



From “Investors and Exploiters in Ecology and Economics: Principles and Applications,”
Luc-Alain Giraldeau, Philipp Heeb, and Michael Kosfeld, eds. 2017. Strüngmann Forum Reports, vol. 21,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03612-2.

3

The Ecological and Economic Conditions of Exploitation Strategies

Maxwell N. Burton-Chellew, Alex Kacelnik, Michal Arbilly,
Miguel dos Santos, Kimberley J. Mathot, John M. McNamara,
Friederike Mengel, Joël van der Weele, and Björn Vollan

Abstract

In many situations across biology and economics, there is often one individual, or “agent,” that invests effort into a beneficial task and also one individual that, in contrast, foregoes the effort of investing, and instead simply exploits the efforts of another. What makes an individual choose to invest in production versus exploiting the efforts of another? If everyone invests, then exploitative strategies become very profitable; however if everyone is exploitative, there will be no investments to exploit. How does natural selection resolve this dilemma? What can economic institutions do to encourage investment? Can biologists and economists learn from the approach of each other’s discipline? This chapter outlines the commonalities and differences in approach of the two disciplines to the general problem of investment versus exploitation. It develops a model to encapsulate the general features of many scenarios (“games”) involving potential exploitation and explores the benefits of a unified approach, outlining current limitations and important areas for future investigation.

Introduction

Picture one of your distant ancestors, eking out a living during the Paleolithic era, spending many hours carefully and patiently crafting some flint stone into

Group photos (top left to bottom right) Max Burton-Chellew, Alex Kacelnik, Michal Arbilly, Joël van der Weele, Friederike Mengel, John McNamara, Miguel dos Santos, Alex Kacelnik, Max Burton-Chellew, Björn Vollan, John McNamara, Joël van der Weele, Kim Mathot, Alex Kacelnik, Michal Arbilly, Max Burton-Chellew, Kim Mathot, Friederike Mengel, Miguel dos Santos and Joël van der Weele, Björn Vollan, Michal Arbilly

a usable hand ax that he (or she) can use later to chop firewood, butcher meat, and dig out nourishing tubers. After several hours work, his axe is now ready, but just as he stands back to admire his handiwork, another man rushes by and takes the hand ax for himself. Your ancestor has spent valuable time and effort, using skills that took years to acquire, in crafting a usable tool, which is now only going to benefit another man—a man who may have spent the entire day resting or wooing potential partners, before exploiting the efforts of your industrious ancestor. In the evolutionary struggle for survival, it would appear that your ancestor is at a distinct disadvantage.

Alternatively, you may prefer to consider a New Caledonian crow, *Corvus moneduloides*, spending considerable time stripping and bending a twig into a tool fine enough to extract a nourishing but stubborn beetle grub hiding within the apparent safety of tree bark. After several minutes investigating the specific problem and the length of the local twigs, the crow then selects the most appropriate twig and shapes it accordingly. The twig is now ready to be used, to provide a beneficial return on the crow's investment of time and energy, but just as it digs out the recalcitrant grub another crow appears—a crow that had perhaps been resting or mating nearby—and helps itself to the tasty snack. Again, the more industrious individual appears to be at a disadvantage.

In both of these examples there is one individual that invests effort into a task with the aim of reaping a later benefit, and one individual that, in contrast, foregoes the effort of investing and instead simply exploits the efforts of another. Here there is a puzzle, for it would seem that the one that competes for instruments without toiling to make them is at an advantage to the one that produces them, for parasitizing upon the efforts of the producer saves time and energy. Surely, therefore, life's winners will adopt this parasitic strategy. However, if everyone is parasitic, there will be neither tools nor food to steal. So maybe everyone should choose to produce. Likewise in this case, if everyone is productive, a potential exploiter will have lots of potential victims and so being an exploiter will be more tempting. Our aims here are to explore the similarities between the biological and economic approaches to studying the above problem concerning investment versus exploitation and, more importantly, to investigate whether the empirical findings and theoretical concepts from one field can inform the other.

Investment versus Exploitation

The above examples show that individuals often face a decision between producing and taking, between relying on themselves to find and produce resources or relying on others to do so, between investing in production versus exploiting the investments of others. An adult individual can generally either invest in being self-sufficient, producing their own necessary resources, or they can adopt “parasitic” or “predatory” strategies that take advantage of, and thus “exploit,” the efforts of others. Although the best terms to describe

these alternatives are debatable, and although many individuals can and will also cooperate to find, produce, or trade resources, it is hard to argue that the decision to invest or exploit is not a recurring feature of life that shapes much of the biological world. For instance, many species predate on others, and it is possible that every multicellular species is at times exploited by multiple parasites, with the number of parasite species estimated to outnumber nonparasite species, perhaps by even 4 to 1. The presence of exploitative parasites has likely been a major driver of evolutionary change in host species (Price 1980; Moore 2002; Agosta et al. 2010).

Biologically, all organisms can be considered as either autotrophic or heterotrophic. Autotrophs, such as some plants and some microbes, obtain all the energy they need from their abiotic environment, deriving their energy from sunlight or inorganic chemicals. Heterotrophs, in contrast, rely upon the investments of autotrophs. Imagine, for example, a zebra that is grazing on the savannah: This zebra is a heterotroph that is exploiting the production of autotrophic grasses in a way that benefits the zebra but harms the grasses it consumes. Now imagine that this grazing zebra is suddenly attacked and eaten by a lion. This lion is likewise clearly exploiting the investments of the zebra, in a way that benefits itself, but harms the zebra.

The lion that eats the zebra that eats the grass is far from the end of the chain, for as explained above, both the lion and the zebra will likely be playing host to several small harmful parasitic species, which could be considered to be exploiting their host. Of course, such predation and parasitism are not the only available interactions. For example, the lion and the zebra will also be playing host to several microbiotic species that provide benefits as part of a cooperative mutualism. In fact, any and all actions that organisms perform (e.g., pigs depositing nitrogen-rich dung or earthworms moving nutrients through soil) may well provide a mix of unintended positive and negative effects to other species. However, the point remains that once there are investors, of any sort, opportunities for exploitative strategies follow. This exploitation is not restricted to predation and parasitism of the flesh. For example, lions will exploit the searching efforts of vultures by following them and monopolizing any carrion they find (Schaller 1972), and when the cuckoo lays its egg in the nest of another species, it does so to exploit the parental investment of another species, tricking them into caring for their own young, a trick referred to as “brood parasitism” (Davies et al. 2012).

When one species evolves to harm or exploit the investments of another, an evolutionary “arms race” typically occurs, where neither species gains a large relative advantage despite increasing adaptations for winning the conflict (Davies et al. 2012). This is because as one species evolves better defenses against exploitation, the other species evolves better countermeasures in response. For example, prey species will often evolve to run faster, and species that suffer from brood parasitism tend to evolve distinctive egg patterning which helps them identify and reject distinctive cuckoo eggs. However, as

prey become faster and harder to catch, the selection pressure upon predators to become faster increases, and vice versa, creating the phenomenal speed of the cheetah. Both sides are running, but neither is getting ahead by much, in evolutionary terms. Likewise, as the victims of cuckoos make their eggs more distinctive, the selection pressure increases upon cuckoos to effectively mimic the egg patterning of their hosts (Dawkins and Krebs 1979; Davies et al. 2012). Again, both sides are adapting to each other, making life difficult for each other, but neither is ever winning by much.

Most of the above examples describe interspecific interactions, that is, interactions between species. However, the decision to invest or to exploit the efforts of others is central to many intraspecific interactions too, where members of the same species compete for the same resources, as in our opening human and crow examples. Arguably the most informative intraspecific example for biologists is in the foraging behaviors of gregarious bird species that move in flocks. As we explain more fully below, many birds face the choice to search for their own food (investment) or to save on energy and perhaps time by watching and exploiting the searching efforts of others before rushing in at the last moment to compete for a share of the bounty. Biologists have often termed these behaviors as “producing” and “scrounging,” respectively, and considered these behaviors as separate strategies or tactics that individuals can employ while playing the “producer–scrounger game” (Barnard and Sibly 1981; Barnard 1984; Giraldeau and Caraco 2000; Giraldeau and Dubois 2008). In this case, to scrounge does not usually mean to beg or rely on the goodwill of producers but rather to compete for the food that has been discovered by another; hence we consider it to be behavior that exploits the investments of others. The producer–scrounger (PS) game has been used to conceptualize the range of foraging behaviors within many species of birds and mammals, including several primate species such as mangabeys, baboons, gorillas, and chimpanzees (see references within Arbilly et al. 2014). In general, we refer below to the behaviors in the PS game as an interaction between “investors” and “exploiters.”

What is interesting in these intraspecific interactions is that rather than having two sides that evolve countermeasures to each other, we have similar individuals that can behave as either an investor (“producer”) or an exploiter (“scrounger”). This means that the ratio of investors to exploiters and/or the decision rules employed to choose between acting as an investor or exploiter may change over time. These changes in behavior can theoretically occur either within individuals over time, through processes such as learning or plasticity, or within populations through genetic evolution. Biologists have therefore studied how various ecological and behavioral factors, such as the distribution of food or group size, affect the ratio of producers to exploiters in both wild and laboratory populations, and what information individuals (primarily birds) use to adjust their choice of strategy (Vickery et al. 1991; Giraldeau and Caraco 2000; Giraldeau and Dubois 2008).

Of course we would not be human if we did not consider our own species when observing such behavior in the natural world, and to wonder how people respond to the decision to invest or to exploit. The problem of investment versus exploitation has also been studied extensively in the various social sciences, such as economics. Not only does the tension between investing and exploiting affect individuals, it can also be seen as a decision facing rival companies and even rival countries. For example, should a firm invest time and money on developing new products, such as a technologically advanced mobile phone, or perhaps a new pharmaceutical product, when it could simply wait for others to do the innovating and then copy them? Should governments order their ships to refrain from overfishing, to allow the replenishment of nearby stocks, at the risk of seeing the fish harvested by ships from another country? When comparing these scenarios to biological examples, one sees a key difference: such decisions play out on a different stage to the ecological stage of biology, and thus are affected by the decisions and institutions of governments and other interested parties. It may be that many of our laws and modern sociopolitical institutions have been shaped by a recognition of the fears of would-be investors (e.g., a permanent, paid police force to protect our bodies, families, and homes, or patent offices to protect our intellectual “property”).

In this chapter, we present a general overview of how the problem of investment versus exploitation is conceived, analyzed, and empirically tested, by both biologists and economists. We also consider how the problem can be advanced and what we still need to know. The following discussion investigates if and how the concept of investment and exploitation can be applied to real-world problems. But first, we need to clarify what kind of exploitation we are focusing on here.

If interested in discussing exploitation, one could be concerned with instances where one party exploits another in the sense that they take advantage of another’s misfortune, desperate situation, or current weakness. For example, a payday loan company that lends moderately small amounts of money in exchange for very high fees could be considered to exploit desperate people that arguably have no choice. The morality of such behavior is far from clear, as reflected by the mix of laws that try simultaneously to allow but limit the severity of such practice. However, here we are not analyzing the morality of usury, or even investment and exploitation. Instead, we are interested in what affects the “choice” to be an investor or an exploiter, and our analytical approach requires us to be able to clarify different situations.

There are at least two key distinctions that allow us to clarify the difference between the payday loan example and those examples that interest us. First, the customer, legally speaking at least, always has the choice to accept the deal on offer or to walk away. Such interactions are thus perhaps best viewed as a form of negotiation during bargaining, with successful agreements only applying to outcomes that are perceived as beneficial by both parties. However, the reality

that the customer will often be desperate for money to pay for primary needs (e.g., food, energy, heating) can make this black and white distinction a blurry gray. Second, the customer does not have the option to exploit the payday loan company. In this way, the asymmetry between the two agents means the situation more resembles an interspecies interaction than an intraspecies one (a zebra cannot become a lion). The effect of asymmetries between individual agents is a topic we explore below when considering how differences in factors, such as foraging ability and dominance, affect an individual's decision to invest or to exploit, but we restrict ourselves to examples where all individuals can, theoretically at least, choose between being an investor or exploiter in the relevant scenario.

The Biology of Individuals versus Groups

If a group of individuals all choose to exploit, there will be no production from which to draw an advantage, and such a group may be outcompeted by a group with more investors. It is therefore tempting to conclude that exploitation will not be favored. Such an explanation is, however, unsatisfactory because selection or competition does not only occur among groups but also within them. Therefore, within-group dynamics, where exploitation may provide a relative advantage, must also be considered (Williams 1966).

Darwin realized that natural selection favors the heritable components of physiology and behavior (the phenotype) that increase an individual's survival and reproduction. This process creates an appearance of design in organisms as they become increasingly well fitted (adapted) to their environment. This process of adaptation is why, within evolutionary biology, many organisms can be reasonably modeled as agents approximately maximizing their own survival and reproduction (Grafen 2006; Gardner et al. 2011). However, the heritable basis of an individual's phenotype is encoded by genes, and thus a complete understanding of evolution requires an appreciation of genetic success and survival.

Adopting a gene-centered approach to adaptation has allowed biologists to show that a gene is selected depending on both how the phenotype it encodes affects the bearer's reproductive success and how it affects the reproductive success of the bearer's relatives (Hamilton 1964). This is because a gene can increase its frequency within the population through two ways: a direct route whereby it produces a phenotype that helps its bearer to have more or better offspring and an indirect route whereby the phenotype helps other individuals (if they also contain identical copies of the gene) to produce more or better offspring. The simplest and most effective way to achieve this indirect route is by helping close relatives. However, often an adaptation that increases the success of one of these routes will come at the cost of reducing the success of the other. In these cases, selection will favor the optimal trade-off between the two, depending on the cost and benefits of helping,

and the relatedness between the actor and the recipient of the behavioral act. This is because relatedness describes, for any gene that encodes behaviors with social consequences, its above-chance probability of affecting a copy of itself (the indirect route).

The evolutionary fitness of an individual's behavior can therefore be described or measured as the sum of their direct and indirect reproduction, referred to as "inclusive fitness." In the context of social interactions, this concept has allowed biologists to retain a legitimate sense of agency by adopting the perspective of individual actors trying to maximize their inclusive fitness (Hamilton 1964; West and Gardner 2013). Inclusive fitness theory has been very successful in identifying situations where individuals will be selected to reduce or increase their exploitation of others, depending upon their relatedness to them (Gardner et al. 2011). Evolutionary biologists have also demonstrated that models which focus on group-level benefits to explain cooperative behaviors still only get cooperation to evolve when the genetic costs and benefits of the behavior to the individual provide an inclusive fitness advantage (Lehmann et al. 2007; Marshall 2011). This means that natural selection will only favor individual behaviors that serve to benefit their group if they also serve to benefit either the individual or the individual's relatives (Gardner et al. 2011). There have, however, been some notable disagreements with this view (Nowak et al. 2010; Allen and Nowak 2015), especially with regard to human evolution (Henrich 2004).

Modeling contemporary human behavior is not so straightforward, largely because of two factors. First, because our environment has undergone rapid change, we may no longer be as well adapted to our environment as we once were. For example, there is an evolutionary lag between our cravings for energy-rich fats and sugars and the modern glut of such foods in certain countries. Sugary and fatty foods were highly beneficial, but rare, in our ancestral environment, so the costly consequences of such cravings were minimal, unlike today where they cause obesity and heart disease. It is likely that the change in our available diet is not the only domain where we are lagging behind. Other potential areas are the sexual domain and our relationship with pornography and reproductive technologies, such as contraception and abortion. Likewise, it is highly unlikely that we are well adapted to handle money and the complexity and longevity of modern financial products, such as pensions.

Second, the human capacities for problem solving, learning, forward planning, and communication indicate that we are, in principle at least, capable of organizing collective behaviors that benefit society. There is debate about how we should model the evolution of group-beneficial behaviors in humans, and whether the fundamental predictions of evolutionary theory still apply (Henrich 2004). If our social traits are transmitted through cultural routes, such as imitating others, more than through genetic routes, then cultural evolutionary models will be needed to explain and predict the distribution of social behaviors (Henrich 2004; Richerson and Boyd 2005). Contrary to the predictions

of genetic models, the outcomes of cultural models are less certain and depend largely on the precise nature of any genetically based biases we exhibit in learning and copying (Aoki et al. 2011; Mesoudi 2011; El Mouden et al. 2014; Acerbi and Mesoudi 2015).

The Economics of Individuals versus Groups

The situations we discuss in this chapter often involve individuals making a choice between what is best for them and what is best for the group. For example, if more people invested in giving blood or contributing to Wikipedia, then society would be better off, but this means that individuals would have to give up some of their time. If more parents had their children vaccinated, then the overall immunity of the society would increase, but parents may feel that vaccinations expose their children to undue risks. If individuals falsely report their earnings, they will pay less tax, but then society has less revenue to use for communal services. If fishermen catch fewer fish, they earn less money, but the shared stocks have a greater chance of replenishing themselves. In all these cases, the interests of the individuals making the decisions diverge from the interests of the larger group, or society, in general.

How then should such decisions be modeled? Traditional economic models of decision making have often assumed that people know what they want and make their choices accordingly. More specifically, traditional models assume that people consider the consequences of their choices, rank the possible outcomes by how desirable they are, and then choose an action that leads to their most preferred outcome, or act as if they were doing so. Such models have the advantage of not requiring the modeler to assume to know what people want. Instead, one can infer an individual's preferences from their choices (the principle of revealed preferences). This principle does not allow us to question whether someone is making a decision in line with their preferences or not, because of the assumption that individuals prefer the outcome they choose, and choose the outcome they prefer (Kreps 1988; Kacelnik 2006).

The above principle of revealed preferences relies on *rational choice theory*, which assumes that people make choices consistent with their preferences. This assumption can only be tested by examining the pattern of choices individuals make (Allingham 2002). More formally, given an individual's preference ordering, which satisfies properties of consistency and completeness (the preference ordering covers all permutations of options), this model can be conveniently formulated as the maximization of a utility function. In this general framework, *utility* is merely a numerical ordering of outcomes; it does not make assumptions on the specific preferences of an individual other than the consistency and completeness requirements.

While rational choice does not imply that people cannot care about others' outcomes, it has often been understood to mean that people prefer to maximize their own monetary and/or material gains. This assumption greatly simplifies

models of decision making and has proved to be a reasonable approximation in many economic domains, such as market competition. This assumption, however, has not fared so well in social domains, where people have to make decisions that affect both themselves and others. Hundreds if not thousands of experiments have repeatedly shown that people in laboratory conditions will often make decisions that are suboptimal for their own finances but beneficial for the finances of others, suggesting that people do not only value their own welfare (Chaudhuri 2011; Engel 2011; Guth and Kocher 2014).

Can groups be modeled the same way? If individuals make rational choices, then surely groups of rational individuals will do so as well. However, as Mancur Olson has shown, it is illogical to consider groups to be rational just because they are comprised of rational individuals (Olson 1965). Olson's logic is that if the individuals within a group are rational, they will follow their own interests; thus the group will collectively appear irrational whenever individual interests conflict with group interests. In other words, groups may well be rational, not because they are made up of rational individuals, but because group and individual interests are aligned, which is rarely the case. This parallels the lesson from inclusive fitness theory in evolutionary biology: individual behaviors only evolve to serve the group when they also provide sufficient benefits to the individual or to the individual's relatives.

Much research has therefore focused on the idea of increasing "efficiency" in scenarios where individual and group interests are in conflict. Efficiency concepts in economics relate to the aggregation and favorable trading off of the utility of different people. Such trade-offs are central to the discipline, as evidenced by the pervasive applications of cost-benefit analysis in policy applications. One problem is that standard utility concepts are based on preference rankings; they only provide (ordinal) rankings over alternatives. Thus, while (numerical) utility values describe the relative valuations for a given individual, the absolute level can be arbitrarily rescaled, making interpersonal utility comparisons impossible. To overcome this problem, economists have developed several different efficiency concepts to help theoreticians and policy makers compare outcomes.

The most commonly accepted notion of efficiency is "Pareto efficiency," because this concept obviates the need for interpersonal comparisons of outcomes. Instead, Pareto efficiency says that an increase in efficiency, a Pareto improvement, occurs when a new allocation of resources makes at least one individual better off while keeping every other individual at least as well off. When no Pareto improvements are possible, the outcome distribution is Pareto efficient. Note that Pareto efficiency does not imply equality; it is perfectly possible for a distribution of resources to be unequal and Pareto efficient, as long as any changes toward greater equality leave at least one person worse off.

Thus, many economic analyses focus on "correcting" deviations from (Pareto) efficiency that would result from individuals pursuing their own interests. This can be accomplished through various institutions designed to shift

the relative payoffs of different choices for individuals. In this way, they attempt to incentivize more desirable behavior in rational individuals (Hurwicz and Reiter 2006). Examples include monetary incentives such as taxes, subsidies, and fines but also manipulations of social incentives such as prestige, shame, and guilt (Kosfeld and Neckermann 2011).

Furthermore, people often appear motivated by a concern for the welfare of others (other-regarding preferences) that may act to reduce the conflict between the individual and the group (Fehr and Schmidt 1999; Bolton and Ockenfels 2000). Thus, efficiency analysis may be improved by including people's sense of fairness when considering trade-offs between economic growth and inequality (Durante et al. 2014). It is also possible that traditional policies have been too pessimistic in assuming people's preferences are entirely self-interested. By making salient the idea that people are selfish, such policies may have provided a justification for selfish behavior and reduced people's faith in each other. If such policies are self-fulfilling, it may instead be better to focus on increasing people's confidence and trust in each other (Gaechter 2007; Bowles and Hwang 2008).

Given all of the above, how then can we analyze the situations of investors and exploiters in which we are interested? It would seem that we have to consider the costs and benefits of the decision to invest or to exploit at the individual and not the group level. However, if people are rational and self-interested, then why will there ever be any investors? If exploiters gain a higher payoff, then surely everyone will choose to exploit?

Let us return to our original examples of one of our ancestors laboring to make a hand ax that ultimately benefitted another man, and a crow laboring to extract a grub that ultimately benefitted another crow. In these extreme cases, there is no benefit to the individual producer, but imagine that instead of the ax being stolen straight away, our ancestor was able to use the ax for a day, or a week, or a year, before it was stolen. In this case, the costs of making the ax would have been recouped, provided that the benefits were large enough to have outweighed the costs. Likewise, imagine that the crow, instead of extracting one beetle grub, had extracted several and gotten to eat some of them before a competitor came and took some for herself. Here, the investments of the industrious crow returned a "finder's share" or "advantage" (Vickery et al. 1991). More generally, if the average finder's share sufficiently outweighed the costs, then we could argue that it is sensible to sometimes be an investor. In human societies, creating or enlarging such a finder's share is the aim of regulations like patent law, to which we will return to below. The effect of a finder's share in regulating fisheries is also discussed by Valone et al. (this volume).

Of course, even with a large finder's share, it will still sometimes be sensible to be an exploiter, for as we have outlined, whenever there are investors, there are opportunities and benefits for exploiters. Therefore, the predicted ratio of investors to exploiters will depend on the circumstances, the behavior of others, and the associated costs and benefits of the different behaviors, which

will depend on how costly it is to invest in a resource, how many potential investors there are to exploit, the cost of exploiting in terms of defense or retaliation by investors, and the value of the resource being produced (Barnard 1984). Calculating exactly how the incentives, and therefore the ratio of investors to exploiters, will change requires mathematical models.

Applying the Concepts of Investment and Exploitation to Biology

In biology, examples of investor–exploiter relationships are widespread across all scales of life. For example, in bacteria, some individuals invest in production while others exploit their efforts. Here, investors make and release siderophores, molecules which sequester valuable and metabolically essential iron from the local environment (Ratledge and Dover 2000). Siderophores, however, are costly to produce and are not guaranteed to return to the original investor. Therefore, other bacterial strains in the same neighborhood can still survive even if they do not invest in siderophore production by harvesting the siderophores produced by others (West and Buckling 2003). These microbial examples are useful because we expect them to conform to the laws of the relentless evolutionary process. Furthermore, they show that cognition and intentions are not necessary to replicate investor–exploiter dynamics. Microbial examples can thus be used to test the predictions of theoretical models without having to worry about the role of intentions or the effects of either imperfect or sophisticated cognition.

Investors do not only produce food and metabolites. The various strategies that animals use to obtain sexual partners can also be viewed as alternative mating tactics within an investor–exploiter framework. In many frog species, some males will invest in calling to attract females, putting themselves at increased risk of predation, while other males will exploit these investments, thus avoiding the increased predation risk, by staying silent, and intercepting females as they approach calling males (Lucas et al. 1996). Likewise, brood parasitism is another major class of reproductive tactics that can be effectively categorized into investor and exploiter roles. In many species of birds, females will surreptitiously lay their eggs in the nest of another female, exploiting the incubation and provisioning efforts of the nest-attending female (Davies et al. 2012).

Arguably the best-studied example of the use of investor–exploiter tactics comes from the field of social foraging (Afshar et al. 2015). When animals forage in groups, individuals can invest in searching for food, or they can exploit the food discoveries of others. This investor–exploiter scenario has been extensively studied using models that typically can only be applied to groups of more than two individuals and contain simplifying assumptions that may not be true of all such scenarios. Typically, researchers make three assumptions in these investor–exploiter models (termed PS models in much of the literature):

1. Tactics are mutually exclusive; that is, individuals cannot simultaneously search for resources and search for opportunities to exploit the resources produced by others. For example, in certain species of fish, males can either invest in building/maintaining a nest to attract females or they can try to exploit the nest building of others and sneak paternity as females approach another male's nest. By definition, a male cannot do both of these things simultaneously, although of course they can invest in doing a bit of both over time, but this means they have to allocate each unit of effort into investment or exploitation. Similarly, in many birds, the nature of their visual system means that searching for food patches on the ground requires a different head orientation than searching for congeners that have discovered food patches (Coolen et al. 2001).
2. The resource is finite and can be depleted, so that any part of the resource consumed or used by one individual is not available for other individuals. For example, food that is consumed by one individual is not available to others, or eggs fertilized by one male can no longer be fertilized by another male. In economics, as explained below, such goods are considered "rivalrous," because different individuals can be considered as rivals competing for the same goods.
3. Investors, on average, gain some advantage for having produced a resource. For example, they get to eat some amount of food they discover before another individual joins them. This is known as the finder's share (discussed above).

Several empirical studies have tested these assumptions (Giraldeau and Dubois 2008; Dubois, this volume; Barta, this volume). The predicted effects of group size, patch size, and finder's advantage on the frequency of exploitative behaviors have been experimentally tested in laboratory experiments with nutmeg mannikins (*Lonchura punctulata*), a small ground-feeding passerine, and the results qualitatively support the predictions. For example, doubling flock size from three to six individuals resulted in a 50% increase in the frequency of exploitation (Coolen 2002). Increasing the number of seeds available per food patch also resulted in increased exploitation (Coolen et al. 2001; Coolen 2002). In a very elegant experiment, Mottley and Giraldeau (2000) experimentally manipulated the finder's advantage by devising an apparatus where individuals had to pull on a string to access seeds in a food patch. The seeds then fell into a collecting dish, and the finder's advantage was manipulated by partially covering the dish in such a way so as to limit the number of seeds that could be consumed by the producer. As predicted, when the finder's advantage was reduced, the frequency of exploitation increased.

Field experiments also provide support for the main predictions of investor-exploiter models. Experimenters allowed free-living Carib grackles (*Quiscalus lugubris*) to invest effort in making dry pieces of dog food more palatable by

dunking them in puddles, or to attempt to exploit the efforts of others by stealing food from group members after it has been softened. Field observations by Morand-Ferron et al. (2007) have shown that the proportion of exploitation attempts changes with natural variation in wild Carib grackles group size (more exploitation in larger groups). In their study, Morand-Ferron et al. also experimentally manipulated the expected finder's advantage for investing in dunking dog food. To do this, they created experimental puddles of equal surface area but varying in shape, which changed the average distance to exploitation opportunities. Larger distances to exploitation opportunities meant that investors were more likely to consume the food before an exploiter could attempt to steal it (i.e., their average finder's share increased). As a result, the frequency of exploitation declined under these experimental conditions.

The above experiments show that natural selection does not have to act upon genes for being an investor or exploiter, but instead can select for behavioral changes based upon genetically encoded learning rules. If investor-exploiter scenarios are sufficiently common and important, then animals will evolve mechanisms to improve their dynamic performance in repeated instances of such "games."

Applying the Concepts of Investment and Exploitation to the Social Sciences

The social sciences often study the conflict between individual and group interests by examining either investment in the production of goods that are freely available and benefit society (public goods), or the private consumption of shared resources (common-pool resources). In these scenarios, individuals can pay a cost (invest) to produce directly a good that benefits everyone, or they can show restraint in their consumption and "invest" in future resources by avoiding overconsumption of a public resource. In both cases, individuals typically benefit more by acting against the interest of the group, and thus outcomes are predicted to be suboptimal at the group level. The scale of this individual versus group dynamic can vary: sometimes the "group" is one local group, or a country, or even the global population, and the conflicting "individual" could be a single person, a region, or even a country, respectively. In all cases, however, the "individual" is a subunit of the larger "group."

Various mechanisms have been proposed to increase the group welfare when individual- and group-level interests are not aligned. Not all scenarios are, however, the same, and because the efficacy of different mechanisms will depend on the nature of the goods being "produced," the social sciences have found it useful to categorize public goods into different types.

Goods can be conceptualized along two dimensions, according to their degree of "rivalry" and "excludability." Rivalrous goods are those which can be used by only one person simultaneously (e.g., a candy bar or a chair), over which people may wish to compete. In contrast, nonrivalrous goods are those

that can be used by many at the same time, and are therefore often intangible (e.g., national defense or scientific knowledge). Excludable goods are those for which one can limit the benefits to only those that contribute to the costs. For instance, one can exclude noncontributors from a mail service by requiring people to pay for a stamp. In contrast, one cannot prevent noncontributors from using or benefiting from nonexcludable goods. For example, it is not easy to prevent people from listening to a radio or television broadcast, or restrict access to many natural resources (e.g., fish or timber stocks).

Taken to their extremes, the combination of rivalry/nonrivalry and excludability/nonexcludability allows one to distinguish between four types of goods. In no particular order, there are private goods, which are often provided commercially, whereby people pay to use/consume an excludable and rivalrous good, such as a sandwich: here production and consumption are in the self-interests of both parties. Second, there are club goods, which are those that are excludable but effectively nonrivalrous because one additional user does not significantly affect overall costs, although too many users may lead to congestion. For example, users of a golf club can exclude nonmembers and benefit from the membership fees of each additional member, but may want to limit the number of members to avoid overcrowding at busy times. Third, there are public goods, whereby society pays for goods that are freely available to everyone (nonexcludable and nonrivalrous). Public goods are expected to be underprovisioned because self-interested individuals will be reluctant to invest to provide public benefits. Finally, there are common goods, also referred to as common-pool resources, which like public goods are also nonexcludable but, unlike public goods, are rivalrous. Fishing grounds provide a good example, because it is difficult to exclude boats from fishing, but any fish caught by one boat reduces the catch available for others. Therefore a recurring problem with common goods is how to prevent overuse. Rational individuals are predicted to maximize consumption because any restraint they show can be exploited by others increasing their consumption. Thus, well-defined property rights for common goods can theoretically solve the problem of overuse but can be costly to monitor and enforce, and may be unpopular.

How does the PS game relate to this classification of goods? If we imagine a social forager that chooses to invest in searching for food or to exploit the searching efforts of others, then it is apparent that the food produced is clearly rivalrous, but not fully excludable, if at all. The finder's share sits in the continuum between full and nonexcludability; thus, this classification of goods may not be the most useful to use in framing the PS game. Instead, the idea of rent seeking from economics may be a better fit to the PS game. Rent seeking is a broad concept used to describe efforts devoted to the capture of resources by diverting them away from other agents (Tullock 1967, 1974; Murphy et al. 1993; Congleton et al. 2008). While the definition is somewhat fluid, it crucially refers to actions that, in contrast to this chapter's use of the term production, do not create new value or surplus. The term rent seeking originates with

Adam's Smith's classification of income into profit, wages, or rent, and refers to the attempt to gain control of (natural) resources.

Gordon Tullock designed a flexible game theoretic framework where participants exert efforts to compete for the possession of a resource, called the Tullock contest or rent-seeking game introduced above (see also King et al., this volume). This framework has been applied to a range of activities, such as lobbying for government protection via restrictions to foreign imports or limiting access to lucrative industries like notaries and pharmacies. Dechenaux et al. (2015) review experimental literature that has studied behavior in this game in the laboratory, and Murphy et al. (1993) discuss an application to economic growth and development, where rent seekers target surplus-generating activities like innovative companies or high-yield farming. The authors distinguish between private and public rent seeking. The former consists of theft or piracy, whereas the latter includes government corruption in the issuance of licenses or other government services. The authors suggest that innovative activities are especially vulnerable to public rent seeking, as innovators are dependent on government services and do not have established lobbies to obtain administrative favors for themselves. As a consequence many companies prefer to operate in the informal sector (and not become too big) and to operate under the radar of the corrupt government.

Murphy et al. (1993) include a discussion of the frequency dependence of the returns to productive and rent-seeking activities. As one may expect, the return to productive activities causes a decline in the amount of rent seekers in the economy. However, the authors propose that there are complementarities in rent seeking which may cause the returns to increase with the amount of rent seekers, at least for low levels of rent seeking. Sources of complementarities may be strength in numbers among rent seekers, as embodied by gangs or mafias, or the escalation of arms races in the fight over resources. Depending on the exact nature of the frequency dependence, multiple equilibria may exist that are characterized by either high levels of rent seeking and low levels of production or vice versa. Murphy et al. (1993) suggest that such multiplicity may help explain variances in the development of different countries.

One application of the rent-seeking model in the context of developing countries is to the phenomenon of forced interfamily solidarity, where the requirement to share within the extended family network is an integral part of everyday life, one exacerbated by the lack of formal insurance systems (this was first reported by anthropologists in the 1960s; Belshaw 1965). The literature on "forced solidarity" or the "dark side of social capital" shows that acts of giving are not always voluntary and are often "demanded" from the network members, thereby deterring investment and effort of the producer (di Falco and Bulte 2011). Entrepreneurs must resist normative pressures to support their extended families if they wish to reinvest in their firms (Belshaw 1965; Nafziger 1969; Bloch 1973; Hart 1975). This has led to the hypothesis that high demand for financial support by members of an entrepreneur's family hinders savings

and investments (in capital and labor) and thus may hinder long-term economic growth. The above similarities between the concepts and models of the PS game and rent seeking are another example of parallel developments in the fields of evolutionary biology and economics. Both fields utilize a cost-benefit approach with the concept of equilibrium thinking, but alas their use of different terms for similar phenomena can create a potential barrier to interdisciplinary collaboration between like-minded researchers.

There are scenarios that are still not captured by either investment in public goods, the PS game, or the rent-seeking game. These include situations where the investments of several members are required to produce a good, and situations where the investments of one or a minority of members are sufficient to benefit all members of a group. Here, the benefits of investment are not directly proportional to the number of investors. Imagine a group of lions hunting for a buffalo. More than one individual lion is required to kill the buffalo, so a team of lions is required, but each lion is better off letting other lions take the risk of attacking the buffalo. A modern human example is when people profit from open-source software. More than one producer is required to generate it, but everybody has access to the software. Alternatively, if several people witness an accident or crime, it only takes one person to phone the emergency services; this “diffusion of responsibility” can create the possibility that everyone leaves the task to someone else (Darley and Latane 1968). Situations requiring a certain threshold of investors before a potential good can be realized have been studied in economics as threshold public goods games, and there is a considerable experimental literature on such games (Croson and Marks 2000).

It has been argued that this type of scenario, the so-called volunteer’s dilemma (Diekmann 1985, 1993), also known as the snowdrift game, chicken game, or hawk–dove game (Maynard-Smith and Price 1973; Maynard-Smith 1982), is more relevant to the study of exploitation, especially within biology, than the classical linear public goods game (Archetti and Scheuring 2011). As in the PS game, there can be an equilibrium point between strategies that invest and exploitative strategies in the volunteer’s dilemma. The idea of requiring a threshold number of investors to produce the actual good is a potential extension of the PS game for future research. It is reasonable to assume that one investor might not be sufficient to produce any goods. Alternatively, the presence of more than one investor at a patch (e.g., by chance, in a world where the number of patches is finite) might be required to attract exploiters. Such modification would probably decrease overall resource use because all lonely investors would not be able to profit from the resource (or attract exploiters), which would decrease the population’s average payoff.

Finally, there has sometimes been concern expressed that if individuals receive help from willing investors, they are then less likely to become investors themselves. This has been termed the Samaritan’s dilemma, whereby an altruistically motivated person has to decide between leaving a helpless person to suffer and providing help at the risk of reducing the recipient’s self-reliance

(Buchanan 1975). If a potential recipient knows that the Samaritan will help, then his/her incentive to work may, theoretically speaking, be reduced. Such arguments can be used as a rationale for reducing publicly funded welfare or social security for people who face unemployment. However, such an argument may rely on the benefits of employment being less than those of unemployment, unless one is arguing that some people are of a different disposition to others with regard to their desire to work. A meaningful analysis of such a dilemma requires more understanding of the psychological and sociological forces that impact an individual's ability, desire, and opportunities to work.

Model Behavior

If decisions are based upon the costs and benefits of different outcomes, then such decisions have to be modeled mathematically. Sometimes a necessary trade-off of such modeling is that complex real-world phenomena have to be studied with simple "games" that abstract away much of the important complexity. This can generate valid concern that the subsequent analyses are overly simplistic and reductionist. However, a cost-benefit approach is necessary to model evolutionary outcomes because the process of natural selection blindly operates on differential success in terms of reproduction. In economics, as outlined above, the cost-benefit approach can be justified by the conservative expectation that people prefer to have more resources/money. While this assumption can be falsified in certain cases, it vastly simplifies the mathematical models and makes the analysis feasible. If we relax this assumption, but maintain the idea that individuals maximize something, as defined by their own personal utility function, then whatever is being maximized can be mathematically modeled (Kacelnik 2006).

Sometimes the best behavior or "strategy" depends on what others are doing. For example, the costs and benefits to a crow that uses a tool to extract grubs depend on how many other nearby crows are being exploiters. This decision can be thought of as stemming from a game, whereby it would be beneficial for an individual to know what other individuals are doing before making a decision on how to respond strategically. One way to model such situations is through game theory, which can be used to calculate what the stable outcome of such games will be, provided the "players" can either anticipate the moves of other players or change their strategy accordingly in response to success or failure. The adjustment of strategy choice can be through natural selection acting upon the differential success of competing versions of the same gene (alleles), through various economic processes that favor organizations using more successful strategies, or through various learning processes within individuals and organizations. However, if the success of a strategy relies upon it being rare, then as successful strategies become more common they will enjoy increasingly less success, which will limit their growth in the population. This dynamic is known as negative frequency dependence (Figure 3.1).




Figure 3.1 The payoffs (net benefit) of investing or exploiting are frequency dependent. When there are many investors, exploiters obtain a larger payoff than investors because they invest no effort in discovering resources, and there are a large number of investors to exploit. In contrast, when there are many exploiters, the average payoff of exploiting decreases because investors are harder to find and there is more competition to exploit them. At some intermediate mixture of investors and exploiters, both tactics receive equal payoffs (dashed vertical line). This is known as the stable equilibrium frequency, and no one individual can use a shift in tactic to increase their own payoff. ESS: evolutionarily stable strategy.

One crucial concept used to predict the outcome of negative frequency dependence, in both biology and economics, is the idea of equilibrium, which can be thought of as the solution to, or the predicted outcome of, the above types of game. This is the idea that when a population of individuals reaches a certain state, whereby there is a certain proportion of individuals adopting a certain behavioral strategy, no one individual can improve their situation by changing their strategy. This is key for economists because it means that no rational individual will be predicted to change their behavior unless they want to incur a loss, and that other rational individuals can use this assumption to anticipate the behavior of other players, leading to a stable outcome if players are rational and have correct beliefs; this is known as a Nash equilibrium (Nash 1950). For biologists, the concept is equally crucial, because it identifies the evolutionary resting state of the population where individuals will behave according to an evolutionarily stable strategy (ESS) (Maynard-Smith and Price 1973; Maynard-Smith 1982).

A strategy is evolutionarily stable when no plausible mutant strategy can improve upon it. The process also applies to the short-term changes in animal behavior when animals use learning to respond to changes in their own payoffs and/or in the behavior of others. We therefore expect animals to be well adapted to their social environment, even if we observe different behaviors in different individuals of the same population, or different behaviors over time within the same individual. Thus a game theoretical mathematical approach,

while limited, offers potentially large benefits to the analysis of social behaviors such as investment and exploitation. One may wonder if it is too hard for human rationality to compete with the rationality of the evolutionary process (Kacelnik 2006). Nash himself, however, provided two interpretations of his equilibrium concept: one based on rational individuals that calculate the best action, and the other based on a large population of individuals in “mass action” that use their experience and limited knowledge to gravitate toward the equilibrium (Kuhn et al. 1996).

A Mathematical Model of Exploitation

Here we outline a mathematical model of exploitation that will form the basis of our subsequent analyses. Consider G group members engaged in joint production. The basic situation we are interested in is one where individuals can either invest in a productive activity or in an exploitative activity that parasitizes the investments and efforts of producers, such as in the PS game (Giraldeau and Caraco 2000).

We refer to the amount invested by an individual i into the productive activity x_i and into exploitation y_i . In the PS game, individuals either invest all their resources into production or into exploitation: $x \in [0, 1]$; $y = 1 - x$.

If an individual invests into production, the individual produces productive effort for T periods and produces an amount F with success rate λ . Parameters that affect who benefits from production, β and δ , allow a direct comparison between the PS and public good models. If a proportion $(1 - p)$ of the individuals in the group produces, then the payoff (fitness) of any individual producer ($x = 1$) is:

$$W_x(p) = [1 - (1 - \beta)pG]T\lambda \left(a + \frac{F - \beta a}{\delta G + \beta} \right). \quad (3.1)$$

The payoff of an individual using the exploitation strategy ($y = 1$) is:

$$W_y(p) = (1 - p)GT\lambda \left(\frac{F - \beta a}{\delta G + \beta} \right). \quad (3.2)$$

Under the PS model, $a > 0$ represents the “finder’s share” discussed above in our crow example, a small reward to the producer for being the first to enjoy the benefits of production. Further, $\beta = 1$ and $\delta = p$; that is, the results of production are shared among all the exploiters as well as the producer of F , but not the other producers. While an exploiter in the PS model can benefit from the investments of all other group members, producers only reap the benefits of their own investments.

The Nash equilibrium of the game, where both strategies coexist and have equal payoffs, can be determined by equating Equation 3.1 and Equation 3.2, and solving for p . This yields the equilibrium frequency of scroungers \hat{p} , which

is also the ESS since for any p above or below \hat{p} , the direction of selection will go toward \hat{p} ; that is, $W_y(p) > W_x(p)$ if $p < \hat{p}$ and $W_x(p) > W_y(p)$ if $p > \hat{p}$ (Giraldeau and Caraco 2000):

$$\hat{p} = 1 - \frac{a}{F} - \frac{1}{G}. \quad (3.3)$$

In economics, socioeconomic interactions similar to the PS setting are typically modeled via rent-seeking models (Tullock 1967). In rent seeking, individuals can allocate any amount $x \in [0, 1]$ to production, while again $y = 1 - x$. The payoffs (fitness) of any individual i is then given by:

$$W(x, y) = \frac{y_i}{\sum_j y_j} \sum_j x_j, \quad (3.4)$$

where the sum of the productive efforts of all individuals gives the overall amount shared among all and the share each individual obtains is linearly proportional to the amount invested into exploitation.

The equilibrium amount each individual allocates to exploitation in the unique symmetric equilibrium (there are asymmetric ones as well) is:

$$\hat{y} = \frac{G(G-1)-1}{G(G-1)}. \quad (3.5)$$

This is only well defined for groups of at least size 2, for which individuals devote exactly half their resources to production. For larger groups, the amount devoted to production (x) decreases, just as it does in the PS model.

A second, and more common way to model the organization of trade-off between private and public interests in the social sciences is the so-called public goods game. Here, producers make personally costly investments into products, resources, or services that benefit the entire group. Typically these group-level benefits are assumed to be larger than the individual costs but not so much as to outweigh the individual costs of production. Therefore, the public goods game is different to the PS game in that $a < 0$: producers do not get a net reward for producing, but instead have to pay a net cost for the effort of producing. A second difference is that $\beta = 0$ and $\delta = 1$: all individuals in a group share the benefits of production equally, regardless of whether they are a producer or not. In other words, the costs of production are entirely private, but the benefits are entirely public. Thus, the income-maximizing strategy is generally to not produce, unless the benefits of production are sufficiently large for a given group size. As a result, under commonly assumed parameters, the unique ESS is $\hat{p} = 1$.

Here we outline a formal representation of this game that shows the similarities and differences to the PS paradigm and notation. In behavioral economic experiments, linear public goods games are implemented slightly differently and the customary payoff function is:

$$W_i(x, y) = \frac{\alpha}{G} \sum_j x_j + y_i, \quad (3.6)$$

where individuals can allocate any amount x between 0 and 1 to production, while again $y = 1 - x$. The parameter α is an efficiency factor that determines how large the group-level benefits are, and is often chosen as 1.6 or 2.0 in economic experiments using groups of around four or five individuals (Isaac et al. 1984; Andreoni 1995; Fehr and Gaechter 2000; Fischbacher and Gaechter 2010). As long as α is not greater than G , the unique equilibrium is again $y = 1$. For example, in a group of four players with $\alpha = 1.6$, each dollar invested in production returns 1.6 dollars to the group, meaning that each individual gains 0.4 dollars and the contributor makes a net loss of 0.6 dollars.

Model Extensions

Models are advantageous in that they provide tractable, analyzable, but simplified versions of the real world. However, the wrong model, or an overly simple model, may not only be of limited use, it may actually provide results that are at odds with those from a more complicated model. Thus it is necessary to push for richer game theoretic models that incorporate more of the actual behavior involved (McNamara 2013). Here we consider the effects of several possible extensions to the PS game modeled above.

Limitations of the Game Theoretic Framework

Simple game theoretic models run the risk of deducing universal solutions that may not fit the problem at hand. Thus, defining ecological and socio-economic conditions under which more or less scrounging emerges is key to understanding systems dynamics, and potentially increasing the efficiency of these systems. The social-ecological system (SES) framework, proposed by Elinor Ostrom (2007) in the context of common-pool resources management, is a comprehensive way of describing the major variables that need to be accounted for in more realistic interactions between (human) agents in a given environment. The SES framework is used to explain the management of a resource according to rules and procedures determined by a governance system. On the most fundamental level, Ostrom distinguished between four fundamental properties of social ecological systems: the resource system, the resource units, the governance system, and the users of that system. Within these “first-tier” categories, specific “second-tier” variables are listed that have been tested empirically to influence the stability of social-ecological systems (Table 3.1).

Laboratory or field evidence of investment and exploitation in the PS game comes from very specific situations, which, of course, limit the ability to generalize these findings. The most common experimental setup involves socially foraging birds as the users competing for immobile resources. The interaction

Table 3.1 Variables in the social-ecological system (SES) framework, based on Ostrom (2007).

Resource System	Governance System
Sector (water, forests, pastures, fish)	Government organizations
Clarity of system boundaries	Nongovernment organizations
Size of resource system	Network structure
Human-constructed facilities	Property rights systems
Productivity of system	Operational rules
Equilibrium properties	Collective choice rules
Predictability of system dynamics	Constitutional rules
Storage characteristics	Monitoring and sanctioning processes
Location	
Resource Units	Users
Resource unit mobility	Number of users
Growth or replacement rate	Socioeconomic attributes of users
Interaction among resource units	History of use
Economic value	Location
Size	Leadership/entrepreneurship
Distinctive markings	Norms/social capital
Spatial and temporal distribution	Knowledge of SES/mental models
	Dependence on resource
	Technology used
Interactions	Outcomes
Harvesting levels of diverse users	Social performance measures (e.g., efficiency, equity, accountability)
Information sharing among users	
Deliberation processes	Ecological performance measures (e.g., overharvested, resilience, diversity)
Conflicts among users	
Investment activities	Externalities to others SESs
Lobbying activities	

among agents is limited to investing in producing food or exploiting the investments of others. Adding to this setup, behavioral ecologists have studied the variable influence of different factors, including the number of users, the size of the finder's share, the resource value and group size. Below we aim to highlight some of the evidence related to these second-tier variables or make projections on how they might shift the fraction of investors and exploiters in a given population.

Resource Heterogeneity

The basic PS game does not incorporate complex ecological conditions, such as variable amounts of resources and how they are distributed. We illustrated earlier that exploitative rent seeking might explain low growth in developing

countries. However, some countries may be more prone to rent seeking because of the amount or type of resources they have. Sachs and Warner (2001) discuss that rent seeking and corruption may be more likely in countries that are rich in natural resources because such resources are more concentrated and appropriable by government officials. Thus, the productivity of the resource system, the predictability of the resource dynamics, and the temporal-spatial distribution of resources, to name just three of the key factors mentioned by Ostrom, might influence the amount of exploitation in a society. Here we aim to present some predictions on the effect of different resource characteristics on the proportion of exploitation.

Resource heterogeneity can arise when patches vary in the amount of resource they contain (i.e., when their quality varies). Recent simulations by Afshar and Giraldeau (2014) predicted that varying patch quality results in an increase in the frequency of exploiters in the population, a result they also confirmed empirically (Afshar et al. 2015). A potential extension of their model is to vary both the mean and variance in patch richness between groups and to allow for (potentially costly) migration between them, so as to determine their effects on the overall frequency of exploiters in the population. Another interesting extension of the baseline model of the PS game is to add a dynamical component to the resource. Barta and Giraldeau (2001) investigate a situation in which food patches are ephemeral, such as flying insects. In their model, once an investor has found a resource patch, both investors and exploiters do not have the time to deplete it completely. Under such conditions, the presence of exploiters might not decrease the investors' intake and resource use might not depend on the frequency of exploiters, because individuals do not have the time to compete for the entire resource and can only consume a fraction of it. Resource heterogeneity might be introduced in this model by varying the time available to consume the resource between groups and to allow for different efficacies of resource consumption (or searching efficacies) between individuals. In this case, individuals who are efficient in collecting the resource might be selected for, which would increase the population-level resource use.

In reality, resource patches are spatially structured in a finite three-dimensional space. Hence, the distance between resource patches is expected to have considerable effects on the population's average intake (Beauchamp 2008). Large distances between patches are likely to result in both lower encounter rates for investors (i.e., λ in our mathematical model), but also lower detection rates for exploiters. Both theoretical and empirical studies have shown that lower probabilities of finding patches result in lower frequencies of exploitation in the population (Beauchamp 2008; Afshar and Giraldeau 2014). The rate at which exploiters detect or join investors will impact directly on the finder's share: if it is low, investors will have more time to profit alone from the patch, or will share the remaining resources with fewer exploiters than the total number of exploiters in the group. This, in turn, results in more investors

at equilibrium than otherwise (Caraco and Giraldeau 1991; Vickery et al. 1991; Hamilton 2002; Afshar and Giraldeau 2014).

Alternatively, with small distances between patches, population density will be higher. This will result in higher encounter rates for investors and hence fewer exploiters (Hamilton 2002), but also higher detection rates for exploiters. In addition, the relative ratio of investors to the number of patches will affect how many investors are present on a patch, in contrast to the basic game where the number of patches is assumed to be considerably larger than group size. Under such conditions, the finder's share will be decreased, and the population's equilibrium will be shifted toward more exploiters than otherwise. Finally, resources might be characterized by both spatial structure and temporal variability in quality. Here, the resulting dynamics are likely to be complex, and how the population's equilibrium would be affected by such characteristics remains to be investigated.

Population Structure

In many models of the PS game, interactions between individuals are assumed to occur at random. Real-life populations are, however, often characterized by some degree of structure, either spatially or in terms of strength of ties between interacting agents. Therefore, interactions do not occur at random. Structured interactions are expected to have a considerable impact on population-level outcomes. For example, Mathot and Giraldeau (2010) showed that in groups of related individuals, the average higher relatedness led to higher proportion of exploiters at the equilibrium, because investors tolerated exploitation from relatives but imposed costs (via aggressiveness) on unrelated scroungers. Likewise, if different groups have different sizes, which can change over time, this can alter the equilibrium frequency of investors. More research is needed, however, to understand fully how dynamic group sizes affect the PS game. Furthermore, it has been shown that group augmentation, whereby groups grow until reaching some maximum size for their environment (carrying capacity), can actually enhance cooperation (Kokko et al. 2001). Therefore, an increased frequency of investors could be predicted in populations that are still growing or have not reached their maximum carrying capacity.

Interactions within groups are also not random because different individuals occupy different positions within groups and social networks. Individuals within the center of groups have reason to be more exploitative, because there are more potential investors for them to exploit, and more potential exploiters to take advantage of their own investment. In contrast, peripheral individuals are more likely to benefit from investing in their own production (Barta et al. 1997). To test these predictions, Flynn and Giraldeau (2001) trained a subset of individuals within groups to be investors in an experimental study of captive ground-feeding nutmeg mannikins (*L. punctulata*). In support of the theory,

they found that individuals with the ability to invest were more likely to adopt peripheral locations, whereas exploiters tended to adopt more central positions.

Defining the Individual: When Teams Compete

The interests of individuals and groups, or more specifically, among the individuals within groups, are often in conflict. For example, an individual benefits if others pay more taxes. Conflicts of interest also occur between subgroups of a larger group, as in when business tries to enforce patents over generic pharmaceuticals, or when different nation-states negotiate a treaty to limit anthropogenic climate change. In the tax example, individual interests compete; in the latter examples, business or international interests compete. Therefore, as there are opportunities for exploitation in both, it may at times be appropriate to consider multiple individuals acting together as a single entity within a game played at a higher level. Conversely, such higher-level individuals may be weakened by an internal game between investors and exploiters. It would be useful to know if the dynamics of the PS game apply equally well regardless of the scale of the competing units.

The cognitive abilities of humans to enter into contracts means that groups of individuals may be very effective in their collective strategy (investing or exploiting) within a higher-level game. For example, Milinski et al. (2016) show that individuals within experimental groups that mimic nations involved in negotiating issues related to climate change will select leaders that exploit on their behalf at the international level. Another example is members of street gangs, who like a pride of lions cooperating to hunt large game, may collaborate to rob other individuals. However, even within the gang, there may be greater investors, who put in most effort or perform the most risky part of the jobs, and exploiters who pocket a share of the gang's proceeds without having shared in the risks. This shows that whether an individual should be classified as an investor or exploiter depends on the level of competition upon which we focus. A Mafioso may be seen as an investor on the lower level where he contributes to the activities of the gang, which as a group constitutes a bunch of exploiters at the higher level of society as a whole (see Foster et al. and King et al., both this volume).

The forces that determine equilibrium at one level effect its impact at higher levels. This is because within-level dynamics lead to the group being less effective at whatever it is doing than it could be in the absence of such an internal game. For instance, fishermen on the same boat have to cooperate to maximize their harvest, but one individual may save energy or avoid danger by shirking work. At the individual level, this fisherman is exploiting the efforts of his crewmates. However, if the boat is competing with other boats to harvest rapaciously a shared patch of sea, then any boat that restrains from maximal captures of fish actually invests in the replenishment of fish stocks. In other words, the unmotivated fisherman who was an exploiter at the individual

level becomes an investor in the game between boats (see Valone et al., this volume). It is thus important to always be clear about the level of competition under analysis and who is responsible for the decisions in the “game.”

Cognition and Learning Behavior

Game theoretical models of social interactions are vastly simplified depictions of animal behavior. The robustness of any results deriving from such simplified depictions needs to be tested by using models that incorporate more mechanistic aspects of behavior (McNamara 2013). For instance, are individuals able to learn over time whether it is better to be an investor or exploiter, or do they have a fixed disposition (Afshar and Giraldeau 2014)? If individuals (or teams) learn, then what information do they use to update their strategies? Do they rely on personal experience (individual learning), or do they learn by observing the behaviors, and perhaps success, of others (social learning)? Here we explore how flexible investor–exploiter tactics are as well as how animals attend to uncertainty, and consider the implications of both individual and social learning in humans and nonhumans.

Individual Consistency

One issue that arises from social foraging experiments and observations is whether investors and exploiters are different types of individuals, or whether these are just different tactical roles that an individual may adopt under different circumstances. The answer appears to be a bit of both. Experiments have shown that individuals appear to shift from being an investor to an exploiter, and vice versa, when the relative payoffs (or risks involved) change accordingly (Koops and Giraldeau 1996; McCormack et al. 2007; Morand-Ferron et al. 2007, 2011). There is also evidence for some individual consistency in strategy use. However, we lack a clear understanding of what makes some individuals be producers and others scroungers.

The corresponding evolutionary question is: Are all individuals in a population genetically endowed with a similar, but flexible, behavioral strategy that can adjust to local circumstances, or do different individuals have different genetically encoded strategies that are maintained by a process of frequency-dependent selection? One way to answer this would be to investigate if there is any genetic basis to these differences in propensity to adopt an investing or exploiting behavior. For example, a genetic basis has been discovered for the alternative mating strategies of male ruffs, *Philomachus pugnax*, whereby three different, genetically encoded, strategies coexist in the population (Lank et al. 1995; Kupper et al. 2016): males can be aggressive “independents,” semi-cooperative “satellites,” or “faeders” that mimic female appearance. The complex interactions between these strategies have at times been considered

analogous to the PS game, whereby some males exploit the efforts of other males to attract and arouse females (Barnard and Behnke 1990), similar to the male frogs mentioned above that avoid the costs of predation by remaining silent and exploiting the calling efforts of other male frogs.

There is some evidence for individual consistency in strategy use, but it is unclear whether individual consistency reflects a stable individual disposition or is an artifact of a flexible behavior that appears stable in a stable social environment. For example, individuals may adopt a suitable tactic considering their relative social rank, with weaker individuals being exploited by dominant individuals. Evidence that an individual's "type" may arise in response to their position in their social world comes from Morand-Ferron et al. (2011), who showed that simply moving individuals to new groups could erase apparent individual consistency in nutmeg mannikins. Furthermore, McCormack et al.'s (2007) work on Mexican jays showed that although many individual jays consistently used one strategy more than the other, many actually used a mixture of strategies, opportunistically choosing to stop searching for food when a subordinate to themselves, from whom they could seize the food, was searching nearby.

Alternatively, it may be true that investors and exploiters belong to different distinct types, but that these types may be better characterized by other, associated, behavioral qualities. For example, Katsnelson et al. (2011) showed that young house sparrows (*Passer domesticus*) were individually consistent when choosing to invest or exploit, but that consistent use of the strategy was predicted by better performance in a prior, foraging-related, learning task. Here, investing may be an associated behavior of superior foragers, because investing is relatively more beneficial for them than it is for inferior foragers. Katsnelson et al.'s (2011) results support, therefore, Arbilly et al.'s (2010) theoretical analysis: over evolutionary time, a tendency for investor behavior may become coupled with sophisticated, but costly, learning behaviors, and likewise simple but cheap learning mechanisms may become coupled with exploitative behaviors.

Much of the literature in psychology and economics implicitly or explicitly claims that players in games have mutually exclusive "types," with some valuing social concerns more than others (Fischbacher and Gächter 2010). In social psychology and economics, different methods have been developed to assess how the social value orientations of different individuals vary (Grzelak et al. 1988; Liebrand and McClintock 1988; Fischbacher et al. 2001; Charness and Rabin 2002; Murphy et al. 2011). One of the most common methods involves so-called decomposed games, which remove any strategic concerns from social decision making and aim to measure an individual's concern for others (Murphy and Ackermann 2014). Recent evidence, however, suggests that a nonnegligible share of people may embrace seemingly mutually exclusive dispositions. Studies in Mexico with fishermen and in Namibia with pastoralists found that a large share of individuals were both prosocial as well as

antisocial toward their fellow villagers, and that contextual environmental factors at the group level explained variation of this behavioral pattern (Prediger et al. 2014; Basurto et al. 2016).

Addressing Uncertainty in Investor–Exploiter Games

Organisms often have to attend to uncertainty in variable environments. Under a range of scenarios, individuals are expected to be sensitive to the average (mean) gain they can achieve from a behavioral action as well as to respond to variance around that average (Stephens 1981). In contrast, the standard biological formulation of the PS game models foraging returns based on a unique value for each model parameter. For example, how often an animal encounters a patch that contains food (the encounter rate) is modeled as λ , which represents the average encounter rate with food patches. However, producers will sometimes be in an environment where they discover food with encounter rates above the average as well as below the average. Thus, for any given searching bout, an animal experiences uncertainty in the exact time it will take to discover a food patch.

Does this type of uncertainty alter investor–exploiter interactions? The nature of the response to uncertainty will depend on whether the negative consequences of a deviation below the average are greater or less than an equal deviation above the average (Stephens 1981). More generally, several stochastic dynamic models demonstrate that stochastic variation, either in patch richness or patch encounter rates, exerts stronger effects on the variance in intake rates for investors compared with exploiters (Caraco and Giraldeau 1991; Barta and Giraldeau 2000; Afshar et al. 2015). These predictions have been confirmed experimentally (Lendvai et al. 2004; Wu and Giraldeau 2005). Therefore, even though both investor and exploiter tactics receive equal payoffs at the equilibrium frequencies, they can differ in the variance of payoffs they experience.

Although uncertainty is likely to have important consequences for the outcomes of investor–exploiter interactions (e.g., stable frequency of tactics, individual differences in tactic use), there are only a handful of empirical studies that have investigated how investors and exploiters respond to uncertainty. This remains an uncharted area in need of exploration (Lendvai et al. 2004; Wu and Giraldeau 2005; Mathot et al. 2009; Afshar et al. 2015).

Learning in Nonhumans

The evidence discussed above suggests that, at least in some species, individuals learn from experience to adjust their use of the investor and exploiter strategies, rather than relying on a fixed strategy, and that populations will reach (or arrive close to) the ESS through learning. This is similar to the idea of Nash’s large population of individuals in “mass action,” who use their experience and limited knowledge to gravitate toward the equilibrium (Kuhn et al. 1996).

However, as learning presumably entails costs, a strategy that uses learning would have to be superior to a fixed or randomly mixed strategy, otherwise it will not be favored by natural selection.

If individual agents face variable environments, they may do better to employ learning behaviors that update tactics, depending on both individual and/or social experience. In the scenario depicted by the PS game, the payoffs of investing and exploiting depend on two quantities: the outcome of each strategy when interacting with each possible alternative, and the frequency with which the strategies are present in the population. This raises the question of how the behavior of agents can be tuned to these two categories of information. The solution depends on the nature of the agents being considered.

Simple psychological mechanisms of reinforcement that reward profitable behaviors can lead to an increase in the actions that are, here and now, best. For instance, as rich patches are depleted, or a large number of individuals are driven to exploit rather than invest in the search for food, the experienced payoff by each individual will shift. Thus, if individuals respond appropriately to the changes in their experienced payoffs, the incidence of each kind of action in the population will shift dynamically in the direction that an outside observer with full knowledge might predict.

Several experiments in birds have shown that individual experience affects strategy choice. Nutmeg mannikins have been shown to adjust their use of the exploiter strategy in response to the distribution of food, but previous experience affected how quickly and accurately they adjusted to the new condition (Morand-Ferron and Giraldeau 2010). Perhaps the most direct evidence of strategy-use learning comes from an analysis of strategy choice in house sparrows. Belmaker et al. (2012) found that individuals were more likely to use a previously experienced strategy that had yielded a higher success rate (Belmaker et al. 2012). Similarly, Katsnelson et al. (2008) showed that individual experience of different social environments can affect strategy choice in socially foraging house sparrows. Specifically, hand-reared house sparrows that experienced a “productive mother” (a stuffed female sparrow that frequently made food available by only pecking in places where there was food) were more likely to later adopt an exploitative strategy.

The mathematical description of how animals update memory with their experience of alternative strategies (learning rule), and how they choose between strategies based on these memories (decision rule), has been debated over recent decades. A number of models have tried to identify the evolutionarily stable learning rule: the mathematical rule which, much like the ESS, once fixed in the population, cannot be invaded and replaced by any other rule (Harley 1981; Tracy and Seaman 1995). By implication, this rule should allow learners to reach the ESS or at least approximate it. More recent work has used agent-based simulations to find the evolutionarily stable learning rule in an investor–exploiter framework (Beauchamp 2000; Hamblin and Giraldeau 2009). In these models, learners continuously modified their strategy based on

previous experience, and their strategy choice at each time step affected the experience of others in the population. These studies identified some learning rules as more evolutionarily stable than others. Still, since the poor performance of some of these rules can be resolved through coupling with flexible decision rules (Arbilly 2015), the evolutionarily stable learning rule in the PS game has yet to be determined. Since there are likely various learning processes that may converge to the ESS (Selten and Hammerstein 1984), fitting into different learning models detailed data of the behavioral choices animals make following experience might be the best way to identify these rules.

While learning may seem like the best way to approach the problem of a changing ESS, evolutionary models of the PS game that compare strategy learning with innately fixed strategies suggest that the advantage of learning is not straightforward. The social, rapidly changing environment presents a serious challenge: when everyone adjusts their behavior based on previous experience, previous experience may become irrelevant. Explicit modeling of the learning process in the PS game, for example, revealed that learning can be favored only in a fast-changing (physical) environment or when individuals have some preexisting trait that makes them perform better in one of the two strategies (Katsnelson et al. 2011). Furthermore, since learning is presumably costly, it may be disfavored once the population reaches a stable equilibrium, because learning is no longer needed (Dubois et al. 2010).

Learning in Economics

In many of the scenarios under discussion, individuals interact repeatedly for some indefinite amount of time. Such repeated interaction gives a potentially important role to cognition, as human actors can theoretically anticipate the future behavior of their “opponent.” By reasoning what their opponent will do in the final round of a series of interactions, individuals can work backward to choose their current actions accordingly, using “backward induction” (Von Neumann and Morgenstern 1953; Aumann 1995; Binmore 1996). For example, imagine a scenario where there is only one round of interactions between two players, and that not cooperating is individually the most attractive option even though mutual cooperation is more attractive than mutual exploitation (the prisoner’s dilemma) (Rapoport and Chammah 1965; Axelrod and Hamilton 1981; Tucker 1983). Here, the Nash equilibrium, the best response by any one player to the actions of the other, is to not cooperate (to exploit the other). If this scenario is repeated indefinitely, it can pay to invest in cooperating as the benefits of a long-run cooperative relationship can outweigh the short-term gains of exploitation, and there can be many Nash equilibria in repeated games (Kreps et al. 1982). However, if the number of interactions is common knowledge, then individuals can use backward induction to reason that their partner will exploit them in the final round. Thus, they may as well exploit them in the

preceding round rather than invest in cooperating, and so on all the way back to the opening interaction (although see Pettit and Sugden 1989).

For another example of backward induction undermining cooperative strategies in repeated games, consider the strategy of grim trigger (Axelrod 2000). This strategy starts out by choosing cooperative actions, such as not using nuclear missiles, but threatens to switch permanently to noncooperative actions after the opponent has been observed using exploitation (e.g., using nuclear missiles). The threat of this strategy serves to deter exploitative strategies and implies that mutual cooperation can be maintained. However, if there is a final round to the game, perhaps due to nuclear annihilation, then there is either no opportunity to retaliate or no incentive to follow through on a costly deterrent. This means the deterrent is no longer a credible threat, and thus exploitative strategies should be used in the final round of the game. Consequently, there is no reward for cooperating in earlier rounds of the game, meaning cooperation is disfavored.

Empirical evidence on whether or not humans use backward induction is mixed (Binmore et al. 2002). Typical evidence from the prisoner's dilemma suggests that backward-induction reasoning kicks in only as the end of an interaction becomes near (Andreoni and Miller 1993; Embrey et al. 2016). This pattern has been explained by learning models that view human actors as adaptive but capable of displaying some degree of foresight (Jehiel 2001; Heller 2015).

Whether or not nonhuman animals should be viewed as consistent with backward induction will largely depend on whether evolutionary selective pressures select for outcomes consistent with backward induction (Noldeke and Samuelson 1993). Learning and cultural transmission also play an important role for humans in acquiring strategies, consistent with approaches such as backward induction (see Mengel and van der Weele, this volume).

Models used to describe learning in economics can require little in terms of cognitive resources from agents. Take, for example, reinforcement or stimulus-response learning, where actions that have led to good outcomes in the past are more likely to be repeated in the future. Agents have a probability distribution over possible actions. When an action is chosen, the probability of that action being taken again rises in proportion to the realized payoff. The action has been "reinforced." Note the very low level of information or processing ability necessary to implement such an algorithm. In the context of game playing, an agent does not need to know the structure of the game to calculate best responses or even to know that a game is being played (Foster and Young 2006; Pradelski and Young 2012; Nax et al. 2016).

More sophisticated models are based on best response behavior and imitation or may involve Bayesian rationality and forward-looking behavior (Fudenberg and Levine 1998). Which of these learning models describes human behavior best under a specific condition is as yet unresolved (Cheung and Friedman 1997; Camerer and Ho 1999). While this issue has mostly been addressed within the context of a single game, some authors have recently started

to acknowledge the importance of understanding learning beyond the context of a single game (Jehiel 2005; Mengel 2012).

When trying to understand how these models extend to nonhuman actors, we would like to know how well the outcomes of these models could be described as outcomes of evolutionary processes. In the case of some of these learning models, such as reinforcement learning, it is well known that close links to the evolutionary replicator dynamics exist (Taylor and Jonker 1978; Börgers and Sarin 1997; Hopkins 2002). Hence, not only can the outcomes of these models of differing cognitive demands be well described by evolutionary models, we should also expect Nash equilibria to be played. However, as some humans arguably use more complex learning rules than nonhumans, it is less clear whether we should expect nonhuman animals to reach the same outcomes as humans (cf. Lange and Dukas 2009).

Innovation and Technology

Individual learning is required for innovation but can be costly, requiring investments of time, energy, and perhaps physical materials. A fundamental problem for the innovator is thus whether these investments can be recouped through the advantages that stem from the innovation. Consequently, there is a strong incentive for other agents to learn socially and to copy the innovation of others, rather than invest in innovations. When the benefits of an innovative behavior stem from increasing relative competitive ability, copying of such behavior can quickly erode any advantage, decreasing the incentives for innovation in the first place. While social learning can reduce the incentives for innovation and the number of potential innovators, it can also facilitate the transmission of innovative behaviors through social networks. Thus populations with a blend of individual and social learning may exhibit a high level of cultural developments, despite the negative aspects of copying (Rogers 1988; Boyd and Richerson 1995; Kendal et al. 2009; Rendell et al. 2010).

The temptation to copy is central to the economics of innovation. Societies aiming to promote innovation have therefore mitigated the costs of being copied by granting temporary monopolies in exchange for the dissemination of innovative technologies. The most prominent economic institutions are patents and copyrights; however, these institutions can have the undesired effect of granting the innovator a monopoly position that can be used to charge high prices, temporarily reducing the availability of new technologies. The trade-off between the incentives for innovation and the ensuing market distortions has been the topic of a large literature, and many different solutions are proposed and observed in practice. For example, the Creative Commons is a nonprofit organization that simplifies the process of copyright and reduces the costs for copyright owners and licensees, facilitating the sharing of useful intellectual property.

Models investigating the potential evolution of innovation using the PS game, where investors are innovators and exploiters are copying their innovations, reaffirm the essential role of patents and copyrights in sustaining innovative behavior. If we assume that innovations may be advantageous as well as disadvantageous, the presence of innovators in the population may greatly fluctuate; their innovations can lead to a substantial increase in population size, but can also result in the extinction of the population altogether (Lehmann and Feldman 2009). Computer simulations of evolution suggest that innovators are not likely to persist for more than a few generations in a population of copiers, unless there is some penalty to copying; for instance, when behavior is not replicated faithfully or some social reward (e.g., prestige or royalties) is given to innovators. Altogether these models suggest that the ability to enforce social reward through institutions may be fundamental to the high rate of innovations in humans compared to other animals.

Technology may serve different purposes for investors and exploiters. Investors may use technology to obtain resources or to protect themselves from exploiters, whereas exploiters may employ technology to track the behavior of investors. For example, some fishermen may invest in technology for locating shoals of fish while others may invest in technology to locate those fishermen's boats and exploit whatever they find. This situation may become unstable. Evolutionary simulations have shown that if the competition between strategies leads to escalation, where each strategy is investing increasingly more in technology to outwit the other, the emerging arms race proves, in the long run, to be unstable. It is likely to result in either the extinction of exploiters, if the investors manage to race ahead and open a large enough gap, or in a slow backward arms race if the exploiters race too far ahead so that investors are better off investing less in their technology (Arbilly et al. 2014).

Conclusion

The benefits of investment into the production of a “good” are undermined by the risk of others exploiting such investments. This risk creates a dilemma for potential investors in a host of real-life situations that are faced by many animal species, including humans. In this chapter, we have demonstrated that the essence of this dilemma is well captured by the so-called PS game, a simple but powerful biological framework for investigating situations of conflicts of interests between individuals. We have identified numerous biological examples, such as foraging behavior in birds and primates, as well as examples in various human societies: from small-scale communities managing resources, to firms investing in research and development, and nations competing over endangered natural resources. Crucially, in all these examples, we have highlighted how the proportion of resources, from which investors can benefit before exploiters deplete the resource (i.e., the finder's share), is key to the population level of exploitation. In both animals and humans, interested parties

have developed elaborate mechanisms that increase the value of the finder's share (e.g., patents and copyrights in human societies, or resource defense in animals).

Finding cooperative solutions to the type of dilemma captured by the PS game has attracted much attention from both the fields of biology and economics. Although the two approaches often tackle this problem from different perspectives, with specific empirical and theoretical tools, we find that they share commonalities. In both approaches, for example, individuals are modeled as agents seeking to maximize a certain quantity (e.g., biological fitness, economic payoffs, or general utility). Consequently, groups of rational agents cannot often be modeled in the same way as individuals, since the interests of agents within a group are not necessarily aligned with those of the group. In addition, we find that the PS game could be encapsulated within a more general mathematical framework, and have provided a single model to unify the well-known games from economics, specifically the public goods game, the volunteer's dilemma, and the rent-seeking model. Our general framework can be used to extend the range of dilemmas that can be studied with the PS game, for example, in situations where a certain number of investors is required to produce a good or to attract exploiters.

In an attempt to highlight current limitations of the current PS framework, we demonstrated how Elinor Ostrom's well-known classification of variables in the social ecological framework can help identify important aspects of the game theoretic framework that remain unexplored (Ostrom 2007). For example, the value of the finder's share determines the degree of excludability, and thus the typical classification of goods in economics needs to reflect this continuum and be less discreet. In general, characteristics of the resource of interest (rivalry, excludability), the type of population structure of the agents, as well as the different types of agents at play (individuals, teams, firms, or nations) are all likely to have important implications for the dynamics of exploitation within populations. Thus, incorporating these features will extend the value of the PS framework to a wider range of biological and economic scenarios.

The PS game is affected by both the resource characteristics and the nature of the agents involved; however, the implications of certain key factors of the agents remain to be resolved. Whether the agents, both human and nonhuman, are capable of learning and how they attend to uncertainty in more realistic biological scenarios is still unclear but important, because learning can potentially alter the predicted equilibrium between investors and exploiters. The role of learning is particularly important in humans because of our unusual ability for cumulative cultural development, innovation, and sophisticated technology, which allows us to find, develop, and keep novel solutions to the problem of investment versus exploitation. These features, however, also allow potential exploiters to increase their ability to benefit from the investments of others.