Contests for Experimentation

Marina Halac   Navin Kartik   Qingmin Liu

September 2014

Introduction (1)

- Principal wants to obtain an innovation whose feasibility is uncertain
- Agents can work on or experiment with innovation
- Probability of success depends on state and agents' hidden efforts

Introduction (1)

- Principal wants to obtain an innovation whose feasibility is uncertain
- Agents can work on or experiment with innovation
- Probability of success depends on state and agents' hidden efforts

→ How should principal incentivize agents to experiment?

→ **This paper**: What is the optimal contest for experimentation?

Introduction (2)

- Long tradition of using contests to achieve specific innovations
  - more broadly, intellectual property and patent policy discussion
- Examples:
  - 1795 Napoleon govt offered a 12,000-franc prize for a food preservation method (winning idea: airtight sealing 1809).
  - Netflix contest: $1M to improve recommendation accuracy by 10%
  - Increased use in last two decades

Details

Introduction (2)

- Long tradition of using contests to achieve specific innovations
  - more broadly, intellectual property and patent policy discussion
- Examples:
  - 1795 Napoleon govt offered a 12,000-franc prize for a food preservation method (winning idea: airtight sealing 1809).
  - Netflix contest: $1M to improve recommendation accuracy by 10%
  - Increased use in last two decades Details
- Contests:
  - Not initially known if target attainable; contestants learn over time
  - Contestants' effort is unobservable $\Longrightarrow$ private learning
  - Contest architecture affects contestants' incentives to exert effort

Introduction (2)

- Long tradition of using contests to achieve specific innovations
  - more broadly, intellectual property and patent policy discussion
- Examples:
  - 1795 Napoleon govt offered a 12,000-franc prize for a food preservation method (winning idea: airtight sealing 1809).
  - Netflix contest: $1M to improve recommendation accuracy by 10%
  - Increased use in last two decades Details
- Contests:
  - Not initially known if target attainable; contestants learn over time
  - Contestants' effort is unobservable $\implies$ private learning
  - Contest architecture affects contestants' incentives to exert effort
- What contest design should be used?
  - Posit fixed budget and aim to max. prob. of one success
  - Propose tractable model based on exponential-bandit framework

Contest design

- Should Netflix award full prize to first successful contestant?
  - Intuit: Yes (under risk neutrality), sharing lowers expected reward

Contest design

- Should Netflix award full prize to first successful contestant?
  - Intuit: Yes (under risk neutrality), sharing lowers expected reward
- Should Netflix publicly announce when a first success is obtained?
  - Intuit: Yes, values only one success, hiding lowers expected reward

Contest design

- Should Netflix award full prize to first successful contestant?
  - Intuit: Yes (under risk neutrality), sharing lowers expected reward
- Should Netflix publicly announce when a first success is obtained?
  - Intuit: Yes, values only one success, hiding lowers expected reward

$\rightarrow$ Intuition says "public winner-takes-all" contest is optimal

$\rightarrow$ Indeed, dominates "hidden winner-takes-all" and "public shared-prize"

Contest design

- Should Netflix award full prize to first successful contestant?
  - Intuit: Yes (under risk neutrality), sharing lowers expected reward
- Should Netflix publicly announce when a first success is obtained?
  - Intuit: Yes, values only one success, hiding lowers expected reward

→ Intuition says "public winner-takes-all" contest is optimal

→ Indeed, dominates "hidden winner-takes-all" and "public shared-prize"

But will show that it is often dominated by "hidden shared-prize"

Main results

- Optimal info. disclosure policy (within a class) and prize scheme
- Conditions for optimality of Public WTA and Hidden Shared-Prize
  - Tradeoff: $\uparrow$ agent's reward for success versus $\uparrow$ his belief he will succeed
- More generally, a Mixture contest is optimal

Main results

- Optimal info. disclosure policy (within a class) and prize scheme
- Conditions for optimality of Public WTA and Hidden Shared-Prize
  - Tradeoff: $\uparrow$ agent's reward for success versus $\uparrow$ his belief he will succeed
- More generally, a Mixture contest is optimal
- Other issues
  1. Social planner may also prefer hidden shared-prize to public WTA
  2. Why a contest? Optimal contest dominates piece rates

Literature

Contest design (no learning)

- *Research contests*: Taylor 95, Krishna-Morgan 98, Fullerton-McAffee 99, Moldovanu-Sela 01, Che-Gale 03
- *Innovation contests*: Bhattacharya et al. 90, Moscarini-Smith 11, Judd et al. 12

Strategic experimentation games

- *Only info. externality*: Bolton-Harris 99, Keller et al. 05, ...
- *WTA contests*: Choi 91, Malueg-Tsutsui 97, Mason-Välimäki 10, Moscarini-Squintani 10, Akcigit-Liu 13
- *Other payoff externalities*: Strulovici 10, Bonatti-Hörner 11, Cripps-Thomas 14

Mechanism design for experimentation

- *Single-agent contracts*: Bergemann-Hege 98, 05, ...
- *Multiple agents & info. disclosure*: Che-Hörner 13, Kremer et al. 13

Model

Model (1)

Build on exponential bandit framework

- Innovation feasibility or state is either good or bad
  - Persistent but (initially) unknown; prior on good is $p_0 \in (0,1)$
- At each $t \in [0,T]$, agent $i \in \mathcal{N}$ covertly chooses effort $a_{i,t} \in [0,1]$
  - Instantaneous cost of effort is $ca_{i,t}$, where $c > 0$
  - $\mathcal{N} := \{1, \ldots, N\}$ is given; $T \geq 0$ will be chosen by principal
- If state is good and i exerts $a_{i,t}$, succeeds w/ inst. prob. $\lambda a_{i,t}$
  - No success if state is bad
  - Successes are conditionally independent given state

Model (2)

- Project success yields principal a payoff $v > 0$
  - Agents do not intrinsically care about success
  - Principal values only one success (specific innovation)
- Success is observable only to agent who succeeds and principal
  - Extensions: only agent or only principal observes success
- All parties are risk neutral and have quasi-linear preferences
  - Assume no discounting

Belief updating

- Given effort profile $\{a_{i,t}\}_{i,t}$, let p_t be the public belief at t, i.e. posterior on good state when no-one succeeds by t:

$$p_t = \frac{p_0 e^{-\int_0^t \lambda A_s ds}}{p_0 e^{-\int_0^t, \lambda A_s ds} + 1 - p_0}$$

where $A_t := \sum_j a_{j,t}$

Belief updating

- Given effort profile $\{a_{i,t}\}_{i,t}$, let p_t be the public belief at t, i.e. posterior on good state when no-one succeeds by t:

$$p_t = \frac{p_0 e^{-\int_0^t \lambda A_s ds}}{p_0 e^{-\int_0^t, \lambda A_s ds} + 1 - p_0}$$

  where $A_t := \sum_j a_{j,t}$

- Evolution of p_t governed by familiar differential equation

$$\dot{p}_t = -p_t \left(1 - p_t\right) \lambda A_t$$

First best

- Efficient to stop after success; hence, social optimum maximizes

$$\int_0^\infty (vp_t\lambda - c)\, A_t \overbrace{e^{-\int_0^t p_s \lambda A_s ds}}^{\text{Prob. no success by } t} dt$$

First best

- Efficient to stop after success; hence, social optimum maximizes

$$\int_0^\infty (vp_t\lambda - c)\, A_t \overbrace{e^{-\int_0^t p_s \lambda A_s ds}}^{\text{Prob. no success by } t} dt$$

- Since p_t decreasing, an efficient effort profile is $a_{i,t} = 1$ for all $i \in \mathcal{N}$ if $p_t\lambda v \geq c$ and no success by t; $a_{i,t} = 0$ for all $i \in \mathcal{N}$ otherwise

- Assume $p_0\lambda v > c$. First-best stopping posterior belief is

$$p^{FB} := \frac{c}{\lambda v}$$

Principal's problem

- Principal has a budget $\overline{w}$; assume $p_0 \lambda \overline{w} > c$
- Maximizes amount of experimentation:

$$p_0 \left(1 - e^{-\int_0^T \lambda A_t dt}\right)$$

Principal's problem

- Principal has a budget $\overline{w}$; assume $p_0 \lambda \overline{w} > c$
- Maximizes amount of experimentation:

$$p_0 \left(1 - e^{-\int_0^T \lambda A_t dt}\right)$$

- Mechanisms: payment rules and dynamic disclosure policies
  - s.t. limited liability & (ex-post) budget constraint

▸ Mechanisms

- Contests: **Subclass** of mechanisms

Contests

- A contest specifies

1. Deadline: $T \geq 0$
2. Prizes: $w(s_i, \boldsymbol{s}_{-i}) \geq 0$, where s_i is time at which i succeeds, s.t.
   (i) Anonymity: $w(s_i, \boldsymbol{s}_{-i}) = w(s_i, \sigma(\boldsymbol{s}_{-i}))$ for any permutation σ
   (ii) Wlog, 0 prize for no success: $w(\emptyset, \cdot) = 0$
3. Disclosure: $\mathcal{T} \subseteq [0, T]$ where outcome-history is publicly disclosed at each $t \in \mathcal{T}$ and nothing is disclosed at $t \notin \mathcal{T}$
   - Salient cases: public ($\mathcal{T} = [0, T]$) and hidden ($\mathcal{T} = \emptyset$)

Contests

- A contest specifies
  1. Deadline: $T \geq 0$
  2. Prizes: $w(s_i, \boldsymbol{s}_{-i}) \geq 0$, where s_i is time at which i succeeds, s.t.
     (i) Anonymity: $w(s_i, \boldsymbol{s}_{-i}) = w(s_i, \sigma(\boldsymbol{s}_{-i}))$ for any permutation σ
     (ii) Wlog, 0 prize for no success: $w(\emptyset, \cdot) = 0$
  3. Disclosure: $\mathcal{T} \subseteq [0, T]$ where outcome-history is publicly disclosed at each $t \in \mathcal{T}$ and nothing is disclosed at $t \notin \mathcal{T}$
     - Salient cases: public ($\mathcal{T} = [0, T]$) and hidden ($\mathcal{T} = \emptyset$)
- Strategies & Equilibrium
  - Wlog, $a_{i,t}$ is i's effort at t conditional on i not having succeeded by t
  - (Symmetric) Nash equilibria; refinements would not alter analysis

Public WTA Contest

Public winner-takes-all contest

- Let $A_{-i,s}$ be (i's conjecture of) total effort by agents $-i$ at s given no success by s. Then i's problem reduces to

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T (\overline{w}p_{i,t}\lambda - c)\, a_{i,t} \overbrace{e^{-\int_0^t p_{i,s}\lambda(a_{i,s}+A_{-i,s})ds}}^{\text{prob. no one succeeds by } t} dt$$

where

$$p_{i,t} = \frac{p_0 e^{-\int_0^t \lambda(a_{i,s}+A_{-i,s})ds}}{p_0 e^{-\int_0^t \lambda(a_{i,s}+A_{-i,s})ds} + 1 - p_0}$$

Public winner-takes-all contest

- Let $A_{-i,s}$ be (i's conjecture of) total effort by agents $-i$ at s given no success by s. Then i's problem reduces to

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T (\overline{w}p_{i,t}\lambda - c)\, a_{i,t} \overbrace{e^{-\int_0^t p_{i,s}\lambda(a_{i,s}+A_{-i,s})ds}}^{\text{prob. no one succeeds by } t}\, dt$$

where

$$p_{i,t} = \frac{p_0 e^{-\int_0^t \lambda(a_{i,s}+A_{-i,s})ds}}{p_0 e^{-\int_0^t \lambda(a_{i,s}+A_{-i,s})ds} + 1 - p_0}$$

- $p_{i,t} \downarrow \implies$ unique solution: $a_{i,t} = \begin{cases} 1 & \text{if } p_{i,t} \geq p^{PW} \\ 0 & \text{otherwise} \end{cases}$

where $p^{PW} := \dfrac{c}{\lambda \overline{w}}$

Public winner-takes-all contest

- For any T, unique equilibrium: all agents exert $a_{i,t} = 1$ until either a success occurs or public belief reaches p^{PW} (or T binds), then stop

Public winner-takes-all contest

- For any T, unique equilibrium: all agents exert $a_{i,t} = 1$ until either a success occurs or public belief reaches p^{PW} (or T binds), then stop
- Deadline T is optimal iff $T \geq T^{PW}$, where

$$\frac{p_0 e^{-N\lambda T^{PW}}}{p_0 e^{-N\lambda T^{PW}} + 1 - p_0} = \frac{c}{\lambda \overline{w}}$$

Public winner-takes-all contest

- For any T, unique equilibrium: all agents exert $a_{i,t} = 1$ until either a success occurs or public belief reaches p^{PW} (or T binds), then stop

- Deadline T is optimal iff $T \geq T^{PW}$, where

$$\frac{p_0 e^{-N\lambda T^{PW}}}{p_0 e^{-N\lambda T^{PW}} + 1 - p_0} = \frac{c}{\lambda \overline{w}}$$

- *Remark*: Amount of experimentation is invariant to N

Hidden WTA Contest

Hidden winner-takes-all contest

- Now i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \Big(\overline{w} p_{i,t}^{(1)} \lambda \underbrace{e^{-\int_0^t \lambda A_{-i,s} ds}}_{\substack{\text{prob. all } -i \text{ fail} \\ \text{until } t \text{ given } G}} - c\Big) a_{i,t} \overbrace{e^{-\int_0^t p_{i,s}^{(1)} \lambda a_{i,s} ds}}^{\substack{\text{prob. } i \text{ does not} \\ \text{succeed by } t}} dt,$$

where $p_{i,t}^{(1)}$ is i's private belief given he did not succeed by t:

$$p_{i,t}^{(1)} = \frac{p_0 e^{-\int_0^t \lambda a_{i,s} ds}}{p_0 e^{-\int_0^t \lambda a_{i,s} ds} + 1 - p_0}$$

Hidden winner-takes-all contest

- Unique solution: $a_{i,t} = \begin{cases} 1 & \text{if } \overline{w} p_{i,t}^{(1)} \lambda e^{-\int_0^t \lambda A_{-i,s} ds} \geq c \\ 0 & \text{otherwise} \end{cases}$

- Unique equilibrium is symmetric

Hidden winner-takes-all contest

- Unique solution: $a_{i,t} = \begin{cases} 1 & \text{if } \overline{w} p_{i,t}^{(1)} \lambda e^{-\int_0^t \lambda A_{-i,s} ds} \geq c \\ 0 & \text{otherwise} \end{cases}$
- Unique equilibrium is symmetric
- The stopping time T^{HW} is given by

$$\frac{p_0 e^{-N\lambda T^{HW}}}{p_0 e^{-\lambda T^{HW}} + 1 - p_0} = \frac{c}{\lambda \overline{w}}$$

Hidden winner-takes-all contest

- Unique solution: $a_{i,t} = \begin{cases} 1 & \text{if } \overline{w} p_{i,t}^{(1)} \lambda e^{-\int_0^t \lambda A_{-i,s} ds} \geq c \\ 0 & \text{otherwise} \end{cases}$

- Unique equilibrium is symmetric

- The stopping time T^{HW} is given by

$$\frac{p_0 e^{-N\lambda T^{HW}}}{p_0 e^{-\lambda T^{HW}} + 1 - p_0} = \frac{c}{\lambda \overline{w}} = \frac{p_0 e^{-N\lambda T^{PW}}}{p_0 e^{-N\lambda T^{PW}} + 1 - p_0}$$

- Hence, $T^{HW} < T^{PW}$ $\rightarrow$ Strictly dominated by public WTA

Public Shared-Prize Contests

Public shared-prize contests

- Now i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \left[\left(w_{i,t} p_{i,t}\lambda - c\right) a_{i,t} + p_{i,t}\lambda A_{-i,t} u_{i,t}\right] \overbrace{e^{-\int_0^t p_{i,s}\lambda(a_{i,s}+A_{-i,s})ds}}^{\text{prob. no one succeeds by } t} dt$$

  where $w_{i,t}$ is i's expected reward if he succeeds at t and $u_{i,t}$ is his continuation payoff if some $-i$ succeeds at t

  - dependence on strategies suppressed

Public shared-prize contests

- Now i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \left[\left(w_{i,t} p_{i,t}\lambda - c\right) a_{i,t} + p_{i,t}\lambda A_{-i,t} u_{i,t}\right] \overbrace{e^{-\int_0^t p_{i,s}\lambda(a_{i,s}+A_{-i,s})ds}}^{\text{prob. no one succeeds by } t} dt$$

  where $w_{i,t}$ is i's expected reward if he succeeds at t and $u_{i,t}$ is his continuation payoff if some $-i$ succeeds at t

  - dependence on strategies suppressed

- Since $u_{i,t} \geq 0$

$$a_{i,t} > 0 \implies p_{i,t} \geq \frac{c}{w_{i,t}\lambda}$$

Public shared-prize contests

- Now i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \left[\left(w_{i,t} p_{i,t}\lambda - c\right) a_{i,t} + p_{i,t}\lambda A_{-i,t} u_{i,t}\right] \overbrace{e^{-\int_0^t p_{i,s}\lambda(a_{i,s}+A_{-i,s})ds}}^{\text{prob. no one succeeds by } t} dt$$

  where $w_{i,t}$ is i's expected reward if he succeeds at t and $u_{i,t}$ is his continuation payoff if some $-i$ succeeds at t

  - dependence on strategies suppressed

- Since $u_{i,t} \geq 0$ and $w_{i,t} \leq \overline{w}$,

$$a_{i,t} > 0 \implies p_{i,t} \geq \frac{c}{w_{i,t}\lambda} \geq \frac{c}{\overline{w}\lambda} = p^{PW}$$

Public shared-prize contests

- Now i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \left[\left(w_{i,t}p_{i,t}\lambda - c\right) a_{i,t} + p_{i,t}\lambda A_{-i,t}u_{i,t}\right] \overbrace{e^{-\int_0^t p_{i,s}\lambda(a_{i,s}+A_{-i,s})ds}}^{\text{prob. no one succeeds by } t} dt$$

  where $w_{i,t}$ is i's expected reward if he succeeds at t and $u_{i,t}$ is his continuation payoff if some $-i$ succeeds at t

  - dependence on strategies suppressed

- Since $u_{i,t} \geq 0$ and $w_{i,t} \leq \overline{w}$,

$$a_{i,t} > 0 \implies p_{i,t} \geq \frac{c}{w_{i,t}\lambda} \geq \frac{c}{\overline{w}\lambda} = p^{PW}$$

$\rightarrow$ Dominated by public WTA (strictly if different)

Hidden Shared-Prize Contests

Hidden shared-prize contest

Proposition

Among hidden contests, an optimal prize scheme is equal sharing: for any number of successful agents $n \in \mathcal{N}$, $w_i = \frac{\overline{w}}{n} \; \forall i \in \{1, \ldots, n\}$.

Hidden shared-prize contest

Proposition

Among hidden contests, an optimal prize scheme is equal sharing: for any number of successful agents $n \in \mathcal{N}$, $w_i = \frac{\overline{w}}{n}\ \forall i \in \{1, \ldots, n\}$.

- Idea of Proof:
  - Without loss to consider a prize regime that induces full effort equilibrium
  - Equal sharing implies constant sequence of expected rewards and stopping time T^{HS} s.t. agent's IC constraint binds at each $t \in [0, T^{HS}]$
  - Hence, cannot induce more experimentation with non-constant reward sequence (if $T > T^{HS}$, IC constraint is violated at some $t \leq T$)

Hidden equal-sharing contest

- Under **equal sharing**, i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \left(w_i p_{i,t}^{(1)} \lambda - c \right) a_{i,t} \overbrace{e^{-\int_0^t p_{i,s}^{(1)} \lambda a_{i,s} ds}}^{\text{prob. } i \text{ does not succeed by } t} dt$$

Hidden equal-sharing contest

- Under **equal sharing**, i's problem is

$$\max_{(a_{i,t})_{t\in[0,T]}} \int_0^T \left(w_i p_{i,t}^{(1)} \lambda - c \right) a_{i,t} \overbrace{e^{-\int_0^t p_{i,s}^{(1)} \lambda a_{i,s} ds}}^{\text{prob. } i \text{ does not succeed by } t} dt$$

- An optimal strategy is $a_{i,t} = 1$ if $w_i p_{i,t}^{(1)} \lambda \geq c$ and $a_{i,t} = 0$ otherwise

- Consider symmetric eqa characterized by stopping time T^{HS}

Hidden equal-sharing contest

- Given T^{HS}, the expected reward for success is

$$w = \overline{w}\mathbb{E}_n \left[\frac{1}{n} \middle| n \geq 1, T^{HS} \right]$$

Hidden equal-sharing contest

- Given T^{HS}, the expected reward for success is

$$
\begin{aligned}
w &= \overline{w}\mathbb{E}_n\left[\frac{1}{n} \,\middle|\, n \geq 1, T^{HS}\right] \\
&= \overline{w}\sum_{m=0}^{N-1}\left(\frac{1}{m+1}\right)\binom{N-1}{m}\underbrace{\left(1-e^{-\lambda T^{HS}}\right)^m}_{\substack{\text{Prob. } m \text{ opponents} \\ \text{succeed by } T^{HS} \text{ in } G}}\ \underbrace{e^{-(N-1-m)\lambda T^{HS}}}_{\substack{\text{Prob. } N-1-m \\ \text{opponents fail in } G}}
\end{aligned}
$$

Hidden equal-sharing contest

- Given T^{HS}, the expected reward for success is

$$w = \overline{w}\mathbb{E}_n\left[\frac{1}{n}\,\middle|\, n \geq 1, T^{HS}\right]$$
$$= \overline{w}\sum_{m=0}^{N-1}\left(\frac{1}{m+1}\right)\binom{N-1}{m}\underbrace{\left(1-e^{-\lambda T^{HS}}\right)^m}_{\text{Prob. } m \text{ opponents succeed by } T^{HS} \text{ in } G}\ \underbrace{e^{-(N-1-m)\lambda T^{HS}}}_{\text{Prob. } N-1-m \text{ opponents fail in } G}$$

- Equilibrium T^{HS} solves

$$\underbrace{\overline{w}\ \frac{1-e^{-\lambda N T^{HS}}}{(1-e^{-\lambda T^{HS}})N}}_{\text{exp. reward}}\ \underbrace{\frac{p_0 e^{-\lambda T^{HS}}}{p_0 e^{-\lambda T^{HS}} + 1 - p_0}}_{\text{stop. private belief}}\ \lambda = c,$$

which has a unique solution; hence essentially unique symmetric eqm

- *Remark*: Amount of experimentation can be non-monotonic in N

Public WTA vs. Hidden Equal-Sharing

Public winner-takes-all versus hidden equal-sharing

- T^{PW} and T^{HS} satisfy respectively

$$\frac{p_0 e^{-N\lambda T^{PW}}}{p_0 e^{-N\lambda T^{PW}} + 1 - p_0} = \frac{c}{\lambda \overline{w}}$$

$$\frac{p_0 e^{-\lambda T^{HS}}}{p_0 e^{-\lambda T^{HS}} + 1 - p_0} \mathbb{E}_n \left[\frac{1}{n} \middle| n \geq 1, T^{HS} \right] = \frac{c}{\lambda \overline{w}}$$

Public winner-takes-all versus hidden equal-sharing

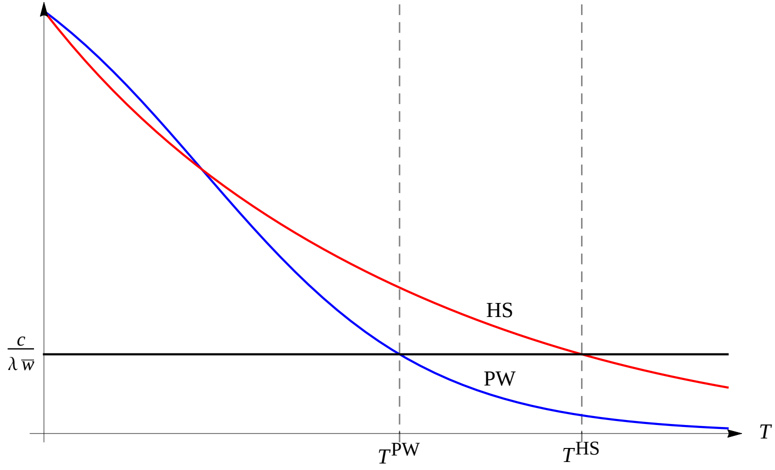

Result for public vs. hidden

Proposition

Among public and hidden contests, if

$$\frac{p_0 e^{-\lambda T^{PW}}}{p_0 e^{-\lambda T^{PW}} + 1 - p_0} \frac{1 - e^{-\lambda N T^{PW}}}{(1 - e^{-\lambda T^{PW}})N} > \frac{c}{\lambda \overline{w}}$$

then a hidden equal-sharing contest is optimal.

Otherwise, a public winner-takes-all contest is optimal.

Result for public vs. hidden

Proposition

Among public and hidden contests, if

$$\frac{p_0 e^{-\lambda T^{PW}}}{p_0 e^{-\lambda T^{PW}} + 1 - p_0} \frac{1 - e^{-\lambda N T^{PW}}}{(1 - e^{-\lambda T^{PW}})N} > \frac{c}{\lambda \overline{w}}$$

then a hidden equal-sharing contest is optimal.

Otherwise, a public winner-takes-all contest is optimal.

Note: If principal can choose N, HS can replicate PW by setting $N = 1$

Intuition: Necessary and sufficient conditions

- Condition for $N = 2$ is
$$\frac{\overline{w}}{2}\lambda > c$$
$\rightarrow$ i would continue experimenting to earn half prize if he knew state is good, or equivalently, if he knew opponent succeeded

Intuition: Necessary and sufficient conditions

- Condition for $N = 2$ is
$$\frac{\overline{w}}{2}\lambda > c$$
$\rightarrow$ i would continue experimenting to earn half prize if he knew state is good, or equivalently, if he knew opponent succeeded

- A sufficient condition for any $N > 2$ is
$$\frac{\overline{w}}{N}\lambda \geq c$$

Intuition: Discussion

- Relative to public WTA, why can hidden shared-prize help but neither public shared-prize nor hidden WTA can?
  - Want to hide info. to bolster agent's belief when no-one has succeeded
  - But hiding is counter-productive if WTA
    $\implies$ to harness benefits of hiding info., must share prize
  - Public shared-prize no help: only $\uparrow$ effort when it does not benefit P
    - and can hurt because of free-riding

Intuition: Discussion

- Relative to public WTA, why can hidden shared-prize help but neither public shared-prize nor hidden WTA can?
  - Want to hide info. to bolster agent's belief when no-one has succeeded
  - But hiding is counter-productive if WTA

    $\Longrightarrow$ to harness benefits of hiding info., must share prize
  - Public shared-prize no help: only $\uparrow$ effort when it does not benefit P
    - and can hurt because of free-riding

- Public WTA optimal if $p_0 = 1$ or arms uncorrelated
  - no learning **from others** $\Longrightarrow$ no benefit to hiding info
  - most patent design papers assume $p_0 = 1$ — hence "patent"

Other Disclosure Policies

Simple disclosure policies

- Principal specifies $\mathcal{T} \subseteq [0, T]$ such that outcome-history publicly disclosed at each $t \in \mathcal{T}$ and nothing disclosed at any $t \notin \mathcal{T}$

Simple disclosure policies

- Principal specifies $\mathcal{T} \subseteq [0, T]$ such that outcome-history publicly disclosed at each $t \in \mathcal{T}$ and nothing disclosed at any $t \notin \mathcal{T}$

Proposition

An optimal contest is a mixture contest that implements public winner-takes-all from 0 until t_S and hidden equal-sharing from t_S until T.

- Idea of Proof: Take arbitrary contest with disclosure $\mathcal{T}$ and let $t' = \sup\{t : t \in \mathcal{T}\}$. Dominated by mixture contest with $t_S = t'$

Simple disclosure policies

- Principal specifies $\mathcal{T} \subseteq [0, T]$ such that outcome-history publicly disclosed at each $t \in \mathcal{T}$ and nothing disclosed at any $t \notin \mathcal{T}$

Proposition

An optimal contest is a mixture contest that implements public winner-takes-all from 0 until t_S and hidden equal-sharing from t_S until T.

Moreover:

1. *If $\overline{w}\lambda/N > c$ then $t_S = 0$ (hidden equal-sharing).*
2. *If $\overline{w}\lambda/2 < c$ then $t_S = T$ (public WTA).*

- Idea of Proof: Take arbitrary contest with disclosure $\mathcal{T}$ and let $t' = \sup\{t : t \in \mathcal{T}\}$. Dominated by mixture contest with $t_S = t'$

Example: Optimal mixture contest

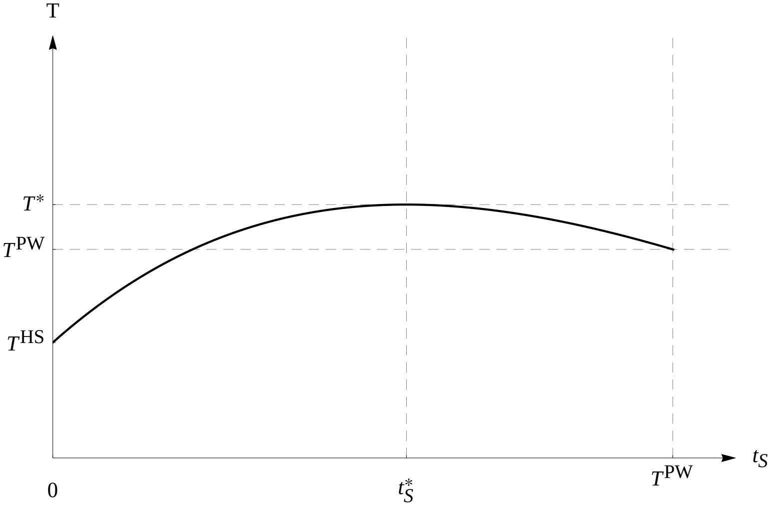

- $t_S \uparrow \Longrightarrow$ from t_S on, belief $\downarrow$ but expected reward $\uparrow$

Conclusions

- Tradeoff in incentivizing experimentation:
  $\uparrow$ agent's reward for success versus $\uparrow$ his belief that he will succeed

- Hidden equal-sharing often dominates public WTA (even for planner)
  - Only hiding info. or dividing prize hurts, but together can help

- Conditions for optimality of these contests

- Broader contributions
  1. contest design in an environment with learning
  2. mechanism design—payments and info. disclosure—to multi-agent strategic experimentation

Thank you!

Contests for experimentation

- R&D competition, patent races
- Increased use of contests to achieve specific innovations
  - McKinsey report: huge increase in large prizes in last 35 years. 78% of new prize money since 1991 is inducement for specific goals
  - New intermediaries such as Changemakers, Idea Crossing, X Prize
  - America Competes Reauthorization Act signed by Obama in 2011
- Many examples
  - British Parliament's longitude prize,
  - Orteig prize
  - X Prizes: Ansari, Google Lunar, Progressive Automotive
  - Methuselah Foundation: Mouse Prize, NewOrgan Liver Prize

Back

Mechanisms

- Principal has budget $\overline{w} > 0$ to incentivize agents' effort
  - Assume $p_0 \lambda \overline{w} > c$
- In general, a (limited-liability) mechanism specifies
  1. Deadline $T \geq 0$
  2. Vector of payments $(w_1, \ldots, w_N) \in \mathbb{R}_+^N$ that are made at T
     $\rightarrow$ as function of principal's info at T and subject to $\sum_{i \in \mathcal{N}} w_i \leq \overline{w}$
  3. Information disclosure policy (signal of history for each agent at each t)
- Strategy for i specifies $a_{i,t}$ for each t given i's information at t

◂ Contests

Observability of success

- If principal observes success but not agent, results readily extend
  - A will condition on failure; P has no reason to hide success from A

Observability of success

- If principal observes success but not agent, results readily extend
  - A will condition on failure; P has no reason to hide success from A
- More subtle: principal does not observe success directly; any agent who succeeds can choose when to verifiably reveal his success
  - Winner-takes-all: dominant for agent to reveal when succeeds
  - Hidden success: equal sharing optimal, outcome unchanged
  - Thus, under same condition, hidden ES dominates public WTA

Back

Implications for the planner's problem

- Hidden shared-prize contest can be optimal for principal who does not internalize effort costs. How about social planner?

Implications for the planner's problem

- Hidden shared-prize contest can be optimal for principal who does not internalize effort costs. How about social planner?
- Suppose social planner has only $\overline{w} < v$ to reward agents
  - Likely if social value of discovery high, e.g. medical innovations
- Then even social planner will sometimes prefers hidden equal-sharing, as public winner-takes-all induces less than efficient experimentation
  - Ex post, planner induces wasteful experimentation after discovery made

Back

Why contests?

- Instead of a contest, suppose principal uses piece rates
  - Payment to i, w_i, independent of others' outcomes, with $\sum_i w_i \leq \overline{w}$
  - Assume independent of time (just a bonus for success)

Why contests?

- Instead of a contest, suppose principal uses piece rates
  - Payment to i, w_i, independent of others' outcomes, with $\sum_i w_i \leq \overline{w}$
  - Assume independent of time (just a bonus for success)

Proposition

1. *Optimal piece rate has hidden success and pays $\frac{\overline{w}}{k^*}$ to each of $1 \leq k^* \leq N$ agents; zero to all others.*
2. *This piece rate dominates public winner-takes-all contest,*
   But is dominated by hidden equal-sharing contest if principal can choose N.

- Domination statements strict if $k^* > 1$

Back

Intuition: Contests versus piece rates

- A piece rate can implement the public winner-takes-all outcome
  - Pay $\overline{w}$ for success to one agent
- But gives less experimentation than hidden equal-sharing with k^*:
  - Stopping rule in optimal piece rate: $p_{i,T}^{(1)} \lambda \frac{\overline{w}}{k^*} = c$
  - Stopping rule in hidden equal-sharing: $\frac{1-e^{-k^* \lambda T}}{1-e^{-\lambda T}} p_{i,T}^{(1)} \lambda \frac{\overline{w}}{k^*} = c$

Back