Sharing is not erring: Pseudo-reciprocity in collective search

Imen Bouhlel[1*] (imen.bouhlel@gredeg.cnrs.fr)**, Charley M. Wu**[2*]**, Nobuyuki Hanaki**[1]**, & Robert L. Goldstone**[3]

[1]Université Côte d'Azur, CNRS, GREDEG, Nice, France
[2]Center for Adaptive Rationality, Max Planck Institute for Human Behavior, Berlin, Germany
[3]Department of Psychological and Brain Sciences, Indiana University, Bloomington, USA
[*]These authors contributed equally to the following work

Abstract

Information sharing in competitive environments may seem counterintuitive, yet it is widely observed in humans and other animals. Companies often release technologies under open source license and scientists openly publish their research, often sharing free-to-access versions of their papers, rather than only publishing in restricted access journals. What drives this behavior and in which environments can it be beneficial? Using simulations in both static and dynamic environments, we show that sharing information can lead to individual benefits through the mechanisms of pseudo-reciprocity, whereby shared information leads to by-product benefits for an individual without the need for explicit reciprocation. Crucially, imitation with a certain level of innovation is required to avoid Roger's paradox, while the novel mechanism of an local visibility radius allowed for the coordination of self-organizing collectives of agents.When these two mechanisms are present, we find robust evidence for the benefits of sharing—even when others do not reciprocate.

Keywords: Collective search; information sharing; pseudo-reciprocity;

Introduction

In a competitive environment, it may seem counterintuitive to share information about a discovered solution to a problem (which could be copied) or the location of resources (which could be stolen). Yet unrestricted sharing of information is widely observed in both human and animal behavior. Private companies often release their technologies under an open source license, rather than keeping them as trade secrets to gain advantage over competitors. Scientists openly publish their research, often sharing free-to-access versions of their papers, rather than only publishing into restricted access journals. What drives this behavior and how can we understand it through the mechanisms of social cooperation and collective behavior?

In ecology, it has been observed that many species share information via mass recruitment systems when foraging for resources, whereby successful foragers send signals (e.g., sounds or pheromone trails) to help locate resources. For example, the American Cliff Swallow has a unique call used only when it finds air-borne food (Brown, Brown, & Shaffer, 1991). The call is an unrestricted social signal to other Cliff Swallows which does not only benefit kin (i.e., no kin selection mechanisms). Rather, the individual welfare of each Cliff Swallow is improved by the recruitment of peers towards the search effort, since the collective performs better at tracking prey than any lone individual. More generally, this behavior—common in different animal populations—is known as *pseudo-reciprocity* (Connor, 1986; Clutton-Brock, 2009), meaning altruistic behavior, such as sharing, can generate "by-product" benefit that improve individual fitness without the need for explicit reciprocation. Thus, pseudo-reciprocity in Cliff Swallows is thought to be mainly driven by the sparse and dynamic distribution of resources in the environment (Brown et al., 1991), which makes individual search very difficult, although in general the mechanisms are not well understood.

Forms of Cooperation

The study of human cooperation has revealed different levels of mechanisms that can lead to altruistic behavior (Rand & Nowak, 2013), such as direct or indirect reciprocity. *Direct reciprocity* requires repeated interactions and conditional strategies (i.e., cooperate or defect), where the evolution of cooperation depends on the likelihood of future interactions (Axelrod, 2006; Rand & Nowak, 2013). Thus, altruistic actions are explained by the expectation of reciprocity from peers in future interactions, conditional on one's own behavior. *Indirect reciprocity* extends this expectation of future reciprocity further, through reputation-based systems, whereby cooperative behavior is viewed as a type of social signaling to third-parties (Nowak & Roch, 2007). Thus, maintaining a positive reputation can have "downstream" effects, making it more likely for third-parties to cooperate, even if there have been no previous history of interactions.

Both direct and indirect reciprocity can be understood as forms of *impure altruism* (Andreoni, 1989), whereby altruistic behavior is explained through the expectation of future reciprocation within a self-interested, utility-maximization framework, and has been applied to game theory in public good games (Andreoni, 1990; Falk & Fischbacher, 2006; Nowak & Sigmund, 2005), but less so to collective search contexts. In contrast, *pseudo-reciprocity*, can be understood as an investment in other individuals in order to acquire or enhance by-product benefits, without the need for explicit reciprocation (Leimar & Connor, 2003). The sharing of information about the solution to a problem or the location of resources can benefit the sharer, because it can lead others to copy and improve on the original information, which can aid the original sharer as a by-product benefit.

For instance, in the case of the Cliff Swallow (Brown et al., 1991), the individual who signals for others gains *by-product benefits* through the recruitment of peers, regardless of whether or not they also signal in the future (Connor, 1995), since recruitment immediately aids in the tracking and

acquisition of food[1]. Importantly, this entails that pseudo-reciprocity can operate in systems without repeated interactions or reputation systems, which is not the case for direct or indirect reciprocity. While both forms of impure altruism have been greatly studied in game-theory (Dufwenberg & Kirchsteiger, 2004; Fehr & Gächter, 2002; Milinski, Semmann, & Krambeck, 2002; Lotem, Fishman, & Stone, 1999; Nowak & Sigmund, 1998), there exist relatively little study of pseudo-reciprocity in human behavior (although considerably more in animal behavior; Brown et al., 1991; Connor, 1986; Wilkinson, 1992).

The Role of the Environment in Collective Search

In a collective search task, one must learn to balance the exploration of yet unknown solutions, while also exploiting existing knowledge in order to yield immediate returns. This frames the *exploration-exploitation* dilemma, where different environments can demand a different ratio of exploration to exploitation (Barkoczi, Analytis, & Wu, 2016). The relative effectiveness of different search strategies can be drastically altered by changing one of many aspects of the task environment, such as the complexity of the problem space (Mason, Jones, & Goldstone, 2008), the connectivity of the communication network (Barkoczi & Galesic, 2016; Lazer & Friedman, 2007; Mason et al., 2008; Mason & Watts, 2012), the type of social information being communicated (Wisdom, Song, & Goldstone, 2013), and the learning strategies of each individual agent (Barkoczi & Galesic, 2016). However, one aspect taken for granted in these previous collective-search paradigms, is that information is always freely given, whereas in many real world situations agents make an explicit decision to either share or withhold information. We examine how different task environments influence the adaptiveness of freely sharing information.

Goals and Scope

In which environments will we find stable benefits for freely sharing information, and where will we fail to find it? We present multi-agent simulations of a competitive search task in n-dimensional search spaces, with both static and dynamic reward structures. Agents are given a limited search horizon to search for rewards, with different populations of agents containing different mixtures of either *sharers* or *non-sharers*. We focus on four main mixtures (All sharers, No sharers, Free-rider, and Free-giver) in order to determine for any given agent whether or not it is better to share or to withhold information from others. We use a competitive search context, whereby agents who choose the same location in the search space must split the reward between them. In the static simulation, we use a fitness landscape where the reward is a monotonically decreasing function of the distance from a randomly selected global optimum. In the dynamic simulation, the reward function is similar, but the global optimum follows a random walk, moving on each trial.

[1] A crucial mechanism is that each Cliff Swallow has free, visual information about the location of nearby peers, thus allowing for localized coordination (see Visibility Radius).

Static Simulations

We present a mathematical framework for studying the benefits of sharing information in a competitive search context. We use multi-agent simulations to implement this framework, and present simulation results across a variety of different environmental structures.

The search task involves k agents searching for rewards on an n-dimensional environment, where each dimension d in $1,2,3,...,n$ can take any positive or negative integer value. The task takes place over $T = 100$ trials, where at each time $t \in [1,T]$, all agents individually make a decision about where to search for rewards, and then whether or not to share that information with other agents (see Information Sharing). Altogether, we vary the number of agents $k \in [4,10]$, the number of dimensions of the search space $n \in [5,15]$, the innovation rate $p_c \in \{0, 0.5, 1\}$, the visibility radius $r \in \{0,1,2\}$, and the mixture of sharing strategies (All sharers, No sharers, Free-rider, and Free-giver).

Fitness function

Each search decision is represented as a vector $\mathbf{x_t} = [d_1^t, ..., d_n^t]$, where the resulting payoff $y_t = f(\mathbf{x_t})$, is given by the inverse Manhattan distance between the global optimum M and $\mathbf{x}_t$:

$$f(\mathbf{x_t}) = \frac{2}{1 + \|\mathbf{x}_t - M\|_1}.100 \tag{1}$$

with Laplace smoothing used to prevent division by zero and rescaled by a factor of 100. The global maximum of the environment at time t is represented by $M = [D_1, ..., D_n]$, which is drawn from a uniform distribution $\mathcal{U}(1,10)$.

The competitive nature of the task is represented by *splitting* earned payoffs between agents who make the same search decision on a specific trial (i.e., subtractive resources). Let $\mathbf{P}_t = [\mathbf{x}_{t1}, ..., \mathbf{x}_{tk}]$ represent the matrix containing all k-agents' search decisions at time t. For each agent i, let $C(\mathbf{x}_{ti})$ be the number of rows in $\mathbf{P}_t$ matching $\mathbf{x}_{ti}$,

$$C(\mathbf{x}_{ti}) = \sum_j \delta(\mathbf{x}_{ti}, \mathbf{x}_{tj}), \tag{2}$$

where $\delta(\mathbf{x}_{ti}, \mathbf{x}_{tj}) = 1$ if $\mathbf{x}_{ti} = \mathbf{x}_{tj}$, and 0 otherwise. Thus, $C(\mathbf{x}_{ti}) = 1$ when $\mathbf{x}_{ti}$ is unique, otherwise, $C(\mathbf{x}_{ti}) > 1$. Thus, each agent earns reward $R(\mathbf{x}_{ti}) = \frac{y_{ti}}{C(\mathbf{x}_{ti})}$, where $R(\mathbf{x}_{ti}) = y_{ti}$ when no splitting of rewards occurs.

Individual Search

At trial $t = 1$, all agents start with a random location sampled from a uniform distribution $\mathcal{U}(1,10)$. For all subsequent trials, each agent i has access to the history of previous observations O_{ti}, which includes both *individually* acquired observations $o_{ti} = \{\mathbf{x}_{li}, R(\mathbf{x}_{li})\}_{l=1}^{t-1}$ and *socially* observed data s_{ti}, where $\{\mathbf{x}_{lj}, R(\mathbf{x}_{lj})\}$ is contained in s_{ti} if agent i observed

agent j's ($j \neq i$) information at trial $l < t$ (see Information Sharing). Contingent on O_{ti}, we use a local search strategy where the agent selects the search decision that had previously yielded the largest reward, $\mathbf{x}^*_{ti} = \text{argmax}_{\mathbf{x}_{li}, l \in [1,t-1]} R(\mathbf{x}_{li})$, and then innovates on that solution by modifying each value in $\mathbf{x}^*_{ti}$ by a discrete value in $\{-1, 0, 1\}$.

We define the *Innovation rate* as the probability $p_c \in \{0, 0.5, 1\}$ with which the agent innovates on the search decision $\mathbf{x}^*_{ti}$. On each trial the agent innovates with probability p_c, where otherwise $\mathbf{x}^*_{ti}$ is copied verbatim. If the agent innovates, then we draw from a Binomial distribution centered on zero $c^t_{im} \sim \text{B}\left(2, \frac{1}{2}\right) - 1$ for each dimension of the search problem $m = 1, ..., n$. Intuitively, half of the time $c^t_{im} = 0$, meaning there is no change along that dimension, while changes of both -1 or $+1$ are equally likely, each with probability 0.25. Thus, when innovation occurs, the search decision at $t+1$ is defined as:

$$\mathbf{x}_{(t+1)i} = \mathbf{x}^*_{ti} + [c^t_{i1}, ..., c^t_{in}] \tag{3}$$

Information Sharing

We consider four different mixtures of sharers and non-sharers in the collective. *No sharers* and *All sharers* are cases where there is a homogeneous mixture of strategies, with either all agents or no agents sharing information. In the *Free-rider* and *Free-giver* conditions, there is a single heterodox agent who acts in contradiction to the rest of the collective. The Free-rider agent is one who withholds information while all others share, and the Free-giver agent is one who shares information, even though all others withhold information. In all cases, sharing information broadcasts both the location $\mathbf{x}_{ti}$ and the earned reward $R(\mathbf{x}_{ti})$ to all other agents.

We use differences in performance across these four conditions to examine the individual benefits of sharing when the behaviors of all others are fixed. By comparing the performance differences between the Free-rider agent and any given agent of the All sharer agents, we observe if there exist individual benefits of sharing when others also share. Similarly, we can compare individual performance differences between the Free-giver agent and any of the No sharer agents to observe the benefit of sharing when others don't share. Environments where sharing is beneficial at both levels provide evidence of pseudo-reciprocity, whereby sharing confers benefits without the need for explicit reciprocation.

Visibility Radius

In addition to explicitly shared information, we also implement a *Visibility Radius* defined as $r \in \{0, 1, 2\}$ where localized clusters of agents have free access to information about each other's current solution and rewards. For any two agents $i \neq j$, agent j is visible to agent i at trial t if the *Chebyshev* distance between $\mathbf{x}_{ti}$ and $\mathbf{x}_{tj}$:

$$D_{Chebyshev}(\mathbf{x}_{ti}, \mathbf{x}_{tj}) = \max_m |d^t_{im} - d^t_{jm}| \tag{4}$$

is smaller than the Visibility Radius r. Intuitively, Chebyshev distance is also referred as "chessboard distance", since it is the minimum number of moves needed by a king to go from one square to another. If agent j is visible to agent i at trial t, then information about the current search decision and reward is automatically shared: $\{\mathbf{x}_{tj}, R(\mathbf{x}_{tj})\} \in s_{(t+1)i}$. The visibility radius is an important coordination mechanism that allows for localized transmission of information, whereas the explicit decision to either share or withhold information operates for all distances greater than r. Crucially, given the high dimensionality and size of the search space, it is very unlikely for any two agents to fall within the same visibility radius without explicit information sharing.

Results

Across each combination of environmental parameters, we conducted 10,000 simulations and computed the mean payoff over trials for each individual agent. In Figure 1 we show heat maps of the simulation results, where the color of each tile represents the net individual benefit of sharing. In Figure 1a, we compute the net benefit of sharing (when others share) as the difference between the mean payoff for a random agent from the All sharers condition and the mean payoff of a Free-rider agent. Thus, red values indicate a net benefit of sharing, while blue values indicate a net cost of sharing. In Figure 1b we run the same analysis, but compute the net benefit of sharing (when others don't share) as the difference between the mean payoff for a Free-giver agent and a random agent from the No sharers condition. Again, red values indicate a net benefit of sharing, while values indicate a net cost. In both cases (when others share and when they do not), we find benefits for sharing information.

Benefits of sharing. Overall, we found evidence that sharing information can be a beneficial strategy, so long as peers copy with some level of innovation and in the presence of a visibility radius. Benefits were found both with and without reciprocation. Interestingly, the relative benefits of sharing when others don't share (Free-giver vs. No sharers) were far larger than when others also share (All sharers vs. Free-rider), likely due to diminishing returns as the collective becomes more saturated with social information (see Fig 2). Thus, sharing information without explicit reciprocation leads to the largest net benefit, whereas any potential disadvantages to sharing when others share is entirely reversed by the mechanisms of innovation and the visibility radius.

Number of agents and size of the environment. As long as there are innovators ($p_c \neq 0$), we find that larger collectives reap larger benefits of sharing when exploring larger environments. Moreover, larger visibility increases these benefits in larger environments. Hence, the potential disadvantages of overcrowding are washed out by the advantages of social coordination through pseudo-reciprocity.

Innovation and visibility. Sharing information for others to imitate can be beneficial, as long as imitation also involves some level of innovation. Thus, Roger's Paradox (Rendell

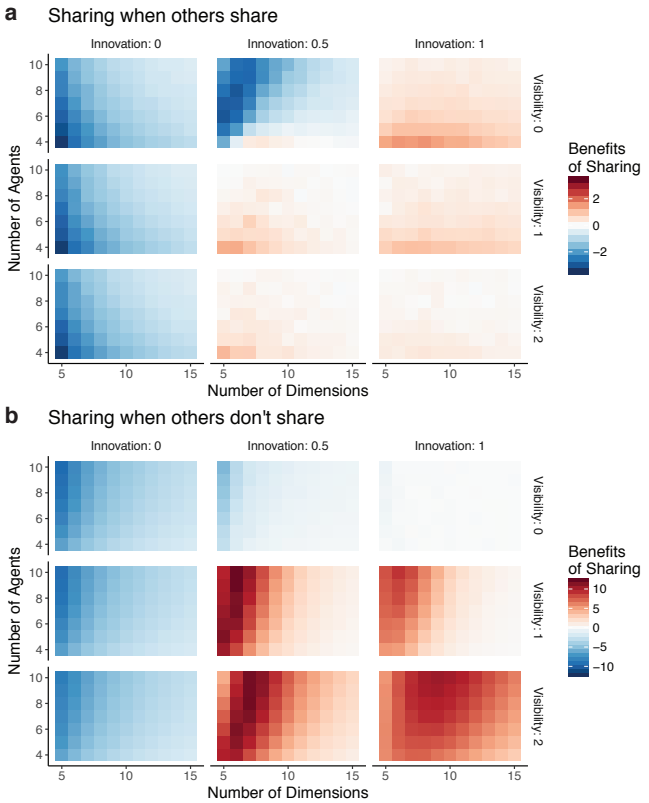

Figure 1: Static simulation results. We simulated 10,000 replications for each combination of environmental parameters (number of agents, number of dimensions, innovation rate and visibility radius) crossed with each of the four scenarios (all sharers, no sharers, free-rider, and free-giver). **a**) Sharing when others share. The color of each tile shows the difference in the mean agent payoffs, comparing an agent who shares when others also share (all sharers) and an agent who withholds information when others share (free-rider). Red values indicate a net benefit of sharing, white indicates no difference, and blue indicates a net disadvantage. We see that all negative side-effects of sharing disappear when both innovation and visibility are present. **b**) Sharing when others don't share. Colors show the difference in mean agent payoffs, comparing an agent who shares when no one else shares (free-giver) and an agent who withholds information along with everyone else (no sharers). We see that both innovation and visibility support the benefits of sharing, although here we see stronger benefits, with a linear relationship between the number of agents and number of dimensions that leads to maximal benefits.

& Laland, 2010), whereby social information can sometimes reduce mean fitness, only holds in our simulations when innovation is zero (Fig. 1). Additionally, Figure 2 shows the effect of the visibility radius on trial-by-trial performance, where we fix the number of agents ($k = 10$), the size of the environment ($n = 7$), and set the innovation rate to $p_c = 0.5$. Here we see that by introducing a visibility radius, the maladaptiveness of sharing when others also share (comparing free-rider to all sharers) becomes non-existent, creating no benefit to withholding information (i.e., free-riding). Moreover, the visibility radius even more dramatically reverses the performance difference between sharing while others don't (free-giver) and when no social learning takes place (no sharers). Whereas without a visibility radius, being a free-giver is a maladaptive strategy, we find a stark reversal when the visibility radius is introduced, where free-giving dramatically

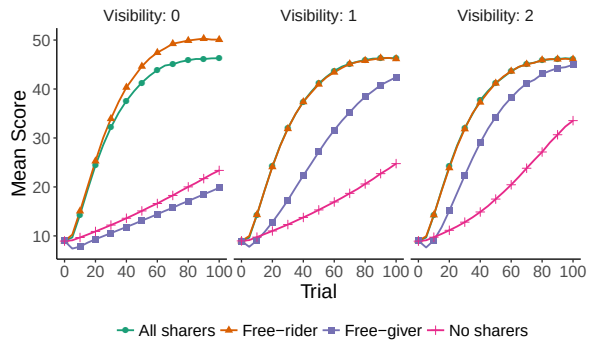

Figure 2: Learning curves. Mean score for 10,000 replications using $k = 10$ agents and $n = 7$ dimensions, and with innovation rate $p_c = 0.5$. Each line is either a random agent drawn from the all sharers or no sharers condition, or specifically the heterodox agent in the free-rider and free-sharer conditions. Across different visibility radii, we see how the maladaptiveness of sharing (left panel) disappears, and transforms into an adaptive benefit. With a visibility of 1, the benefits of free-riding are reduced to zero, while the reversal for free-giving is much more dramatic and confers much larger rewards than the counter-factual case of withholding information (i.e,. no sharer).

increases individual fitness, and even begins to approach the performance levels of the all-sharers condition. In our simulations, we find that both innovation and visibility are essential mechanisms that facilitate the adaptiveness of pseudo-reciprocity by creating conditions where sharing information confers by-product benefits to the individual—without the need for reciprocation.

Dynamic Simulations

We adapted the static simulations so that the global maximum changes over time. The global maximum at time t is then defined by $M_t = [D_1^t, ..., D_n^t]$. We define *Environmental change* as the probability $p_e \in \{0.25, 0.5, 0.75\}$ with which the global maximum changes at each time t. If the environment changes, then each dimension D_m^t is updated using:

$$D_m^t = D_m^{t-1} + e_m^t \tag{5}$$

where $e_m^t \sim \mathrm{B}\left(2, \frac{1}{2}\right) - 1$, is a Binomial distribution centered on zero. Thus, environmental change operates on the same mechanism as used by agents in individual search.

In order to account for the decreasing validity of past observations in a dynamically changing environment, we introduce a *Discount rate* parameter $\gamma \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ to discount the memory of previous rewards as a function of elapsed time, where smaller values of γ result in harsher discounting. Similar to temporal discounting, observations of past rewards R_{ti} are discounted by :

$$\hat{R}(\mathbf{x}_{ti}) = \gamma^{\tau} \cdot R(\mathbf{x}_{ti}) \tag{6}$$

where $\hat{R}(\mathbf{x}_{ti})$ is the discounted reward observation and τ is the elapsed time between the observation and the current time. Thus, in the dynamic simulations the local search strategy innovates on the search decision that has the largest discounted reward, $\hat{\mathbf{x}}_{ti}^* = \mathrm{argmax}_{\mathbf{x}_{li, l \in [1, t-1]}} \hat{R}(\mathbf{x}_{li})$.

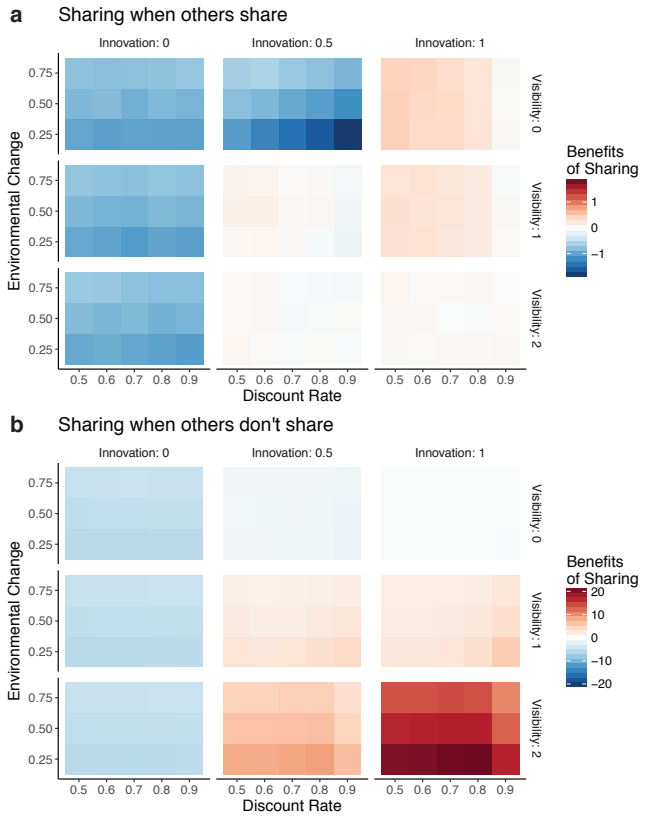

Figure 3: Dynamic simulation results. Here we fixed the number of agents ($k = 10$) and the number of dimensions ($n = 7$) to more freely explore differences in the dynamic environment parameters. We simulated 10,000 replications for each combination of the remaining environmental parameters (innovation rate, visibility radius, environmental change, and discount factor), crossed with each of the four scenarios (all sharers, no sharers, free-rider, and free-giver). **a**) Sharing when others share. Both innovation and visibility are robust to differences in rates of environmental change, and mediate net benefits of sharing when others also share. A strong environmental change rate also generally amplifies the net benefit of sharing. **b**) Sharing when others don't share. The benefits of sharing are observed as soon as innovation and visibility are positively valued, and are amplified by small discount rates (resulting in faster decay rates). Moreover, with a high level of innovation ($p_c = 2$) and visibility ($r = 2$), we see the strongest net benefits for sharing out of all simulations presented.

In our dynamic simulations, we fix the number of agents and the size of the environment ($k = 10$; $n = 7$) to the combination that yielded one of the highest benefits of sharing in the static analysis, while exploring the influence of environmental change p_e and the discount rate γ, along with innovation rate and visibility radius. Our goal is to find whether the sharing benefits observed in the static environment are robust to temporal dynamics in a changing environment.

Results

Across each combination of environmental parameters and each of the four mixtures of sharers, we conducted 10,000 simulations and computed the mean payoff over trials for each individual agent, identical to the static simulations. Again, we find evidence of benefits for sharing information (Fig 3), both with and without reciprocation. While we see

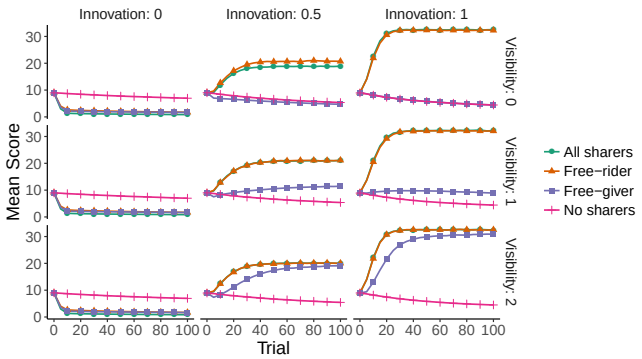

Figure 4: Learning curves. Mean score for 10,000 replications using $k = 10$ agents and $n = 7$ dimensions, and with environmental change $p_e = 0.25$ and discount rate $\gamma = 0.8$. Each line is either a random agent drawn from the all sharers or no sharers condition, or specifically the heterodox agent in the free-rider and free-sharer conditions. Innovation increases the performance, and visibility increases the benefits of free giving. With positive values for both innovation and visibility radius, we find that there is never any disadvantage to sharing, either when others also share (all sharers vs. free-rider) or when others don't (free-giver vs. no sharers).

smaller benefits of sharing when others also share compared to the static simulations (Fig 3a), we still find that any net disadvantages of sharing are reversed into either a net benefit or a cost-free outcome when innovation and the visibility radius are present. In the case of sharing without reciprocation (Fig 3b), we find even stronger benefits of sharing than in the static simulations for particularly high innovation and visibility.

Environmental change and discount rate. We observe the largest benefits of sharing in slower changing environments. Indeed, coordination of search through sharing social information is even more crucial in tracking changing rewards. However, rapidly changing environment can quickly become too difficult for any mixture of sharing strategies. Interestingly, smaller—and thus stronger—discount rates seem to be the most adaptive for populations with at least one sharer.

Innovation and visibility. We replicate the main finding from the static simulations, that both innovation and visibility are key ingredients for finding pseudo-reciprocity. Each ingredient has a distinct directed effect: high innovation increases the mean performance of all agents as soon as the population contains a majority of sharers, whereas high visibility increases the individual benefits of sharing in populations where only one agent shares (Fig. 4). Thus, when these two conditions are met, the benefits of sharing can be even larger than in the static case.

General Discussion

In this study, we investigate the rationality of sharing information in a competitive search context using agent-based modeling. Agents searched for rewards through n-dimensional environments using a local search strategy, while information about observations could either be shared with the collective or withheld. The competitive context was induced by agents being forced to equally split rewards if they chose the same location. Across both static and dynamic

reward environments, we found evidence that sharing information can enhance individual performance *without the need for reciprocation*, operating on the as yet poorly-understood mechanisms of pseudo-reciprocity. Across our environmental analysis of 2,700 different environments, we found two essential mechanisms that facilitated the benefits of sharing. Firstly, we found that sharing information for others to imitate can prove beneficial so long as imitation also involves some level of innovation. Secondly, we found that a visibility radius, which allowed agents free access to information from peers within a limited distance, greatly improved the benefits of sharing without reciprocation. Crucially, this localized visibility acts as a coordination device, where sharing information acts as a recruitment mechanism, allowing for small localized collectives to cooperatively find better rewards—yet still without the need for reciprocation through unrestricted sharing.

One limitation of our study is that we only investigate environments with a single global maximum, whereas differences in the search environments can induce different demands on the exploration-exploitation trade-off (Barkoczi et al., 2016). Indeed, highly rugged fitness landscapes may cause shared information to be a burden rather than boon, misleading individuals or collectives to explore sub-optimal local maxima of the search space. In the future, we would also look at the influence of different types of social networks governing the flow of shared information (Lazer & Friedman, 2007; Mason & Watts, 2012), along with a wider variety of individual search strategies that more closely resemble human search behavior (Wu, Schulz, Speekenbrink, Nelson, & Meder, 2017).

Conclusion

While there have been many studies on collective search, they all typically assume free access to social information, without considering the decision to either share or withhold information. In this study, we find that even in a competitive search context, sharing information can prove beneficial, in both static and dynamic reward environments. This is one of the first studies to look at the environmental conditions that support *pseudo-reciprocity* whereby shared information can lead to by-product benefits for an individual without the need for reciprocation. This work is one of the first steps towards understanding the ecological rationality of the open source movement through the lens of a coordination system that does not require either reputation or repeated interactions.

Acknowledgments

CMW is supported by the International Max Planck Research School on Adapting Behavior in a Fundamentally Uncertain World.

References

Andreoni, J. (1989). Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of political Economy*, *97*(6), 1447–1458.

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The economic journal*, *100*(401), 464–477.

Axelrod, R. M. (2006). *The evolution of cooperation*. Basic books.

Barkoczi, D., Analytis, P. P., & Wu, C. M. (2016). Collective search on rugged landscapes: A cross-environmental analysis. In *Proceedings of the annual conference of the cognitive science society.*

Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature communications*, *7*.

Brown, C. R., Brown, M. B., & Shaffer, M. L. (1991). Food-sharing signals among socially foraging cliff swallows. *Animal Behaviour*, *42*(4), 551–564.

Clutton-Brock, T. (2009). Cooperation between non-kin in animal societies. *Nature*, *462*(7269), 51–57.

Connor, R. C. (1986). Pseudo-reciprocity: investing in mutualism. *Animal Behaviour*, *34*(5), 1562–1566.

Connor, R. C. (1995). Altruism among non-relatives: alternatives to the prisoner's dilemma. *Trends in Ecology & Evolution*, *10*(2), 84–86.

Dufwenberg, M., & Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and economic behavior*, *47*(2), 268–298.

Falk, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and economic behavior*, *54*(2), 293–315.

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, *415*(6868), 137–140.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*(4), 667–694.

Leimar, O., & Connor, R. C. (2003). By-product benefits, reciprocity, and pseudoreciprocity in mutualism. *Genetic and cultural evolution of cooperation*, 203–222.

Lotem, A., Fishman, M. A., & Stone, L. (1999). Evolution of cooperation between individuals. *Nature*, *400*(6741), 226–227.

Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, *137*(3), 422.

Mason, W. A., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, *109*(3), 764–769.

Milinski, M., Semmann, D., & Krambeck, H.-J. (2002). Reputation helps solve the tragedy of the commons. *Nature*, *415*(6870), 424–426.

Nowak, M. A., & Roch, S. (2007). Upstream reciprocity and the evolution of gratitude. *Proceedings of the Royal Society of London B: Biological Sciences*, *274*(1610), 605–610.

Nowak, M. A., & Sigmund, K. (1998). The dynamics of indirect reciprocity. *Journal of theoretical Biology*, *194*(4), 561–574.

Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, *437*(7063), 1291–1298.

Rand, D. G., & Nowak, M. A. (2013). Human cooperation. *Trends in cognitive sciences*, *17*(8), 413–425.

Rendell, F. L., L., & Laland, K. N. (2010). Rogers paradox recast and resolved: population structure and the evolution of social learning strategies. *Evolution*, *64*(2), 534–548.

Wilkinson, G. S. (1992). Information transfer at evening bat colonies. *Animal Behaviour*, *44*, 501–518.

Wisdom, T. N., Song, X., & Goldstone, R. L. (2013). Social learning strategies in networked groups. *Cognitive Science*, *37*(8), 1383–1425.

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2017). Exploration and generalization in vast spaces. *bioRxiv*, 171371.