Contests for Experimentation

Marina Halac   Navin Kartik   Qingmin Liu

November 2015

Introduction

- Principal wants to obtain an innovation whose feasibility is uncertain
- Agents can work on or experiment with this project
- 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?

Contests for experimentation

- Long tradition of using contests to achieve specific innovations
  - More broadly, intellectual property and patent policy debates
- Increased use in last two decades
  - Accounts for 78% of new prize money since 1991 (McKinsey)
  - America Competes Reauthorization Act signed by Obama in 2011
- Many examples
  - British Parliament's longitude prize
  - Napoleon's food preservation prize
  - Orteig prize
  - X Prizes: Ansari, Google Lunar, Progressive Automotive
  - Methuselah Foundation: Mouse Prize, NewOrgan Liver Prize

Contest design

- Netflix contest: \$1M to improve recommendation accuracy by 10%
  - Not initially known if target attainable; contestants learn over time
  - Contestants' effort is unobservable $\implies$ learning is private
  - Contest architecture affects contestants' incentives to exert effort
- What contest design should be used?
  - Given a prize, principal aims to maximize probability of success
  - Propose tractable model based on exponential-bandit framework

Contest design: Payments and info disclosure

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

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

→ Indeed, dominates *any other* public and *any other* winner-takes-all

Contest design: Payments and info disclosure

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

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

→ Indeed, dominates *any other* public and *any other* winner-takes-all

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

Main results

- Optimal info disclosure policy and prize scheme

- Conditions for optimality of hidden shared-prize and public WTA
  - Tradeoff: $\uparrow$ agent's reward for success *vs* $\uparrow$ his belief he will succeed

- More generally, sharing the prize with cutoff disclosure is optimal

Related literature

Contest design without learning

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

Innovation contests with learning

- *WTA*: Choi 91, Malueg-Tsutsui 97, Mason-Välimäki 10, Moscarini-Squintani 10
- *Contest design*: Bimpikis et al. 14, Moroni 15

Multi-agent strategic experimentation

- *Games*: Keller et al. 05, Keller-Rady 10, Bonatti-Hörner 11, Cripps-Thomas 14
- *Info. disclosure*: Bimpikis-Drakopoulos 14, Che-Hörner 14, Heidhues et al. 14, Kremer et al. 14, Akcigit-Liu 14

Model (1)

Build on exponential bandit model (Keller, Rady, and Cripps, 2005):

- 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 $c > 0$
  - $\mathcal{N} \equiv \{1,\ldots,N\}$ is given; $T \geq 0$ will be chosen by principal
- If state is good and i exerts $a_{i,t}$, succeeds with 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:

$$p_t = \frac{p_0 e^{-\int_0^t \lambda A_z dz}}{p_0 e^{-\int_0^t \lambda A_z dz} + 1 - p_0}$$

where $A_t \equiv a_{1,t} + ... + a_{N,t}$

- 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

$$\int_0^\infty (p_t \lambda v - c)\, A_t \overbrace{e^{-\int_0^t p_z \lambda A_z dz}}^{\text{Prob. no success by } t} dt$$

- p_t decreasing $\implies$ an efficient effort profile is, for all $i \in \mathcal{N}$,

$$a_{i,t} = \begin{cases} 1 & \text{if } p_t \lambda v \geq c \text{ and no success by } t \\ 0 & \text{otherwise} \end{cases}$$

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

$$p^{FB} \equiv \frac{c}{\lambda v}$$

Contests

A contest specifies:

1. Deadline: $T \geq 0$

2. Prizes: $\overline{w}$ and prize-sharing scheme $(w_i(\boldsymbol{s}))_{i \in \mathcal{N}}$ such that
   - (i) $w_i(\boldsymbol{s}) = w(s_i, \boldsymbol{s}_{-i})$, where $w(s_i, \boldsymbol{s}_{-i}) = w(s_i, \sigma(\boldsymbol{s}_{-i}))$ for any perm. σ
   - (ii) $w(\varnothing, \cdot) = 0$
   - (iii) $\boldsymbol{s} \neq (\varnothing, \ldots, \varnothing) \implies \sum_{i=1}^{N} w_i(\boldsymbol{s}) = \overline{w}$

$\rightarrow$ Salient cases: WTA and equal-sharing

3. Disclosure: $(M_t, \mu_t)_{t \in [0,T]}$, at each t agents observe $m_t = \mu_t(\boldsymbol{o}^t) \in M_t$

$\rightarrow$ Salient cases: public and hidden

Principal's problem

- Principal designs contest to maximize her expected payoff gain

$$(v - \overline{w})p_0 \left(1 - e^{-\lambda A^T}\right)$$

where $A^T \equiv \int_0^T A_z dz$

Principal's problem

- Principal designs contest to maximize her expected payoff gain

$$(v-\overline{w})p_0\left(1-e^{-\lambda A^T}\right)$$

where $A^T \equiv \int_0^T A_z dz$

- Decompose problem into two steps
    1. For any given $\overline{w}$, solve for optimal contest
    2. Use solution to step 1. to solve for optimal prize $\overline{w}$

- Strategies & Equilibrium:
    - Wlog, $a_{i,t}$ is i's effort at t conditional on i not succeeding by t
    - Symmetric Nash equilibria; refinements would not alter analysis

Principal's problem: Step 1

- For any given $\overline{w}$, solve for optimal prize scheme and info disclosure
- Given $\overline{w} \leq v$, principal's objective is to maximize prob. of a success
- Study public and hidden contests, then general info disclosure

Public winner-takes-all contest

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

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

where

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

Public winner-takes-all contest

- Unique equilibrium is symmetric: for all $i \in \mathcal{N}$,

$$a_{i,t} = \begin{cases} 1 & \text{if } p_{i,t} \geq \frac{c}{\lambda \overline{w}} \equiv p^{PW} \text{ and no success by } t \\ 0 & \text{otherwise} \end{cases}$$

- Implies deadline T 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 1: Implements first-best solution iff $\overline{w} = v$

Remark 2: Probability of success is invariant to N

Hidden winner-takes-all contest

- Now i's problem is

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

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

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

Hidden winner-takes-all contest

- Unique equilibrium is symmetric: for all $i \in \mathcal{N}$,

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

- Under non-binding T, stopping time T^{HW} is then 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 contest

- Now i's problem is

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

where (suppressing dependence on strategies):

- $w_{i,t} \equiv i$'s expected reward if he succeeds at t
- $u_{i,t} \equiv i$'s continuation payoff if some $-i$ succeeds at t

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

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

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

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:
  - Wlog to consider prize scheme that induces full effort from 0 to T
  - Equal sharing $\implies$ constant sequence of expected rewards
  - Stopping time T^{HS} s.t. each agent's IC binds at each $t \in [0, T^{HS}]$
  - Thus, no hidden contest can induce more experimentation
    - If $T > T^{HS}$, IC violated at some $t \leq T$

Hidden equal-sharing contest

- Under equal sharing, i's problem is

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

- An optimal strategy is

$$a_{i,t} = \begin{cases} 1 & \text{if } p_{i,t}^{(1)}\lambda w_i \geq c \\ 0 & \text{otherwise} \end{cases}$$

- In a symmetric equilibrium, expected reward w^{HS}, stopping time T^{HS}

Hidden equal-sharing contest

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

$$w^{HS} = \overline{w}\ \frac{1 - e^{-\lambda N T^{HS}}}{(1 - e^{-\lambda T^{HS}})N}$$

- Under non-binding T, stopping time T^{HS} solves

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

which has a unique solution; hence essentially unique symmetric eqm

Remark: Increase in N can increase or decrease probability of success

Public or hidden?

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

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

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

Public winner-takes-all versus hidden equal-sharing

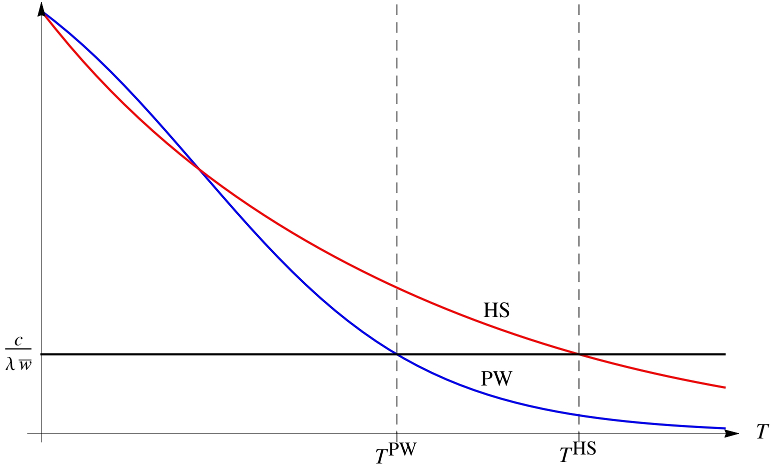

Result for public and hidden contests

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.

Intuition: Interpreting the condition

- We can rewrite condition as follows:

$$\lambda\overline{w}\sum_{m=1}^{N-1}\frac{\Pr[m \text{ opponents succeed by } T^{PW} \mid G]}{\Pr[\text{at least one opponent succeeds by } T^{PW} \mid G]}\left(\frac{1}{m+1}\right) > c$$

- At T^{PW}, if all $-i$ failed, i is indifferent over exerting effort
- So, i strictly prefers to continue iff he does when some $-i$ succeeded

Intuition: Necessary and sufficient conditions

- Condition for $N = 2$ is
$$\frac{\overline{w}}{2}\lambda > c$$
$\rightarrow i$ would experiment to earn half prize if he knew $-i$ succeeded

- If $N > 2$, above condition necessary, and simple sufficient condition is
$$\frac{\overline{w}}{N}\lambda \geq c$$

$\Longrightarrow$ HS dominates (is dominated by) PW if $c/\lambda\overline{w}$ sufficiently small (large)

Intuition: Discussion

- Why can hidden shared but not hidden WTA/public shared dominate?
  - Want to hide info to bolster agent's belief when no-one succeeded
  - But hiding info is counter-productive under WTA
  - And public shared can only $\uparrow$ effort when not beneficial ($+$ free-riding)

- Hiding information can be beneficial because agents learn from others
  - If $p_0 = 1$ or arms uncorrelated $\implies$ public WTA always optimal
  - Higher $p_0 \implies$ hidden equal-sharing optimal for smaller parameter set

Implication: Number of contestants

- If principal can choose N, HS does always at least as well as PW
  - HS can replicate PW by setting $N = 1$
- Our results imply it can be strictly optimal to have multiple agents
  - Despite no exogenous forces such as heterogeneity and discounting
- $N > 1$ allows to harness benefits from hiding info and sharing prize

General disclosure policies

- Rank monotonicity: for any $\boldsymbol{s}$, $s_i < s_j \implies w(s_i, \boldsymbol{s}_{-i}) \geq w(s_j, \mathbf{s}_{-j})$
- Cutoff disclosure: $M_t = \{0, 1\}$, $\mu_t(\boldsymbol{o}^t) = 1$ iff n or more succeed by t
- Define

$$n^* \equiv \max\left\{n \in \{1, \ldots, N\} : \lambda \frac{\overline{w}}{n} \geq c\right\}$$

Proposition

A cutoff-disclosure equal-sharing contest with cutoff n^* is optimal among rank-monotonic contests.

- Intuition:
  - Rank monotonicity $\implies$ reward for success bounded by equal share
  - Exert effort given G and equal-sharing iff share w/less than n^* agents
  - Increase (reduce) effort incentive if reveal $n \geq n^*$ ($n < n^*$) successes

Optimal cutoff disclosure equal-sharing contest

- Note that $n^* = 1$ when $\lambda \frac{\overline{w}}{2} < c$, whereas $n^* = N$ when $\lambda \frac{\overline{w}}{N} > c$
- Since agents stop exerting effort when n^* successes are announced,
  - Cutoff-disclosure equal-sharing with $n^* = 1$ is equivalent to PW
  - Cutoff-disclosure equal-sharing with $n^* = N$ is equivalent to HS

Corollary

Among rank-monotonic contests, public WTA is optimal if $\lambda\overline{w}/2 < c$ and hidden equal-sharing is optimal if $\lambda\overline{w}/N > c$.

- Finally, given salience and widespread use of WTA contests, we note:

Proposition

A public contest is optimal among WTA contests.

Principal's problem: Step 2

- Given optimal contest as function of $\overline{w}$, principal solves for optimal $\overline{w}$

Proposition

Fix any parameters (p_0, λ, c, N) *and consider rank-monotonic contests.*

- v *large enough* $\Longrightarrow$ *principal chooses* $\overline{w} \in (0, v)$ *and hidden equal-sharing*
- v *small enough* $\Longrightarrow$ *principal chooses* $\overline{w} \in (0, v)$ *and public WTA*

- Intuition: For v large (small) enough, optimal $\overline{w}$ s.t. $\frac{\lambda \overline{w}}{N} > c$ $\left(\frac{\lambda \overline{w}}{2} < c\right)$

Extensions and discussion

- Social planner: May also prefer hidden equal-sharing over public WTA
  - If budget constrained ($\overline{w} < v$), which is likely if value of discovery high
  - Ex post, planner induces wasteful experimentation after discovery made
- Observability of success: Results robust to different assumptions
  - If only P observes success, no reason to hide it from successful A
  - If only A observes success and can verifiably reveal it, main result holds
  - If A can verifiably reveal success to opponents, main result holds
- Discounting, convex costs: Main insight is robust
  - With discounting, benefit of public WTA: can use success immediately
  - But hidden equal-sharing can yield higher probability of success

Applications

First-to-file vs. first-to-invent rules in patent law

- FTF seen as beneficial because it induces earlier filing, more disclosure (e.g. Scotchmer-Green 90)
- Our results: FTI beneficial because it limits disclosure! (and induces sharing)

Optimal task allocation in organizations

- Principal assigns two tasks of uncertain and indep. difficulty to two agents
- Our results: benefits of making agents jointly responsible for the two tasks

Design of contract awards in government procurement

- In challenge-based acquisitions, no disclosure until evaluation date
- Moreover, multiple contractors, often to have stable supply and competition
- Our results: contract sharing beneficial beyond these considerations

Conclusions

- Tradeoff in incentivizing experimentation:

  $\uparrow$ agent's reward for success *vs* $\uparrow$ his belief that he will succeed

- Hiding info and sharing prize often dominates public WTA
  - Only hiding info or dividing prize hurts, but both together can help

- Broader contributions
  - Contest design in an environment with learning
  - Mechanism design approach—over payments and info disclosure—to multi-agent strategic experimentation