Award Structure in Collaborative Contests

James Siderius and Mohamed Mostagir

(Preliminary Work)

INFORMS 2019
Seattle, WA, October 22, 2019

An Example: Netflix Contest

Motivation/Literature Review

What We Do

- Understand incentives for researchers to share progress, and how it shapes societal outcomes.
  - how does a collaborative society use resources to solve complex problems?
  - how do rewards influence whether agents hoard preliminary results?
  - how should society structure rewards to promote collaborative behavior?

- Try to understand strategic incentives for agents to work on similar problems and keep breakthroughs private.

- How should a designer better align private incentives and societal goals for solving a complex problem?

- Today:
  - understand the collaborative solution
  - characterize the equilibria with private research efforts
  - understand optimal design of "partial-progress" rewards

What We Do

- Understand incentives for researchers to share progress, and how it shapes societal outcomes.
  - how does a collaborative society use resources to solve complex problems?
  - how do rewards influence whether agents hoard preliminary results?
  - how should society structure rewards to promote collaborative behavior?

- Try to understand strategic incentives for agents to work on similar problems and keep breakthroughs private.

- How should a designer better align private incentives and societal goals for solving a complex problem?

- Today:
  - understand the collaborative solution
  - characterize the equilibria with private research efforts
  - understand optimal design of "partial-progress" rewards

What We Do

- Understand incentives for researchers to share progress, and how it shapes societal outcomes.
  - how does a collaborative society use resources to solve complex problems?
  - how do rewards influence whether agents hoard preliminary results?
  - how should society structure rewards to promote collaborative behavior?

- Try to understand strategic incentives for agents to work on similar problems and keep breakthroughs private.

- How should a designer better align private incentives and societal goals for solving a complex problem?

- Today:
  - understand the collaborative solution
  - characterize the equilibria with private research efforts
  - understand optimal design of "partial-progress" rewards

What We Do

- Understand incentives for researchers to share progress, and how it shapes societal outcomes.
  - how does a collaborative society use resources to solve complex problems?
  - how do rewards influence whether agents hoard preliminary results?
  - how should society structure rewards to promote collaborative behavior?

- Try to understand strategic incentives for agents to work on similar problems and keep breakthroughs private.

- How should a designer better align private incentives and societal goals for solving a complex problem?

- Today:
  - understand the collaborative solution
  - characterize the equilibria with private research efforts
  - understand optimal design of "partial-progress" rewards

Model: Society's Problem

- Society is presented with some complex problem. Complexity of problem has several dimensions:
  - Value to society $K > 0$ (e.g., cure for cancer vs. marginal technological improvement).
  - Time sensitivity of the problem $\beta \in (0, 1)$ (e.g., Apollo-11 mission vs. twin-prime conjecture).
  - Difficulty of problem (likelihood p of a breakthrough per unit effort).

- Profitability of effort can be measured as the expected value of the problem's solution given the effort today. Decompose into contribution and tractability:
  - If value to society (K) is large, solution has a sizable contribution.
  - If breakthroughs occur frequently (p) relative to time sensitivity (β), problem is tractable.

- Whether society exerts effort will depend on whether the solution has a large contribution and/or reasonable tractability.

Model: Society's Problem

- Society is presented with some complex problem. Complexity of problem has several dimensions:
  - Value to society $K > 0$ (e.g., cure for cancer vs. marginal technological improvement).
  - Time sensitivity of the problem $\beta \in (0, 1)$ (e.g., Apollo-11 mission vs. twin-prime conjecture).
  - Difficulty of problem (likelihood p of a breakthrough per unit effort).

- Profitability of effort can be measured as the expected value of the problem's solution given the effort today. Decompose into contribution and tractability:
  - If value to society (K) is large, solution has a sizable contribution.
  - If breakthroughs occur frequently (p) relative to time sensitivity (β), problem is tractable.

- Whether society exerts effort will depend on whether the solution has a large contribution and/or reasonable tractability.

Model: Society's Problem

- Society is presented with some complex problem. Complexity of problem has several dimensions:
  - Value to society $K > 0$ (e.g., cure for cancer vs. marginal technological improvement).
  - Time sensitivity of the problem $\beta \in (0, 1)$ (e.g., Apollo-11 mission vs. twin-prime conjecture).
  - Difficulty of problem (likelihood p of a breakthrough per unit effort).

- Profitability of effort can be measured as the expected value of the problem's solution given the effort today. Decompose into contribution and tractability:
  - If value to society (K) is large, solution has a sizable contribution.
  - If breakthroughs occur frequently (p) relative to time sensitivity (β), problem is tractable.

- Whether society exerts effort will depend on whether the solution has a large contribution and/or reasonable tractability.

Model: Solving Complex Problems

- Society is composed of n agents, who can each choose whether or not to work on a given problem.

- Problem requires a sequence of breakthroughs.
  - There are m stages of the problem; society is at stage $k \in \{1, 2, \ldots, m\}$.
  - Can only progress to stage $k + 1$ if society currently knows the solution to stage k.

- Time is discrete $t = 1, 2, \ldots$
  - Each agent i chooses to exert effort $e_{i,t} \in \{0, 1\}$ at each time t on the problem (i.e., either the agent works on the problem or not).
  - If $e_{i,t} = 1$, with probability p agent i advances society from stage s_t to stage $s_t + 1$.

- Society maximizes $\sum_{t=0}^{\infty} \beta^t \left(K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} e_{i,t} \right)$, where $\mathbb{1}_{\text{sol}}$ is the *first* period where society has advanced to stage m.

Model: Solving Complex Problems

- Society is composed of n agents, who can each choose whether or not to work on a given problem.

- Problem requires a sequence of breakthroughs.
  - There are m stages of the problem; society is at stage $k \in \{1, 2, \ldots, m\}$.
  - Can only progress to stage $k + 1$ if society currently knows the solution to stage k.

- Time is discrete $t = 1, 2, \ldots$
  - Each agent i chooses to exert effort $e_{i,t} \in \{0, 1\}$ at each time t on the problem (i.e., either the agent works on the problem or not).
  - If $e_{i,t} = 1$, with probability p agent i advances society from stage s_t to stage $s_t + 1$.

- Society maximizes $\sum_{t=0}^{\infty} \beta^t \left(K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} e_{i,t} \right)$, where $\mathbb{1}_{\text{sol}}$ is the *first* period where society has advanced to stage m.

Model: Solving Complex Problems

- Society is composed of n agents, who can each choose whether or not to work on a given problem.

- Problem requires a sequence of breakthroughs.
  - There are m stages of the problem; society is at stage $k \in \{1, 2, \ldots, m\}$.
  - Can only progress to stage $k+1$ if society currently knows the solution to stage k.

- Time is discrete $t = 1, 2, \ldots$
  - Each agent i chooses to exert effort $e_{i,t} \in \{0, 1\}$ at each time t on the problem (i.e., either the agent works on the problem or not).
  - If $e_{i,t} = 1$, with probability p agent i advances society from stage s_t to stage $s_t + 1$.

- Society maximizes $\sum_{t=0}^{\infty} \beta^t \left(K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} e_{i,t} \right)$, where $\mathbb{1}_{\text{sol}}$ is the *first* period where society has advanced to stage m.

Model: Solving Complex Problems

- Society is composed of n agents, who can each choose whether or not to work on a given problem.

- Problem requires a sequence of breakthroughs.
  - There are m stages of the problem; society is at stage $k \in \{1, 2, \ldots, m\}$.
  - Can only progress to stage $k + 1$ if society currently knows the solution to stage k.

- Time is discrete $t = 1, 2, \ldots$
  - Each agent i chooses to exert effort $e_{i,t} \in \{0, 1\}$ at each time t on the problem (i.e., either the agent works on the problem or not).
  - If $e_{i,t} = 1$, with probability p agent i advances society from stage s_t to stage $s_t + 1$.

- Society maximizes $\sum_{t=0}^{\infty} \beta^t \left(K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} e_{i,t} \right)$, where $\mathbb{1}_{\text{sol}}$ is the *first* period where society has advanced to stage m.

Comparative Statics

- Let $\ell^*(k)$ denote the total amount of effort exerted at stage k (i.e., the number of agents working toward the solution).

Proposition

Optimal effort $\ell^*(k)$ is non-decreasing in k and K.

- Intuition: Sprint to the finish. As society gets closer to solving the problem in its entirety, devote more resources to finishing the project.
  - Holds even though the feasibility of the problem is unaffected by earlier stages' progress.
  - Time-value of the solution: When solution is close, effort today will translate into contribution soon.

Comparative Statics

- Let $\ell^*(k)$ denote the total amount of effort exerted at stage k (i.e., the number of agents working toward the solution).

Proposition

Optimal effort $\ell^*(k)$ is non-decreasing in k and K.

- Intuition: Sprint to the finish. As society gets closer to solving the problem in its entirety, devote more resources to finishing the project.
  - Holds even though the feasibility of the problem is unaffected by earlier stages' progress.
  - Time-value of the solution: When solution is close, effort today will translate into contribution soon.

Comparative Statics

- Let $\ell^*(k)$ denote the total amount of effort exerted at stage k (i.e., the number of agents working toward the solution).

Proposition

Optimal effort $\ell^*(k)$ is non-decreasing in k and K.

- Intuition: Sprint to the finish. As society gets closer to solving the problem in its entirety, devote more resources to finishing the project.
  - Holds even though the feasibility of the problem is unaffected by earlier stages' progress.
  - Time-value of the solution: When solution is close, effort today will translate into contribution soon.

Non-Monotonicity in p

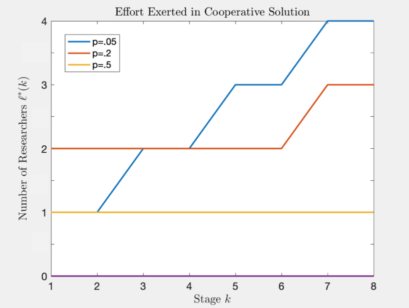

- How does effort vary with tractability (as measured by p)?
  - Depends on how far away you are from the solution.
- Low p: breakthroughs are infrequent.
  - Many remaining stages $\Longrightarrow$ intractable. Exert little to no effort.
  - Few remaining stages, exert a lot of effort.
- High p: redundant breakthroughs are common.
  - Do not waste resources leading to multiple (but the same) breakthroughs in each period.

Non-Monotonicity in p

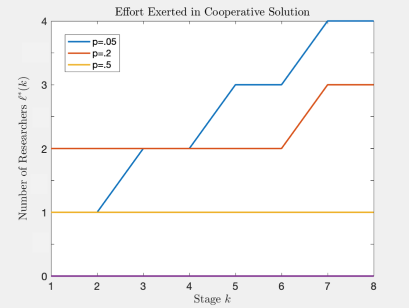

- How does effort vary with tractability (as measured by p)?
  - Depends on how far away you are from the solution.
- Low p: breakthroughs are infrequent.
  - Many remaining stages $\Longrightarrow$ intractable. Exert little to no effort.
  - Few remaining stages, exert a lot of effort.
- High p: redundant breakthroughs are common.
  - Do not waste resources leading to multiple (but the same) breakthroughs in each period.

Non-Monotonicity in p

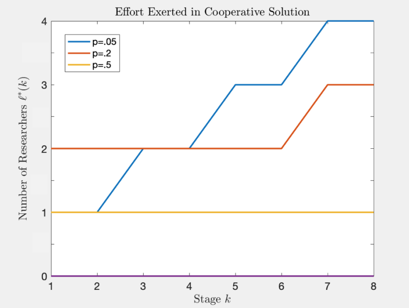

- How does effort vary with tractability (as measured by p)?
  - Depends on how far away you are from the solution.
- Low p: breakthroughs are infrequent.
  - Many remaining stages $\Longrightarrow$ intractable. Exert little to no effort.
  - Few remaining stages, exert a lot of effort.
- High p: redundant breakthroughs are common.
  - Do not waste resources leading to multiple (but the same) breakthroughs in each period.

Model: Individual Efforts

- Each of the n agents works (or not) on the problem at their own pace.
- Agent i chooses an effort level $e_{i,t} \in \{0, 1\}$. If $e_{i,t} = 1$, then with probability p agent i advances from step $s_{i,t}$ to step $s_{i,t} + 1$; that is, $s_{i,t} = s_{i,t-1} + 1$.
- Agents may publish "new" results. If last publication was s^*, agent i may publish any stage s^{**} such that $s^* < s^{**} \leq s_{i,t}$.
- After the intermediate progress has been published, all other agents catch up to this stage; that is, $s_{i,t} \leftarrow \max\{s_{i,t}, s^{**}\}$.
- Suppose there is a reward for solving the problem. Will any agent voluntarily publish intermediate results?
    - No, this creates additional competition. Inefficient because agents develop the same intermediate progress in parallel.

Model: Individual Efforts

- Each of the n agents works (or not) on the problem at their own pace.
- Agent i chooses an effort level $e_{i,t} \in \{0, 1\}$. If $e_{i,t} = 1$, then with probability p agent i advances from step $s_{i,t}$ to step $s_{i,t} + 1$; that is, $s_{i,t} = s_{i,t-1} + 1$.
- Agents may publish "new" results. If last publication was s^*, agent i may publish any stage s^{**} such that $s^* < s^{**} \leq s_{i,t}$.
- After the intermediate progress has been published, all other agents catch up to this stage; that is, $s_{i,t} \leftarrow \max\{s_{i,t}, s^{**}\}$.
- Suppose there is a reward for solving the problem. Will any agent voluntarily publish intermediate results?
  - No, this creates additional competition. Inefficient because agents develop the same intermediate progress in parallel.

Model: Individual Efforts

- Each of the n agents works (or not) on the problem at their own pace.
- Agent i chooses an effort level $e_{i,t} \in \{0, 1\}$. If $e_{i,t} = 1$, then with probability p agent i advances from step $s_{i,t}$ to step $s_{i,t} + 1$; that is, $s_{i,t} = s_{i,t-1} + 1$.
- Agents may publish "new" results. If last publication was s^*, agent i may publish any stage s^{**} such that $s^* < s^{**} \leq s_{i,t}$.
- After the intermediate progress has been published, all other agents catch up to this stage; that is, $s_{i,t} \leftarrow \max\{s_{i,t}, s^{**}\}$.
- Suppose there is a reward for solving the problem. Will any agent voluntarily publish intermediate results?
  - No, this creates additional competition. Inefficient because agents develop the same intermediate progress in parallel.

Model: Individual Efforts

- Each of the n agents works (or not) on the problem at their own pace.
- Agent i chooses an effort level $e_{i,t} \in \{0, 1\}$. If $e_{i,t} = 1$, then with probability p agent i advances from step $s_{i,t}$ to step $s_{i,t} + 1$; that is, $s_{i,t} = s_{i,t-1} + 1$.
- Agents may publish "new" results. If last publication was s^*, agent i may publish any stage s^{**} such that $s^* < s^{**} \leq s_{i,t}$.
- After the intermediate progress has been published, all other agents catch up to this stage; that is, $s_{i,t} \leftarrow \max\{s_{i,t}, s^{**}\}$.
- Suppose there is a reward for solving the problem. Will any agent voluntarily publish intermediate results?
  - No, this creates additional competition. Inefficient because agents develop the same intermediate progress in parallel.

Model: Individual Efforts

- Each of the n agents works (or not) on the problem at their own pace.
- Agent i chooses an effort level $e_{i,t} \in \{0, 1\}$. If $e_{i,t} = 1$, then with probability p agent i advances from step $s_{i,t}$ to step $s_{i,t} + 1$; that is, $s_{i,t} = s_{i,t-1} + 1$.
- Agents may publish "new" results. If last publication was s^*, agent i may publish any stage s^{**} such that $s^* < s^{**} \leq s_{i,t}$.
- After the intermediate progress has been published, all other agents catch up to this stage; that is, $s_{i,t} \leftarrow \max\{s_{i,t}, s^{**}\}$.
- Suppose there is a reward for solving the problem. Will any agent voluntarily publish intermediate results?
  - No, this creates additional competition. Inefficient because agents develop the same intermediate progress in parallel.

Model: Designer's Problem

- Suppose instead the designer can commit to offering reward r^k for publication of stage k to incentivize publication.

- Other extreme: set $r^k = R^k \gg R^{k+1}$ for every stage k. Then progress evolves just as in the cooperative solution, where all agents publish every stage immediately.

- Agents care only about extrinsic rewards. For simplicity (and largely WLOG), assume agents discount payoffs at the same rate.
  - They choose to exert effort and whether to publish in order to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t (r_{i,t} - e_{i,t})\right]$, where $r_{i,t}$ is the reward received by agent i in period t.

- Designer (e.g., a social planner) reaps the societal rewards K, but must pay out the intermediate rewards r^k.
  - Chooses reward structure to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} r_{i,t}\right)\right]$.

Model: Designer's Problem

- Suppose instead the designer can commit to offering reward r^k for publication of stage k to incentivize publication.
- Other extreme: set $r^k = R^k \gg R^{k+1}$ for every stage k. Then progress evolves just as in the cooperative solution, where all agents publish every stage immediately.
- Agents care only about extrinsic rewards. For simplicity (and largely WLOG), assume agents discount payoffs at the same rate.
  - They choose to exert effort and whether to publish in order to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t (r_{i,t} - e_{i,t})\right]$, where $r_{i,t}$ is the reward received by agent i in period t.
- Designer (e.g., a social planner) reaps the societal rewards K, but must pay out the intermediate rewards r^k.
  - Chooses reward structure to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(K\mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} r_{i,t}\right)\right]$.

Model: Designer's Problem

- Suppose instead the designer can commit to offering reward r^k for publication of stage k to incentivize publication.

- Other extreme: set $r^k = R^k \gg R^{k+1}$ for every stage k. Then progress evolves just as in the cooperative solution, where all agents publish every stage immediately.

- Agents care only about extrinsic rewards. For simplicity (and largely WLOG), assume agents discount payoffs at the same rate.
  - They choose to exert effort and whether to publish in order to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t (r_{i,t} - e_{i,t})\right]$, where $r_{i,t}$ is the reward received by agent i in period t.

- Designer (e.g., a social planner) reaps the societal rewards K, but must pay out the intermediate rewards r^k.
  - Chooses reward structure to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} r_{i,t}\right)\right]$.

Model: Designer's Problem

- Suppose instead the designer can commit to offering reward r^k for publication of stage k to incentivize publication.

- Other extreme: set $r^k = R^k \gg R^{k+1}$ for every stage k. Then progress evolves just as in the cooperative solution, where all agents publish every stage immediately.

- Agents care only about extrinsic rewards. For simplicity (and largely WLOG), assume agents discount payoffs at the same rate.
  - They choose to exert effort and whether to publish in order to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t (r_{i,t} - e_{i,t})\right]$, where $r_{i,t}$ is the reward received by agent i in period t.

- Designer (e.g., a social planner) reaps the societal rewards K, but must pay out the intermediate rewards r^k.
  - Chooses reward structure to maximize $\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(K\mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} r_{i,t}\right)\right]$.

Two-Stage Equilibria

- Suppose $m = 2$, so there is only one intermediate stage. Take as given reward structure $r^1, r^2 > 0$.

Theorem

All pure-strategy, symmetric perfect Bayesian equilibria are of the form:

(i) *No agent exerts effort at any point in time.*

(ii) *Agents exert effort in every period of stage 1 and publish the stage 1 result immediately. No agent exerts effort in period 2.*

(iii) *Agents always exert effort in stage 1 until some time T^*. There exists some interval $[\underline{\Delta}, \overline{\Delta}]$ such that any agent who has the stage 1 result publishes at (and only at) times $\mathcal{T} = \{\tau_1, \tau_2, \ldots\}$, where $\tau_j - \tau_{j-1} = \Delta_j^*$ for some $\Delta_j^* \in [\underline{\Delta}, \overline{\Delta}]$. Agents publish the stage 2 result immediately.*

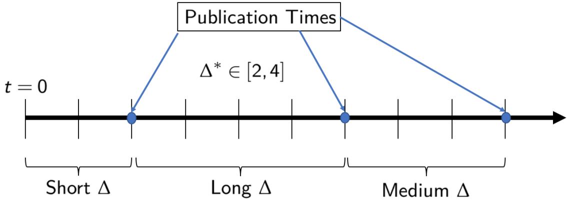

Two-Stage Equilibria

- Suppose $m = 2$, so there is only one intermediate stage. Take as given reward structure $r^1, r^2 > 0$.

Theorem

All pure-strategy, symmetric perfect Bayesian equilibria are of the form:

(i) *No agent exerts effort at any point in time.*

(ii) *Agents exert effort in every period of stage 1 and publish the stage 1 result immediately. No agent exerts effort in period 2.*

(iii) *Agents always exert effort in stage 1 until some time T^*. There exists some interval $[\underline{\Delta}, \overline{\Delta}]$ such that any agent who has the stage 1 result publishes at (and only at) times $\mathcal{T} = \{\tau_1, \tau_2, \ldots\}$, where $\tau_j - \tau_{j-1} = \Delta_j^*$ for some $\Delta_j^* \in [\underline{\Delta}, \overline{\Delta}]$. Agents publish the stage 2 result immediately.*

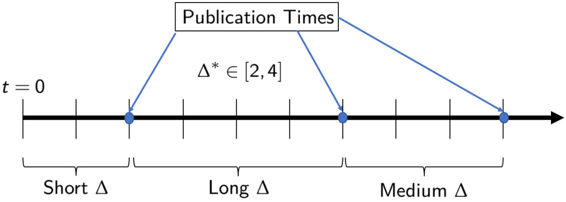

Two-Stage Equilibria

- Suppose $m = 2$, so there is only one intermediate stage. Take as given reward structure $r^1, r^2 > 0$.

Theorem

All pure-strategy, symmetric perfect Bayesian equilibria are of the form:

(i) *No agent exerts effort at any point in time.*

(ii) *Agents exert effort in every period of stage 1 and publish the stage 1 result immediately. No agent exerts effort in period 2.*

(iii) *Agents always exert effort in stage 1 until some time T^*. There exists some interval $[\underline{\Delta}, \overline{\Delta}]$ such that any agent who has the stage 1 result publishes at (and only at) times $\mathcal{T} = \{\tau_1, \tau_2, \ldots\}$, where $\tau_j - \tau_{j-1} = \Delta_j^*$ for some $\Delta_j^* \in [\underline{\Delta}, \overline{\Delta}]$. Agents publish the stage 2 result immediately.*

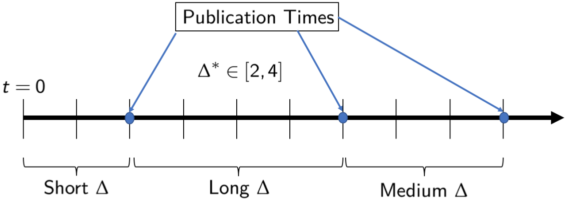

Two-Stage Equilibria

- Suppose $m = 2$, so there is only one intermediate stage. Take as given reward structure $r^1, r^2 > 0$.

Theorem

All pure-strategy, symmetric perfect Bayesian equilibria are of the form:

(i) *No agent exerts effort at any point in time.*

(ii) *Agents exert effort in every period of stage 1 and publish the stage 1 result immediately. No agent exerts effort in period 2.*

(iii) *Agents always exert effort in stage 1 until some time T^*. There exists some interval $[\underline{\Delta}, \overline{\Delta}]$ such that any agent who has the stage 1 result publishes at (and only at) times $\mathcal{T} = \{\tau_1, \tau_2, \ldots\}$, where $\tau_j - \tau_{j-1} = \Delta_j^*$ for some $\Delta_j^* \in [\underline{\Delta}, \overline{\Delta}]$. Agents publish the stage 2 result immediately.*

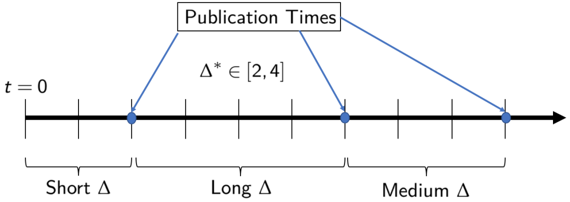

Intuition: Publication Cycles

- Shortest cycle, $\underline{\Delta}$, corresponds to most collaborative equilibrium and longest cycle, $\overline{\Delta}$, corresponds to most secretive.
  - Multiple equilibria because of *strategic complementarities*.
  - Suppose today is Monday and agent *i* has the stage 1 result. No one will publish until Thursday. When is the earliest agent *i* will publish? What if agent *i* believes her competitor might publish tomorrow?

- Two effects which jointly determine the range of Δ supportable in equilibrium:
  - Fear of scooping: Wait an extra period, risk too many publications at time τ_j who split r^1 (or someone finishes stage 2 and gets $r^1 + r^2$). Instead could publish today and guaranteed entire r^1.
  - Marginal competition: Long publication cycles mean most agents have (independently) solved stage 1. Publication is not helping as many competitors catch-up.

- After time τ_j has passed, all agents know that no one has solved stage 1 problem. Environment resets.

Intuition: Publication Cycles

- Shortest cycle, $\underline{\Delta}$, corresponds to most collaborative equilibrium and longest cycle, $\overline{\Delta}$, corresponds to most secretive.
  - Multiple equilibria because of *strategic complementarities*.
  - Suppose today is Monday and agent i has the stage 1 result. No one will publish until Thursday. When is the earliest agent i will publish? What if agent i believes her competitor might publish tomorrow?

- Two effects which jointly determine the range of Δ supportable in equilibrium:
  - Fear of scooping: Wait an extra period, risk too many publications at time τ_j who split r^1 (or someone finishes stage 2 and gets $r^1 + r^2$). Instead could publish today and guaranteed entire r^1.
  - Marginal competition: Long publication cycles mean most agents have (independently) solved stage 1. Publication is not helping as many competitors catch-up.

- After time τ_j has passed, all agents know that no one has solved stage 1 problem. Environment resets.

Intuition: Publication Cycles

- Shortest cycle, $\underline{\Delta}$, corresponds to most collaborative equilibrium and longest cycle, $\overline{\Delta}$, corresponds to most secretive.
  - Multiple equilibria because of *strategic complementarities*.
  - Suppose today is Monday and agent i has the stage 1 result. No one will publish until Thursday. When is the earliest agent i will publish? What if agent i believes her competitor might publish tomorrow?

- Two effects which jointly determine the range of Δ supportable in equilibrium:
  - Fear of scooping: Wait an extra period, risk too many publications at time τ_j who split r^1 (or someone finishes stage 2 and gets $r^1 + r^2$). Instead could publish today and guaranteed entire r^1.
  - Marginal competition: Long publication cycles mean most agents have (independently) solved stage 1. Publication is not helping as many competitors catch-up.

- After time τ_j has passed, all agents know that no one has solved stage 1 problem. Environment resets.

Intuition: Publication Cycles

- Shortest cycle, $\underline{\Delta}$, corresponds to most collaborative equilibrium and longest cycle, $\overline{\Delta}$, corresponds to most secretive.
  - Multiple equilibria because of *strategic complementarities*.
  - Suppose today is Monday and agent i has the stage 1 result. No one will publish until Thursday. When is the earliest agent i will publish? What if agent i believes her competitor might publish tomorrow?

- Two effects which jointly determine the range of Δ supportable in equilibrium:
  - Fear of scooping: Wait an extra period, risk too many publications at time τ_j who split r^1 (or someone finishes stage 2 and gets $r^1 + r^2$). Instead could publish today and guaranteed entire r^1.
  - Marginal competition: Long publication cycles mean most agents have (independently) solved stage 1. Publication is not helping as many competitors catch-up.

- After time τ_j has passed, all agents know that no one has solved stage 1 problem. Environment resets.

Intuition: Publication Cycles

- Shortest cycle, $\underline{\Delta}$, corresponds to most collaborative equilibrium and longest cycle, $\overline{\Delta}$, corresponds to most secretive.
  - Multiple equilibria because of *strategic complementarities*.
  - Suppose today is Monday and agent *i* has the stage 1 result. No one will publish until Thursday. When is the earliest agent *i* will publish? What if agent *i* believes her competitor might publish tomorrow?

- Two effects which jointly determine the range of Δ supportable in equilibrium:
  - Fear of scooping: Wait an extra period, risk too many publications at time τ_j who split r^1 (or someone finishes stage 2 and gets $r^1 + r^2$). Instead could publish today and guaranteed entire r^1.
  - Marginal competition: Long publication cycles mean most agents have (independently) solved stage 1. Publication is not helping as many competitors catch-up.

- After time τ_j has passed, all agents know that no one has solved stage 1 problem. Environment resets.

Simplified Designer's Task

- Recall, designer wishes to maximize expectation of $\sum_{t=0}^{\infty} \beta^t (K \mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} r_{i,t})$. Optimal reward structure (r^1, r^2) often difficult to solve in general.
- Consider the heuristics: all-or-nothing contest and partial-progress contest.
  - All-or-nothing: only reward for final contribution (i.e., $r^1 = 0$) and choose r^2 optimally.
  - Partial-progress: choose r^2 smallest so agents still exert effort in period 2, choose r^1 strategically.

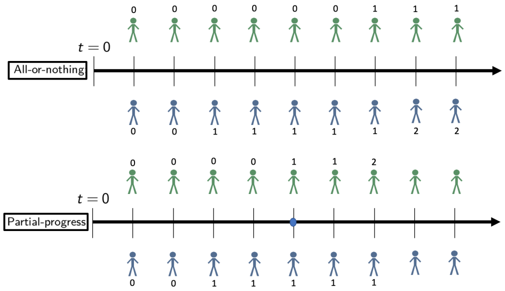

Simplified Designer's Task

- Recall, designer wishes to maximize expectation of $\sum_{t=0}^{\infty} \beta^t (K\mathbb{1}_{\text{sol}} - \sum_{i=1}^{n} r_{i,t})$. Optimal reward structure (r^1, r^2) often difficult to solve in general.

- Consider the heuristics: all-or-nothing contest and partial-progress contest.
  - All-or-nothing: only reward for final contribution (i.e., $r^1 = 0$) and choose r^2 optimally.
  - Partial-progress: choose r^2 smallest so agents still exert effort in period 2, choose r^1 strategically.

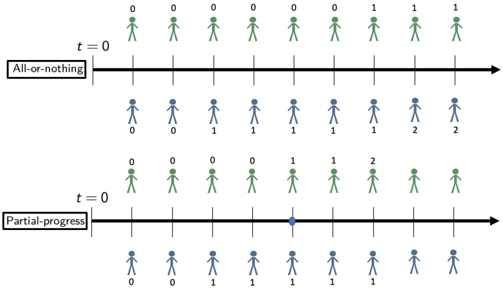

Simulated Designs

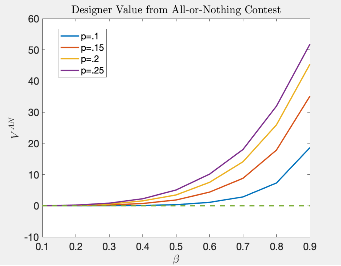

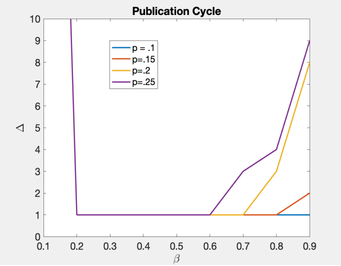

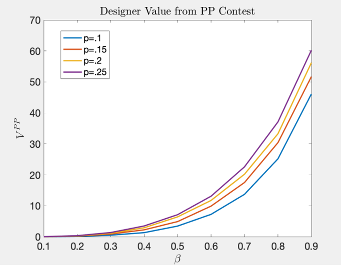

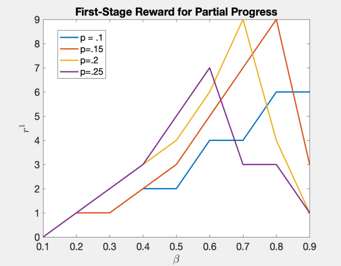

Optimal Reward Structure: Results

- Recall tractability (π) is the measure of the rate of progress (via p) relative to the time sensitivity of the problem (via β). Formally,

$$\pi = \frac{\beta(1-p)^n}{(1-\beta)(1-(1-p)^n)}$$

Theorem

As $\pi \to 0$ or $\pi \to \infty$, the all-or-nothing contest is the optimal partial-progress contest.

- Publication of partial progress is most important when the problem is somewhat tractable.
  - As tractability increases, optimal reward structure induces the most secrecy and longest publication cycles.

Theorem

There exists π^* such that the optimal r^1 is increasing for all $\pi < \pi^*$ and decreasing for all $\pi > \pi^*$.

- If a tractable problem becomes more tractable or intractable problem becomes more tractable, increase the intermediate progress rewards.

Optimal Reward Structure: Results

- Recall tractability (π) is the measure of the rate of progress (via p) relative to the time sensitivity of the problem (via β). Formally,

$$\pi = \frac{\beta(1-p)^n}{(1-\beta)(1-(1-p)^n)}$$

Theorem

As $\pi \to 0$ or $\pi \to \infty$, the all-or-nothing contest is the optimal partial-progress contest.

- Publication of partial progress is most important when the problem is somewhat tractable.
  - As tractability increases, optimal reward structure induces the most secrecy and longest publication cycles.

Theorem

There exists π^* such that the optimal r^1 is increasing for all $\pi < \pi^*$ and decreasing for all $\pi > \pi^*$.

- If a tractable problem becomes more tractable or intractable problem becomes more tractable, increase the intermediate progress rewards.

Optimal Reward Structure: Results

- Recall tractability (π) is the measure of the rate of progress (via p) relative to the time sensitivity of the problem (via β). Formally,

$$\pi = \frac{\beta(1-p)^n}{(1-\beta)(1-(1-p)^n)}$$

Theorem

As $\pi \to 0$ or $\pi \to \infty$, the all-or-nothing contest is the optimal partial-progress contest.

- Publication of partial progress is most important when the problem is somewhat tractable.
  - As tractability increases, optimal reward structure induces the most secrecy and longest publication cycles.

Theorem

There exists π^* such that the optimal r^1 is increasing for all $\pi < \pi^*$ and decreasing for all $\pi > \pi^*$.

- If a tractable problem becomes more tractable or intractable problem becomes more tractable, increase the intermediate progress rewards.

Optimal Reward Structure: Results

- Recall tractability (π) is the measure of the rate of progress (via p) relative to the time sensitivity of the problem (via β). Formally,

$$\pi = \frac{\beta(1-p)^n}{(1-\beta)(1-(1-p)^n)}$$

Theorem

As $\pi \to 0$ or $\pi \to \infty$, the all-or-nothing contest is the optimal partial-progress contest.

- Publication of partial progress is most important when the problem is somewhat tractable.
  - As tractability increases, optimal reward structure induces the most secrecy and longest publication cycles.

Theorem

There exists π^* such that the optimal r^1 is increasing for all $\pi < \pi^*$ and decreasing for all $\pi > \pi^*$.

- If a tractable problem becomes more tractable or intractable problem becomes more tractable, increase the intermediate progress rewards.

Optimal Reward Structure: Results

- Recall tractability (π) is the measure of the rate of progress (via p) relative to the time sensitivity of the problem (via β). Formally,

$$\pi = \frac{\beta(1-p)^n}{(1-\beta)(1-(1-p)^n)}$$

Theorem

As $\pi \to 0$ or $\pi \to \infty$, the all-or-nothing contest is the optimal partial-progress contest.

- Publication of partial progress is most important when the problem is somewhat tractable.
  - As tractability increases, optimal reward structure induces the most secrecy and longest publication cycles.

Theorem

There exists π^* such that the optimal r^1 is increasing for all $\pi < \pi^*$ and decreasing for all $\pi > \pi^*$.

- If a tractable problem becomes more tractable or intractable problem becomes more tractable, increase the intermediate progress rewards.