Given Certain Conditions, Even Selfish Prisoners Can Learn to Cooperate

Michael Cho

Seoul International School

INTRODUCTION

"Under what conditions will cooperation emerge in a world of egoists without a central authority?" - Robert Axelrod

In a predominantly selfish society where all are "free to cooperate or not," the absence of a central authority seems intimidating because everyone works for themselves, and thus, no one works for the whole. In fact, in such an absence of central authority, it may even be said that we live in a state of general anarchy in consideration of the more excellent picture. On pretty much any scale imaginable — whether it be within a community, state, or on a global scale — all of us are but ""prisoners"," profoundly affected in every moment by the choices of others; indeed, this is the context behind the game theory as a field. The idea that humans are purely motivated by greed, ego, and profit was initially a sound explanation on how competition and productivity improved under capitalism, and the magical invisible hand promotes the best for all. However, it is also an obsession with self-interest that causes capitalism's most fatal flaws, namely growth in economic inequality and moral hazard (Piketty 304).

A solution to growing inequality, climate change, and a remarkable lack of "morality," all rapidly increasing under capitalism, requires an understanding of how cooperation can be elicited in a capitalistic society without a central authority. The author presents Game Theory as a viable tool of study to reach an understanding of the merits of cooperation among selfish competitors, as well as predict the outcomes of various situations of conflict in political, social, and global backgrounds. In particular, Robert Axelrod's work provides insight into how cooperative strategies can rise out of selfishness in his iterated Prisoner's dilemma computer tournament. The present paper aims to extend his work by using it as the foundation for a concept within the economic sphere of competition within capitalism, which the author will term "cooperative capitalism." This term revolves around the concept of creating empowerment amongst people to do things for others but due to personal incentives. In addition to this, cooperative capitalism brings a motive for compassion and charity by focusing on voluntary acts. Furthermore, this paper will come up with specific conditions under which participants in greed-driven capitalism can cooperate than defect to yield the most productive outcome for themselves and the rest of society. Ultimately, the author presents cooperative capitalism as an ideal within modern anarchist society, as well as a solution to many current capitalist flaws.

BASIC CONCEPTS OF GAME THEORY.

To briefly explain, Game Theory is the study of strategic interactions between decision-makers, or agents, in particular situations, or "games,", particularly through the mathematical modeling of such strategies and games. The 'Prisoner's Dilemma, the focus of 'Axelrod's study, is a very popularly researched game between two agents, each with two identical choices: to cooperate or to defect.

The Game Theory isn't a conclusive study developing towards winning a scenario, but in explanatory terms, the Game Theory is a study of how strategies are determined, situationally acted out by the players. This results in surprising yields around the insights of human nature through the selection of options based on their incentives and assessed outcomes. The results provided, therefore, create the availability of assessable information regarding each player or agent and the factors present, which lead to cooperating or defecting as a result.

Since there are two options for both agents, there are four possible outcomes with payoffs for each agent modeled in numerical points. In the case where both agents choose to cooperate, both receive 3 points, the **reward for cooperation**. If one agent cooperates and the other defects, the defector gets 5 points, the **temptation to defect**, while the other gets 0 points, the **sucker's payoff**. If both defects, they receive 1 point each, the **punishment for mutual defection.** The hallmark of the Prisoner's Dilemma, and the reason for its name, is that no matter what the other chooses, defection yields a higher payoff than cooperation; if the other agent cooperates, you will receive 5 points for defecting and 3 points for cooperating.

On the other hand, if the other defects, you will receive 1 point for defecting and 0 points for cooperating. In both cases, defecting yields more a more substantial payoff. However, if both agents defect as a result of the above logic, both will

be worse off than if they had both cooperated. Furthermore, the total number of points earned is the highest when both players cooperate; thus, it may even be said that cooperation is the best choice for the benefit of the "whole". This failure of individual, seemingly rational decisions to maximize payoff (both for the individual and the whole) marks the reason for the "dilemma" in the game's name.

The Iterated Prisoner's Dilemma, or IPD, is a variation on the Prisoner's Dilemma. While PD only described a singled-out interaction or situation, the Iterated Prisoner's Dilemma describes a series of multiple interactions between agents under the assumption that agents retain the memory of previous interactions. Such a distinction is important because of the ability of previous outcomes to influence future decisions, creating another dimension to strategy. If each agent does not need to consider future interactions, then there is no incentive to cooperate; thus, it can be safely anticipated that both will defect. This line of reasoning not only holds for one-time interactions but also in where there is a known, finite number of iterations of the game. This is because the last move, both agents knowing that there will be no next interaction to consider, can be predicted to have both defecting (as described above). Both agents won't have the incentive to cooperate on the second to last move either since it is known that the outcome of the last move may be safely anticipated to be mutual defection. As a result of this logic, if there is a definite and finite number of interactions between the agents, the game will ultimately devolve to both agents always defecting. Therefore, it is necessary to introduce an element of uncertainty to the game through the number of iterations being indefinite, unknown to the agents. This causes both to be unsure of when the last interaction will take place, and thus cannot reason accordingly (and thus the game does not merely become pure defection). The Iterated Prisoner's Dilemma has become commonplace in terms of modeling corporate strategies and game theories, being used more by investment strategists as globalization occurs.

Finally, we also take into account the decreasing importance of payoffs as we look further and further into the future. Players tend to value the future less than the present, and in reality, there is always the chance that the players will not have any future interactions; keeping these factors in mind, we may find that the payoff of each iteration decreases in relative value further in the future. Thus, if the chance of the players interacting again is low enough, we once again reach a situation in which the players have nothing to lose from defecting. However, assuming that the chance of more iterations is sufficiently high, we may find that there is no best strategy for any player, independent of that used by the other player. This is because, though it is easy to engineer a "best strategy" in the knowledge of the other 'player's strategy, this counter-strategy is entirely variable, thus removing the possibility for anyone's strategy to be the best against every possible opposing strategy.

'AXELROD'S TOURNAMENT

In his study of the IPD, Axelrod uses a tournament as an experimental case. He hosts a computer tournament in which every entrant to the tournament submits a program to choose whether to cooperate or defect on each move. Each program has, available to it, the history of its interactions with the other player so far, which it may use to aid its decision-making. It was structured as a round-robin, meaning that each program would be paired once with every other program, as well as with itself and a RANDOM program that randomly chooses to cooperate or defect on every move. Each game consisted of 200 moves, and the payoffs used were the same as the ones described above. Though the fact that each program is pained once with itself could theoretically allow for programs that do well against itself to have an advantage in the tournament, we see that there are enough submissions to the tournament for this effect to be negligible in comparison to how the strategy fares with the other submissions. Any concerns about pairing each program with RANDOM are also similarly negligible; instead, it could even be said that this allows for a more thorough investigation, since often in social situations, we may find agents acting without consideration (RANDOM) or agents with similar thinking as ours (mirror). If there are any methodological flaws, it would have to be how each game consisted of precisely 200 moves; this contradicts our earlier theoretical finding that, if a game is finite, each player only has the incentive to defect. In theory, it would have been possible for submission to meta-analyze the number of passing moves, thus deviating from the intention of the tournament and "tainting" the results. However, because the results centered around strategies that best-utilized cooperation (rather than defection) and the entrants were experts in related fields that crafted strategies unrelated to such meta-analysis, this does not seem to have seriously affected the final results. Furthermore, for the Prisoner's Dilemma, deducing the best opted strategy isn't possible as the strategy is entirely based upon perspective and what strategy and results it is matched towards. This is because succeeding in the Prisoner's Dilemma can be as simple as strategy selection based off what suits as a count strategy to the opponent's strategy.

Moving on to the final results of 'Axelrod's investigation, we

find that a strategy named TIT FOR TAT fared the best in every round of the tournament. TIT FOR TAT was a straightforward rule; it cooperated on the first move, then in every consecutive move would make the same choice that its opponent made on the previous move. The TIT FOR TAT a less hostile strategy and produces a defect resulting from an opponent defecting more than once in correspondence. Furthermore, it was always cooperative and benefited both the cooperating agents by rendering the full benefits towards them. The TIT FOR TAT resulted in a defect if it was matched with a defecting opponent. The drawbacks of the TIT FOR TAT were based off of individual score maximizing assumptions and when this was assumption was implemented with a strategy backed up by mindless and haphazard selective approaches, like RANDOM, the opponent's level would also host the TIT FOR TAT strategy that rendered ineffective if not less effective. Similarly, if two of the TITS FOR TATS strategies were implemented, they would result in the punishment of every defection with two of its own. Another rule that was very successful in the first round, earning second, was TIDEMAN AND CHIERUZZI.

Observing these rules, we see that one property of rules that seems to differentiate between well-faring and less-successful rules is the property of being *nice*. We define being nice as never defecting first, and observe that each of the eight top entries are nice, while the rest aren't. Nice entries overall got average points between 472 and 504 per round, while non-nice entries got max 401. In part, this is because each nice entry got 600 points from continual cooperation with both its twin and the other nice rules. Even accounting for minor variations in end-game strategies (some of the "nice" strategies actually may defect after move 198), they tended to cooperate for the vast majority of the game, thus not producing many variations in the actual points received.

Furthermore, we also see that *forgiveness*, or the continued willingness for a rule to cooperate even after instances of defection from the opponent is an important property for a well-performing rule. We see that similarly to how the property of being nice largely differentiates the top rules, the property of forgiveness determines the rankings within the nice rules as well. As mentioned above, the top two rules are both forgiving, while an extreme example of unforgiveness marks the lowest ranking of the nice rules. FRIEDMAN, which scored the worst out of all of the nice rules, used a strategy of permanent retaliation, which in essence represents the property of pure un-forgiveness. Overall we find that this lack of forgiveness ties back to the success of nice (and forgiving) rules since rules that defect unprovoked will elicit varying levels of punishment from other rules, more so because most of the submissions were not very forgiving (note that only two out of the nice rules were forgiving to DOWNING). Even when a rule that will defect unprovoked is matched with a nice, forgiving rule like TIT FOR TAT or TIDEMAN AND CHIERUZZI, we see that retaliation between the rules makes such defection not worth it in the long-term. Looking specifically at the interaction between JOSS, an imitation of TIT FOR TAT that has a chance to defect even when the opponent cooperated the last turn, and TIT FOR TAT, we see that once JOSS decides to defect, there is a chain of retaliation between the two rules. In other words, once JOSS defects, TIT FOR TAT defects the next turn in retaliation, prompting JOSS to defect the turn after that (it is an imitation of TIT FOR TAT after all), which causes TIT FOR TAT to defect after that-- and so on. Eventually, JOSS defects in the middle of the chain again, causing an alternating double-chain that, although with only one "echo" each rule would alternate cooperation and defection, with two "echoes" each rule would defect every turn. Thus, the point output after this second "echo" is minimal for both rules, making such unprompted defection not worth it for JOSS in the long run. Thus we observe that in the first tournament environment, although it may appear to be beneficial for a rule to defects occasionally in the short-term, underlying "echoes" emerging as a result of such behavior cause the deterioration of relationships with other rules in the long run. Thus, not being nice as a strategy appeared not to punish their opponents, but punish themselves in every relationship as a result, while nice rules did not have that tendency.

Thus, we find that these qualities of being "nice" and "forgiving", clearly essential qualities for cooperation, are often necessary for agents to maximize their payoff; not only individually, but also for the whole (as established previously, cooperation benefits the whole group/society). However, in the current economic system, these qualities are often missing or face a lack of incentive. This is because, over time, the current economic system of capitalism has developed, through properties that have observed to have been superior (though these properties are, through comparison to the 'author's findings, are far less than ideal), to a form of capitalism that will henceforth be referred to as "competitive capitalism", contrary to the 'author's proposed "cooperative capitalism". As the name suggests, competitive capitalism is, while still being centered around free-acting, fundamentally self-interested agents, most encouraging of competitive behavior over the cooperative. However, competitive capitalism is at a notable lack of the properties above. In particular, one of competitive 'capitalism's most notable characteristics is its tendency for inequality, monopoly, and monopsony. For example, take a monopoly; in an established monopoly,

the one at the top is in a constant state of defection with its competitors. Regardless of their payoff in such interactions, as long as their competitors are also prevented from obtaining high payoffs, the monopoly is sustained; thus, competitive capitalism **discourages** "nice" behavior, for both sides of the monopolize-competitor interaction. A similar phenomenon is observed for monopsonies and situations of inequality; competitive capitalism tends to encourage the ones at the top to maintain the relative status quo (in other words, keep the situation in which they have more "points" than their opponents), which is done through constant defection. This is not specific to the companies or individuals at the very top; rather, competitive capitalism encourages such behavior at every level of inequality. As such, not only are the qualities of "niceness" and "forgiveness" completely ignored, the concept of cooperation itself seems to be undermined by the current economic system. This concept,' 's reinvigoration, is one of the leading aspects towards creating a cooperative behavior at each level of inequality.

THE POSITION OF THE CURRENT SYSTEM

The current system is, as a result of this lack of cooperation, detrimental in many ways, both individual and for society. For society, there is undoubtedly much that is lost; many social benefits are ignored (harming, for example, living standards that benefit the majority, if not all), there is much innovation and advancement lost that 'would've taken place with cooperation, and especially looking long-term towards issues like climate change, a lack of constant cooperation is worrying. Not to mention the diminished collective happiness due to inequality, as well as the diminishing marginal utility of wealth as it is funneled towards the upper echelons of the economy. Due to the social benefits not being considered through decision making (defect or cooperate), inequality and the overall wealth utility having a weaker effect whilst being funneled towards the upper echelons of the economy, the system leads towards a short-term mindset which aims towards grasping on to any nature of the benefit, ignoring the morality of defection and the consequences that could be resultant of a highly defective society as compared to a highly cooperative society.

The author tentatively proposes, as an outline of a potential system for cooperative capitalism, more in agreement with the most successful qualities found in 'Axelrod's experiment, a couple of changes. Firstly, there is much to be gained through the establishment of long-term cooperative relationships between corporations, encouraging an economy built upon relationships rather than pure self-interest. One way of accomplishing this/providing incentive for such a shift in the economic system is for companies to own share in other companies and their particular shares in each 'other's corporations. This would result in a self-interested incentive to have lasting cooperative relations with other companies. Furthermore, such long-term relationships would be encouraged on the individual level as well; if workers, even in the lowest classes of work, were to establish long-term relationships with their employers — for instance, through offering social and financial security to a greater and longer extent in exchange for greater loyalty and work, similarly to the system found in Japan — the economic system would shift to be more cooperation-focused. The resulting economic system would be the sketch outlines of cooperative capitalism.

N-PERSON IPD

The N-person prisoner's Dilemma (NPD) is the variant of the Prisoner's Dilemma with the significant difference of hosting more than two agents or players. It became popular amongst social theorists and economists after emergence. The Prisoner's Dilemma involved with multi-player scenarios have a common strategically implemented structure between them, but for the NPD, there is more to be gained by defecting that cooperating, but this effect is negated by the ability to gain more through cooperation rather than mutual defection.

As an example of the use of the N-person IPD is shown through the market for labor where the prices of goods or services rendered feature increase leaving every individual in a worse position than if everyone displayed and exercised restraint. This scenario was the result of trade unions standing up for wages that catered to their self-interest rather than on the individual level of negotiation and self-interest for inflation rate toppling wage rates. This problem could only be solved by encouraging collective rationality in bargaining. Another type of NPD that can be easily assessed within the international marketplace is the NPD that stimulates scenarios where there is a shortage of resources already happening or predicted. In exemplary terms, if any resource is short, there is a call for conversation, but if only everyone restraints from unnecessary resource usage, then individuals' benefit. On the other hand, if only one individual restrains, then the conservation is futile. Therefore, the self-interest of every individual would be NOT to conserve, but if everyone acts with an individualistically strategic mindset, all are worse off.

Furthermore, the unstable equilibrium of the system displayed that the larger number of cooperators were rewarded slightly, and the small number of defectors are awarded greatly. The

stable equilibrium displayed a low number of cooperators with a high number of defectors, where the unstable equilibrium displayed the opposite. This aspect of the NPD referred to the potential issues within the payoff structure.

The conclusions of the NPD differed from Axelrod significantly due to the NPD› ‹s possibility of a conclusion more based around defection. In addition to this, Axelrod›s recognition of communication between the agents proves to be the key determining factor. This may be more accurate as opposed to Axelrod›s deductions. However, the author maintains his stance towards supportive a ‹cooperative capitalistic› ideology. Lessons from ‹Axelrod›s Tournament:

However, cooperation cannot work in every environment; in fact, the competitive capitalism system is seen in 'today's economies is proof of this. In that case, the question arises: what conditions are necessary for cooperation to be possible? The author observes both 'Axelrod's tournament and literature on the N-player IPD to identify a couple of conditions of cooperation.

From 'Axelrod's theory, we first note that cooperation is mostly profitable only when the number of interactions is infinite or practically predicted to be sufficiently large to discourage the short-term benefit of defection. In a similar vein, the probability of a repeated interaction must also be sufficiently high for cooperation to work. The author notes that this further supports the need to establish long-term cooperative relationships between corporations and individuals for a working system of cooperative capitalism.

Furthermore, we find that, especially in the N-player IPD, the initial preferences of the agents have a significant bearing on the practicality of cooperation in the long-term. When considering a community of Pavlovian agents (agents that change their probability for each action by an amount proportional to its reward or penalty) that each interacts with every other agent, a majority-cooperating environment may only be reached in the long-term if the ratio of initially cooperating agents to initially defecting agents is sufficiently high. Otherwise, the community settles at a largely defection-based environment, which we may liken to competitive capitalism. Then, we find that a certain level of initial encouragement of individuals and corporations to cooperate is necessary to sustain a system of cooperative capitalism in the long-term. Thus, in an environment of agents that base their behavior solely on previous personal benefit (in other words, a simplified model of the behavior of self-interested agents), we find that the initial preferences largely determine the viability of cooperation; a sufficiently high number of cooperators is necessary for cooperation.

However, the willingness of agents to change their behavior also has a large bearing on the viability of cooperation. If there is a sufficiently large number of agents that prefer defection and are not willing to change behavior to cooperation, cooperation can't be beneficial even to agents that initially preferred it. For a very simplistic example, if TIT FOR TAT interacts with a player that only defects, cooperation is impossible to result. In other words, there is no reason to cooperate with an entity that continuously defects; this is the equivalent of a particularly uncooperative country receiving, in return, no cooperation from its neighbors as a result of its constant defection. The willingness of agents to change their behavior also revolves around the potential of long-term benefits significantly outweighing short-term advantages, which will be easy to assess. This is a barrier to cooperation, which can only be overcome through constructive efforts with each significant aspect, individually, translating towards a stronger collective effect.

Finally, the viability of cooperation also depends on how the agents decide to update their action preferences. If the agents base their preferences purely on their payoffs, as described above, then cooperation may occur in specific situations; however, if the agents base their behavior based on the success of other agents, likely, cooperation cannot exist. This is because, as long as there is one initial defector, the defector will earn more points than their neighboring cooperators; thus, other agents will copy the 'defector's behavior, eventually leading to complete defection within the community of interacting/visible agents.

Often in the real world, however, economic and political agents are too short-sighted to maximize their *long-term payoff* in egoist interest. On the broadest scale, we see that the condition of infinite interactions for cooperation to be possible is met; for the foreseeable future, most of humanity will remain on Earth, and within its confines, we will continue to interact. However, then, consider an issue like climate change that affects the Earth itself, the medium on which our interactions happen. Modeling the choice of going "green" versus not going "green" in favor of the economic benefit of utilizing CO2 emissions like the 'Prisoner's Dilemma options of cooperation and defection, the above study shows that follow climatic agreements and sacrificing the immediate payoff of defection is the superior egoist option.

However, governments and corporations are often seen to defect anyway, both out of greed for the immediate economic payoff and out of fear that other entities will betray them

(notice the clear similarities to the 'Prisoner's Dilemma). Similarly, to how competitive capitalism, as described above, creates a diminishing marginal utility of wealth as inequality proliferates, so does the economic payoff of not going green at the expense of the Earth. Over time, a repeated betrayal of environmental treaties produces diminishing returns by increasing costs (for example, things like water or fertile land may increase in price as supply decreases). This could even apply to personal expenses like labor if overall health decreases due to phenomena like extreme temperatures. Combined with the likelihood that the international community of corporations/governments will eventually reach complete defection due to the unwillingness of some to change preference to cooperation, such consistent defection will yield lower payoffs over time, all the way down to zero. Thus, in the long-term, it is in the best egoist interest to cooperate rather than a defect. This applies to almost all current issues of cooperation, on the international and national levels; the general atmosphere of competitive capitalism, then, is also heavily supported by the short-sightedness of realistic agents in calculating the optimal strategy for the maximal payoff.

The caveats present through a cooperative capitalist system revolve around the implementation of long-term relationship importance, which would drive individuals and corporations to think past short-term decisions, which are the reason for a large number of defects throughout the assessments. This is one area that would result in higher levels of cooperation and a lower chance of defects resulting from decisions derived by greed and short-term insights. The high reliance on long-term relationship insight management points towards how necessary and effective this aspect can be for the sustainability of such a system.

CONCLUSION

In this study, the author used Game Theory as a tool of study to real-life model interactions between self-serving entities in economics and politics. It is found that under certain conditions, individual decision-makers can act on pure self-interest to cooperate. It is found that cooperation, in such conditions, is the better choice for the agents to maximize long-term payoff. Analysis of 'Axelrod's IPD tournament and N-player IPD environments further supports the benefits of cooperating for individuals. As a result, the author contends that states and corporations in the real world should follow these conclusions rationally.

However, the current economic system does not support such cooperation in the real world; as a result, the system of cooperative capitalism is proposed as a solution. Individuals unaware of this study or seeking conditional benefits still stray away from the concepts and results deduced from the discussed models. Cooperative capitalism will fundamentally be the same as current capitalism, in that the means of production will still be owned by independent corporations and not the state; however, it will establish long-term cooperative relationships between entities (corporations and/or individuals) through systems like companies owning a share in each 'other's companies. This is an improvement over the current system of competitive capitalism, based on the lessons learned from 'Axelrod's tournament and related 'literature's findings that long-term, linked interactions (and each 'entity's recognition of this longevity) are necessary for cooperation and bringing down the level of defects which are more likely to be caused through short-term insight Cooperative Capitalism aims towards bringing better results in terms of cooperation through imposing an environment that is more inviting towards long-term relationships as well as simultaneously constructing an environment which makes defecting difficult to cater.

It can, however, be argued that the outlined flaws of modern capitalism above are not actual flaws; once the payoff of defecting becomes low enough, corporations and governments could cooperate of their own accord. Thus, it could be said that the proposed system of "cooperative capitalism" is unnecessary and perhaps even in danger of introducing too much state control over the economy. However, the author still contends that the initial state of and continued cooperation, as seen in 'Axelrod's tournament and N-player IPDs, yields superior payoffs for agents, especially considering short-term benefits are weighed higher than those in the future. Thus, cooperative capitalism still bears a significant improvement over the current system for the good of the "whole".