

Re-incentivizing Discovery: Mechanisms for Partial-Progress Sharing in Research

SIDDHARTHA BANERJEE, Stanford University
ASHISH GOEL, Stanford University
ANILESH K. KRISHNASWAMY, Stanford University

An essential primitive for an efficient research ecosystem is *partial-progress sharing* (PPS) – whereby a researcher shares information immediately upon making a breakthrough. This helps prevent duplication of work; however there is evidence that existing reward structures in research discourage partial-progress sharing. Ensuring PPS is especially important for new online collaborative-research platforms, which involve many researchers working on large, multi-stage problems.

We study the problem of incentivizing information-sharing in research, under a stylized model: non-identical agents work independently on subtasks of a large project, with dependencies between subtasks captured via an acyclic subtask-network. Each subtask carries a reward, given to the first agent who publicly shares its solution. Agents can choose which subtasks to work on, and more importantly, when to reveal solutions to completed subtasks. Under this model, we uncover the strategic rationale behind certain anecdotal phenomena. Moreover, for any acyclic subtask-network, and under a general model of agent-subtask completion times, we give sufficient conditions that ensure PPS is incentive-compatible for all agents.

One surprising finding is that rewards which are approximately proportional to perceived task-difficulties, are sufficient to ensure PPS in all acyclic subtask-networks. The fact that there is no tension between local fairness and global information-sharing in multi-stage projects is encouraging, as it suggests practical mechanisms for real-world settings. Finally, we also characterize the efficiency of PPS – we show that PPS is necessary, and in many cases, sufficient, to ensure a high rate of progress in research.

Categories and Subject Descriptors: H.4 [Information Systems Applications]: Miscellaneous

Additional Key Words and Phrases: Collaborative Research, Information Sharing, Makespan Minimization

1. INTRODUCTION

Academic research has changed greatly since the Philosophical Transactions of the Royal Society were first published in 1665. However, the process of disseminating research has remained largely unchanged, until very recently. In particular, journal publications have remained to a large extent the dominant mode of research dissemination. However, there is growing belief among researchers that existing systems are inefficient in promoting collaboration and open information-sharing among researchers [Nosek and Bar-Anan 2012]. This is evident in the increasing use of platforms such as ArXiv, StackExchange, the Polymath project, etc., for collaboration and research dissemination; Nielsen [2011] provides an excellent survey of this development.

Open information sharing in research *reduces inefficiencies due to researchers duplicating each other's work*. A researcher who has made a breakthrough in a part of a large problem can help speed up the overall rate of research by sharing this partial-

Siddhartha Banerjee and Ashish Goel (sidb,ashishg@stanford.edu) are with the Department of Management Science and Engineering, and Anilesh K. Krishnaswamy (anilesh@stanford.edu) is with the Department of Electrical Engineering at Stanford University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'14, June 8–12, 2014, Stanford, CA, USA.

ACM 978-1-4503-2565-3/14/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2600057.2602888>

progress – however withholding this information is often in her self-interest. Partial-progress sharing forms the basis of new online platforms for collaborative research like the *Polymath project* [Gowers 2013], where mathematicians work on open problems with discussions occurring online in full public view. Polymath and similar platforms, though successful, are currently small, and participants often do not receive any formal credit for their contributions. However mechanisms for allocating credit for partial-progress are increasingly being recognized as an important primitive for scaling such systems. For example, the results of the first polymath were published under a generic pseudonym – subsequently, the Polymath wiki started publishing details regarding individual contributions, with grant acknowledgements.

Incentivizing *partial-progress sharing* in research forms the backdrop for our work in this paper. In particular, we formally justify the following assertions:

- *Partial-progress sharing is critical for achieving the optimal rate of progress in projects with many researchers.*
- *Existing reward-structures often disincentivize researchers from publicly sharing their partial progress.*
- *However, it is possible to reward subtasks of a project so as to make partial-progress sharing incentive-compatible for all researchers.*

Though these assertions appear intuitive, there are few formal ways of reasoning about them. We develop a model of research dynamics with which we can study these topics.

1.1. A Model for Strategic Behavior in Collaborative Projects

We develop a stylized model for dynamics and information-sharing in research. For brevity, we refer to it as the *Treasure-hunt game*, as it is reminiscent of a treasure-hunt where agents solve a series of clues towards some final objective, and solving certain sets of clues ‘unlocks’ new clues.

We consider n non-identical agents working on a large project, subdivided into m inter-dependent subtasks. The dependencies between subtasks are captured via a *subtask-network* – a directed acyclic graph, where edges represent subtasks. A subtask u becomes *available* to an agent once she knows the solution to *all* the predecessor subtasks – those from which subtask u is reachable in the subtask-network. The entire project is completed once all subtasks are solved.

For example, a project consisting of subtasks arranged in a line requires agents to execute the subtasks in order starting from the first, and moving to the next subtask only after completing the preceding one. On the other hand, a network consisting of two nodes with parallel edges between them represents a project with complementary subtasks – each can be solved independently, but all subtasks need to be solved to complete the project. More examples are given in figure 1.

We assume that *agents work on subtasks independently*. If agent i focusses her complete attention on subtask u , then she solves it after a random time interval, drawn from an exponential distribution $Exp(a_i(u))$ – the rate parameter $a_i(u)$ corresponds to




Fig. 1. Examples of the Treasure-Hunt game: (a) Linear subtask-network – Agent 1 has solved (but not shared) subtask u and can start subtask v . The rest must first finish subtask u . (b) More complex acyclic subtask-network. Solid subtasks (a, b) are completed, while green subtasks (c, d, e) are available for solving.

the aptitude of agent i for task u . While most of our results are derived for a general aptitude matrix, we obtain sharper results in the case of separable-aptitudes (SA):

Definition 1.1 (*Separable-Aptitudes (SA) Model*). The agent-task aptitudes decompose as $a_i(u) = a_i \cdot s_u$, where:

- a_i is the innate *ability* of agent i , which is independent of the tasks.
- s_u is the *simplicity* of task u , which is independent of the agents.

We assume that agents know the complete aptitude matrix – however, as we discuss later, our results hold under much weaker assumptions. If an agent has more than one subtask available for solving, then we assume she can split her effort between them.

Subtasks are associated with rewards, given to the first agent to *publicly share* its solution – this captures the standard norm in academia, sometimes referred to as the *Priority Rule* [Strevens 2003]. We assume that the *rewards are fixed exogenously*.

Upon solving a subtask, an agent can choose when to share the solution. Sharing without delay helps speed up the overall rate of research; we formalize this as follows:

Definition 1.2 (*Partial-Progress Sharing (PPS) Policy*). An agent i is said to follow the *partial-progress sharing* policy if, upon solving a subtask, she immediately shares the solution with all other agents.

Agents can alternately defer sharing solutions to an arbitrary future time, and continue working on downstream subtasks, while others remain stuck behind. For example, an agent may choose to share her partial-progress after a fixed delay, or until she solves the next subtask and can claim rewards for multiple subtasks together.

The Treasure-hunt game captures the strategic behavior of agents with respect to sharing information. The strategy space of agents in the game consists of: (i) the delay in announcing the solution to a subtask, and (ii) choice of subtask when multiple are available. On the other hand, for the social planner, there are two natural objectives:

- (1) To *ensure partial-progress sharing* – choosing reward-vectors in a way that agents are incentivized to publicly share solutions without any delay.
- (2) To *minimize the makespan* – ensuring the project is completed in the least time.

1.2. Two Warmup Examples

To understand the above objectives, we first present two examples. Our first objective is to allocate rewards to incentivize PPS – the following example shows that even in simple subtask-network, the rewards affect agent strategies in a non-trivial way:

Example 1.3. Consider a linear subtask-network under the SA model, with two subtasks $\{p, q\}$ (p preceding q) with simplicities $\{2s, s\}$, and two agents $\{1, 2\}$ with abilities $\{a, a\}$. Suppose the rewards are $R_p = R, R_q < 4R$: then for both agents, following PPS is a dominant strategy. However if instead the rewards are $R_p = R, R_q > 4R$: then for both agents *withholding partial-progress* is a dominant strategy.

PROOF. Suppose the strategy space of an agent i is $S_i = \{P, \bar{P}\}$, where P corresponds to agent i following PPS, and \bar{P} to agent i sharing solutions *only after solving both tasks*¹. For each $(x, y) \in S_1 \times S_2$, let U_{s_1, s_2}^i to be the payoff to agent $i \in \{1, 2\}$ with agent 1 opting for strategy s_1 and agent 2 opting for s_2 . By symmetry (because both agents are identical), we have $U_{P, P}^1 = U_{P, P}^2 = U_{\bar{P}, \bar{P}}^1 = U_{\bar{P}, \bar{P}}^2$, $U_{P, \bar{P}}^1 = U_{P, \bar{P}}^2$ and $U_{\bar{P}, P}^1 = U_{\bar{P}, P}^2$. Further, since it is a constant sum game with total reward $R_p + R_q$, we have: $U_{P, P}^1 + U_{P, P}^2 = U_{\bar{P}, \bar{P}}^1 + U_{\bar{P}, \bar{P}}^2 = U_{P, \bar{P}}^1 + U_{P, \bar{P}}^2 = U_{\bar{P}, P}^1 + U_{\bar{P}, P}^2 = R_p + R_q$.

¹Though sufficient for this example, we note that this is a simplified strategy space – more generally, after solving an open subtask, an agent can choose to share the solution after any arbitrary delay.

Furthermore, considering the strategy profile (\bar{P}, P) , we have:

$$\begin{aligned}
U_{\bar{P}, P}^1 &= \mathbb{P}[1 \text{ solves } p \text{ first}] \cdot \mathbb{P}[1 \text{ solves } q \text{ before } 2 \text{ claims } p | 1 \text{ solved } p \text{ first}] \cdot (R_p + R_q) \\
&\quad + \mathbb{P}[1 \text{ solves } p \text{ first}] \cdot \mathbb{P}[2 \text{ claims } p \text{ before } 1 \text{ solves } q | 1 \text{ solved } p \text{ first}] \cdot \left(\frac{R_q}{2}\right) \\
&\quad + \mathbb{P}[2 \text{ solves } p \text{ first}] \cdot \left(\frac{R_q}{2}\right) \\
&= \frac{1}{2} \cdot \frac{sa}{sa + 2sa} \cdot (R_p + R_q) + \frac{2sa}{sa + 2sa} \cdot \frac{R_q}{2} + \frac{1}{2} \cdot \frac{R_q}{2} = \frac{2R_p + 7R_q}{12}.
\end{aligned}$$

Now we have $\frac{2R_p + 7R_q}{12} > \frac{R_p + R_q}{2} \iff 4R_p < R_q$: thus, if $4R_p < R_q$, then $U_{\bar{P}, P}^1 > U_{P, P}^1$, and by symmetry, $U_{P, \bar{P}}^2 > U_{\bar{P}, P}^2$. Further, since it is a constant sum game, the latter inequality implies $U_{P, \bar{P}}^1 < U_{\bar{P}, P}^1$ (and $U_{P, P}^2 < U_{\bar{P}, P}^2$). Thus \bar{P} , a non-PPS strategy, is a dominant strategy for both players. On the other hand, if $4R_p > R_q$, a similar argument shows that PPS is a dominant strategy for both players. \square

Thus, it is non-trivial to ensure PPS. On the other hand, the following example suggests it is necessary for efficient research, by showing that the makespan in the absence of PPS can be greater by a factor proportional to the number of agents:

Example 1.4. Consider a treasure-hunt game under the SA model on a linear subtask network, with $m \sim \Omega(\log n)$ subtasks, each with simplicity 1, being attempted by n agents, each having ability 1. Suppose all agents agree to follow PPS – then, the expected makespan is $\mathbb{E}[T_{PPS}] = m/n$.

On the other hand, suppose the rewards are distributed so that all subtasks have a reward of 0, except the last subtask which has a reward of 1 – it is easy to see that withholding partial-progress is a dominant-strategy for all agents in this case. Further, the makespan is given by $T_{Non-PPS} = \min_{i \in [n]} \left\{ \sum_{u \in [m]} T_i(u) \right\}$, where $T_i(u) \sim \text{Exp}(1)$. Since $m = \Omega(\log n)$, via a standard Chernoff bound, we have that for all agents, $\sum_{u \in [m]} T_i(u) = \Omega(m)$ with high probability – thus we have that $\mathbb{E}[T_{Non-PPS}] = \Omega(m)$.

Taken together, the examples raise the following two questions: Can one design subtask rewards to ensure PPS under general aptitude matrices and acyclic subgraphs? Is ensuring PPS sufficient to minimize makespan in different settings? Our work provides answers to both these questions.

1.3. Overview of our Results

The Treasure-Hunt Game on a Line: We first focus on the special case of a linear subtask-network (see figure 1(a)). At any time, each agent has only one available subtask; upon solving it, she is free to choose when and with whom to share the solution. In this setting, under the separable-aptitudes model, we have the following result:

MAIN RESULT 1. *On linear subtask-networks under the SA model, suppose rewards $\{R_t\}$ satisfy the following condition – for every pair of subtasks (u, v) s.t. u precedes v :*

$$\frac{R_u s_u}{R_v s_v} \geq \alpha,$$

where $\alpha = \max_{i \in [n]} \left\{ a_i / \sum_{j \in [n]} a_j \right\}$. Then PPS is a Perfect Bayesian Equilibrium.

Some comments and extensions (for details, refer Section 2.3):

- *Sufficient condition for general aptitude matrix:* The above result follows from more general conditions that ensure PPS under *any aptitude matrix* (Theorem 2.1). However, under the SA model, the condition admits the following nice interpretation:
- *The Proportional-Allocation Mechanism:* Under the SA model, for any agent, the ratio of expected completion times for any two tasks is inversely proportional to their simplicities. Thus, a natural ‘locally-fair’ mechanism is to set subtask rewards to be inversely proportional to their simplicity – the above result shows that this surprisingly also ensures global information sharing.
- *Approximate Proportional-Allocation:* Moreover, the above result shows that PPS is incentive compatible even under *approximate* proportional-allocation – as long as rewards for earlier subtasks are at least an α fraction of rewards for later subtasks. Note that α corresponds to the fraction of the total ability possessed by the best agent – in particular, $\alpha < 1$.
- *PPS payoffs in the core:* Our results extend to transferable-utility cooperative settings, where agents can form groups – in particular, we get conditions for ensuring that PPS is in the core. Moreover, the modified conditions admit the proportional-allocation mechanism under the SA model.
- *Uniqueness of equilibrium with Stackelberg agents:* Finally, we modify the above result to show that following PPS can be shown to be the *unique* Perfect Bayesian Equilibrium by admitting Stackelberg strategies – wherein a group of agents ex-ante commit to following PPS. In particular, we show that a sufficient condition for uniqueness is that for any task, the Stackelberg agents together dominate any other individual agent.

The Treasure-Hunt Game on Acyclic Subgraph-Networks: Analyzing the treasure-hunt game is more complicated in general acyclic subtask-networks where, at any time, an agent may have several available subtasks, and can split her effort between them. In particular, given available subtasks M , an agent i can choose a subtask according to any distribution $\{x_i(u)\}_{u \in M}$. Despite this added complexity, under the SA model we obtain an extension to our previous result for the line:

MAIN RESULT 2. *On an acyclic subtask-network under the SA model, suppose rewards $\{R_t\}_{t \in [m]}$ satisfy the following condition: for every pair of subtasks (u, v) such that v is reachable from u in the subtask-network, we have:*

$$R_u s_u \geq R_v s_v.$$

Then PPS is a Perfect Bayesian Equilibrium.

Some comments and extensions regarding this setting (see Section 3 for details):

- Unlike linear subtask-networks, we can no longer admit approximate proportional-allocation because earlier tasks need to be weighted more – however, exact proportional allocation still incentivizes PPS.
- We again get conditions under which PPS is in the core, and sufficient conditions on Stackelberg agents to ensure PPS is the unique equilibrium.
- For a *general aptitudes matrix*, we provide a constructive way to set rewards so as to ensure PPS in any acyclic subtask-network (Theorem 3.2).
- We also show that ensuring PPS indirectly results in all agents working on the same subtask at any time. Such behavior, referred to as *herding* [Strevens 2013], has been reported in empirical studies [Boudreau and Lakhani 2013a].

The Efficiency of PPS in the Treasure-Hunt Game: Finally, we explore the connections between PPS and minimum makespan. Although PPS intuitively appears to be a pre-requisite for minimizing makespan, this may not be true in general acyclic networks. To quantify the efficiency of PPS, we study the ratio of the *optimal makespan*

to that under any reward-vector which incentivizes PPS. Note that under conditions wherein PPS is the unique equilibrium, this is a natural notion of *price of anarchy* (POA) for the Treasure-Hunt game.

In linear subtask-networks, it is not hard to see that under any aptitude-matrix, *ensuring PPS is necessary and sufficient for minimizing the makespan*. For a general acyclic network however, we show that the ratio of the makespan under PPS to the optimal can be unbounded. Significantly, however, we show that the ratio scales *only in the number of tasks, and not the number of agents*. Furthermore, in the case of the SA model, we have the following result:

MAIN RESULT 3. *Consider the Treasure-hunt game on any acyclic subtask-network under the SA model. Suppose rewards are chosen to ensure that all agents follow PPS. Then any resulting equilibrium also minimizes the expected makespan.*

Note that the above result does not reference details of how agents choose which subtask to solve – PPS alone is sufficient to minimize makespan.

Interpreting our model and results: Collaboration/competition in research involves many factors – trying to incorporate all of these in a model may make the strategy-space so complicated as to not allow any analysis. On the other hand, with a stylized model like ours, there is a danger that some phenomena arise due to modeling artifacts. We address this concern to some degree in Section 5, where we argue that our qualitative results are not sensitive to our assumptions.

More significantly, our model captures the main tension behind information sharing in research – partial-progress in a problem by one agent creates an externality in the form of information asymmetry. To promote PPS, the reward for sharing must compensate the agent for removing this externality. Our model lets us quantify this ‘price of information’ – in particular, we show that proportional-reward mechanisms indirectly help ensure partial-progress sharing. We also show that PPS, while necessary, is also often sufficient, for efficiency in large-scale research.

Our work suggests that an effective way to conduct research on a big project is to first break the project into many small subtasks, and then award credits for subtasks proportional to their difficulty. The former recommendation matches the evidence from the Polymath project [Cranshaw and Kittur 2011; Gowers 2013], which is based largely on this idea of aggregating small contributions. The latter suggests that in order to scale such platforms, it is important to have mechanisms for rewarding contributions. The fact that that simple mechanisms suffice to ensure PPS also points the way to potentially having these rewards defined endogenously by the agents.

Technical Novelty: Our model, though simple, admits a rich state space and agent strategies, which moreover affect the state transitions. Thus, a priori, it is unclear how to characterize equilibria in such settings. Our results follow from careful induction arguments on appropriate subgames. We define a class of subgames corresponding to nested pairs of connected subgraphs of the network. In general acyclic networks, in addition to PPS, we also show that agents follow a herding strategy – this turns out to be crucial for characterizing equilibria. Finally, our efficiency result under the SA model is somewhat surprising. In a related one-shot game where agents need to choose one among a set of available tasks, Kleinberg and Oren [2011] show that under a similar separable-aptitudes model, though there exist reward-structures achieving the social optimum (in their case, maximizing the number of completed tasks), finding these rewards is NP-hard. In contrast, we show that ensuring PPS alone is sufficient to minimize makespan, and give simple sufficient conditions for the same. Thus, optimizing for the makespan in dynamic settings appears to circumvent the hardness of static settings under the SA model.

1.4. Related Work

Understanding the strategic aspects of contests for limited resources has a long history [Konrad 2009]. Recently, these models have been adapted to study distribution of credit among researchers. Ackerman and Brânzei [2012] study the impact of authorship order on research; Bachrach et al. [2013] study collaboration across multiple projects under simple local-sharing rules; Kleinberg and Oren [2011] study the effect of the priority rule on agents choosing between open problems, and show how redistributing rewards can make the system more efficient. Our work shares much with these works, in particular, in the idea of engineering rewards to achieve efficiency. A fundamental difference however is that these works consider static/one-shot settings.

Another relevant line of work focusses on dynamics in R&D settings. Taylor [1995] and Harris and Vickers [1987] model R&D contests as a one-dimensional race. They focus only on understanding the equilibrium strategies of the contestants, assuming fixed rewards. Moreover, information-sharing dynamics are ignored in these models.

The information-sharing aspect of our work shares similarities with models for crowdsourcing and online knowledge-sharing forums. Chawla et al. [2012] model crowdsourcing platforms as all-pay auctions, where agents pay a cost to participate – this however is not the case in research, where agents derive utility only by solving open problems in their chosen field. Information-sharing also forms the backdrop for models for online Q&A forums. Jain et al. [2012] and Ghosh and Hummel [2012] study a setting where rewards are given to the top answers to a question. Each user has some ex ante knowledge, and by submitting an answer later, can aggregate earlier answers. In research however, knowledge does not exist ex ante, but is created while working on problems, as is captured in our model. Finally, information dynamics are studied by Ghosh and Kleinberg [2013] in the context of online education forums – though differing in topic, our work shares much of their modeling philosophy.

Finally, there is a growing body of empirical work studying open information in collaborative research. Cranshaw and Kittur [2011] study the Polymath projects, and conclude that better incentive structures are required to make such endeavors sustainable and scalable. The studies of Boudreau and Lakhani [2013a,b] help quantify the value of open information-sharing in collaborative research, via experiments in Topcoder contests. By offering incentives for partial-progress sharing, they observe that it increases the efficiency of problem solving. They also observe that open information leads to herding in these settings – both these findings correspond to behavior we observe in our models.

Outline: The remaining paper is organized as follows: In Section 2, we consider projects with linear subtask-networks – here, we obtain a complete characterization of conditions under which PPS is a Perfect Bayesian equilibrium. In Section 3, we extend these results to general acyclic subtask-networks. Next, in Section 4, we characterize the efficiency, with respect to the makespan, with and without PPS in different settings. We conclude with a discussion in Section 5. Section 6 presents proofs for some of our results, which illustrate our main ideas and techniques; due to lack of space, complete proofs are deferred to our technical report [Banerjee et al. 2014].

2. THE TREASURE-HUNT GAME ON LINEAR SUBTASK-NETWORKS

2.1. Extensive-Form Representation of the Treasure-Hunt Game

We denote $[n] = \{1, 2, \dots, n\}$. The treasure-hunt game on a line is defined as follows:

- **Task Structure:** A project is divided into a set $[m]$ of smaller subtasks, arranged in a linear subtask-network. Subtask u becomes available to an agent only when she knows solutions to all preceding subtasks; the overall project is completed upon

solving the final subtask. A set of n non-identical agents are committed to working on the project. Each agent wants to maximize her own expected reward.

- **Rewards:** Subtask u has associated reward $R_u \geq 0$, awarded to the first agent to publicly share its solution. We assume that *the rewards are fixed exogenously*.
- **System Dynamics:** We assume that *agents work on subtasks independently*. Agent i , working alone, solves subtask u after a delay of $T_i(u) \sim \text{Exp}(a_i(u))$, where we refer to $a_i(u)$ as the aptitude of agent i for subtask u . For the special case of separable-aptitudes (see Definition 1.1), we have $a_i(u) = a_i s_u$, where we refer to $\{a_i\}_{i \in [n]}$ as the agent-ability vector, and $\{s_u\}_{u \in [m]}$ as the subtask-simplicity vector.
- **Strategy Space:** The game proceeds in continuous time – for ease of notation, we suppress the dependence of quantities on time. If an agent i solves subtask u at some time t , then she can share the solution with any set of agents $A \subseteq [n]$ at any subsequent time – if the solution is shared with all other agents, we say it is *publicly shared*. When agent i shares the solution to subtask u with an agent j , then j can start solving the subtask following u , bypassing all intermediate subtasks.
- **Information Structure:** We assume all agents ex ante know the aptitude-matrix $\{a_i(t)\}_{i \in [n], t \in [m]}$ – later we show how this can be relaxed. During the game, we assume agents only know their own progress, and what is shared with them by others.

2.2. Progress Sharing in the Treasure-Hunt Game

Our main result in this section is a sufficient condition on rewards to incentivize PPS:

THEOREM 2.1. *Consider the Treasure-hunt game on a linear subtask-network G , and a general agent-subtask aptitude matrix $\{a_i(u)\}$. Suppose the rewards-vector $\{R_u\}$ satisfies the following: for any agent i and for any pair of subtasks u, v such that u precedes v , the rewards satisfy:*

$$\frac{R_u a_{-i}(u)}{R_v a_{-i}(v)} \geq \frac{a_i(v)}{a_i(v) + a_{-i}(v)}, \quad (1)$$

where for any task w , we define $a_{-i}(w) \triangleq \sum_{j \neq i} a_j(w)$. Then for all agents, partial-progress sharing (PPS) is a Perfect Bayesian Equilibrium.

Note that the first scenario in Example 1.3 satisfies the above condition, while the second violates it. The result follows from a careful backward-induction argument – refer Section 6 for the proof. Though the condition in Theorem 2.1 is somewhat hard to interpret, it admits several interesting extensions, as we discuss next.

2.3. The SA model, and stronger equilibrium concepts:

- *Approximate Proportional-Allocation in the SA model:* Applying the theorem to the SA model gives the following corollary (presented as Main Result 1 in Section 1.3):

COROLLARY 2.2. *Consider the Treasure-hunt game on linear subgraph-network G under the SA model. Suppose $A \triangleq \sum_{i \in [n]} a_i$, and $\alpha_{NE} \triangleq \max_{i \in [n]} \frac{a_i}{A}$. Then the PPS policy is a Nash equilibrium for all agents if the rewards-vector satisfies:*

$$\frac{R_u s_u}{R_v s_v} \geq \alpha_{NE}, \quad \forall (u, v) \text{ s.t. } u \text{ precedes } v \text{ in } G$$

Note that under the SA model, for any agent i , the ratio of expected completion time for subtasks u and v obeys $\frac{\mathbb{E}[T_i(u)]}{\mathbb{E}[T_i(v)]} = \frac{s_v}{s_u}$. Thus we can reinterpret the above condition in terms of an *approximate proportional-allocation mechanism*, which rewards subtasks proportional to the expected time they take to solve. Observe also that α_{NE} is the fraction of the total ability possessed by the top agent. In a sense, if all agents

act selfishly, then the top agent gains the most by deviating from PPS; the above result captures that the rewards should be calibrated so as to prevent this.

- *PPS is in the Core of Cooperative Settings:* Theorem 2.1 can be extended to derive a sufficient condition for PPS to be in the core, i.e., to ensure that no coalition of agents, who share partial-progress and rewards amongst themselves, have an incentive to deviate from the PPS policy. We state this result for the SA model, but it can easily be extended to a general aptitudes matrix:

COROLLARY 2.3. *Consider the Treasure-hunt game on linear subtask-network G , under the separable-aptitudes model. Define $A \triangleq \sum_{i \in [n]} a_i$ and $\alpha_C \triangleq 1 - \min_{i \in [n]} \frac{a_i}{A}$. Then the PPS policy is in the core if the rewards-vector satisfies:*

$$\frac{R_u s_u}{R_v s_v} \geq \alpha_C, \quad \forall (u, v) \text{ s.t. } u \text{ precedes } v \text{ in } G$$

The corollary follows from Theorem 2.1 from observing that under exponentially distributed completion-times, if a group of agents form a coalition, then their total aptitude for a subtask is the sum of their aptitudes. Note that $\alpha_{NE} \leq \alpha_C < 1$ – thus an approximate proportional-allocation mechanism is in the core, but with a stricter approximation factor.

- *Stackelberg leaders and dominant strategies:* One subtle point regarding Theorem 2.1 is that the resulting equilibrium need not be unique. For two agents, an argument similar to Example 1.3 can be used to show uniqueness – this approach does not easily extend to more players. However, we can circumvent this difficulty and find conditions under which PPS is the unique Perfect Bayesian equilibrium, by considering settings with *Stackelberg leaders* – agents who *ex ante* commit to following the PPS policy. We state this here for the SA model:

COROLLARY 2.4. *Consider the Treasure-hunt game on linear subtask-network G under the SA model, with $A \triangleq \sum_{i \in [n]} a_i$. Suppose a subset of agents $S \subseteq [n]$ ex ante commit to PPS. Define $\alpha_S = \max_{i \in [n]} \frac{a_i}{a_S} \frac{a_{-i}}{a_i + a_{-i}}$, where $a_S = \sum_{i \in S} a_i$. Then PPS is the unique Perfect Bayesian Equilibrium if the rewards satisfy:*

$$\frac{R_u s_u}{R_v s_v} \geq \alpha_S, \quad \forall (u, v) \text{ s.t. } u \text{ precedes } v \text{ in } G$$

Note that if $a_S \geq a_i$, then $\alpha_S < 1$ – thus if the Stackelberg agents dominate any individual agent, then proportional allocation results in PPS. Such behavior mirrors real-life settings, wherein when top researchers commit to open information sharing, then many others follow suit. An example of this can be seen in the Polymath project, where a core group of mathematicians first committed to open discussion, and this led to others joining the project.

- *Weaker Requirements on Ex Ante Knowledge:* We assumed that all agents are aware of the entire aptitude matrix, and use this to determine their strategy. Note however that the conditions in Theorem 2.1 depend only on an agent’s own aptitude vector, and the *aggregate* aptitudes of the remaining agents. In particular, in the SA model, an agent i needs to only know $\alpha_i = a_i/A$, the ratio of her ability to the total ability of all agents, in order to determine if PPS is a best response or not. On the other hand, to correctly design exogenous rewards, a social planner needs to know the simplicities of all subtasks. In case these estimates are not perfect, but are within a $1 \pm \epsilon$ factor of the true values, then as long as $\epsilon > (1 - \alpha_{NE})/(1 + \alpha_{NE})$, then proportional-allocation using the noisy estimates still incentivizes PPS.

3. THE TREASURE-HUNT GAME ON GENERAL ACYCLIC SUBTASK-NETWORKS

We now extend the results from the previous section to general acyclic subtask-networks. The subtask-network now encodes the following dependency constraints: a subtask u with source node v_u can be attempted by an agent i only after i knows the solutions to *all* subtasks terminating at v_u (i.e., all the in-edges of node v_u). Note that at any time, an agent may have multiple available subtasks. We assume that agent i chooses from amongst her available subtasks M according to some strategically-chosen distribution $\{x_i(u)\}_{u \in M}$ – as a result, subtask $u \in M$ is solved by i in time $Exp(x_i(u)a_i(u))$. The remaining model is as described in Section 2.1.

A technical difficulty in acyclic subtask-networks is that even assuming agents follow PPS, there is no closed-form expression for agent’s utility in a subgame – consequently, at any intermediate stage, there is no simple way to determine an agent’s choice of available subtask. In spite of this, we can still derive a sufficient condition on rewards to ensure PPS – this also gives a constructive proof of existence of reward-vectors that ensure PPS in any acyclic subtask-network.

We first need some additional definitions. For any task u and any agent $i \in [n]$, we define $a_{-i}(u) = \sum_{j \in [n] \setminus \{i\}} a_j(u)$ – we introduce the following notion of a virtual reward:

Definition 3.1 (Virtual rewards). For any agent $i \subseteq [n]$ and subtask $u \in [m]$, the virtual-reward $\gamma_i(u)$ is defined as:

$$\gamma_i(u) \triangleq \frac{R_u a_i(u) a_{-i}(u)}{a_i(u) + a_{-i}(u)}.$$

Using this definition, we have the following theorem:

THEOREM 3.2. *Consider the Treasure-hunt game on a directed acyclic subtask-network G , and a general agent-subtask aptitude matrix $\{a_i(u)\}$. Suppose the reward-vector $\{R_u\}$ satisfies the following conditions:*

- (Monotonicity) *There exists a total ordering \prec on subtasks $[m]$ such that for every agent i and pair of subtasks u, v , we have $v \prec u \iff \gamma_i(v) < \gamma_i(u)$.*
- *For all pairs of subtasks (u, v) such that v is reachable from u in G , we have $v \prec u$.*

Then the following strategies together constitute a Perfect Bayesian equilibrium:

- *Every agent implements the PPS policy.*
- *At any time, if $M_o \subseteq [m]$ is the set of available subtasks, then every agent i chooses to work on the unique subtask $u^* = \arg \max_{u \in M_o} \gamma_i(u)$*

The proof of Theorem 3.2, which is provided in Section 6, is again based on a careful backward induction argument. The induction argument, however, is more involved than in Theorem 2.1 – in particular, when multiple subtasks are available to each agent, the utility of the subgame for an agent depends on the strategies of all other agents. We circumvent this by observing that under the given conditions, besides following PPS, *agents also have an incentive to choose the same subtask as other agents.* We henceforth use the term *herding* to refer to this behavior.

Theorem 3.2 can be summarized as stating that as long as virtual rewards are monotone and have the same partial-ordering as induced by the subtask-network, then *PPS + herding is a Nash equilibrium.* As with the condition in Theorem 2.1, this result may be difficult to interpret. Note however that given *any acyclic subtask-network, and any aptitudes matrix* with non-zero aptitudes, it gives a constructive way to choose $\{R_u\}$ to ensure PPS. The observation that ensuring PPS may also result in herding is an interesting outcome of our analysis – such behavior has in fact been observed in research settings, for example, Boudreau and Lakhani [2013a] report such behavior in contests with incentives for partial-progress sharing.

In the SA model, note that for every agent i , the virtual reward $\gamma_i(u) \propto R_u s_u$. Thus for any rewards vector, the monotonicity property is satisfied. Moreover, we obtain the following sharper characterization of sufficient conditions for ensuring PPS:

THEOREM 3.3. *Consider the treasure-hunt game under the SA model, on a directed acyclic subtask-network with m subtasks, associated simplicities $\{s_k\}_{k \in [m]}$ and exogenous (non-zero) rewards $\{R_k\}_{k \in [m]}$, and n agents with associated abilities $\{a_i\}_{i \in [n]}$. If the rewards satisfy:*

- v is reachable from $u \implies R_v s_v \leq R_u s_u$.

Then the following strategies jointly constitute a Perfect Bayesian Equilibrium:

- *Every agent follows the PPS policy.*
- *At any time, if M_o denotes the set of open subtasks, then for any agent i , the support of her choice distribution $\{x_i(u)\}_{u \in M_o}$ is contained in $\arg \max_{u \in M_o} \{R_u s_u\}$.*

Note that in case of perfect proportional-allocation, agents are indifferent between all available subtasks – thus herding no longer occurs in this setting.

Finally, as in linear subtask-networks, we can extend the above result to get conditions for PPS to be in the core, and also, for it to be a dominant strategy for agents in the presence of appropriate Stackelberg agents. Details of these extensions are provided in our technical report [Banerjee et al. 2014].

4. THE EFFICIENCY OF PPS IN THE TREASURE-HUNT GAME

Having obtained conditions for incentivizing PPS, we next turn to the problem of minimizing the *expected makespan* of a multi-stage project. The importance of PPS is clear in linear subtask-networks, where *ensuring PPS is necessary and sufficient for minimizing the expected makespan*. Essentially, under PPS, all agents work together on a single current subtask, and thus each subtask gets solved at a rate equal to the sum of rates of the agents. On the other hand, as we show in Example 1.4, in settings without PPS, each subtask is solved at a rate close to that of the maximum single-agent rate.

An ideal result would be if ensuring PPS alone is sufficient to minimize makespan in all acyclic subtask-networks. However the following example shows this is not possible:

Example 4.1. Consider a treasure-hunt game with m parallel subtasks and m agents. We index agents and subtasks with unique labels from $[m]$. Suppose the aptitudes are given by: $a_i(u) = 1$ if $i = u$, else $1/(m-1)$. From Theorem 3.2, we know that we can set rewards so as to ensure that all agents follow PPS + herding, i.e., they solve the subtasks as a group in a fixed order. Thus we have that $\mathbb{E}[T_{PPS}] = \frac{m}{1+(m-1)/(m-1)} = m/2$. On the other hand, suppose agent i starts on task $u = i$, and stops after finishing it – clearly this is an upper bound on the optimal makespan, and so we have $\mathbb{E}[T_{OPT}] = \frac{1}{m} + \frac{1}{m-1} + \dots + 1 = \Theta(\log m)$.

Thus, in the equilibrium wherein all agents follow PPS, the expected makespan is greater than the optimal by a multiplicative factor of $\Omega(m/\log m)$. Conversely, however, we can show that this is in fact the worst possible ratio up to a logarithmic factor:

THEOREM 4.2. *Given an acyclic subtask-network G with edge set $[m]$, suppose the rewards-vector satisfy the conditions in Theorem 3.2. Then the expected makespan of the corresponding equilibrium where all agents follow PPS is no worse than m times the optimal makespan, irrespective of the number of agents.*

Recall in Example 1.4, we show that the expected makespan *without PPS* can be greater by an $\Omega(n)$ factor. Theorem 4.2 shows that under PPS, the makespan can be off by a factor depending on the number of tasks, but *not the number of agents*. This is significant in large collaborative projects, where the number of tasks often far exceeds

the number of agents. The proof of Theorem 4.2 is provided in [Banerjee et al. 2014]. Moreover, under the SA model, we get a surprising optimality result:

THEOREM 4.3. *Given any acyclic subtask-network G under the SA model, with edge set $[m]$ and agents $[n]$. Then any corresponding equilibrium where all agents follow PPS also minimizes the expected makespan.*

In particular, setting rewards to satisfy the conditions in Theorem 3.3 *minimizes the makespan*. The proof of this result is given in Section 6.

5. DISCUSSION

We focus on a setting with *compulsory participation* – agents derive utility only by solving subtasks. A vast majority of research is done by professional researchers, committed to working on their chosen topics, which makes the act of research different from settings such as crowdsourcing [Chawla et al. 2012]. Moreover, our results are not sensitive to the number of researchers engaged in a project – the question of partial-progress sharing arises as soon as there is more than one participant.

We assume that the time an agent takes to solve a task has an *exponential distribution*. Such an assumption is common when considering dynamics in games as it helps simplify the analysis (for example, see [Ghosh and Kleinberg 2013]). We conjecture however that the equilibria in the treasure-hunt game remain qualitatively similar under a much broader class of distributions. One reason for this is that assuming exponentially distributed completion times is, in a sense, *over-optimistic*. For an agent trying to decide whether or not to share a breakthrough, assuming memoryless completion times essentially corresponds to assuming other agents are starting from scratch, discounting any time already invested in the problem.

We assume that agents work on tasks independently. This allows us to focus on information sharing. However, as discussed in Section 2.3, our results extend to give conditions for the payoffs-vector resulting from everyone following PPS to be in the core.

Partial-progress sharing is one of several primitives required for enabling true collaborative research. However, there are very few good formal ways for reasoning about such settings. The model we develop for studying information sharing in collaborative research, though stylized, exhibits a variety of phenomena which correspond to observed behavior among researchers. It captures the uncertain dynamics of research, while retaining analytic tractability. Although we focus on understanding incentives for PPS, our model may prove useful for understanding other related questions.

6. PROOFS FOR SELECTED RESULTS

Due to lack of space, we only provide proofs for Theorems 2.1, 3.2 and 4.3. For complete proofs of all the results, please refer to our technical report [Banerjee et al. 2014]

PROOF OF THEOREM 2.1. We number the subtasks from the end, with the last subtask being denoted as 1, and the first as m . Fix an agent $i \in [n]$, and assume every other agent always follows the PPS policy. We define $\mathcal{G}_{k,l}^i, 0 \leq k \leq l \leq m$, to be the subgame where agent i starts at subtask k , and all other agents start at subtask l . Note that agent i starts either at par or ahead of others – this is because we assume that all other agents follow PPS. When $k < l$, agent i is ahead of the others at the start – however she has not yet revealed the solutions to her additional subtasks. In this setting, we claim the following invariant: under the conditions specified in the theorem, PPS is a best response for agent i in *every* subgame $\mathcal{G}_{k,l}^i, 1 \leq k \leq l \leq m$.

First, some preliminary definitions: given two subgames $\mathcal{G}_{k,l}^i$ and $\mathcal{G}_{p,q}^i$, we say that $\mathcal{G}_{k,l}^i$ is smaller than $\mathcal{G}_{p,q}^i$ if $k < p$ OR $k = p$ and $l < q$. Also, assuming the invariant is true, the expected reward earned by agent i in subgame $\mathcal{G}_{k,l}^i$ is given by:

$$R(\mathcal{G}_{k,l}^i) = \sum_{u=k+1}^l R_u + \sum_{u=1}^k R_u \cdot \left(\frac{a_i(u)}{a_i(u) + a_{-i}(u)} \right),$$

where $a_{-i}(u) = \sum_{j \in [n] \setminus \{i\}} a_j(u)$. Finally, when $k < l$, we refer to the (non-empty) subtasks $\{l, l-1, \dots, k+1\}$ as the *captive subtasks* of agent i .

We prove the invariant by a joint induction on k and l . For the base case, consider the subgames $\mathcal{G}_{0,l}^i$, which correspond to a game where i knows solutions to all subtasks – clearly following PPS (i.e., sharing the solutions at $t = 0$) is a best response. Now we split the remaining subgames into two cases: $k < l$ and $k = l$.

• *Case i.* ($\mathcal{G}_{k,l}^i, 1 \leq k < l$): Assume game $\mathcal{G}_{k,l}^i$ starts at $t = 0$, and the invariant holds for all smaller games. If agent i follows PPS, and publicly shares her captive subtasks at $t = 0$, then $\mathcal{G}_{k,l}^i$ reduces to the smaller subgame $\mathcal{G}_{k,k}^i$. Suppose instead i chooses to defect from PPS. Let τ_1 to be the first time when either agent i solves subtask k OR some agent $j \neq i$ solves subtask l . Via a standard property of the minimum of independent exponential random variables, we have $\tau_1 \sim \min\{\text{Exp}(a_i(k)), \text{Exp}(a_{-i}(l))\} \sim \text{Exp}(a(i, k, l))$, where we define $a(i, k, l) = a_i(k) + a_{-i}(l)$.

Next, we observe that it is sufficient to consider deviations from PPS of the following type: i chooses a fixed delay $\tau > 0$, and publicly shares the solutions to captive subtasks at time $t = \min\{\tau_1, \tau\}$. This follows from the following argument (*):

- Suppose $\tau_1 < \tau$, and assume some agent $j \neq i$ solves subtask l before i solves subtask k . Then agent j immediately shares the solution to subtask l (as other agents follow PPS). This reduces the game to $\mathcal{G}_{k,l-1}^i$, wherein by our induction hypothesis, i immediately shares her remaining captive subtasks, i.e., $\{l-1, l-2, \dots, k+1\}$.
- If $\tau_1 < \tau$, but agent i solves subtask k before any other agent solves subtask l , then the game reduces to $\mathcal{G}_{k-1,l}^i$, wherein again by our induction hypothesis, i immediately shares her captive subtasks, i.e., $\{l, l-1, \dots, k\}$.
- Any randomized choice of τ can be written as a linear combination of these policies. To complete the induction step for $\mathcal{G}_{k,l}^i$, we need to show that $\tau = 0$ is the best response for i in the game $\mathcal{G}_{k,l}^i$. Let $U_{k,l}^i(\tau)$ be the expected reward of agent i if she chooses delay τ . Also, for any subtask u , let $a(u) = a_i(u) + a_{-i}(u)$. Then for $\tau = 0$, we have:

$$U_{k,l}^i(0) = \sum_{u=k+1}^l R_u + \sum_{u=1}^k R_u \cdot \left(\frac{a_i(u)}{a(u)} \right) = R_l + \Delta + \frac{a_i(k)}{a(k)} \cdot R_k + R(\mathcal{G}_{k-1,k-1}^i),$$

where we define $\Delta = \sum_{j=k+1}^{l-1} R_j$. Next we turn to the case of $\tau > 0$. Suppose $\tau_1 < \tau$: conditioned on this, we have that i solves subtask k before any other agent solves subtask l with probability $\frac{a_i(k)}{a(i, k, l)}$ (from standard properties of the exponential distribution). Now for any $\tau > 0$, we can write:

$$\begin{aligned} U_{k,l}^i(\tau) = & \mathbb{P}[\tau_1 \geq \tau] \left(\sum_{u=k+1}^l R_u + R(\mathcal{G}_{k,k}^i) \right) \\ & + \mathbb{P}[\tau_1 < \tau] \left(\frac{a_i(k)}{a(i, k, l)} R(\mathcal{G}_{k-1,l}^i) + \frac{a_{-i}(l)}{a(i, k, l)} R(\mathcal{G}_{k,l-1}^i) \right) \end{aligned}$$

This follows from the argument (*) given above. Expanding the terms, we have:

$$\begin{aligned} U_{k,l}^i(\tau) = & \mathbb{P}[\tau_1 \geq \tau] \left(R_l + \Delta + \frac{a_i(k)}{a(k)} R_k + R(\mathcal{G}_{k-1,k-1}^i) \right) \\ & + \mathbb{P}[\tau_1 < \tau] \frac{a_i(k)}{a(i,k,l)} \left(R_l + \Delta + R_k + R(\mathcal{G}_{k-1,k-1}^i) \right) \\ & + \mathbb{P}[\tau_1 < \tau] \frac{a_{-i}(l)}{a(i,k,l)} \left(\Delta + \frac{a_i(k)}{a(k)} R_k + R(\mathcal{G}_{k-1,k-1}^i) \right) \end{aligned}$$

Simplifying using $a(i,k,l) = a_i(k) + a_{-i}(l)$ and $\mathbb{P}[\tau_1 \geq \tau] + \mathbb{P}[\tau_1 < \tau] = 1$, we get:

$$\begin{aligned} U_{k,l}^i(\tau) = & R_l \left(1 - \mathbb{P}[\tau_1 < \tau] \frac{a_{-i}(l)}{a(i,k,l)} \right) + \Delta \\ & + R_k \left(\frac{a_i(k)}{a(k)} + \mathbb{P}[\tau_1 < \tau] \frac{a_{-i}(k)}{a(k)} \cdot \frac{a_i(k)}{a(i,k,l)} \right) + R(\mathcal{G}_{k-1,k-1}^i) \end{aligned}$$

Subtracting this from our expression for $U_{k,l}^i(0)$, we have:

$$U_{k,l}^i(0) - U_{k,l}^i(\tau) = \frac{\mathbb{P}[\tau_1 < \tau]}{a(i,k,l)} \left(R_l a_{-i}(l) - R_k \frac{a_{-i}(k) a_i(k)}{a(k)} \right).$$

Finally from the condition in Theorem 2.1, we have $\frac{R_l a_{-i}(l)}{R_k a_{-i}(k)} \geq \frac{a_i(k)}{a(k)}$: substituting we get $U_{k,l}^i(0) - U_{k,l}^i(\tau) > 0 \forall \tau > 0$. Thus setting $\tau = 0$ maximizes expected reward.

• *Case ii.* ($\mathcal{G}_{k,k}^i, k > 1$): Assume the invariant holds for all smaller games. In this case, note that the game reduces to either $\mathcal{G}_{k-1,k}^i$ (if i solves subtask k) or to $\mathcal{G}_{k-1,k-1}^i$ (if some agent $j \neq i$ solves subtask l). For both these subgames, we have that PPS is a best response for i via our induction hypothesis. This completes the induction. \square

PROOF OF THEOREM 3.2. We first need some definitions. Recall that in an acyclic subtask-network, a subtask v can be solved by an agent only if she knows the solutions to all subtasks corresponding to edges from which v is reachable – we denote this set of predecessors of v as $L(v)$. Moreover, we define $\mathcal{V} \triangleq \{P \subseteq [m] \mid v \in P \implies L(v) \subset P\}$ as the set of *valid knowledge-subgraphs* – subtasks whose solutions may be jointly known to agents at some time. Finally, given that an agent knows the solution to subtasks $P \in \mathcal{V}$, the complement $P^c = [m] \setminus P$ represents the subtasks still open for the agent. We define the set $\mathcal{S} = \{T \subseteq M \subseteq [m] : M^c, T^c \in \mathcal{V}\}$ – essentially \mathcal{S} contains all nested-pairs of subtask-sets, whose complements are valid knowledge-subgraphs.

Fixing agent i , we define $\mathcal{G}_{M,T}^i, (M, T) \in \mathcal{S}$ to be the subgame where agent i knows the solutions to subtasks T^c , while any agent $j \neq i$ knows the solutions to subtasks M^c . Note that $T \subseteq M$, and $T^c \setminus M^c$ are the captive subtasks of i , i.e., those which i has solved, but not publicly shared. Also, for any set M such that $M^c \in \mathcal{C}$, we define $M_o \subseteq M$ to be the available subtasks in M .

Given this set of subgames, we claim the following invariant: under the conditions specified in the theorem, PPS + herding is a best response for agent i in *every* subgame $\mathcal{G}_{M,T}^i, (M, T) \in \mathcal{S}$. As in Theorem 2.1, we prove this via induction. To do so, we first partition the set \mathcal{S} as follows: given $1 \leq k \leq l \leq m$

$$\mathcal{S}_{l,k} \triangleq \{(M, T) \in \mathcal{S} : |M| = l, |T| = k\}$$

We abuse notation and say subgame $\mathcal{G}_{M,T}^i \in \mathcal{S}_{l,k}$ if $(M, T) \in \mathcal{S}_{l,k}$ – note that this corresponds to agent i having k unsolved subtasks, while the remaining agents have

$l \geq k$ unsolved subtasks. We can restate the above invariant as follows – for all $1 \leq k \leq l \leq m$, and for all games $\mathcal{G}_{M,T}^i \in \mathcal{S}_{l,k}$, PPS + herding is a best response for agent i .

For the base case, consider the set of subgames $\mathcal{G}_{M,T}^i \in \mathcal{S}_{1,1}$, i.e., subgames where all agents have a single open subtask. Clearly PPS + herding is a best response for i in this case. For the inductive step, fixing (l, k) , we assume the invariant is true for all subgames in $\mathcal{S}_{p,q}$ where $q < l$ and $p \leq k$ (i.e., agent i has strictly less open subtasks) or $q = l$ and $p < k$ (i.e., agent i has k open subtasks, but the remaining agents have strictly less open subtasks).

Now, as in Theorem 2.1, we break the induction into two cases: (i) $k < l$, and (ii) $k = l$. Note that $k < l$ corresponds to subgames $\mathcal{G}_{M,T}^i$ with $T \subset M$ – here we need to show that PPS, i.e., i sharing solutions to captive subtasks $T^c \setminus M^c$ at $t = 0$, is a best response. On the other hand, in any subgame $\mathcal{G}_{M,M}^i \in \mathcal{S}_{k,k}$, we need to show herding, i.e., i chooses to work on subtask $v^* = \arg \max_{u \in M_o} \gamma_i(u)$. We note that we do not need to assume/show that i follows herding in subgame $\mathcal{G}_{M,T}^i$ – we show that PPS is a best response regardless of which subtask in T_o that i attempts.

• Case (i) ($\mathcal{S}_{l,k}, 1 \leq k < l$)

Consider subgame $\mathcal{G}_{M,T}^i \in \mathcal{S}_{l,k}$. As in Theorem 2.1, the only deviation strategies consist of agent i choosing a delay parameter $\tau > 0$, and sharing her captive subtasks at time $\min\{\tau, \tau_1\}$, where τ_1 is the time of the first breakthrough in a subtask, either by her, or some agent $j \neq i$. Suppose agent i concentrates her efforts on some subtask $v \in T_o$; further, from our induction hypothesis, we know any agent $j \neq i$ concentrates on subtask $u^* = \arg \max_{u \in M_o} \gamma_j(u)$ – note that this is the same subtask $\forall j \neq i$.

If $\tau = 0$ or $\tau < \tau_1$, then the game reduces to $\mathcal{G}_{T,T}^i$. If $\tau_1 < \tau$ and agent i solves subtask v first, then by our induction assumption, agent i immediately shares her remaining subtasks and the game reduces to $\mathcal{G}_{T-v, T-v}^i$, where we use the notation $T - v = T \setminus \{v\}$. Finally, if $\tau_1 < \tau$ and some agent $j \neq i$ solves subtask u^* , then again agent i follows PPS and the game reduces to $\mathcal{G}_{T-u^*, T-u^*}^i$ – note though that task u^* could be in T^c (i.e., its solution be known to agent i). For the remaining proof, *the only deviation of i from PPS we consider is $\tau = \infty$* . To see why this is sufficient, note that for any τ , in sample paths where $\tau < \tau_1$, agent i shares all her captive subsets (albeit with a delay) before any agent makes an additional breakthrough. The memoryless property of exponential random variables implies that choosing any $\tau < \infty$ only scales the difference in utility with and without PPS, by a factor of $\mathbb{P}[\tau_1 < \tau]^2$.

As before, for subtask u we define $a(u) = \sum_j a_j(u)$, and for any pair of subtasks u, w , we define $a(i, u, w) = a_i(u) + a_{-i}(w)$. Now for any $\mathcal{G}_{M,T}^i \in \mathcal{S}_{l,k}$, defining $U_{M,T}^i$ to be the expected reward earned by agent i by following PPS, we have:

$$U_{M,T}^i = \sum_{u \in T^c \setminus M^c} R_u + \sum_{u \in T} \frac{a_i(u)}{a(u)} R_u.$$

Note also that by our induction hypothesis, this is true for any pair $(P, Q) \in \mathcal{S}$ such that $P \subseteq M, Q \subset T$ or $P \subset M, Q = T$.

Let $U_{M,T}^{i,dev}$ to be the expected payoff of agent i when she deviates from PPS by waiting till τ_1 before sharing solutions. Now if $u^* \in T^c$, we have:

$$U_{M,T}^{i,dev} = \frac{a_{-i}(u^*)}{a(i, u^*, v)} U_{M-u^*, T}^i + \frac{a_i(v)}{a(i, u^*, v)} U_{M, T-v}^i$$

²This fact is somewhat non-trivial, but can be formally derived, following the same steps as in the proof of Theorem 2.1 – we omit it for the sake of brevity.

Subtracting from the expected utility under PPS, we get:

$$\begin{aligned}
U_{M,T}^i - U_{M,T}^{i,dev} &= \frac{a_{-i}(u^*)}{a(i, u^*, v)} (U_{M,T}^i - U_{M-u^*,T}^i) + \frac{a_i(v)}{a(i, u^*, v)} (U_{M,T}^i - U_{M,T-v}^i). \\
&= \frac{a_{-i}(u^*)}{a(i, u^*, v)} (R_{u^*}) - \frac{a_i(v)}{a(i, u^*, v)} \left(\frac{a_{-i}(v)}{a(v)} R_v \right) \\
&\geq \frac{a_{-i}(u^*)}{a(i, u^*, v)} \left(\frac{a_i(u^*)}{a(u^*)} R_{u^*} \right) - \frac{a_i(v)}{a(i, u^*, v)} \left(\frac{a_{-i}(v)}{a(v)} R_v \right) = \frac{\gamma_i(u^*) - \gamma_i(v)}{a(i, u^*, v)}.
\end{aligned}$$

On the other hand, if $u^* \notin T^c$, we have:

$$U_{M,T}^{i,dev} = \frac{a_{-i}(u^*)}{a(i, u^*, v)} U_{M-u^*,T-u^*}^i + \frac{a_i(v)}{a(i, u^*, v)} U_{M,T-v}^i.$$

Again subtracting from $U_{M,T}^i$, we get:

$$\begin{aligned}
U_{M,T}^i - U_{M,T}^{i,dev} &= \frac{a_{-i}(u^*)}{a(i, u^*, v)} (U_{M,T}^i - U_{M-u^*,T-u^*}^i) + \frac{a_i(v)}{a(i, u^*, v)} (U_{M,T}^i - U_{M,T-v}^i), \\
&= \frac{a_{-i}(u^*)}{a(i, u^*, v)} \left(\frac{a_i(u^*)}{a(u^*)} R_{u^*} \right) - \frac{a_i(v)}{a(i, u^*, v)} \left(\frac{a_{-i}(v)}{a(v)} R_v \right) = \frac{\gamma_i(u^*) - \gamma_i(v)}{a(i, u^*, v)}.
\end{aligned}$$

From the induction hypothesis, we have that for all agents $j \neq i$, herding implies that $\gamma_j(u^*) > \gamma_j(w) \forall w \in M_o$. Further, from the monotonicity condition (i.e., all agents have the same ordering on virtual rewards), we have that $\gamma_i(u^*) > \gamma_i(w) \forall w \in M_o$. On the other hand, for any task $w \in M_o$ such that v is reachable w , then from the second condition in the theorem, we have that $\gamma_i(w) \geq \gamma_i(v)$. Putting these together, we get that $\gamma_i(u^*) \geq \gamma_i(v)$; plugging into the above inequalities, we get that $U_{M,T}^i - U_{M,T}^{i,dev} \geq 0$. Thus following PPS is a best response for agent i .

• Case (ii). ($\mathcal{S}_{k,k}, k > 1$)

Consider a subgame $\mathcal{G}_{M,M}^i \in \mathcal{S}_{k,k}$. In this case, we need to show that i follows herding, i.e., i chooses available subtask $u^* = \arg \max_{w \in M_o} \gamma_i(w)$. Let $U_{M,M}^i$ be the expected reward earned by i if she attempts u^* , and $U_{M,M}^{i,dev}$ be her expected reward if she attempts some other subtask $v \in M_o$. Note that due to our induction hypothesis, all other agents concentrate on u^* , and share it once they solve it. Now we have:

$$\begin{aligned}
U_{M,M}^i - U_{M,M}^{i,dev} &= \frac{a_{-i}(u^*)}{a(i, u^*, v)} (U_{M,M}^i - U_{M-u^*,M-u^*}^i) + \frac{a_i(v)}{a(i, u^*, v)} (U_{M,M}^i - U_{M,M-v}^i), \\
&= \frac{a_{-i}(u^*)}{a(i, u^*, v)} \left(\frac{a_i(u^*)}{a(u^*)} R_{u^*} \right) - \frac{a_i(v)}{a(i, u^*, v)} \left(\frac{a_{-i}(v)}{a(v)} R_v \right) = \frac{\gamma_i(u^*) - \gamma_i(v)}{a(i, u^*, v)}.
\end{aligned}$$

Since $u^* = \arg \max_{w \in M_o} \gamma_i(w)$, therefore herding is a best response for agent i . This completes the proof. \square

PROOF OF THEOREM 4.3. As before, given subtask-network G , we define \mathcal{V} as the set of valid knowledge-subgraphs. Further, for any $M \subseteq [m]$ with $M^c \in \mathcal{V}$, we defined $M_o \subseteq M$ to be the set of available tasks.

Let OPT denote the optimal *centralized* agent-subtask allocation algorithm. In particular, for any subgame \mathcal{G}_M (i.e., wherein solutions for all subtasks in knowledge-subgraph $M^c \in \mathcal{V}$ are public), OPT chooses the agent-subtask distribution matrix $\{x_i^{OPT}(u)\}_{i \in [n], u \in M_o}$ that minimizes the expected makespan of the remaining subtasks. Let $T_{OPT}(M)$ denote the expected time achieved by OPT in subgame \mathcal{G}_M , and for each subtask $u \in M_o$, let $\lambda_u^{OPT} = \sum_{i \in [n]} x_i^{OPT}(u) a_i$.

On the other hand, assuming the rewards-vector satisfies the conditions for PPS, let $T_{PPS}(\mathcal{G}_M)$ denote the makespan achieved by some chosen equilibrium state in subgame \mathcal{G}_M . Also, as we did for OPT , we define $\{x_i^{eq}(u)\}_{i \in [n], u \in M_o}$ (and $\lambda_u^{eq} = \sum_{i \in [n]} x_i^{eq}(u) a_i$) to be the agent-subtask distribution matrix for some equilibrium state.

For any arbitrary work-conserving strategy $\{x_i(u)\}_{u \in M_o}$, with corresponding λ_u , we claim that subgame \mathcal{G}_M reduces to the subgame \mathcal{G}_{M-u} with probability $\frac{\lambda_u s_u}{\sum_{u \in M_o} \lambda_u s_u}$, after an expected time of $\frac{1}{\sum_{u \in M_o} \lambda_u s_u}$. This follows from the SA model, and also standard properties of the exponential distribution. Thus, we have:

$$\mathbb{E}[T_{OPT}(M)] = \frac{1}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} + \sum_{u \in M_o} \frac{\lambda_u^{OPT} s_u}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} \mathbb{E}[T_{OPT}(M-u)] \quad (2)$$

$$\mathbb{E}[T_{PPS}(\mathcal{G}_M)] = \frac{1}{\sum_{u \in M_o} \lambda_u^{eq} s_u} + \sum_{u \in M_o} \frac{\lambda_u^{eq} s_u}{\sum_{u \in M_o} \lambda_u^{eq} s_u} \mathbb{E}[T_{PPS}(\mathcal{G}_{M-u})]. \quad (3)$$

The first follows from linearity of expectation, and the property that OPT chooses the optimal agent-subtask distribution in each subgame; the second follows from the subgame-perfect property of the equilibrium.

For a given subgame-network with subtasks M , we define the *herding-time* $T_H(M) = \sum_{u \in M} \frac{1}{A s_u}$, where $A = \sum_{i \in [n]} a_i$. We now claim the following invariant: for every subgraph M with $M^c \in \mathcal{V}$, we have $\mathbb{E}[T_{PPS}(\mathcal{G}_M)] = \mathbb{E}[T_{OPT}(M)] = T_H(M)$. Since this is true for all subgames \mathcal{G}_M , therefore we have that the price-of-anarchy is 1 for all acyclic subtask networks under the SA model.

We now prove the above invariant via induction. Let $S_k = \{\mathcal{G}_M : M^c \in \mathcal{V}, |M| = k\}$ be the set of all valid subgames with k remaining subtasks, for $1 \leq k \leq m$. For $k = 1$, the invariant is trivially, since the subgames in S_1 have only a single subtask left. Now assume the invariant is true for $\forall M \in S_l, l \leq k$. Consider a subgame $M \in S_{k+1}$. Applying the induction hypothesis to equations 2 and 3, we have $\forall u \in M_o, \mathbb{E}[T_{OPT}(M-u)] = \mathbb{E}[T_{PPS}(\mathcal{G}_{M-u})] = T_H(M-u)$. Now from equation 2, we have for T_{OPT} :

$$\begin{aligned} \mathbb{E}[T_{OPT}(M)] &= \frac{1}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} + \sum_{u \in M_o} \left[\frac{\lambda_u^{OPT} s_u}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} \mathbb{E}[T_H(M-u)] \right] \\ &= \frac{1}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} + \sum_{u \in M_o} \left[\frac{\lambda_u^{OPT} s_u}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} \left(\mathbb{E}[T_H(M)] - \frac{1}{A s_u} \right) \right]. \end{aligned}$$

Simplifying this expression, we get:

$$\begin{aligned} \mathbb{E}[T_{OPT}(M)] &= \mathbb{E}[T_H(M)] + \frac{1}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} - \frac{1}{\sum_{u \in M_o} \lambda_u^{OPT} s_u} \left(\frac{\sum_{u \in M_o} \lambda_u^{OPT}}{A} \right) \\ &= \mathbb{E}[T_H(M)]. \end{aligned}$$

Similarly, we can do the same for T_{PPS} , using equation 3 and the induction hypothesis. Combining the two, we get the desired result. \square

ACKNOWLEDGMENTS

The authors were supported in part by the DARPA GRAPHS program and the DARPA XDATA program, via grant FA9550-12-1-0411 from the U.S. Air Force Office of Scientific Research (AFOSR) and the Defense Advanced Research Projects Agency (DARPA).

We thank the anonymous reviewers for their comments, which greatly improved the paper. We also thank the members of the SOAL Lab and Ramesh Johari for their suggestions.

REFERENCES

- ACKERMAN, M. AND BRÂNZEI, S. 2012. Research quality, fairness, and authorship order. *arXiv preprint arXiv:1208.3391*.
- BACHRACH, Y., SYRGKANIS, V., AND VOJNOVIC, M. 2013. Incentives and efficiency in uncertain collaborative environments. *arXiv preprint arXiv:1308.0990*.
- BANERJEE, S., GOEL, A., AND KRISHNASWAMY, A. K. 2014. Re-incentivizing discovery: Mechanisms for partial-progress sharing in research. Tech. rep., Stanford University. available on ArXiv.
- BOUDREAU, K. AND LAKHANI, K. R. 2013a. Cumulative innovation & open disclosure of intermediate results: Evidence from a policy experiment in bioinformatics. Working paper.
- BOUDREAU, K. AND LAKHANI, K. R. 2013b. How disclosure policies impact search in open innovation. Working paper.
- CHAWLA, S., HARTLINE, J. D., AND SIVAN, B. 2012. Optimal crowdsourcing contests. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 856–868.
- CRANSHAW, J. AND KITUR, A. 2011. The polymath project: lessons from a successful online collaboration in mathematics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1865–1874.
- GHOSH, A. AND HUMMEL, P. 2012. Implementing optimal outcomes in social computing: a game-theoretic approach. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 539–548.
- GHOSH, A. AND KLEINBERG, J. 2013. Incentivizing participation in online forums for education. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 525–542.
- GOWERS, T. 2013. The polymath blog.
- HARRIS, C. AND VICKERS, J. 1987. Racing with uncertainty. *The Review of Economic Studies* 54, 1, 1–21.
- JAIN, S., CHEN, Y., AND PARKES, D. C. 2012. Designing incentives for online question-and-answer forums. *Games and Economic Behavior*.
- KLEINBERG, J. AND OREN, S. 2011. Mechanisms for (mis) allocating scientific credit. In *Proceedings of the 43rd annual ACM Symposium on Theory of Computing*. ACM, 529–538.
- KONRAD, K. A. 2009. Strategy and dynamics in contests. *OUP Catalogue*.
- NIELSEN, M. 2011. *Reinventing discovery: the new era of networked science*. Princeton University Press.
- NOSEK, B. A. AND BAR-ANAN, Y. 2012. Scientific utopia: I. opening scientific communication. *Psychological Inquiry* 23, 3, 217–243.
- STREVENSON, M. 2003. The role of the priority rule in science. *The Journal of Philosophy* 100, 2, 55–79.
- STREVENSON, M. 2013. Herding and the quest for credit. *Journal of Economic Methodology* 20, 1, 19–34.
- TAYLOR, C. R. 1995. Digging for golden carrots: an analysis of research tournaments. *The American Economic Review*, 872–890.