To Stand on the Shoulders of Giants: Should We Protect Initial Discoveries in Multi-Agent Exploration?

Extended Abstract

Hodaya Lampert
State Comptroller Office, Jerusalem
Israel
hodaya.lampert@gmail.com

Reshef Meir
Technion—Israel Institute of Technology, Haifa, Israel
reshefm@technion.ac.il

Kinneret Teodorescu
Technion—Israel Institute of Technology, Haifa, Israel
kinnerett@technion.ac.il

ABSTRACT

We present a game theoretic analysis of a competitive search game which clarifies the expected advantages and disadvantages of granting first innovators exclusive rights to enjoy subsequent discoveries. We compare the theoretical predictions with actual behavior of players in the lab and find that the benefit of protection stems from increasing exploration efficiency, rather than encouraging initial exploration efforts. The latter, which contradicts theoretical predictions, can be explained by the cognitive bias of underweighting rare events.

KEYWORDS

Innovation, Exploration, Patents, Decisions from Experience

ACM Reference Format:

Hodaya Lampert, Reshef Meir, and Kinneret Teodorescu. 2025. To Stand on the Shoulders of Giants: Should We Protect Initial Discoveries in Multi-Agent Exploration?: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Stimulating innovation remains a key challenge for policymakers, with debates between proponents of free competition and advocates of rights protection. While competition generally enhances incentives [4, 13, 19], innovation presents a unique case due to the cumulative nature of discoveries [14]. Radical breakthroughs create new research avenues, enabling incremental discoveries at lower costs. However, if knowledge is publicly shared, only the original inventor bears the full cost, making discoveries akin to a public good that may suffer from underinvestment in a competitive environment [8]. A common solution is granting exclusive rights to initial inventors, as seen in patents protecting technological breakthroughs. This type of protection prevent competitors from leveraging the discovery, enhancing incentives for radical innovation [10]. In academia, policies allowing researchers to keep data private increase the reward for data collection. Protection can also promote specialization, as teams refining their own discoveries learn which research directions are most promising. Additionally, preventing redundant parallel discoveries can improve overall search efficiency [1, 6]. However, restricting access may slow down innovation [3, 5, 7, 12]. For example, academic journals increasingly require open data sharing to foster new insights and error detection [15, 20], and open-source platforms thrive by enabling developers to build on shared code [11]. Yet, open policies might lead to insufficient investment in radical innovation and inefficient exploration. This work examines the trade-offs between protection and open knowledge sharing, focusing on how economic competition, with and without protection, shapes exploration behaviors, discovery rates, and efficiency. To this end, we present an abstract theoretical model of sequential discoveries with and without protection, confirming that protection encourages initial discoveries but inhibits followup discoveries. In Section 3 we present a concrete game that simulates such an environment and compare theoretical with behavioral predictions. For proofs and details of the experiment see the full version on https://arxiv.org/abs/2502.14112.

This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. NowÃľ (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

2 A THEORETICAL MODEL

There are n players, each of which chooses how much to invest in exploration for novel knowledge (or *research*), and how much to invest in *exploitation* of existing knowledge towards *application*.

The strategy of each agent is thus composed of two real numbers, $r_i, x_i \geq 0$, representing the effort i invests in research and in exploitation of knowledge provided the opportunity, respectively.

We call the aggregated research product *knowledge*, $K := \sum_{i=1}^{n} r_i$, which can in turn be exploited for applications. As x_i is the effort i invests in applying knowledge, the overall *work* i invests in exploiting knowledge is $w_i := K \cdot x_i$.

The generated knowledge K can be applied by the competing agents, proportionally to their exploitation efforts: $a_i := \frac{x_i}{\sum_j x_j} K$.

Costs and utilities. We associate a fixed reward $R_r, R_a \geq 0$ with *research* and with *application*, respectively, as well as a single convex cost function $c : \mathbb{R}_+ \to \mathbb{R}_+$.

We further assume that ceteris paribus, exploitation is more rewarding than exploration per invested effort, and hence $R_a \geq R_r$.

- The total knowledge generated is $K := \sum_i r_i$;
- The exploitation work of i is $w_i := x_i \cdot K$;
- The knowledge applied by i is $a_i := \frac{x_i}{\sum_j x_i} K$, or just $a_i = x_i K$ if there is no competition (i.e. if $\sum_j x_j < K$);
- the overall utility of i is $u_i(r, x) := r_i R_r + a_i R_a - c(r_i) - c(w_i)$.

When initial research is protected (e.g. by patents), there is no interaction between players. In our model, this essentially means that there is a single player $i = 0$, and $K = r_0$. The optimal strategy then becomes a simple optimization problem.

Proposition 2.1. *The optimal strategy in the protected condition is to play* $x_0^* = 1$*, and* r_0^* *is the unique* r *s.t.* $c'(r) = \frac{R_r + R_a}{2}$.

When there are multiple players with access to the generated knowledge, we have that $K = \sum_i r_i$, and the applications a_i each agent generates depend both on K and the exploitation strategies $x_1, \ldots, x_n$, as explained above.

Proposition 2.2. *There is a symmetric equilibrium, where for every agent i, (a)* $c'(r_i^*) = R_r + \frac{R_a}{n^2}$*; and (b)* $x_i^* c'(n \cdot r_i^* x_i^*) = \frac{n-1}{n^2} R_a$.

Corollary 2.3. *The rate of exploration is* higher with protection *as long as* $\frac{R_a}{R_r} > \frac{n^2}{n^2-2}$*; and the rate of exploiting available knowledge is* lower with protection *as long as* $\frac{R_a}{R_r} > \frac{n^2}{n^2-n-1}$.

Note that the condition on $\frac{R_a}{R_r}$ becomes trivial for large n.

3 COMPETITIVE TREASURE HUNT GAME

In the "Competitive Treasure Hunt" game, each group of 4 players is presented with a hive of hexagons and need to find treasures. 5% of the hexagons are hidden treasures that simulate discoveries in the real world. Every three treasures are arranged in clusters which form a tight triangle. Thus, discovering the first treasure in a cluster increases the probability of finding the next two treasures from 0.05 to at least 0.33. Finding a first treasure simulates a breakthrough discovery and its value is set to 320. The other two treasures simulate sequential discoveries and their value is 80 each. The costs of exploration are distributed between 5 and 35, and are sampled independently for every player in each round. In the "Protection" setting, whenever a player finds the first treasure in a cluster, only he can explore the adjacent hexagons. In the "No Protection" setting, when a player finds a treasure, this does not restrict future searches of other players.

3.1 Simulation Results

We programmed artificial Fully Informed Bayesian Players (FIBP) in both game conditions. In our simulations, the optimal/equilibrium initial search threshold increases from 15 to 20 when applying protection (compared to $\sim 20\%$ increase expected from the theoretical model), and sequential search threshold decreased from 25 to 20 (compared to an expected theoretical decrease of $5\% - 20\%$). This qualitatively confirms the results expected from the abstract model with the appropriate parameters set.

3.2 Theoretical Predictions

Following the theoretical analysis and simulations, we get two clear theoretical predictions under profit maximization and full information assumptions:

Theoretical Prediction 1: Under the Protection condition, initial and sequential search activities should be at a similar rate.

Theoretical Prediction 2: Protection increases exploration activity for first treasures.

Theoretical Prediction 3: Protection decreases exploration activity for subsequent treasures.

3.3 Behavioral predictions

Unlike the assumptions underlying FIBS, in real life (and also in our lab experiment), the a-priori probability of making a new discovery is unknown to the competing players in advance, and they can learn it only throughout ongoing experiences. Learning takes time, but once consistent exploration threshold is formed, under the assumption of rationality, it should be close to the optimal one. However, the Decisions from Experience (DfE) literature suggests that in repeated choice settings, people tend to underweight rare events [2, 9, 16–18], which yields the following predictions:

Behavioral Prediction 1: Sequential search activity should be higher than initial search activity, under both conditions.

Behavioral Prediction 2: we should not expect a difference in initial search activity between the two conditions.

Note that each of these behavioral predictions 1,2 directly contradicts its theoretical counterpart, while theoretical prediction 3 is not affected by the above discussion.

4 EXPERIMENTAL RESULTS

4.1 Initial vs. Sequential search

In contrast to Theoretical Prediction 1 and in line with Behavioral Prediction 1, under protection, search rates for sequential discoveries were significantly higher than search rates for initial discoveries (0.72 vs. 0.6, $p<0.05$). This result was consistent across search costs.

4.2 How protection affects initial search

In line with Behavioral Prediction 2 there was no significant difference between search rates for initial discoveries in the two conditions. Yet, the direction was in line with Theoretical Prediction 2, with higher search rates for initial discoveries under protection than under no-protection (0.6 vs. 0.53, ns). This result is consistent over search costs. An additional analysis supported the behavioral account, by revealing under-sensitity to the reward level, in line with underweighting of rare events in decisions from experience.

4.3 How protection affects sequential search

In line with Theoretical Prediction 3, sequential search rates were lower under protection than under the no-protection condition (0.78 vs. 0.9). This difference becomes more significant for search costs ≤ 15 due to ceiling effect under low search costs.

4.4 Search Efficiency

Despite lower overall search rates under protection, When comparing the number of searches to treasures found, we see a sharp drop from 8.4 searches-per-treasure without protection, to about 7 one protection is applied. Additional analyses suggest this is due to players 'wasting' some of their searches on treasures eventually picked by others. Thus, regardless of its effect on search behavior, protection has the added benefit of *coordinating* players' effort.

5 CONCLUSIONS

The experimental results show that the benefit of protecting subsequent searches around initial discoveries stems from increasing exploration efficiency, rather than encouraging exploration intensity—an insight to consider in the design of IP systems.

Acknowledgements. This project was funded by the Israeli Ministry of Science and Technology (Grant no. 3-15284) and the Israel Science Foundation (ISF; Grant No. 2539/20).

REFERENCES

[1] Justus Baron and Tim Pohlmann. 2013. Who Cooperates in Standards Consortia-Rivals or Complementors? *Journal of Competition Law and Economics* 9, 4 (2013), 905–929.

[2] Greg Barron and Ido Erev. 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making* 16, 3 (2003), 215–233.

[3] James Bessen and Eric Maskin. 2009. Sequential innovation, patents, and imitation. *The RAND Journal of Economics* 40, 4 (2009), 611–635.

[4] Richard Blundell, Rachel Griffith, and John Van Reenen. 1999. Market share, market value and innovation in a panel of British manufacturing firms. *The Review of Economic Studies* 66, 3 (1999), 529–554.

[5] Michele Boldrin and David K Levine. 2005. The economics of ideas and intellectual property. *Proceedings of the National Academy of Sciences* 102, 4 (2005), 1252–1256.

[6] Vincenzo Denicolò. 2000. Two-stage patent races and patent policy. *the RAND Journal of Economics* (2000), 488–501.

[7] Alberto Galasso and Mark Schankerman. 2014. Patents and cumulative innovation: Causal evidence from the courts. *The Quarterly Journal of Economics* 130, 1 (2014), 317–369.

[8] Gene M Grossman and Elhanan Helpman. 1990. *Trade, innovation, and growth.* Technical Report.

[9] Ralph Hertwig, Greg Barron, Elke U Weber, and Ido Erev. 2004. Decisions from experience and the effect of rare events in risky choice. *Psychological science* 15, 8 (2004), 534–539.

[10] Erich Kaufer. 2012. *The economics of the patent system.* Routledge.

[11] Josh Lerner, Parag A Pathak, and Jean Tirole. 2006. The dynamics of open-source contributors. *American Economic Review* 96, 2 (2006), 114–118.

[12] Gaston Llanes and Stefano Trento. 2009. Anticommons and optimal patent policy in a model of sequential innovation. *Harvard Business School Entrepreneurial Management Working Paper* 09-148 (2009).

[13] Stephen J Nickell. 1996. Competition and corporate performance. *Journal of political economy* 104, 4 (1996), 724–746.

[14] Suzanne Scotchmer. 1991. Standing on the shoulders of giants: cumulative research and the patent law. *Journal of economic perspectives* 5, 1 (1991), 29–41.

[15] Stefan Stieglitz, Konstantin Wilms, Milad Mirbabaie, Lennart Hofeditz, Bela Brenger, Ania López, and Stephanie Rehwald. 2020. When are researchers willing to share their data?–Impacts of values and uncertainty on open data in academia. *Plos one* 15, 7 (2020), e0234172.

[16] Kinneret Teodorescu and Ido Erev. 2014. Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological science* 25, 10 (2014), 1861–1869.

[17] Kinneret Teodorescu and Ido Erev. 2014. On the decision to explore new alternatives: The coexistence of under-and over-exploration. *Journal of Behavioral Decision Making* 27, 2 (2014), 109–123.

[18] Kinneret Teodorescu, Ori Plonsky, Shahar Ayal, and Rachel Barkan. 2021. Frequency of enforcement is more important than the severity of punishment in reducing violation behaviors. *Proceedings of the National Academy of Sciences* 118, 42 (2021), e2108507118.

[19] Kenneth A Younge and Tony W Tong. 2018. Competitive pressure on the rate and scope of innovation. *Journal of Economic Behavior & Organization* 150 (2018), 162–181.

[20] Yimei Zhu. 2020. Open-access policy and data-sharing practice in UK academia. *Journal of Information Science* 46, 1 (2020), 41–52.