Optimal Feedback in Contests

George Georgiadis (Northwestern Kellogg)

with

Jeff Ely • Sina Moghadas Khorasani • Luis Rayo

Model (1/4): Players & Timing

- *Players:* A principal and $n \geq 2$ agents
- At $t = 0$, the principal designs a contest comprising
  i. a termination rule τ,
  ii. a rule for allocating a \$1 prize, and
  iii. a feedback policy.
- At every $t > 0$, each agent
  - receives a message per the feedback policy, and
  - chooses effort $a_{i,t} \in [0, 1]$
- The contest ends at τ and prize is awarded according to allocation rule

Model (2/4): Agents' Output & "Who observes what"

- **Each agent's output** takes the form of a Poisson "breakthrough":
  - During $(t, t+dt)$ agent i "succeeds" with probability $a_{i,t}dt$
  - Each agent can succeed at most once
  - Denote $x_{i,t} = 1$ if agent i has succeeded by t, and $x_{i,t} = 0$ otherwise
- **Who observes what:**
  - Principal observes successes but not efforts
  - Each agent observes his effort but not successes
- Denote by $p_{i,t}$ agent i's belief at t that he has succeeded
  - *Note:* Effort is worthwhile for an agent only if he hasn't yet succeeded

Model (3/4): Principal's Choice Variables

- A **termination rule** τ is a stopping time w.r.t $\mathbf{x}_t = \{x_{1,t}, \ldots x_{n,t}\}_{s \le t}$
  - *e.g.*, if $\tau = \inf\{t : x_{i,t} = 1 \text{ some } i\}$, contest ends upon first success
- A **prize allocation rule q** $\in [0,1]^n$ specifies the probability each agent wins the prize as function of $\mathbf{x}_\tau$; *i.e.*, each agent's time of success
  - *e.g.*, if $q_i(\mathbf{x}_\tau) = \mathbb{I}_{\{x_{i,t} \ge x_{j,t} \, \forall j,t\}}$, first agent to succeed wins prize w.p 1
- A **feedback policy** $\mathcal{M}$ specifies the message sent to each agent at every t as a function of $\mathbf{x}_t$ and past messages
  - *e.g.*, if $m_{i,t} = x_{i,t} \ \forall i, t$, ea. agent is told whether he has succeeded
  - *Alternatives:* Random feedback, feedback about others' successes, feedback about feedback, etc

Model (4/4): Payoffs

- Given a contest, each agent's expected utility is

$$u_{i,t} = \max_{a_i \in [0,1]} \mathbb{E}\left[q_i(\mathbf{x}_\tau) - \int_t^\tau c a_{i,s} ds \right]$$

  - *First term:* Probability agent i wins the prize
  - *Second term:* Cost of effort where $c \in (1/n, 1)$
  - *BNE:* Each agent chooses effort optimally anticipating rivals' efforts

- Principal chooses a contest $\{\tau, \mathbf{q}, \mathcal{M}\}$ and effort recommendations to

$$\max_{\tau, \mathbf{q}, \mathcal{M}, \mathbf{a}} \mathbb{E}\left[\sum_{i=1}^{n} \int_0^\tau a_{i,t} dt \right]$$

s.t. $a_{i,t}$ is IC for all i, t.

A Motivating Example

- Consider a manager who uses a promotion, acting as the prize, to motivate a group of employees
- Each agent must clear some “bar” to be eligible for promotion
  - This “bar” is represented by a success in the model (hence agents can succeed only once)
- Agents don’t definitively know whether they have cleared said “bar”, but the principal can disclose this (or other) information
- Manager cares about aggregate effort (not clearing the bar per-se)
- **Question:** How to design contest to get the most effort for $1 prize?

Remarks

i. *No discounting.*
   - Model is equivalent to one in which players discount time at some rate, and the value of the prize appreciates at the same rate

ii. *Agents don't observe their own successes.*
   - Goal is to give the principal full control of the agents' information
   - In optimal contest, each agent is fully appraised of his own success; *i.e.*, main result would be unchanged if agents observed own successes

iii. *Constant hazard rate of success.*
   - Success during $(t, t + dt)$ depends only on effort during this interval
   - *Extension:* Arrival rate of success increases with past efforts

Outline of Results

Proposition 1: Optimal contest *without* feedback

- No messages permitted and contest ends at some deterministic T
- *Egalitarian allocation rule* is optimal: Each agent who succeeds by T wins the prize with equal probability

Outline of Results

Proposition 2: Optimal contest *with* feedback

- *Cyclical structure:*
  - Initially, principal sets provisional deadline T^*
  - If one or more agents succeed by T^*, contest ends at T^*
  - Otherwise, the deadline is extended to $t = 2T^*$
  - If no agent succeeds by $2T^*$, deadline extended until $3T^*$. And so on.
- When contest ends, prize is awarded according to egalitarian rule.
  - *i.e.*, every agent who succeeded wins prize with equal probability
- Agents are fully appraised of their own success. They are informed about their rivals' successes at $T^*, 2T^*, ..$
  - *i.e.*, if deadline is extended, then no one has succeeded yet
- This contest achieves the *first-best payoff* for the principal

Outline of Results

Proposition 3: Optimal contest with increasing hazard rate

- Effort today makes success tomorrow more likely
- Similar structure, except that each provisional deadline has a stochastic duration

No-feedback Contests

- First, we restrict attention to contests without feedback
  - No message transmission permitted (*i.e.*, no direct feedback)
  - Principal chooses a deterministic deadline T (*i.e.*, no indirect feedback)
- Fix a contest and for each agent, define **reward function**

$$R_{i,t} = \mathbb{E}\left[q_i(\mathbf{x}_T) \,\middle|\, dx_{i,t} = 1\right]$$

i.e., agent's expected reward conditional on succeeding at t

- The agent's payoff can thus be expressed as

$$u_{i,t} = \max_{a_{i,s}\in[0,1]} \int_t^T (1 - p_{i,s}) a_{i,s} R_{i,s} - c a_{i,s} \, ds$$

  - During $(t, t + dt)$, succeeds w.p $(1 - p_{i,t}) a_{i,t} dt$, in which case earns $R_{i,t}$,
  - and he incurs cost $c a_{i,t} dt$

Agents' Problem

- Fix an arbitrary deadline T and reward function $R_{i,t}$. Agent solves

$$u_{i,0} = \max_{a_{i,t} \in [0,1]} \int_0^T \left[(1 - p_{i,t}) R_{i,t} - c \right] a_{i,t} dt$$

$$\text{s.t. } \dot{p}_{i,t} = (1 - p_{i,t}) a_{i,t} \text{ with } p_{i,0} = 0$$

- On the constraint:
  - Evolution equation for $p_{i,t}$ follows from Bayes' rule
  - Captures fact that effort today lowers future probability of success
- Std. optimal control problem: Use Pontryagin's maximum principle

Agents' Problem: Incentive Compatibility

- *Today's talk:* Restrict attention to contests with $a_{i,t} = 1$ for all $[0, T]$

Lemma 1.

- Consider no-feedback contest w. deadline T and reward function $R_{i,t}$
- Effort $a_{i,t} = 1$ is incentive compatible for all $t \in [0, T]$ if and only if

$$\underbrace{e^{-t} R_{i,t}}_{MB \text{ at } t} \geq \underbrace{c}_{\text{direct } MC} + \underbrace{\int_t^T e^{-s} R_{i,s} ds}_{\text{strategic } MC} \text{ for all } t.$$

- 1^{st} *term:* Success arrives at rate e^{-t}, and reward is $R_{i,t}$
- 2^{nd} *term:* *(Direct)* marginal cost of effort
- 3^{rd} *term:* Success today eliminates possibility of success in the future

No-feedback Contest: Principal's Problem

- Optimal no-feedback contest solves the following problem:

$$\max_{T,\mathbf{q}} \; n\int_0^T 1\,dt$$

$$\text{s.t. } e^{-t}R_{i,t} \geq c + \int_t^T e^{-s}R_{i,s}ds \;\forall i,t$$

$$T \geq 0,\ \mathbf{q} \text{ is a feasible prize allocation rule}$$

where $R_{i,t} = \mathbb{E}\left[q_i(\mathbf{x}_T) \,\middle|\, dx_{i,t} = 1\right]$.

- The principal chooses
  - a terminal date T, and
  - a prize allocation rule $\mathbf{q}$

  to maximize aggregate effort s.t IC constraint.

- Restriction to symmetric contests with max. effort shown to be wolog

Optimal No-feedback Contest

- *Definition 1:* Egalitarian prize allocation rule (EGA)

$$\mathbf{q}_i^{ega}(\mathbf{x}_T) = \frac{x_{i,T}}{\sum_j x_{j,T}}$$

  - *i.e.*, every agent who succeeds wins the prize with equal probability

- *Definition 2:* $\widehat{T}$ is the unique solution of $1 - e^{-n\widehat{T}} = nc(e^{\widehat{T}} - 1)$
  - Given EGA & no feedback, this is longest max. effort is IC

Proposition 1.

- The optimal no-feedback contest has deadline $\widehat{T}$ and egalitarian prize allocation rule $\mathbf{q}^{ega}$.
- In equilibrium, each agent exerts maximum effort for all $t \in [0, \widehat{T}]$.

Optimal No-feedback Contest: Heuristic Derivation (1/3)

- *Observation #1:* $R_{i,t} = ce^T$ satisfies IC with equality for all t
  - Time-invariant & symmetric $R_{i,t}$ corresponds to EGA allocation rule
- *Observation #2:*
  - Recall $R_{i,t}$ is prob. agent wins prize conditional on succeeding at t
  - Given $R_{i,t}$, agent i wins the prize with probability $\int_0^T e^{-t} R_{i,t} dt$. So

$$\underbrace{\sum_i \int_0^T e^{-t} R_{i,t} dt}_{\Pr\{\text{prize is awarded}\}} \leq \underbrace{1 - e^{-nT}}_{\Pr\{\text{at least one agent succeeds}\}}$$

  - In other words, increasing $e^{-t} R_{i,t}$ entails an opportunity cost, and so the principal wants to minimize $e^{-t} R_{i,t}$ subject to satisfying IC.

Optimal No-feedback Contest: Heuristic Derivation (2/3)

- Consider alternative contest with $e^{-t}\widetilde{R}_{i,t} > e^{-t}R_{i,t}$ on some interval

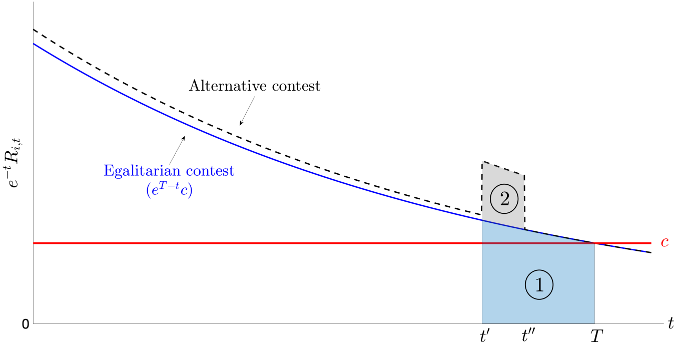

- *Egalitarian contest:* IC at t' requires that $e^{-t'}R_{i,t'} \geq c + (1)$
- *Alternative contest:* IC at t' requires that $e^{-t'}\widetilde{R}_{i,t'} \geq c + (1) + (2)$
- Thus $e^{-t}\widetilde{R}_{i,t} > e^{-t}R_{i,t}$ for all $t < t'$; i.e., $\widetilde{R}_{i,t}$ is more *expensive*

Optimal No-feedback Contest: Heuristic Derivation (3/3)

- Thus, any non-egalitarian contest with deadline T can be replaced by EGA contest with same deadline that is *cheaper* for principal
  - *Cheaper* $\Rightarrow$ Can extend deadline and still satisfy IC for all t
- It remains to pin down the optimal deadline $\widehat{T}$
  - Fix a T. Given the egalitarian allocation rule,

$$\Pr\{\text{agent } i \text{ wins prize}\} = \int_0^T e^{-t} R_{i,t}^{ega} dt = \frac{1 - e^{-nT}}{n}$$

  - Since $R_{i,t}^{ega}$ is time-invariant, we have $R_{i,t}^{ega} = [1 - e^{-nT}]/[n(1 - e^{-T})]$
  - By def. $\widehat{T}$ is largest deadline for which $R_{i,t}^{ega} \geq e^T c$; *i.e.*, max effort IC

Optimality of Egalitarian Contest: Intuition (1/2)

- As an alternative, take "winner-takes-all" contest with deadline T
  i.e., at T, the prize is awarded to the first agent who succeeded
- Assuming max. effort is IC on $[0, T]$, we have reward functions

$$R_{i,t}^{wta} = e^{-(n-1)t}$$

  i.e., if agent i succeeds at t, he is the first to do so w.p. $e^{-(n-1)t}$

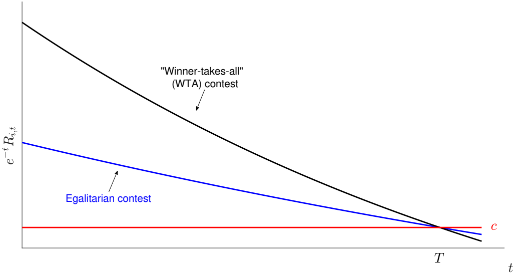

- Notice that $e^{-t}R_{i,t}^{wta} > e^{-t}R_{i,t}^{ega}$; *i.e.*, WTA is more expensive than EGA

Optimality of Egalitarian Contest: Intuition (2/2)

- The problem is that the WTA contest frontloads incentives too much
  - IC is slack for all $t < T$; i.e., incentives excessively strong early on
- In contrast, EGA (maximally) backloads incentives s.t IC binds $\forall t$

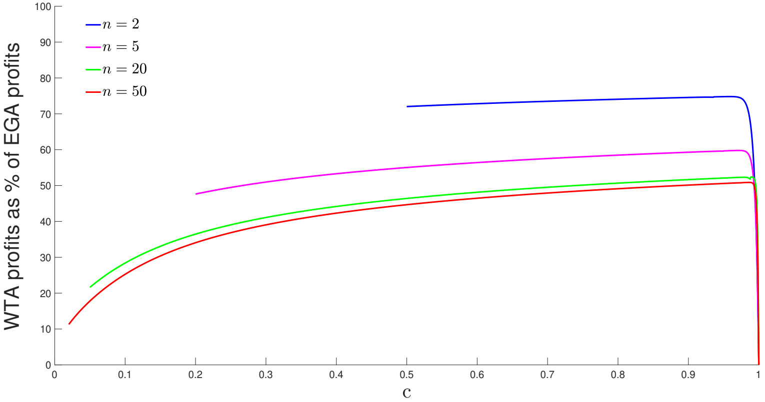

Key Lemma: Sufficient Condition for Optimality

- **Next, we consider contests with an arbitrary feedback policy**

Lemma 2:

A contest is guaranteed to be optimal if in equilibrium:

i. The prize is awarded with probability 1; *i.e.*, $\sum_i \mathbb{E}[q_i(\mathbf{x}_\tau)] = 1$

ii. Each agent earns zero rents; *i.e.*, $u_{i,0} = 0$ for all i

- The principal's object can be rewritten as

$$\mathbb{E}\left[\sum_{i=1}^{n} \int_0^\tau a_{i,t} dt\right] = \frac{1}{c}\left[\underbrace{\sum_i \mathbb{E}\left[q_i(\mathbf{x}_\tau)\right]}_{\Pr\{\text{prize awarded}\} \leq 1} - \underbrace{\sum_i u_{i,0}}_{\text{rents} \geq 0} \right]$$

- If a contest attains those bounds, it must be optimal (and first-best)

Step 1: Constructing a Zero-Rent Contest (1/2)

- We can write each agent's payoff as

$$u_{i,t} = \max_{a_{i,s}\in[0,1]} \int_t^{\tau} [(1 - p_{i,s})R_{i,s} - c]a_{i,s}ds$$

- For a contest to concede no rents to the agents,

$$(1 - p_{i,t})R_{i,t} = c \text{ for all } i, t$$

- *Claim:* Whenever $a_{i,t} > 0$, such a contest must have $p_{i,t} = 0$
  - Suppose there is an interval on which $\dot{p}_{i,t} > 0$ and $(1 - p_{i,t})R_{i,t} = c$
  - Agent can pause effort during first half of interval so $p_{i,t}^{private} < p_{i,t}^{eqm}$
  - Then $(1 - p_{i,t}^{private})R_{i,t} > c$, so agent can earn rents during second half
- Thus feedback policy must keep agents appraised of own success
  - Define the feedback policy $\mathcal{M}^{pronto} = \{m_{i,t} = x_{i,t} \text{ for all } i, t\}$

Step 2: Constructing a Zero-Rent Contest (2/2)

- Since $p_{i,t} = 0$ until each agent succeeds, contest must have

$$R_{i,t} = c \quad \text{for all } i, t$$

- For $R_{i,t}$ to be time-invariant & symmetric, alloc. rule must be EGA
- Suppose prize is awarded according to EGA rule at some fixed T
- *My* reward conditional on succeeding at t, $R_{i,t}$, depends on how many rivals I expect to succeed by T
  - This number $N_T \sim Binom(n-1, 1-e^{-T})$, and

$$R_{i,t}^{ega} = \mathbb{E}\left[\frac{1}{1+N_T}\right]$$

  - If $T \simeq 0$, no rivals will succeed a.s, so $R_{i,t}^{ega} \simeq 1$
  - As $T \to \infty$, all $n-1$ of my rivals will succeed a.s, so $R_{i,t}^{ega} \to 1/n$
  - There is a unique T^* such that $R_{i,t}^{ega} = c$

Step 3: Towards an Optimal Contest

- Consider the contest with:
  i. Deterministic deadline T^*
  ii. Egalitarian allocation rule
  iii. Feedback policy $\mathcal{M}^{pronto}$
- By construction,
  - $R_{i,t}^{ega} = c$ so ea agent exerts max. effort until he succeeds and $u_{i,t} = 0$
  - But the prize is awarded with probability $\sum_i \mathbb{E}\left[q_i(\mathbf{x}_\tau)\right] = 1 - e^{-nT^*} < 1$

i.e., this contest satisfies part (ii) of Lemma 2, but **not** part (i)

- Next, we amend this contest such that $\sum_i \mathbb{E}\left[q_i(\mathbf{x}_\tau)\right] = 1$
  - By Lemma 2, such contest will be optimal.

Step 3: Cyclical Structure

- Consider the (cyclical) termination rule:

$$\tau^* = \inf\left\{ t \,:\, t = kT^* \ ,\ k \in \mathbb{N} \text{ and } \sum_i x_{i,t} \geq 1 \right\}$$

- This contest comprises "cycles" of length T^*, and is terminated at the end of the first cycle in which one or more agents have succeeded
- Within each cycle, $R_{i,t}^{ega} = c$ by construction, so maximum effort is IC, and each agent's instantaneous payoff is 0. Thus, $u_{i,t} = 0$ for all t.
- Since the contest doesn't end until at least one agent succeeds, the prize is awarded with probability 1.
  - i.e., the contest satisfies conditions of Lemma 2, and is hence optimal

Optimal Contest (with Feedback)

Proposition 2.

- The following contest is optimal:
  - i'. termination rule $\tau^* = \inf\{t : t = kT^* \;, k \in \mathbb{N} \text{ and } \sum_i x_{i,t} \geq 1\}$,
  - ii. egalitarian prize allocation rule, and
  - iii. feedback policy $\mathcal{M}^{pronto}$
- In equilibrium, each agent exerts max. effort until he succeeds

- Intuition for cyclical structure:
  - If rivals exert max. effort during $[0, T]$, my expected reward conditional on succeeding $\downarrow T$ (because I will have to share prize with more rivals)
  - By construction, T^* is critical value such that $R_{i,t} = c$
  - Cycles inform agents noone has succeeded, "resetting" incentives

The Value of (optimal) Feedback

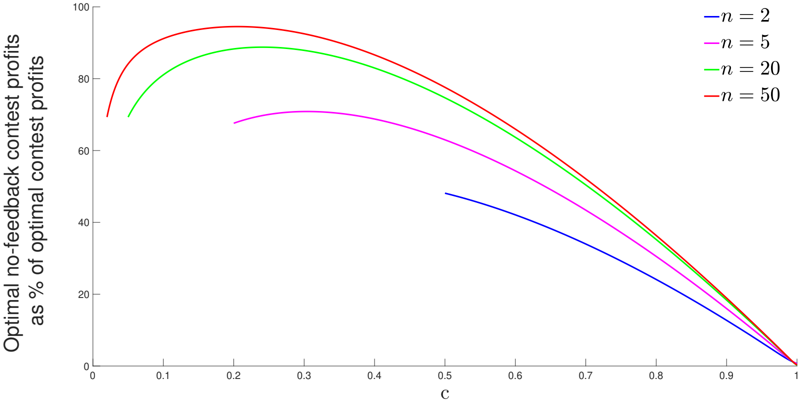

- Optimal feedback is most valuable when
  - Marginal costs c are small or large; *i.e.*, close to $1/n$ or 1, or
  - Number of agents n is small

Increasing Hazard Rate

- So far, we have assumed constant (unit) hazard rate of success
  - *i.e.*, agent succeeds during $(t, t + dt)$ with probability $a_{i,t}dt$
- Suppose instead that success arrives at rate $\lambda_{i,t}a_{i,t}$, and

$$\dot{\lambda}_{i,t} = f(\lambda_{i,t})a_{i,t}dt$$

for some function $f(\cdot)$ and $\lambda_{i,0} = \underline{\lambda}$.

I. Case $f(\lambda) < 0$: Effort today makes future success *less* likely
  - *e.g.*, Halac et al. (2017): "good news Poisson experimentation"

II. Case $f(\lambda) > 0$: Effort today makes future success *more* likely
  - Optimal contest has similar features & properties as in base model: it awards the prize with probability 1 and extracts all rents

Building Blocks

- *Assume:* $f(\lambda) \geq 0$ and satisfies $\lambda_{i,t} \in (c, nc)$
  - Suffices to assume $\underline{\lambda} > c$ and $f(\overline{\lambda}) = 0$ for some $\overline{\lambda} \in (\underline{\lambda}, nc)$
- Let λ_t^* solve $\dot{\lambda}_{i,t} = f(\lambda_{i,t})$ subject to $\lambda_{i,t} = \underline{\lambda}$
  - This is the trajectory of $\lambda_{i,t}$ if agent exerts max. effort
- By an earlier argument, feedback policy $\mathcal{M}^{pronto}$ to extract all rents
- For max. effort to be IC and rents to be 0, we must have

$$\lambda_t^* R_{i,t} = c \quad \text{for all } i, t$$

  - Because λ_t^* increases in t, $R_{i,t}$ must decrease in t

    i.e., incentives should be frontloaded since "earlier" success is "tougher"
- Suffices to find prize allocation and termination rules s.t $R_{i,t} = c/\lambda_t^*$

Optimal Contest

Proposition 3.

- There exists an optimal contest from the following class:
  1. *Cyclical stochastic structure:* Each cycle ends with rate $\gamma(t, \lambda_t)$
  2. At the end of each cycle, if a success has occurred, contest ends and prize is awarded according to EGA; otherwise, a new cycle starts
  3. Feedback policy $\mathcal{M}^{pronto}$; *i.e.*, agents appraised of own success
- In equilibrium, each agent exerts max. effort until he succeeds

- *i.e.*, similar structure to before, except cycles have stochastic length
- If $\gamma = \infty$, contest is "winner-takes-all"
- If $\gamma = 0$ for $t < T$ and $\gamma = \infty$ for $t \geq T$, the contest is egalitarian
- By choosing function $\gamma(t, \lambda_t)$, can fine-tune degree of frontloading

Related Literature

- Static tournaments / contests:
  - Lazear & Rosen ('81), Green & Stokey ('83), Nalebuff & Stiglitz ('83)
  - *Optimal prize allocation:* Moldovanu & Sela ('01), Drugov & Ryvkin ('18, '19), Olszewski and Siegel ('20)
  - *"Turning down the heat":* Fang et al. ('18) and Letina et al. ('20)
- Dynamic contests:
  - Taylor ('95), Benkert & Letina ('20)
  - *Tugs of war:* Moscarini & Smith ('11), Cao ('14)
- Feedback in contests:
  - *"Reveal intermediate progress?":* Yildirim ('05), Lizzeri et al. ('05), Aoyagi ('10), Ederer ('10), Goltsman & Mukherjee ('19)
  - *Contests for experimentation:* Halac et al. ('17)

Discussion

- Contest design with endogenous feedback
  - Cyclical structure
  - Egalitarian prize allocation rule (maximally backloads incentives)
  - Each agent is always appraised of own success, but is informed of rivals' successes periodically
- Future work
  - Continuous effort
  - Decreasing hazard rate
  - Continuous output / more general production functions
  - Asymmetric agents