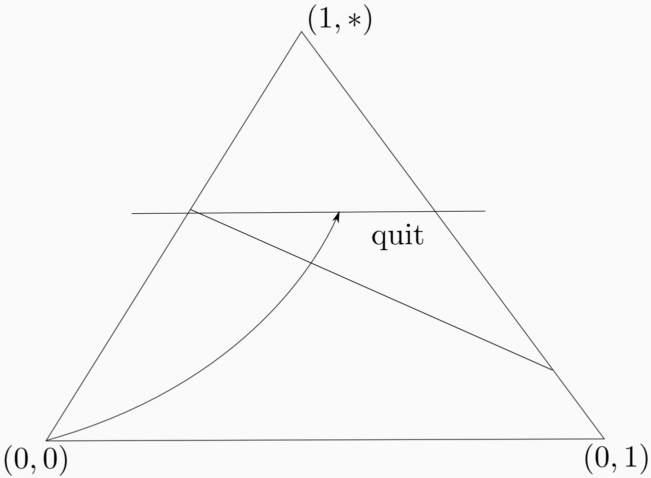

Horizontal level sets of my own success probability.

Intuition

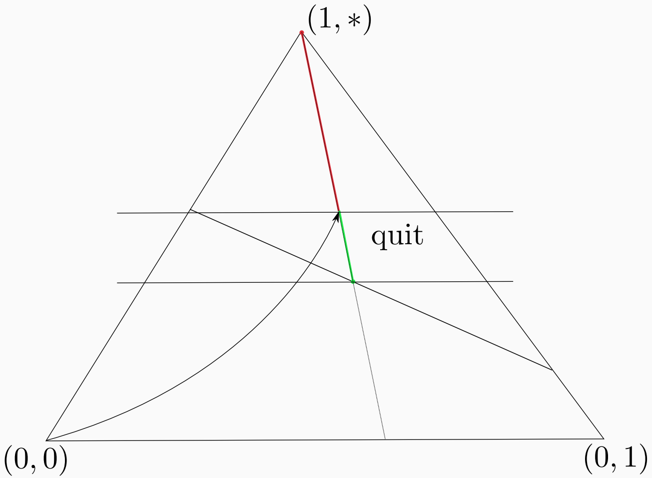

Level set of the relative probability that the other player has already succeeded.

Intuition

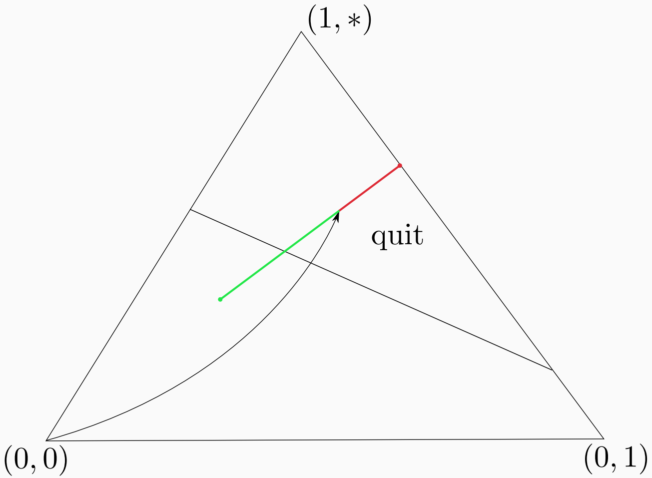

Here is one possible feedback.

Intuition

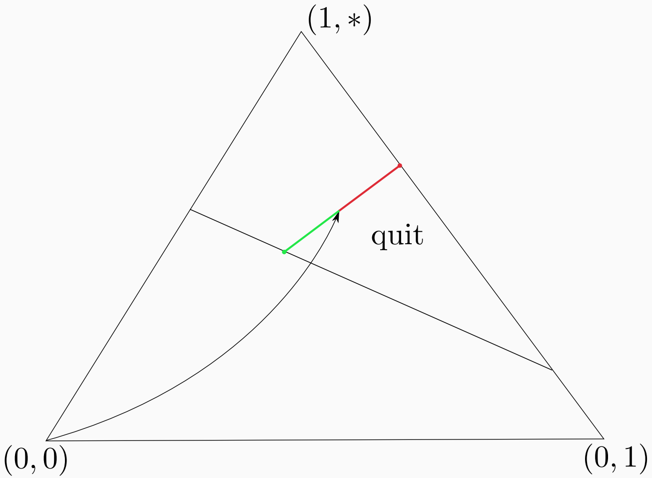

This one is better

Intuition

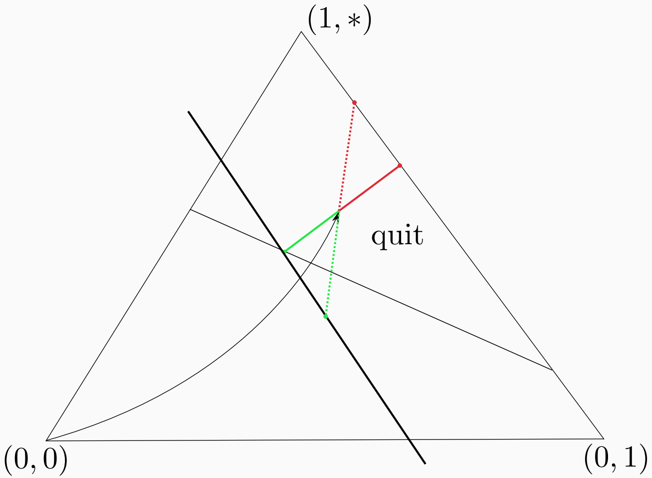

All signals like this are equally good

Intuition

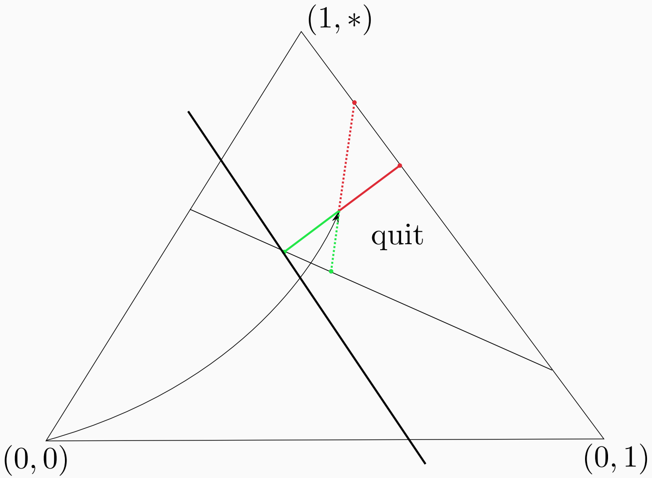

But this one is better

Intuition

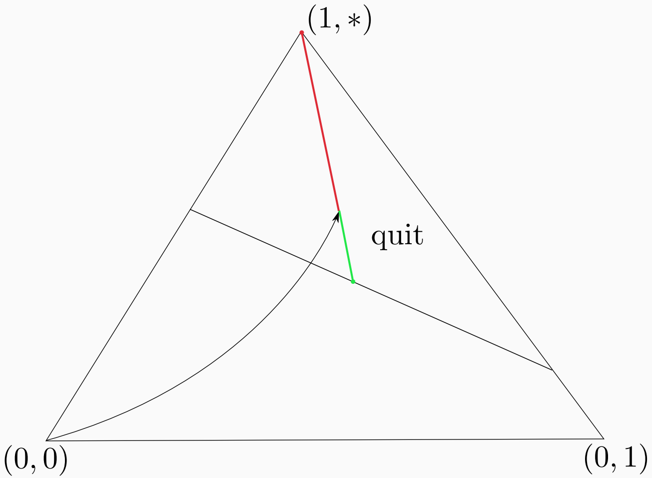

So this one is the best

Intuition

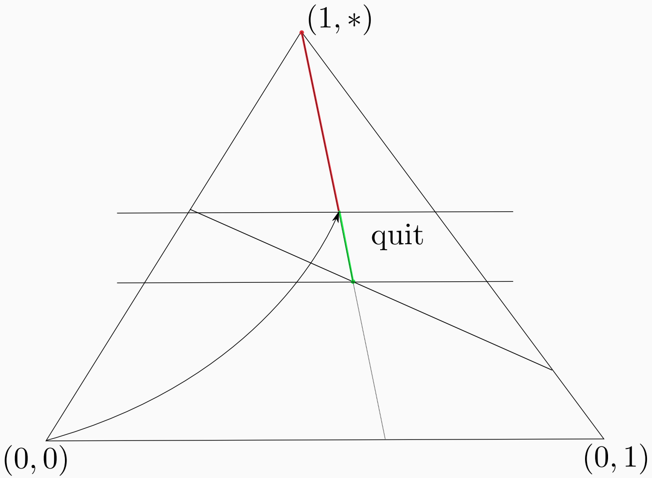

This signal conveys no information about the other's progress.

Last To Succeed

The egalitarian contest works well because its induced minmax marginal benefit is *backloaded.*

Can we backload even more?

LTA: The full prize is awarded to the *last* player to succeed.

Minmax Payoffs

Suppose player 2 succeeds with probability $\bar{p}$. If player 1 succeeds with probability p_1 her payoff is

$$q = p_1 \left[p_1/2 + (1 - \bar{p}) \right]$$

yielding *minmax* marginal benefit

$$\overline{\mathrm{MB}}^{\mathrm{LTA}}(p_1) = (1 - \bar{p}) + p_1$$

Backloading

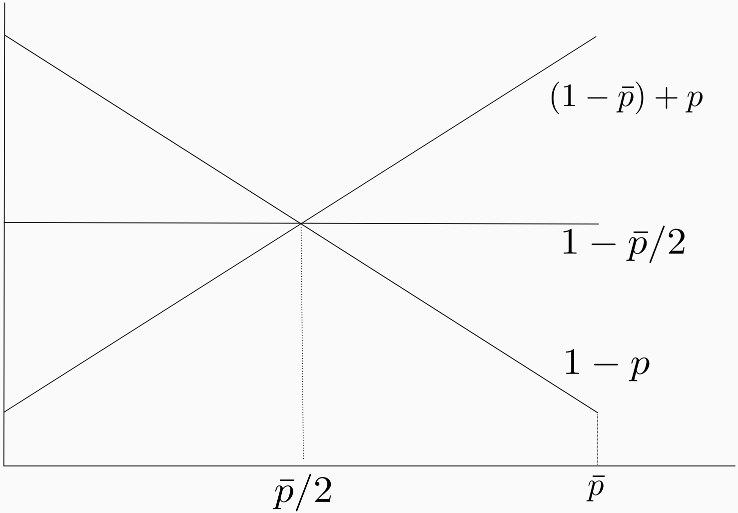

Unfortunately

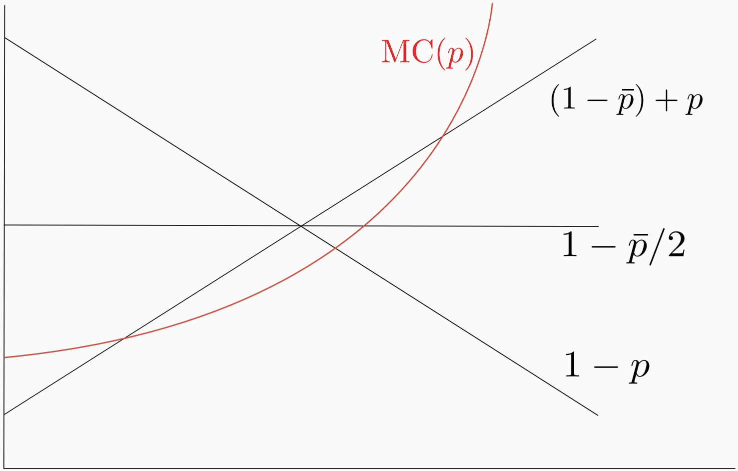

Backloading

Define the *minmax total benefit* of a contest by

$$\mathsf{TB}^{\mathrm{C}} = \int_0^{\bar{p}} \overline{\mathsf{MB}}^{\mathrm{C}}(s)ds$$

Say that contest C is more backloaded than contest D if

$$\mathsf{TB}^{\mathrm{C}} \geq \mathsf{TB}^{\mathrm{D}}$$

and $\overline{\mathsf{MB}}^{\mathrm{C}}(\cdot)$ is larger than $\overline{\mathsf{MB}}^{\mathrm{D}}(\cdot)$ in terms of single-crossing.

The Value of Backloading

If contest C is incentive-compatible and more backloaded than contest D then C induces more expected effort.

The Value of Backloading Illustrated

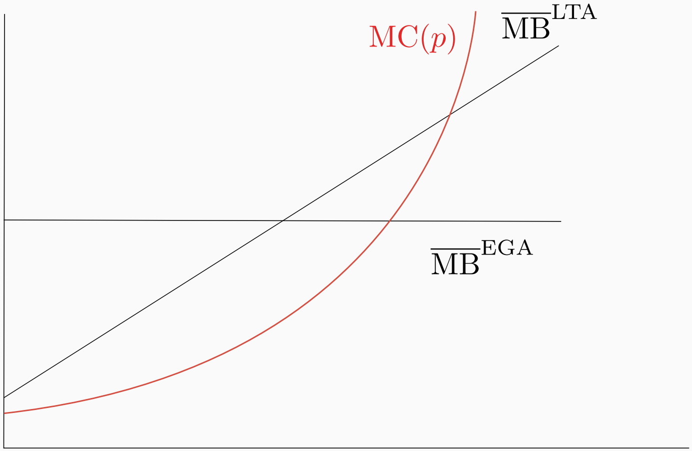

The Value of Backloading Illustrated

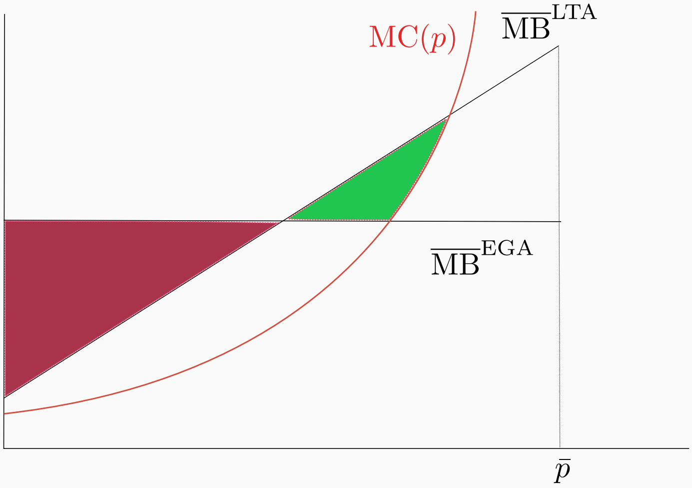

Which is larger, the red or green area?

The Value of Backloading Illustrated

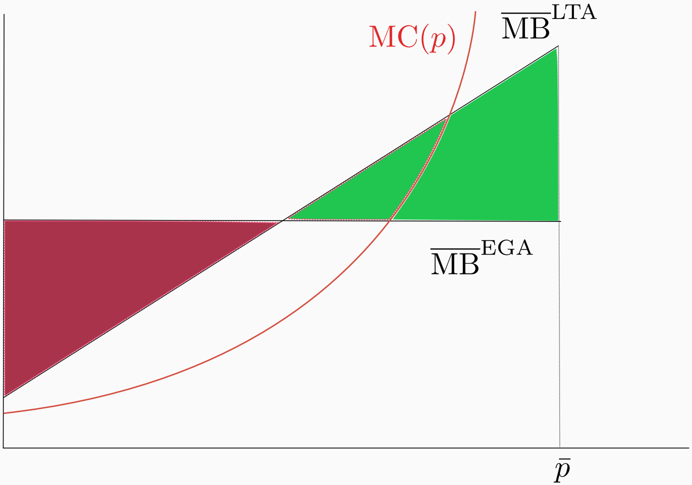

The red area equals this larger green area.

Nonlinear Example

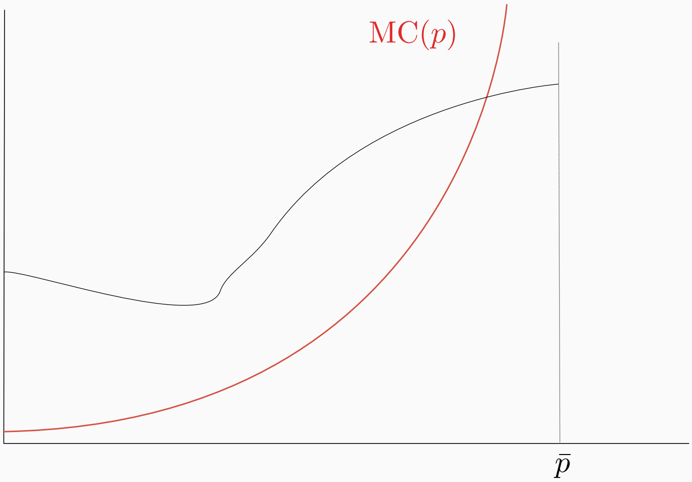

Nonlinear Example

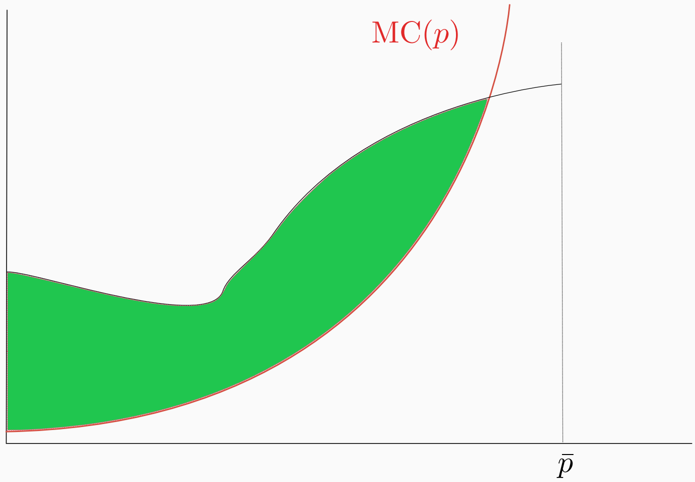

Nonlinear Example

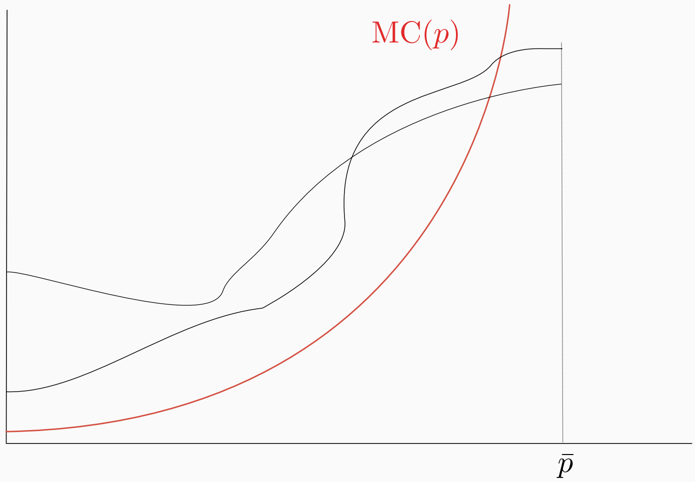

Nonlinear Example

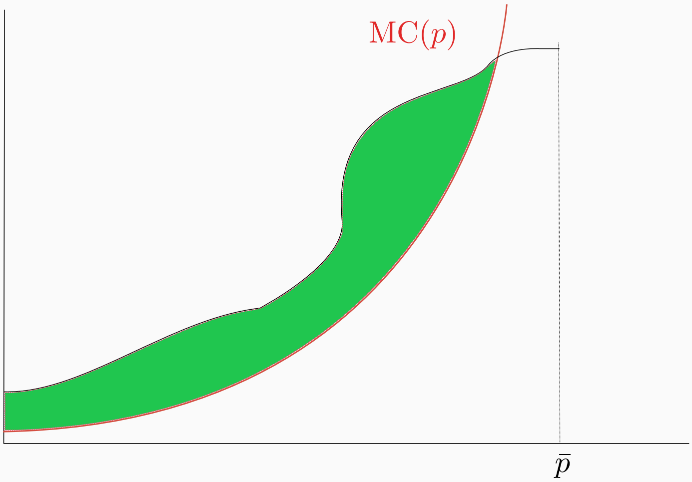

Nonlinear Example

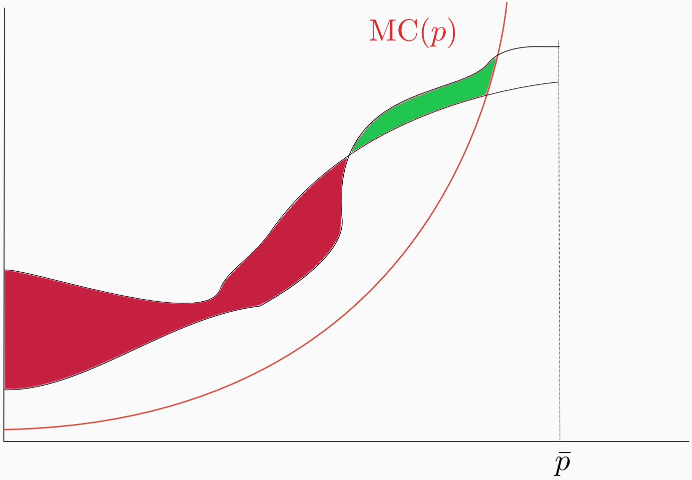

Nonlinear Example

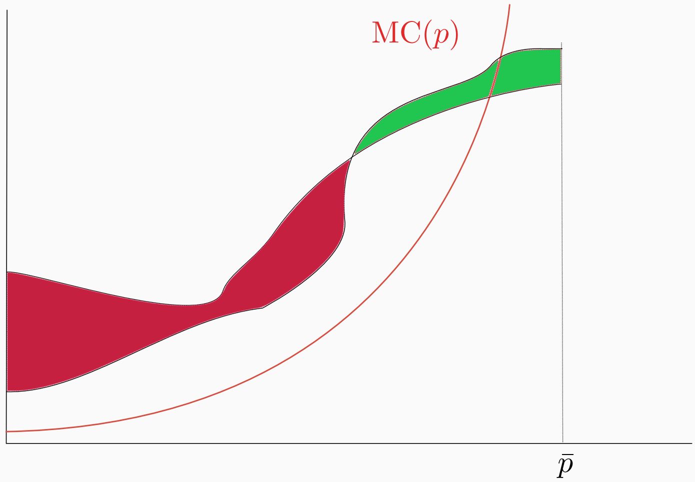

Disclosure Policy for LTA

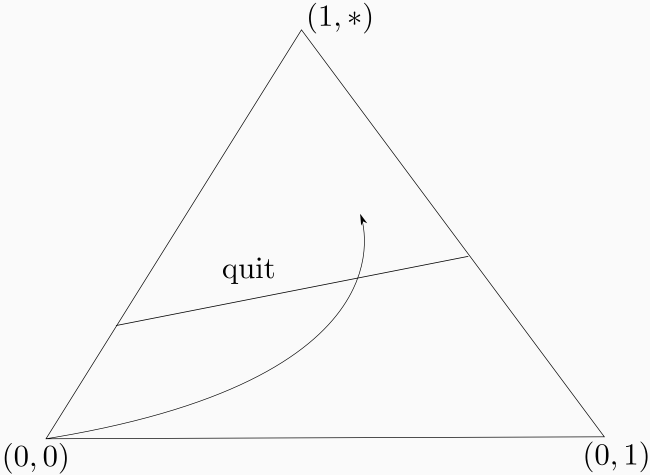

Disclosure Policy for LTA

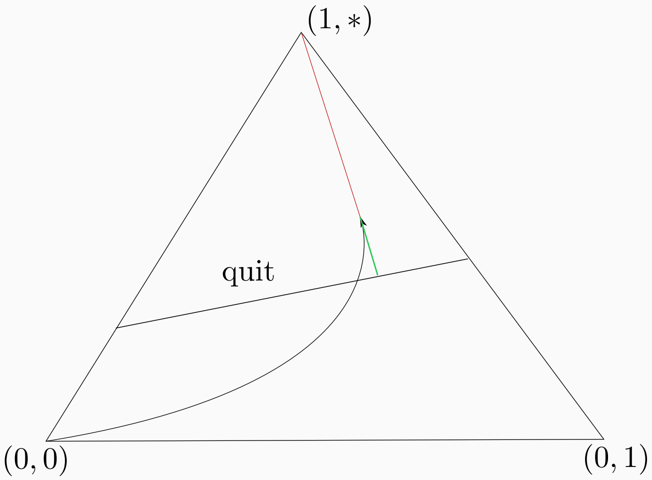

Disclosure Policy for LTA

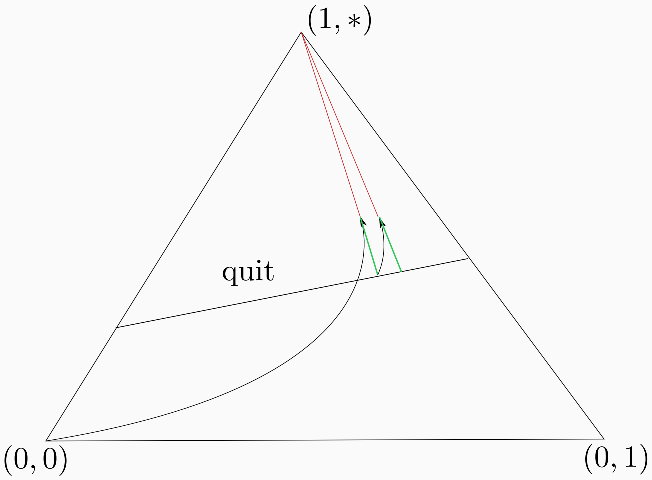

Disclosure Policy for LTA

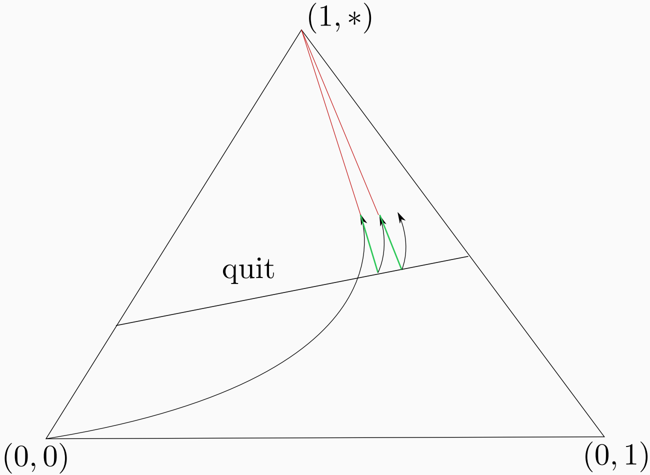

Fine Tuning

When LTA is not incentive compatible, the reason is that it is *too* backloaded.

To fine-tune backloading we can consider mixing a frontloaded contest like WTA with a backloaded contest like LTA.

Hybrid Contests

Suppose the principal randomly chooses beween WTA and LTA and never informs the players.

Hybrid Contests: Marginal Benefit

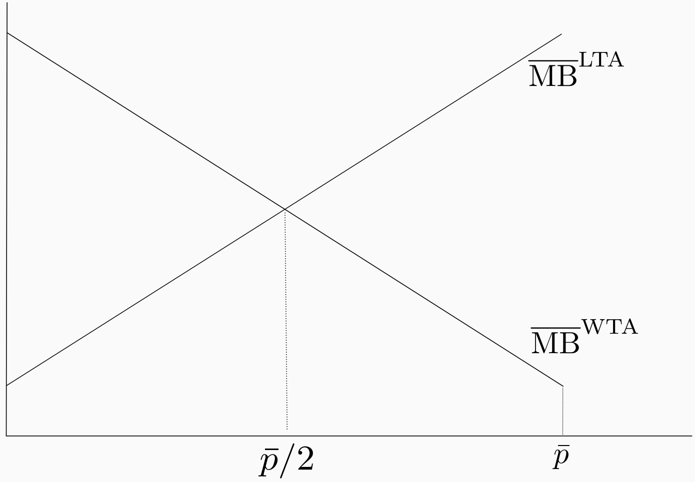

Hybrid Contests: Marginal Benefit

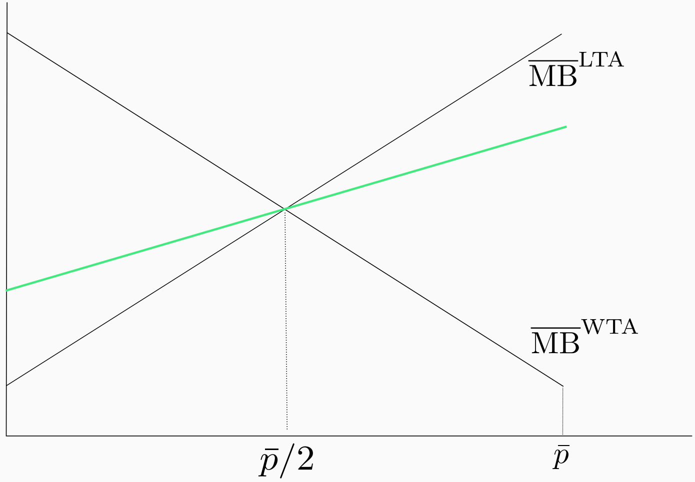

Hybrid Contests: Marginal Benefit

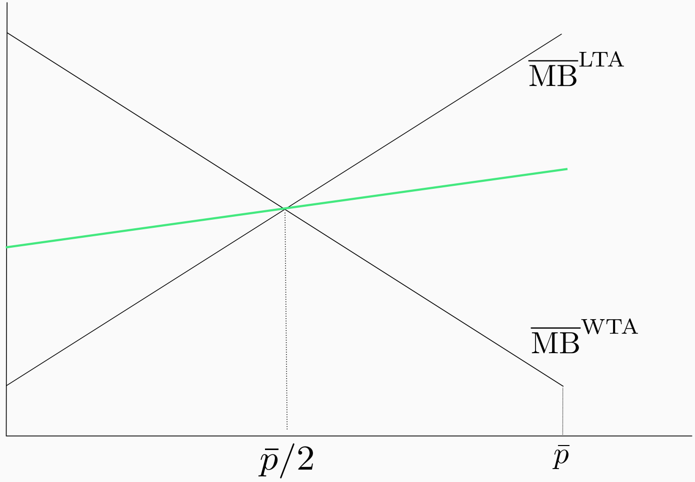

Hybrid Contests: Marginal Benefit

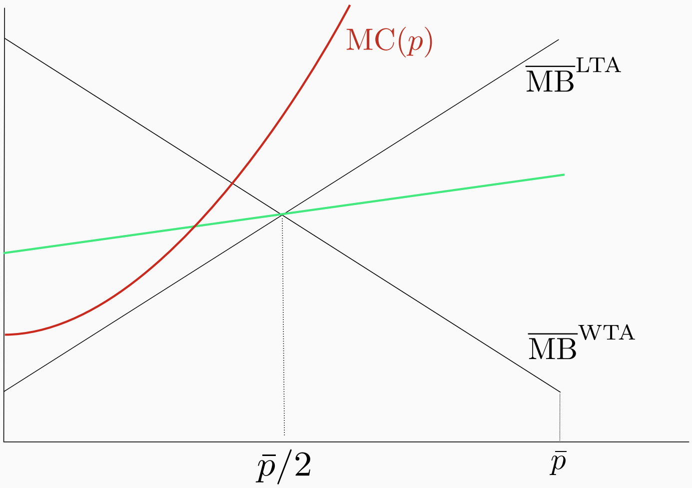

Hybrid Contests: Marginal Benefit

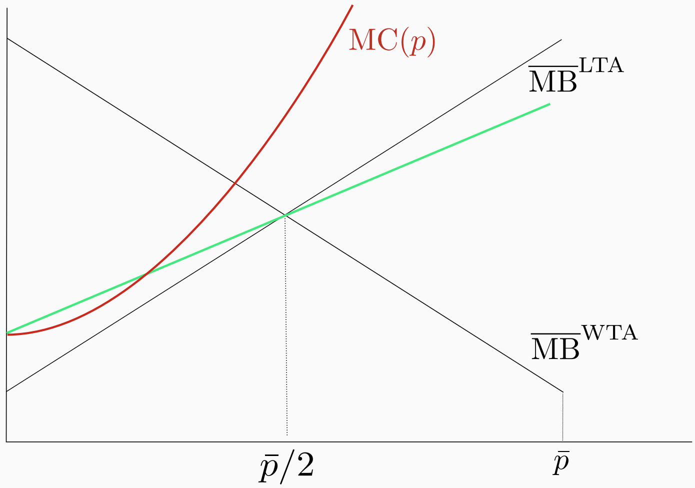

Hybrid Contests: Marginal Benefit

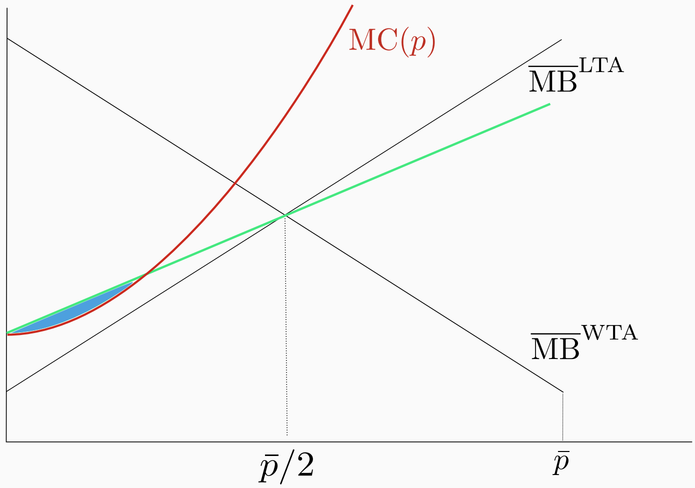

Even Better

The "optimal mixture" of LTA and WTA varies over time. How can we design a time-varying hybrid contest?

Time-Varying Hybrid Contest

Suppose the principal waits until the first success.

At that point she randomly determines whether the contest is LTA or WTA.

Define $g(s)$ to be the probability, if the first success arrives at "time" s, that the contest will be LTA.

Minmax Marginal Benefit for the g-Hybrid Contest.

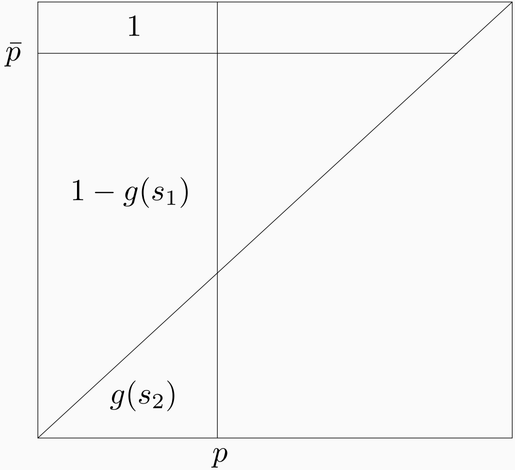

Minmax Marginal Benefit for the g-Hybrid Contest.

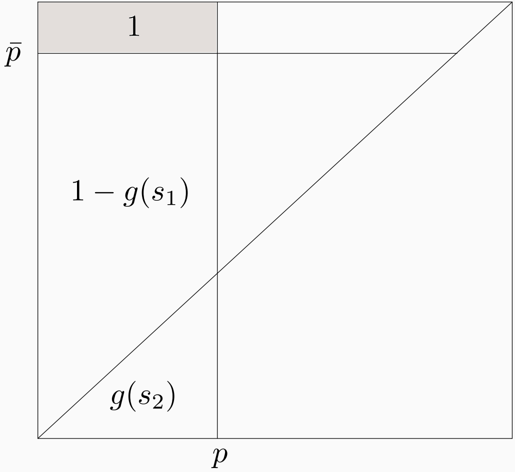

$p(1-\bar{p})$

Minmax Marginal Benefit for the g-Hybrid Contest.

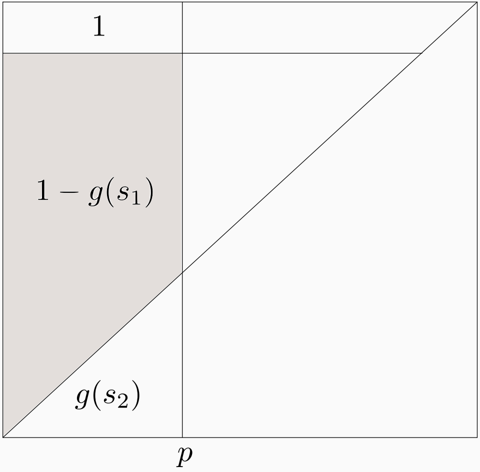

$$p(1-\bar{p}) + \int_0^p (1-s)(1-g(s))ds$$

Minmax Marginal Benefit for the g-Hybrid Contest.

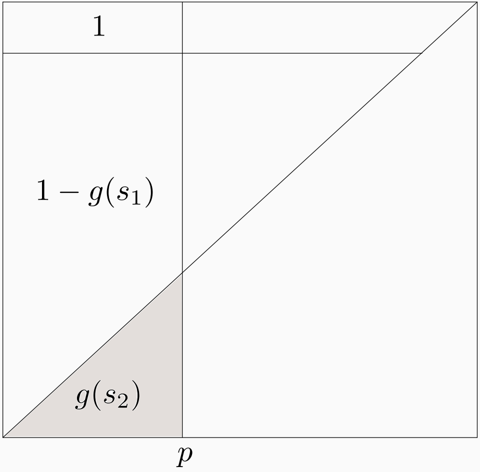

$$p(1-\bar{p}) + \int_0^p (1-s)(1-g(s))ds + \int_0^p \left[\int_0^r g(s)ds\right] dr$$

Minmax Marginal Benefit for the g-Hybrid Contest.

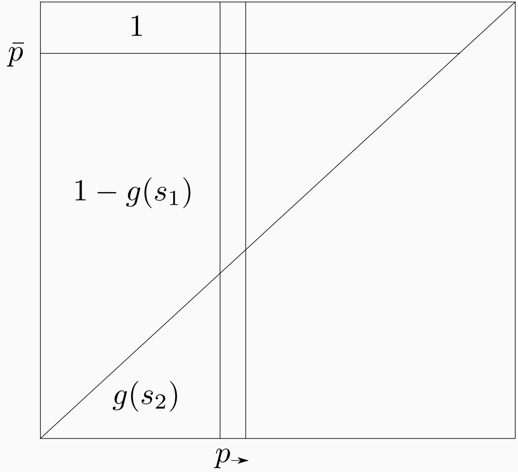

Minmax Marginal Benefit for the g-Hybrid Contest.

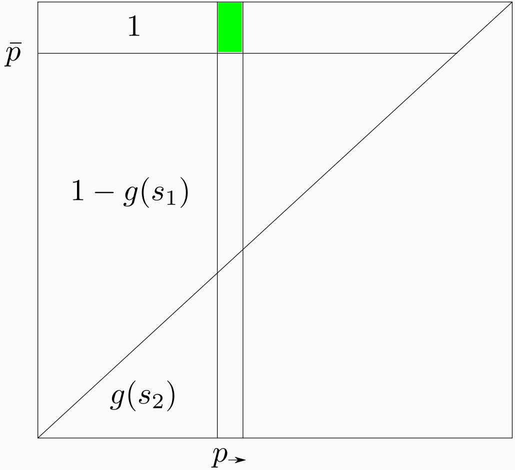

$(1 - \bar{p})$

Minmax Marginal Benefit for the g-Hybrid Contest.

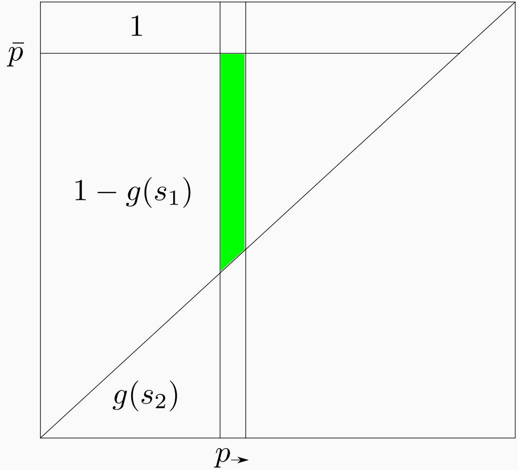

$$(1 - \bar{p}) + (\bar{p} - p)(1 - g(p))$$

Minmax Marginal Benefit for the g-Hybrid Contest.

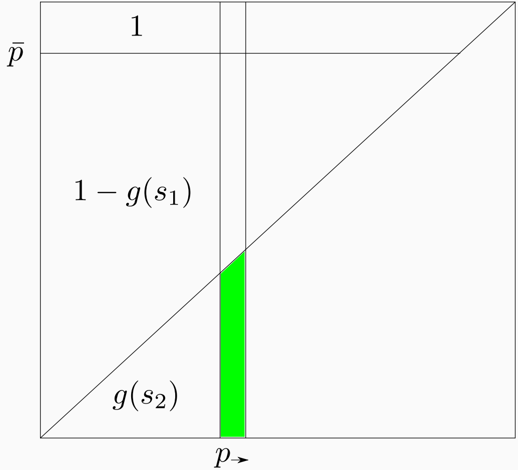

$$(1 - \bar{p}) + (\bar{p} - p)(1 - g(p)) + \int_0^p g(s)ds.$$

Minmax Marginal Benefit for the g-Hybrid Contest.

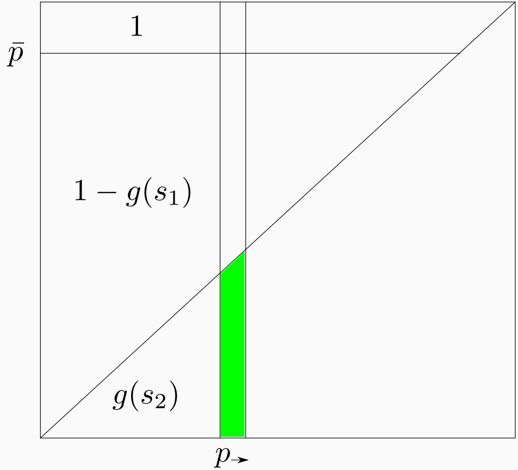

$$(1-\bar{p}) + (\bar{p} - p)(1 - g(p)) + G(p).$$

Stocks and Flows

We cannot perfectly fine-tune the flow of incentives without constraining the range of future incentives.

$$\overline{\text{MB}} = (1 - \bar{p}) + (\bar{p} - p)(1 - g(p)) + G(p)$$

where

$$G(p) = \int_0^p g(s)ds$$

is the stock of incentives and $g(p)$ is the flow.

Differential Equation

We can try to fine-tune the incentives to track marginal cost:

$$\overline{\mathrm{MB}}(p) = (1 - \bar{p}) + (\bar{p} - p)(1 - g(p)) + G(p) = \mathrm{MC}(p)$$

or

$$g(p)\,(1 - \bar{p}) + (1 - g(p))\,(1 - p) = \mathrm{MC}(p) - G(p)$$

Solution

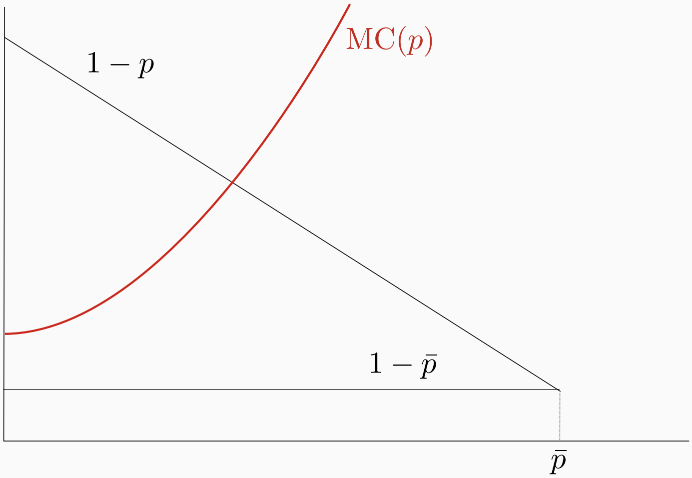

$$g(p)\,(1-\bar{p}) + (1-g(p))\,(1-p) = \mathrm{MC}(p) - G(p)$$

Solution

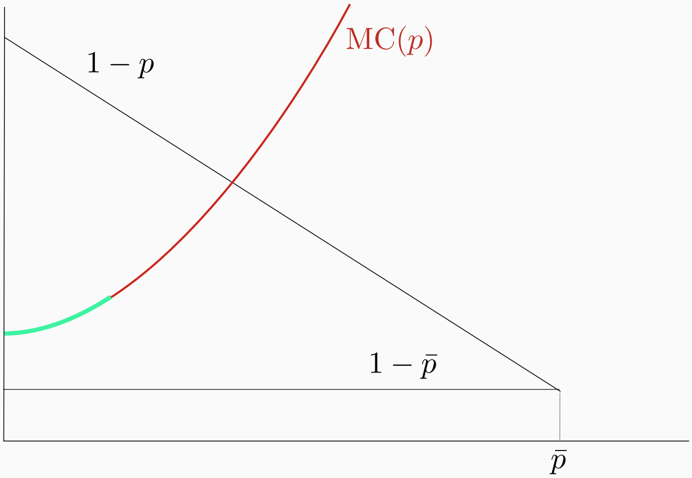

$$g(p)\,(1-\bar{p}) + (1-g(p))\,(1-p) = \mathrm{MC}(p)$$

Solution

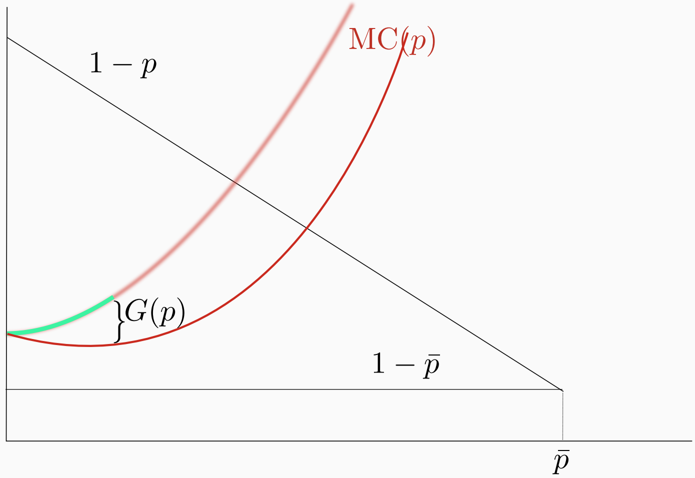

$$G(p) = \int_0^p g(s)ds$$

Solution

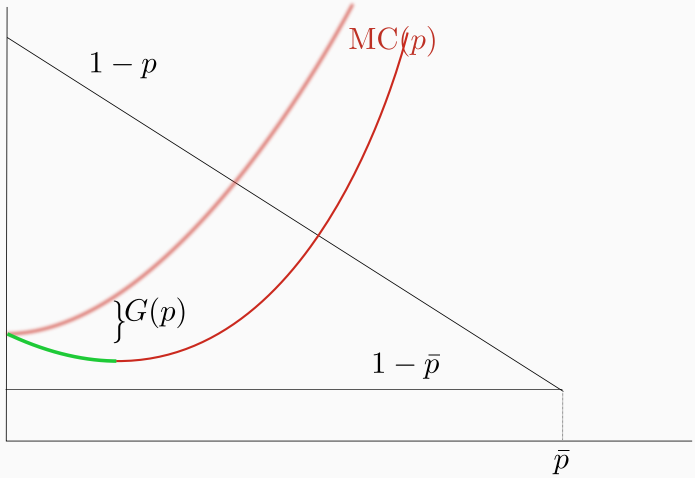

$$g(p)\,(1-\bar{p}) + (1-g(p))\,(1-p) = \mathrm{MC}(p) - G(p)$$

Solution

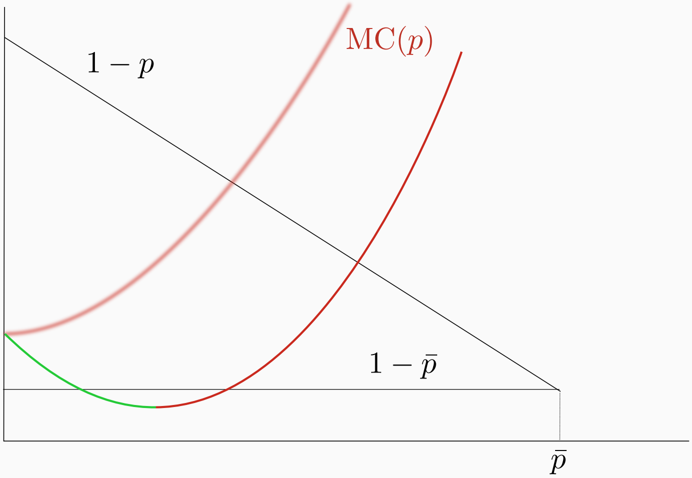

$g(p) > 1$

Solution

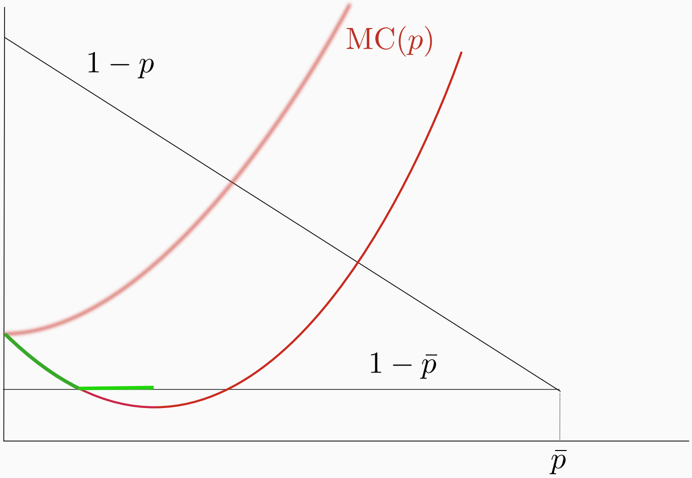

$g(p) \leq 1$ constraint binding.

Solution

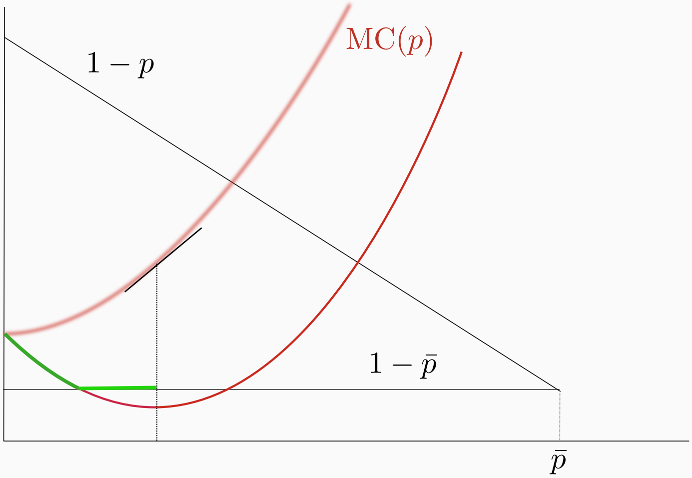

When $\frac{d}{dp}\,\mathrm{MC}(p) = 1$, we bottom out.

Solution

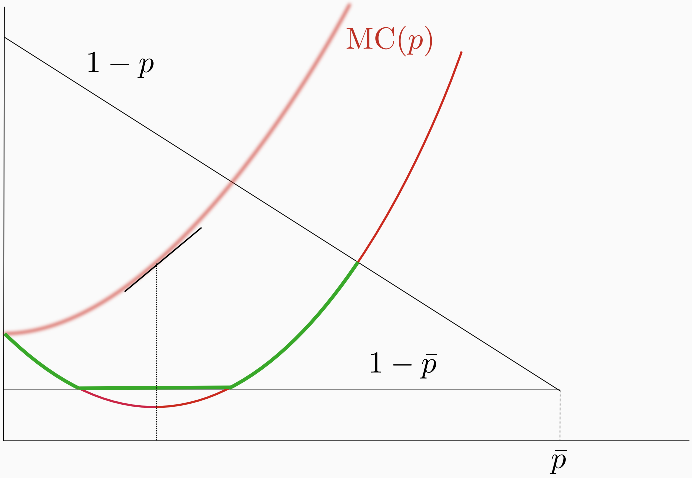

Constraint is slack again.

Solution

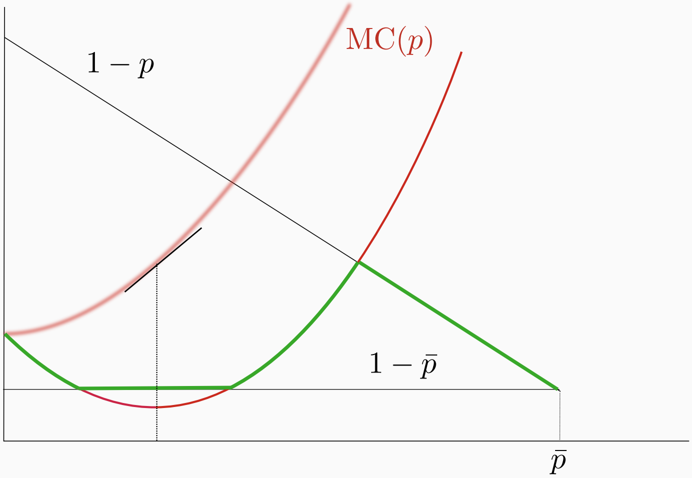

Finally $g(p) \geq 0$ constraint binds and we are done.

Solution

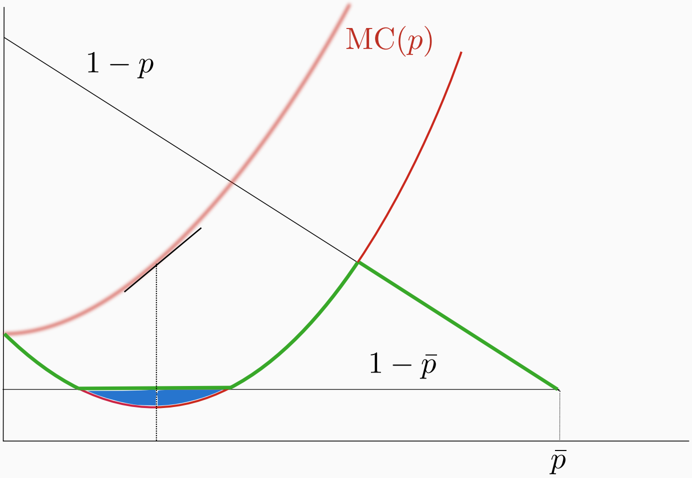

Rents

Solution

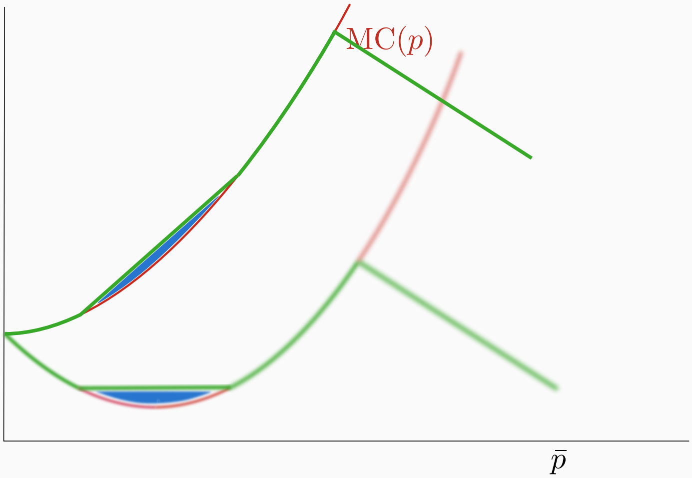

The solution: $\overline{\text{MB}} = (1 - \bar{p}) + (\bar{p} - p)(1 - g(p)) + G(p)$

Summing Up

- The differential equation yields the rent-minimizing hybrid contest (mixing LTA and WTA.)
- Conjecture: You cannot do better than to mix these two contests.
- In the background we have been assuming that the contest is efficient.
- Information Design/Leading on.

Disclosure Policy for LTA

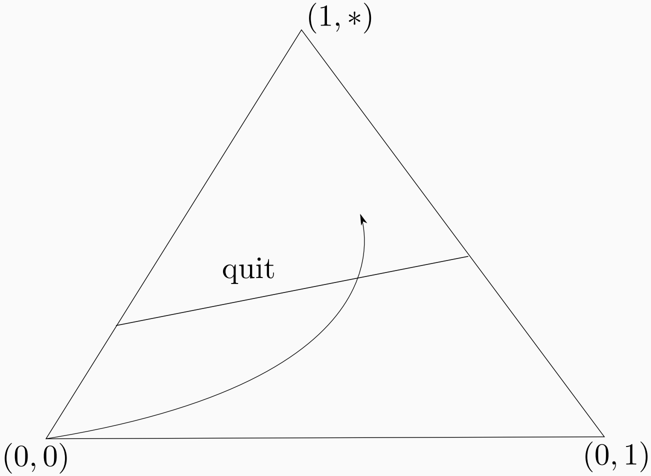

Disclosure Policy for LTA

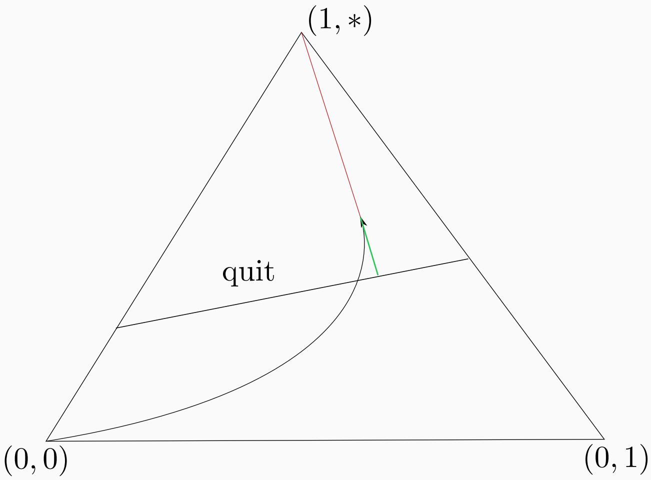

Disclosure Policy for LTA

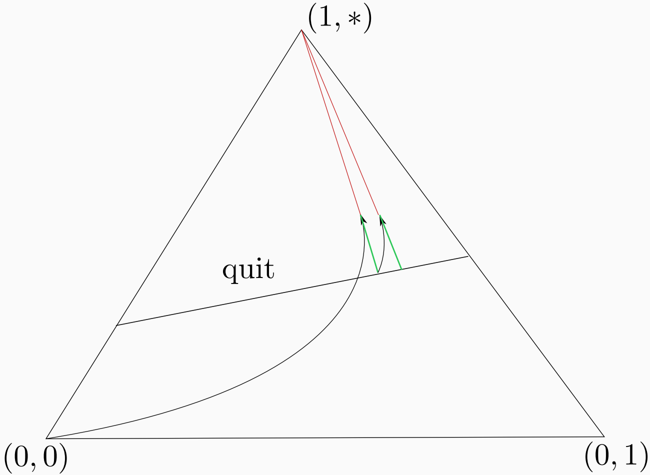

Disclosure Policy for LTA

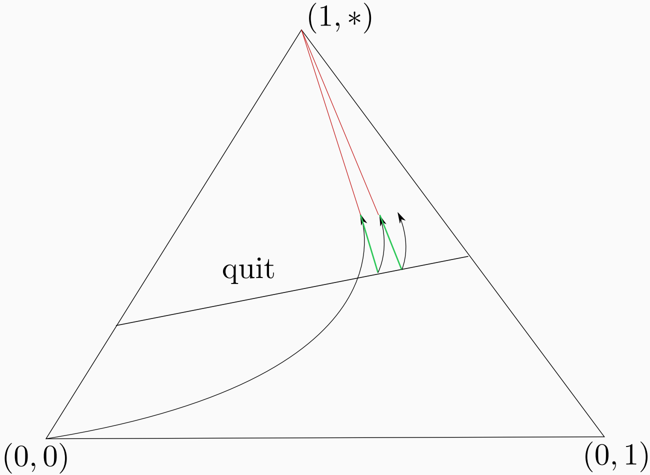

Reversal Example: Independent Contests

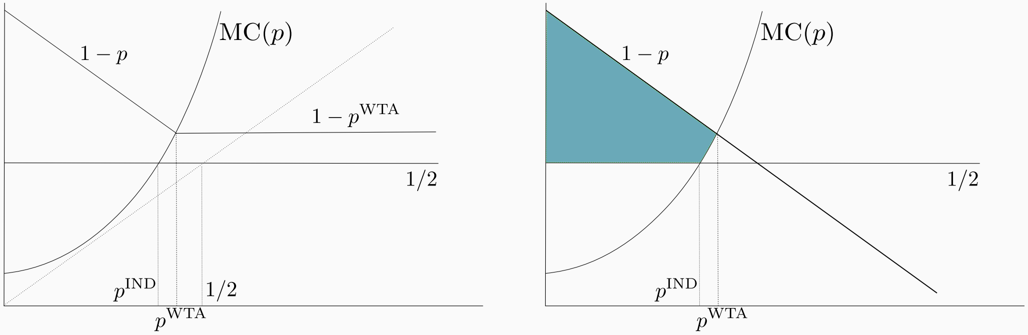

(a) No Disclosure **(b)** Optimal Disclosure