Econ 354

Greg Dow

February 10, 2021

Notes on Elinor Ostrom, "Governing the Commons"

Chapters 1-2 (first set of lecture notes)

Elinor Ostrom started off as a political scientist, but won the Nobel Prize for Economics in 2009 for her research on institutional economics. This is regarded as a classic book.

I will do several things in this introduction. First, I want to define the idea of a common pool resource and give a few examples. Second, I will discuss why CPRs tend to cause economic problems. Finally, I want to describe the research strategy Ostrom uses.

<u>Common pool resources</u> (CPRs).

I will summarize Ostrom's definition of a CPR as follows (see p. 30 for a longer version):

A common pool resource is a resource system sufficiently large to make it costly (but not impossible) to exclude potential users.

The key idea is that there is a valuable resource, which may be natural or human-made, and it is difficult but not impossible to prevent people from using it.

Economists often use the concepts of private goods and public goods. A *private good* is something that can easily be divided up among individuals, and different people can use different amounts. It is also possible to charge people a price for it. A good example is toothpaste. Stores charge a price per tube and you can buy whatever quantity you want.

A *public good* is something that everyone in some area consumes the same amount of, and people cannot easily be excluded from using it. A standard example is defense. If you live in Canada, you are defended by the Canadian armed forces. Everyone gets the same level of protection, and individual people do not pay a price for this service (instead we pay for national defense through taxes).

A common pool resource is not exactly a private good, but also not exactly a public good. Often it may be a natural resource but sometimes we have similar problems with artificial resources. Here are some examples:

Natural resources: grazing areas for animals, fishing areas, groundwater basins (aquifers), bodies of surface water (lakes, rivers).

Artificial resources: irrigation systems, transportation systems (bridges, roads), electronic systems (computer networks, telecommunications networks).

In each of these cases, individuals get some personal benefit from using the resource: fish harvests, water used, roads driven on, etc. But this typically reduces the benefit or raises the cost for other users. For example, the more fish I catch, the fewer there are for you. The more I drive around on the road, the more traffic congestion this creates for others.

It is often useful (especially for renewable natural resources) to think in terms of stocks and flows. Picture the water in a reservoir. People take flows of water out for drinking and other uses. There may also be some evaporation. New water flows in from sources like rainfall or melting snow. If the replacement rate is greater than the withdrawal rate, the amount of water in the reservoir increases. In the opposite case the amount of water in the reservoir decreases, and eventually it may disappear completely.

More formally: economists talk about a *stock* when they want to describe the quantity of some resource at a point in time (like the amount of water in the reservoir today). They talk about a *flow* when they want to describe the rate at which something is going in or out (like the rate at which people take water out for drinking or rate at which the rainfall replenishes the reservoir). The distinction is that a stock exists at a point in time while a flow is always a rate per unit of time (per day, per year, or whatever).

In nature, lakes and rivers tend to reach a stable balance between the flows in and the flows out, so the level of the resource remains relatively constant over time (subject to some random fluctuations). But with human use, excessive withdrawal can deplete the resource, reduce its economic value, and possibly even destroy it.

Johnson and Earle gave numerous examples of resource depletion including destruction of wildlife habitat (like forests), soil erosion or loss of nutrients, overgrazing of pastures, and so on. People often developed complicated systems for using and allocating land in order to avoid making these problems worse.

For artificial CPRs, it is often useful to think of the issue as involving congestion: if too many people try to use the resource, this reduces the value of the resource to others. For example, if everyone tries to cross a bridge simultaneously, it takes a long time for all of them to get across. Or if everyone tries to access the same web site simultaneously, they could overwhelm the server and crash the site. In both cases, it could be technologically possible to charge a price to each user, and this may reduce demand on the resource, but charging a price may involve significant administrative costs.

<u>The tragedy of the commons</u>.

There is a famous short article by Garret Hardin from 1968 called "The Tragedy of the Commons". Hardin tells the following story.

Suppose there is a pasture and everyone can use it. The local people are shepherds and they want their sheep to graze in the pasture. There is an upper limit on the number of sheep that can graze in the pasture without eventually destroying it. Each person gets a direct benefit from grazing their own sheep and ignores the cost this imposes on others. People think: "If no one else will use the pasture, I should definitely use it. On the other hand, if everyone else is going to use it, eventually the pasture will be destroyed, but I may as well use it in the meantime." The problem is that no individual has a noticeable effect on the future quality of the pasture. The result is that everyone uses the resource, and eventually it is destroyed.

Several points about this:

First, this story is a metaphor. Hardin was not actually worried about sheep. Instead, he was thinking about human population growth. His real concern was that each individual or family might believe there is a benefit from having a large number of children. Each person's decision about having kids is independent of everybody else's decision, and no individual has a noticeable effect on global population. However, there are finite natural resources on the planet (water, air, land), and if everyone decides to have a large family, we will wreck the planet. I don't want to go into the details of this argument now. We will come back to the subject of world population growth at the end of the course. But the phrase "tragedy of the commons" is used frequently and you should be aware of the original context.

Second, Johnson and Earle argue that population growth had crucial implications for the evolution of human societies. So the idea that population pressure could arise and might have important consequences should not be totally new at this point in the course.

Third, Ostrom calls her book *Governing the Commons*, not *The Tragedy of the Commons*. She does not believe that common pool resources necessarily lead to a tragedy. What she does believe is that it matters whether people develop good institutions to govern a CPR, or whether they do not. Her book is about what would be required to govern a commons effectively (what good institutions would look like) and identifying the conditions under which people are likely to succeed or fail in developing good institutions.

<u>Ostrom's research strategy</u>.

Ostrom wants to find out how people in local communities manage CPRs when they are allowed to design their own institutions. She also wants to know when these institutions succeed or fail. Her criteria for success involve equity (are the institutions fair in some sense?) and efficiency (this is an economic concept I will define later).

To pursue this agenda, Ostrom needed a broad comparative database. She wanted to look at many individual cases and search for patterns. Before she wrote the book, there were case studies by people in a variety of different disciplines (economics, political science, sociology, history, and so on), with very little attempt to synthesize the lessons from all of them.

Ostrom and her colleagues looked for cases having information about

1. The structure of the CPR (such as the physical environment and technology).
2. The behavior and characteristics of the people using the resource.
3. The institutional rules governing the individual users of the CPR.
4. The outcomes resulting from the institutional rules and the behavior of users.

Other criteria for a good case included the following. Ostrom wanted cases where the community of users was relatively small in scale and located in just one country. She also wanted to focus upon renewable resources (often involving stock and flow issues). There needed to be substantial scarcity, so it was important to people how the resource was allocated. There also needed to be a potential for substantial harm if the resource was governed poorly.

Ostrom's general method was to go back and forth between the empirical evidence from her case studies and relevant types of theoretical analysis (such as game theory). This is useful because theoretical ideas suggest what observations might be relevant, while the observations shed light on the limitations of the theory and suggest ways to improve it.

Ostrom assumes that the distribution of human characteristics is roughly the same across societies or communities. Therefore she does not want to explain cases where people had bad outcomes by saying that this resulted from stupidity or selfishness. These things can be found in every society.

If you believe that the percentage of people who are stupid or selfish is roughly the same in Canada, the U.S., Turkey, Sri Lanka, and so on, then you should not use such factors to explain why the outcomes were different in different places. Instead you need to look for explanations such as differences in the structure of the resource, the constraints facing the designers of the institutions, and so on.

An important general point in social science is that you should not reduce issues about a social system to issues about individual psychology. It is ok to talk about individuals in your model. Economists often do this, and they typically make some assumptions about the preferences of these individuals. But the institutional rules of the game are crucial in determining how people will interact with each other and whether the outcomes are good or bad. It is not just a matter of psychology, although sometimes this is important too.

Economists tend to dislike explanations based on someone's preferences or beliefs. For example, people often say that the price of a good increased because the owners of a firm were greedy. There are two problems with this. First, it is not a good strategy to say that something changed (the price) due to something else that remained constant (the level of greediness of the owners). Changes need to be explained by things that changed, not by things that were constant. You might try to respond by saying that the owners of the firm became greedier than they were before. But now we have a second problem: how would we determine whether the owners were greedier today than they were yesterday? This is generally not a testable hypothesis. An economist would typically look for a change in cost or demand conditions that might help explain the price change, while assuming that the psychology of the firm owners remained constant (e.g., they were maximizing profit yesterday and they are still maximizing profit today).

A good institutional design takes into account that some people are greedy, uninformed, or whatever. Real world institutions cannot be based on the assumption that everyone is a saint or a genius. Even if many people are cooperative, honest, and smart, you have to anticipate that some people will be selfish, some will lie, some will become confused by complex rules, and so on. In short, people are heterogeneous (we will come back to this point when we discuss experimental economics in the Putterman book). Institutions have to be designed so that they will operate successfully even though individuals have a wide variety of personal characteristics.

Econ 354

Greg Dow

February 18, 2021

Notes on Elinor Ostrom, "Governing the Commons"

Chapters 1-2 (second set of lecture notes: game theory)

In this part of my lecture notes, I want to develop some ideas about game theory based on an example Ostrom uses, starting around p. 4 of the book. My way of drawing the game will be different from hers but the underlying numerical example is the same.

Consider the following situation. There are two herders named 1 and 2. They use the same pasture to graze their sheep. Each person gets $1 in profit for each sheep that is well fed.

The payoff matrix from the game is as follows:

		Player 2	
		cooperate (10 sheep)	defect (20 sheep)
Player 1	cooperate (10 sheep)	10, 10	-1, 11
	defect (20 sheep)	11, -1	0, 0

Each player has two possible strategies: cooperate or defect. This implies four possible strategy combinations: (cooperate, cooperate); (cooperate, defect); (defect, cooperate); and (defect, defect). In each case, the first strategy listed in a pair refers to player 1 and the second strategy refers to player 2. The numbers shown in the payoff matrix refer to the profit of each player. So for example if player 1 cooperates while player 2 defects, then player 1 gets -1 and player 2 gets +11.

Now let's interpret the strategies. I will construct a story that is consistent with Ostrom's numerical example. Assume the capacity of the pasture is 20 sheep. If there are no more total sheep than this, they are all well fed. Therefore, if each person puts 10 sheep on the pasture, then each person gets a profit of 10.

If one person puts 10 sheep on the pasture and the other puts 20 sheep there, this leads to a problem. About half of the big flock (in fact, exactly 11 of them) are well fed, and they chase away the smaller flock. Not only are the sheep in the smaller flock poorly fed, but it costs their owner $1 to go find them. So the person who chose 20 sheep gets a payoff of 11, while the person who chose 10 sheep gets -1.

Finally, if both people put 20 sheep on the pasture simultaneously, for a total of 40, then no sheep are well fed and each person gets a payoff of zero.

Both people know what the strategies and the payoffs are (imagine that each person has a copy of the payoff matrix). The two players simultaneously decide what to do (imagine that they each think about the game overnight and then decide in the morning how many sheep they will bring to the pasture). For the moment, assume they do not communicate in advance about what they will do (I'll come back to this issue later).

What predictions can we make about the outcome? First, look at the game from the point of view of player 1.

Suppose player 1 expects that player 2 will bring 10 sheep (cooperate). In this situation player 1 says "if I choose 10 sheep, I can get a payoff of 10. But if I choose 20 sheep, I can get a payoff of 11, which is larger. Therefore I should bring 20 sheep (defect)."

Suppose instead player 1 expects that player 2 will bring 20 sheep (defect). Then player 1 says "if I choose 10 sheep, I will get a payoff of -1. But I choose 20 sheep, I can get a payoff of 0. Since zero is better, I should bring 20 sheep (defect)."

Here is the point: no matter what player 1 expects that player 2 will do, player 1 is better off defecting. So it is a safe prediction that player 1 will defect.

You should be able to see that if you look at the game from the point of view of player 2, you will reach the same conclusion (note: the game is symmetric). Regardless of whether player 2 expects player 1 to cooperate or defect, player 2 is better off defecting.

Based on this logic, the overall outcome is (defect, defect), and the payoffs are (0, 0). Of course, if both people had cooperated, they would have had (10, 10) and both would have been better off.

This is an example of a type of game called the *Prisoner's Dilemma*. Such games always have the same structure: if each player is individually rational and uses the strategy that seems best to them, the resulting social outcome will be bad for everyone.

I won't take the time here to explain the reasons for the name, but if you google prisoner's dilemma (or prisoners' dilemma; some people put the apostrophe before the s while some do it the other way around), you can find out the history behind this.

For a prisoner's dilemma game, it is really only the ranking of the outcomes that matters, not the exact numbers. To see this, consider the payoff matrix

		Player 2	
		cooperate	defect
Player 1	cooperate	3, 3	1, 4
	defect	4, 1	2, 2

The numbers here do not refer to dollar amounts, but just the ranking of the outcomes from best (4) to worst (1). The best outcome is when you defect and the other player cooperates; the next best is when you both cooperate, the third best is when you both defect, and the worst outcome is when you cooperate while the other player defects.

Based on this ranking, you should be able to see that player 1 is better off defecting regardless of what player 2 does, and vice versa. So both defect and they get (2, 2) although it would have been possible for both players to do better with (3, 3).

Now let's define several general concepts.

Dominant strategy. An individual player has a dominant strategy if one of their strategies is best no matter what the other player does (the optimal choice does not change when the other player changes strategies).

Dominant strategy equilibrium. A game has a dominant strategy equilibrium (DSE) if each player has a dominant strategy. The DSE is the strategy combination where each player uses their dominant strategy. For example, (defect, defect) is a DSE above.

Pareto improvement. We say that a Pareto improvement occurs if we move from one situation to another (for example, one strategy combination to another) and everyone becomes better off simultaneously.

For example, in the above game, going from (defect, defect) to (cooperate, cooperate) is a Pareto improvement because we go from (2, 2) to (3, 3).

Pareto efficiency. We say that a situation is Pareto efficient if when we start from that situation, it is impossible to achieve a Pareto improvement.

For example, suppose we start from (3, 3). Then it is impossible to make both people better off simultaneously, because if one person gets 4 the other must get 1 and becomes worse off. The same is true if we start from (4, 1); if we try to make the person with 1 better off, we have to give the other person less than 4 and they become worse off. The same is true for (1, 4). So we conclude that each of the outcomes (4, 1), (3, 3), and (1, 4) is Pareto efficient, but the outcome (2, 2) is not, because a Pareto improvement is clearly possible in that case.

Here is my general definition of a Prisoner's Dilemma game:

A *Prisoner's Dilemma* is a game that has a dominant strategy equilibrium and the DSE is not Pareto efficient.

This definition could be applied to a game with any number of players and any number of strategies for each player. Here are two points to keep in mind though.

First, some games do not have a DSE because at least one of the players does not have a dominant strategy. For example, if everyone else is driving on the right side of the road, you are better off driving on the right side; but if everyone else drives on the left side of the road, you are better off driving on the left side too. In this situation you do not have a dominant strategy, because your optimal behavior depends on what you expect everyone else to do. This is an example of a 'coordination' game. It is not a prisoner's dilemma.

Second, it can be true that a game does have a DSE but the DSE is Pareto efficient. For example, suppose each person decides whether or not to attend a party. It is always fun to be at the party no matter how many other people are there, so attending is a dominant strategy. But if this is true for everyone, then everyone will go. This outcome is Pareto efficient, so again we do not have a prisoner's dilemma. Games that have a DSE and the DSE is Pareto efficient are sometimes called 'invisible hand' games.

Now let's go back to the prisoner's dilemma and the example of the sheep.

You might have been wondering whether communication would help. We are assuming each player has a copy of the payoff matrix, and each has time to think about the problem overnight, so they should be able to understand that there is some danger of getting a bad outcome. Why not pick up the phone and agree that it is better to cooperate, so each will only bring 10 sheep and the outcome will be (10, 10)?

The problem is that as soon as they hang up the phone, each individual shepherd has the following temptation: "if the other person is only going to bring 10 sheep, that's fine, but I'm still better off bring 20 sheep. So I will ignore what we said on the phone and defect anyway." And of course both players face the same temptation. If they think this way, then we are right back to (defect, defect).

In the real world, people may not want to break their promises, and the guilt from doing this might be enough to prevent defection. If so, it means that we did not really provide a complete description of the game. We should have spelled out that there was a first stage of the game in which people have an opportunity to talk, and they can make promises if they want to. At the second stage, they decide how many sheep to put on the pasture. In this larger game, at the second stage each person's payoff depends not just on the profit they get from their sheep, but also the promises each person made, and how guilty they would feel about breaking a promise.

However, if we set aside issues about guilt or similar emotions, and we say that the true payoffs are the ones in the prisoner's dilemma itself, then communication will not solve the problem, because it will not change the fact that each person has a dominant strategy and an individual incentive to use it.

Another possible solution is for the players to sign a contract where each person agrees to bring only 10 sheep to the pasture. If either person violates the contract, the other person can hire a lawyer, bring a lawsuit, and ask a court to impose a penalty on the person who violated the contract (maybe in the form of monetary compensation). This could work in some cases, but in a wide range of real-world situations it would not be practical to create or enforce such a contract. I'll give a few examples of the prisoner's dilemma below, and you can ask yourself whether you think a contractual solution would work.

One other possible solution involves repetition of the game. If the same players are going to meet many times (maybe the shepherds have to decide every day how many sheep to put on the pasture), then people can use more complicated strategies. A player might say, "if you cooperate today, I will cooperate tomorrow. But if you defect today, I will defect tomorrow". Depending on how much people care about the future, strategies of this kind could be useful in preventing defection and escaping from the prisoner's dilemma. We'll talk more later about this approach.

So let's sum up: I have defined a prisoner's dilemma and argued that if the game is only played once, we probably get the result (defect, defect), which is socially undesirable.

The prisoner's dilemma game is useful in modeling a wide range of social situations. It raises a very general issue: what do we do if there is a conflict between individual self-interest and efficient social outcomes? Institutions are often designed to deal with this problem, and they have varying degrees of success in solving it.

Here are some examples where prisoner's dilemma issues arise in the social sciences.

1. The 'tragedy of the commons' (Garrett Hardin's 1968 article discussed earlier).

2. There is a famous book by a political scientist named Mancur Olson called "The Logic of Collective Action". The point was that in politics, group interests are not always well represented. For example: suppose we all agree that we want the BC

government to spend more money on university education. I might decide that my own efforts to lobby the government are unlikely to have a noticeable effect and it is more fun to stay home and watch the hockey game (dominant strategy). And maybe each of you will feel the same way. But if we all stay home, we don't get any additional funding for education. It might have made everyone better off if we had all gotten off the sofa and lobbied the government.

3. The arms race between the U.S. and Soviet Union during the Cold War. In this example, cooperation meant that a country would build fewer missiles and defect meant that it would build more missiles. The outcome where both countries build more missiles is both expensive and dangerous, so both would be better off if they could agree to invest less in weapons. But there is an incentive to cheat because if the other country invests less while you invest more, you will gain an advantage.

4. Collusion in an duopoly. Economists have models of industries where there are only two firms. The firms would like to collude and act like a monopoly because this is profitable. From the point of view of the firms, cooperation means setting a high price and defect means setting a lower price. But if the other firm is going to charge a high price, it could be more profitable for you to charge a slightly lower price and grab a larger market share, which increases your profit.

5. The natural environment. In this case cooperation means reduce your pollution, and defect means emit more pollution. For example, everyone might think it is convenient to drive around in their cars, and if they drive less it will not lead to any noticeable reduction in total pollution (dominant strategy). But if everyone drives around all the time, we have dirty air and global warming (inefficient).

6. Team production. Suppose you are working on a group project. Cooperation means contribute a lot of time and effort. Defection means slack off. It may be that defection is a dominant strategy (you prefer to relax and let other people do the work), but if everyone behaves this way, the project isn't finished on time.

I could give more examples, but that's probably enough.

An idea that is closely related to the prisoner's dilemma is the "free rider problem". This arises in the context of public goods. Recall that a public good is one where everyone consumes the same amount, and no one (at least within some geographic area) can be excluded from consuming it. Earlier I mentioned national defense as an example, and another good example is law enforcement (everyone who lives in a local community generally gets the same level of protection from crime). Back in the old days, a classic example was a lighthouse that warned ships away from some rocks. A similar example today is the global positioning system. Other examples include unscrambled radio or television transmissions that anyone can receive. In general, knowledge, information, and research tend to resemble public goods, because it often costs little or nothing to make the information available to everyone (although it might be possible to exclude some people from gaining access to it).

The *free rider problem* is that when people are making voluntary private contributions to a public good, the total contribution is usually too small from a social point of view.

The reason is that people often base their contributions on the personal benefit that they will get from the good, while ignoring the benefits to everyone else. However, they do take into account the full cost of their own contribution. As a result, everyone does too little, just as when people slack off on group projects.

Each person may have a dominant strategy where they make some voluntary contribution to the public good (something above zero) but the result is not Pareto efficient, because if everyone contributed more, everyone would be better off simultaneously. In practice we often solve problems like this by having the government collect taxes, which are used to pay for things like national defense, law enforcement, research, and so on. We don't rely on voluntary contributions to pay for public goods of this kind.

Now let's go back to the common pool resource (CPR) problem Ostrom is talking about. For convenience I will repeat the earlier payoff matrix here:

		Player 2	
		cooperate (10 sheep)	defect (20 sheep)
Player 1	cooperate (10 sheep)	10, 10	-1, 11
	defect (20 sheep)	11, -1	0, 0

Ostrom talks about two common ways of solving a CPR problem like this. The first is government regulation and the second is private ownership. She thinks both solutions have significant problems. I'll run through the arguments about each of them.

<u>Government regulation</u>.

The idea is that the government could monitor the behavior of each shepherd and punish anyone who defects. This changes the payoffs in the game. For example, suppose every time someone defects, they have to pay a fine of $2. The new payoff matrix is

		Player 2	
		cooperate (10 sheep)	defect (20 sheep)
Player 1	cooperate (10 sheep)	10, 10	-1, 9
	defect (20 sheep)	9, -1	-2, -2

If you go through the same analysis as before, you will find that player 1 has a dominant strategy. However, in the new game, the dominant strategy is to cooperate (10 is better than 9 and -1 is better than -2). Again, the game is symmetric, so the same logic applies to player 2: cooperate is a dominant strategy for the same reasons. Therefore we have a DSE, which is (cooperate, cooperate), and the payoffs are (10, 10).

So we have solved the problem and we don't need to read the rest of the book!

Not quite. There are several potential problems with this approach.

1. We might have self-interested behavior by law enforcers, regulators, politicians, governments, and so on. What we have done is added another player to the game but we haven't said anything about the payoffs or incentives of the new player. If law enforcement is a public good, how do we know enough of it will be supplied? What if the enforcers take naps, take bribes, or steal all of the sheep?

2. Even if the regulators are honest and competent, you generally have to pay them something. There are administrative costs of monitoring the shepherds, keeping records, imposing fines when necessary, etc.

3. The regulators may not have an adequate technical understanding of the situation. For example, they may not know that the maximum number of sheep that can use the pasture is 20, and that the pasture will be damaged if there are more than this.

4. The regulators can make mistakes. There are two kinds of possible mistakes. We have a Type I error if someone was punished even though they were cooperating, and a Type II error if someone was not punished even though they were defecting.

Ostrom runs through a long algebraic discussion to show that if the error rates are too high, it is still a dominant strategy to defect. This is easy to see in an extreme situation where people are punished randomly, regardless of their true behavior. Clearly in this case people have no incentive to change their behavior and start cooperating; they will just continue to defect while paying whatever fines are imposed. In order for regulation to work, the regulators must have reasonably accurate information about what individual people are really doing.

I won't go through all the algebra but in case you are trying to follow Ostrom's logic, here are some clues. A type I error is to punish someone when they choose C (cooperate) and a type II error is to fail to punish someone when they choose D (defect). The probability of a type I error is x and the probability of a type II error is $1-y$. Then set up a new payoff matrix in which each person receives the resulting expected utility under this system. For example, for the strategy combination CC, each player will receive $10(1-x) + 8x$. For the strategy combination DC, player 1 gets $11(1-y) + 9y$ and player 2 gets $(-1)(1-x) + (-3)x$. And so on. This gives the payoffs in Game 3 on p. 11 of the book.

After you crank through all the algebra, what you are left with is this. Given that player 2 chooses C, under what conditions would player 1 prefer to choose C over D? The answer is that player 1 thinks C is at least as good as D when

$$10(1-x) + 8x \geq 11(1-y) + 9y$$

This implies $y \geq x + 1/2$. So given some level of x (type I error), y must be high enough to satisfy this inequality, which means there can't be too much type II error.

By symmetry, the same result applies to the other player. Given that player 1 chooses C, player 2 is willing to choose C rather than D when $y \geq x + 1/2$. Therefore this condition is necessary and sufficient for CC to be a stable equilibrium.

In the example in the book, Ostrom assumes the government is correct 7/10 of the time. This means the probability of both types of error is 0.3, so we have $x = 0.3$ and $y = 0.7$ (recall that the rate of type II error is $1-y = 0.3$).

Is it true that $0.7 \geq 0.3 + 1/2$? No. Therefore, in this numerical example, if player 2 uses C then player 1 is better off choosing D rather than C. Therefore CC is not stable.

In this example there is a DSE and it is still (defect, defect). Ostrom's conclusion is that the error rate is too high to solve the prisoner's dilemma problem in this particular way.

In fact, in Ostrom's example things are even worse for the shepherds than before. Instead of getting (0, 0), they each get -1.4, because they both defect but there is a probability 0.7 of being caught, and when they are caught they have to pay the fine, which results in -2. So we haven't changed anyone's behavior and we are still wrecking the pasture, but the shepherds are now being punished for it.

A minor point: there is a typo in the book. The text says -1.6, but Figure 1.4 on p. 12 has the correct numerical result, which is -1.4.

Of course, this is just an example. If the error rates are lower, this approach will work. In the real world it is often possible to reduce error rates by allocating more resources to the problem (hiring more people to monitor the pasture). So in practice the question is likely to be whether it is worthwhile to spend the money that would be needed to solve the problem. Accurate information is not available for free.

Private ownership.

Again let's go back to the original payoff matrix on p. 1 of these notes. We could use private ownership in two different ways:

(a) Divide up the pasture equally so each shepherd owns half of it.
(b) Have one shepherd own the entire pasture (a monopoly).

In case (a) we have several potential problems.

1. Again there are enforcement costs. These are now borne by individuals. Each shepherd might have to build a fence to keep out the sheep of the other person, watch for possible trespassing, impose penalties for trespassing, and so on (we will see more about this type of thing in the Ellickson book).

2. Again there may be Type I and Type II errors. In this case, one shepherd might make a mistaken accusation (accusing the other shepherd of trespassing when in fact they are innocent), or might fail to notice a real violation when it occurs.

3. There is a possible new problem: risk. What if rain is unpredictable, sometimes falling on person 1's side of the pasture and sometimes person 2's side? When the pasture was freely available to both, then both shepherds could take their sheep to the area where the rain fell and the grass was green. If they could figure out how to cooperate, they could both obtain 10. But now, maybe rain on person 1's side of the fence gives payoffs of (15, 5) while rain on person 2's side gives (5, 15). So having separate territories can result in risk. If the shepherds are risk averse, they might want to find some other solution.

4. Another possible problem with separate territories involves fugitive resources (resources that can move around, like fish). You can assign property rights that allow individuals to fish in certain parts of a lake, but that doesn't mean the fish will decide to be there. Again, this can create risk.

In case (b), suppose person 1 owns the entire pasture and all of the sheep. Person 1 hires person 2 as an employee at a fixed wage w. As before, interpret C as 10 sheep and D as 20 sheep.

Person 1 now gets the total payoff (the sum of the payoffs in the game from p.1 of these notes) minus the constant wage payment w. Person 2 just gets the fixed amount w and doesn't care how many sheep the owner wants to have on the pasture. Assuming that 1 tells 2 what to do, player 1 will choose (cooperate, cooperate), because this gives 1 the total payoff of 20 (minus the wage), while any other outcome gives either 10 (minus the wage) or zero (minus the wage).

Notice that we are assuming that person 2 is indifferent and only cares about the wage (not things like the effort cost associated with different numbers of sheep).

This solves the problem because player 1 bears the cost of having too many sheep on the pasture. Another way to put it is that player 1 owns the entire pasture and therefore has an incentive to preserve it.

Note that we are also avoiding the risk issues raised earlier: player 1 can graze their 20 sheep on whichever part of the pasture had rainfall, where the grass happens to be best.

So what's the problem? There are two main issues.

1. Equity. If the shepherds could have found a way to cooperate, person 2 could have had 10 in the original game. But maybe the wage person 2 is getting as an employee is less than 10. This could make person 2 unhappy. In fact there are two issues: what person 2 actually gets relative to person 1 (if person 1 owns the entire pasture, why is that fair? why shouldn't person 2 own it?) and what person 2 would have had if it had been possible to achieve cooperation without resorting to private ownership at all.

2. Efficiency. Normally economists don't like monopolies. The reason is that they tend to restrict output in order to raise prices to consumers. This might not be a major problem if the shepherds are selling wool on a world market. But if they are selling their output in a local market, giving one person a monopoly over local wool production could be bad for consumers, because the consumers will have to pay a higher price for wool clothes. If there were two shepherds competing with each other, that would keep wool prices low and make consumers better off.

That's all for now. I'll talk about repeated games in my next set of notes. This is often an important way of solving the prisoner's dilemma problem.

Econ 354

Greg Dow

February 22, 2021

Notes on Elinor Ostrom, "Governing the Commons"

Chapters 1-2 (third set of lecture notes: repeated games)

This is a continuation of my previous notes on game theory. Here I focus specifically on repeated games and trigger strategies. These ideas are in the background in chapter 2 of Ostrom's book, as well as chapter 3 and elsewhere.

Ostrom frequently talks about the following concepts:

1. Discounting of future payoffs
2. Self-enforcing agreements
3. Mutual monitoring
4. Credible commitment

The game theory framework I develop here will incorporate each of these concepts. At the end of these notes I will discuss the role of each concept in the model and also some limitations of the model from Ostrom's viewpoint.

Let's start with a general Prisoner's Dilemma (PD) game:

		Player B	
		cooperate (C)	defect (D)
Player A	cooperate (C)	x, x	z, w
	defect (D)	w, z	y, y

where w > x > y > z. Based on previous notes, you should be able to see that this game has a dominant strategy equilibrium (DD) and that the DSE is not Pareto efficient (if we start from DD, the strategy pair CC is a Pareto improvement).

The usual question is whether there is any way to get to the cooperative outcome CC. If the game is only played once, it is not easy to see how this can happen. But suppose we repeat the game. Now people can use more complex strategies where their decisions are contingent on past behavior. For example, in a two period version of the game, a player could use C in period 2 if CC occurred in period 1, but use D otherwise.

The problem is that if we only play the game twice, and everyone knows this will occur, then people can anticipate what will happen in the second round. At that point it will be clear that each player has a dominant strategy (D) and each player will find it optimal to defect, regardless of what happened in the first round.

Now move back to the first round. If you know we will have DD in the second round no matter what happens in the first round, what should you do in the first round? The answer is that you should use your dominant strategy (D) in the first round also, because your choice in the first round will have no effect on the outcome in the second round, and you may as well do the best you can in the first round. If both people think this way then we also get DD in the first round. So repeating the game does not accomplish anything: we just get DD followed by DD again.

The idea of starting from the end of a game and working backwards to the beginning is called *backward induction*. This is a common way for game theorists to figure out how people will behave in a game with a finite number of moves.

Unfortunately, in the PD game, the same argument applies whenever the game will be repeated any finite number of times, and the players know this in advance. In that case backward induction says to start with the last round. Because there is a DSE at that point everyone will anticipate that DD occurs in the last round, regardless of anything that may have happened earlier in the game. Now go to the next to the last round, and ask whether anyone would choose C. The answer is no, because this does not change the fact that DD will occur in the last round. So each player will use D in the next to the last round too. You can keep moving back to earlier periods and applying the same argument back to round one. Thus the entire game unravels and the outcome is DD, DD, . . . DD.

One way to interpret this result is that in a finitely repeated prisoner's dilemma game, promises to cooperate in the future are *not credible*, in the sense that if one player makes such a promise, the other player will not believe it. The reason is that it will not be in the self-interest of the player who makes the promise to carry it out when the time comes.

You might have some doubt about the realism of the backward induction argument. If two people are going to play the PD game 500 times, is it really true that each person will start by thinking about the last period and working back to the beginning? Is it really true that people will never cooperate in any period? I'll postpone these issues for the moment and come back to them later when we talk about the Ellickson book. For the moment, let's assume the backward induction argument is correct in the case of finitely repeated games, and see whether cooperation could still be possible.

Suppose we have an unlimited time horizon, where the PD game will be played infinitely many times. In this situation, there is no last round so the backward induction argument cannot be used. What happens then?

One immediate issue is how we will calculate the payoffs from the repeated game if there are infinitely many periods. You might think this means that the payoffs themselves are infinite. However, I will show that the payoffs can still be finite.

Define a *discount factor* δ (the Greek letter delta), where we assume $0 < \delta < 1$ so delta is a fraction between zero and one. I'll interpret delta in a moment.

Now define the following *present values*, one for each player:

$$V_A = u_{A0} + \delta u_{A1} + \delta^2 u_{A2} + \ldots + \delta^t u_{At} + \ldots$$

$$V_B = u_{B0} + \delta u_{B1} + \delta^2 u_{B2} + \ldots + \delta^t B_{At} + \ldots$$

There are infinitely many time periods $t = 0, 1, \ldots$ Note that we call the first period zero instead of one. This is convenient algebraically but it has no real effect on anything.

Player A gets some payoff or utility u_{A0} from playing the game in period 0. The size of this payoff depends on the strategies used by each player in that round. Then the game is played again and player A gets u_{A1}, etc. In general, player A gets the payoff u_{At} in period t. Similarly player B gets the payoff u_{Bt} in each period $t = 0, 1, \ldots$

The role of the discount factor δ is to determine the weight placed on the future payoffs relative to the present. There is no discounting at $t = 0$. We multiply the payoffs at $t = 1$ by δ. Because this is a fraction, player A gives less weight to u_{A1} than to u_{A0} (a unit of utility tomorrow is worth less than the same unit of utility today). As we go further into the future, we discount payoffs by δ^2, δ^3, and so on, where each expression gets smaller because δ is a fraction. The payoff in period t is discounted by δ^t. When t is a very large number (a period far in the future), δ^t is very small and the corresponding payoff receives very little weight.

There are two extreme cases. As δ approaches zero, player A or B puts zero weight on all future payoffs, and only cares about u_{A0} or u_{B0}. In this case we are back to the one shot game where people ignore the possibility of future interactions.

As δ approaches one, the coefficient δ^t approaches one for any fixed period t. In this case we are putting almost equal weight on the payoff from each period (a player cares almost as much about the payoff 500 rounds in the future as about the payoff today). However, we cannot have $\delta = 1$ exactly, because then we would be adding up the payoffs from an infinite number of periods with equal weights, and the sum will blow up.

Note: people sometimes write $\delta = 1/(1+r)$ where $r > 0$ is an interest rate. As long as $r > 0$ we will have $0 < \delta < 1$ so this works. The terminology sometimes gets confusing though. I will call δ the *discount factor* and r the *discount rate*.

You should try to be clear about this distinction. For example if Ostrom says "person A discounts the future more than person B", what this means is that person A has a lower discount factor δ and a higher discount rate r than person B, so A places less weight on future payoffs compared to B.

As an aside, it would be possible to have an infinitely repeated game where the players have different discount factors δ_A and δ_B. The general results would be the same, and I ignore this possibility in order to simplify the algebra.

Another note (this one is very important): the discount factor δ can be interpreted as the *probability* that the game will be played at least one more time. So if we are in period 0, we get the payoffs u_{A0} and u_{B0} for sure. There is a probability δ that we will also play the game in period 1, and a probability $1-\delta$ that we stop playing (permanently). In the latter case, both players get zero in all future periods.

If we do continue to period 1, we repeat this process: there is a probability δ that we will continue to period 2, and a probability $1-\delta$ that we will stop and get zero in all the future periods (2, 3, 4 and so on). More generally, the probability that the players will still be playing the game in period t is δ^t.

Now you can think about the repeated game in terms of expected utility. The probability that the players get the utilities (u_{At}, u_{Bt}) in period t is δ^t for each t = 0, 1, 2 . . . So when you are adding up all the terms in the present value V_A or V_B, you are just computing the expected utility for that player.

This has a very useful real-world interpretation. If it is likely that the players will keep interacting for a long time (maybe they see each other every day and are unlikely to move to any other community), then they will put a high weight on their future payoffs. But if they are unlikely to meet in the future, they will put a low weight on their future payoffs. This has a major influence on their incentives to cooperate, as I will explain below.

Now that we understand present values and discounting, let's go back to the question of whether it is possible to get the cooperative outcome CC in every period. The problem is that there is a temptation for each player to cheat in the short run by switching from C to D and grabbing the payoff w instead of x (go back and look again at the payoff matrix on the first page of these notes).

The way to prevent this kind of behavior is to punish players who choose D instead of C. The most severe punishment is for the other player to use D forever, from that point on. This leads to the idea of trigger strategies.

We say that a player uses a *trigger strategy* if

1. The player starts by using C in period $t = 0$.
2. The player uses C at any $t \geq 1$ where CC has always occurred in the past.
3. The player uses D at any $t \geq 1$ where anyone has ever used D in the past.

If someone is using a trigger strategy, they start with cooperation. After that, they look at what each player has done in the past. They continue to cooperate as long as both players have always cooperated in the past. But if anyone has ever defected in the past (it doesn't matter who defected, or how long ago it happened), then the player defects now.

If both players are using trigger strategies, then we get CC at $t = 0$ because both people cooperate. When we go to period $t = 1$, both people look back at what happened in $t = 0$, they see that CC has always occurred in the past, and so they both cooperate again. Now the history of the game is CC, CC. When we go to period $t = 2$, this process repeats, and so on. So if both players are using trigger strategies, no one ever defects and the outcome is CC in every period.

This is clearly what we want. The question is whether this is an equilibrium, or whether one player could do better by switching to defection.

To see what is involved, suppose B is using a trigger strategy. A knows this and is trying to decide whether to use a trigger strategy too. To make things simple, suppose A looks at two options: use a trigger strategy, or defect in period $t = 0$.

If A follows a trigger strategy, we know the outcome will be CC in every period. From the payoff matrix at the beginning of these notes, that means A will get the payoff x in every period. Therefore A's present value will be

$$V_A = x + \delta x + \delta^2 x + \ldots + \delta^t x + \ldots \qquad \text{or}$$

$$V_A = x(1 + \delta + \delta^2 + \ldots + \delta^t + \ldots)$$

At this point we need to use an algebraic trick: when δ is a fraction, we have

$$1 + \delta + \delta^2 + \ldots + \delta^t + \ldots = 1/(1 - \delta)$$

Therefore if player A uses a trigger strategy, the resulting present value for A is

$$V_A = x/(1-\delta)$$

Notice that this is a finite number despite the infinite number of periods.

Now suppose instead that player A defects in period $t = 0$. This gives the payoff w at $t = 0$ because B is using a trigger strategy and starts by cooperating (again, look at the payoff matrix on the first page). But after that, B will always defect because B is using a trigger strategy and A has defected in the past. So in each period $t \geq 1$, the best A can do is defect and obtain y (this is better than z, which is what A would get by cooperating).

Therefore A's present value in this case is

$$V_A = w + \delta y + \delta^2 y + \ldots + \delta^t y + \ldots \qquad \text{or}$$

$$V_A = w + \delta y(1 + \delta + \ldots + \delta^t + \ldots) \qquad \text{or}$$

$$V_A = w + \delta y/(1-\delta)$$

where we used the same algebraic trick as before.

Comparing the present values from A's two strategies (trigger strategy versus defect immediately), the trigger strategy is at least as good if

$$x/(1-\delta) \geq w + \delta y/(1-\delta)$$

After some algebraic manipulation this reduces to

$$\delta \geq (w - x)/(w - y).$$

Let's define the right hand side to be δ_{min} and abbreviate this condition as $\delta \geq \delta_{min}$.

Recall that the payoffs in the one shot game satisfy $w > x > y > z$ (this is necessary in order to have a prisoner's dilemma). This implies that the numerator of δ_{min} is positive, the denominator is positive, and the numerator is smaller than the denominator. Hence $0 < \delta_{min} < 1$. So if the discount factor δ is close enough to one, we will have $\delta \geq \delta_{min}$.

If this is true, then player A prefers to use a trigger strategy rather than defect in the first round, given that player B is using the trigger strategy. Because the payoffs in the game are symmetric, the same is true for the other player: B prefers to use the trigger strategy rather than defect in the first round, given that A is using the trigger strategy.

We conclude that if $\delta \geq \delta_{min}$ then it is an equilibrium for both of the players to use trigger strategies (neither can obtain a higher present value by defecting). The outcome is CC in every period, and therefore the payoffs are (x, x) in every period.

The interpretation of this result is simple: if both players place enough weight on their future payoffs relative to the present, we can get cooperation in an infinitely repeated prisoner's dilemma game.

Using the probability interpretation given earlier, we can also say that if people are sufficiently likely to continue interacting in the future, we can get cooperation.

With the probability interpretation, we don't have to assume that people live an infinitely long time in order to have cooperation. As long as people are uncertain about when they will die, and in particular they are always uncertain about whether the game might go on for at least one more period, cooperation is a possibility, provided that the probability of continuation is high enough relative to the threshold δ_{min}.

The reason why this works is that if people care enough about the future, then the threat of defecting forever is a big threat. Although someone can get w in the short run from a defection, after that they will get y forever. If they care a lot about the future, the drop from x to y in all future periods is a big loss, and it outweighs the immediate gain from getting w once in the present. Of course, if they don't care much about future payoffs, then threats of future retaliation don't get much weight and the player prefers to take w rather than x today.

In the framework of the infinitely repeated prisoner's dilemma, defecting forever is the harshest possible punishment a player can impose. If we can't get cooperation through the use of trigger strategies, we will never be able to get it, because there is no way to impose a larger punishment for defection.

You might be worrying about something: when I went through the preceding analysis, I assumed that if player A was going to defect, this happened immediately at $t = 0$. What if A waits a while, uses C for a few periods, and then defects? It turns out that this doesn't change the algebra. A can't gain anything by waiting. If it is ever going to be profitable for A to defect, it is profitable to do this immediately so we only have to check that case.

However, there is a deeper problem, which is that more than one equilibrium can exist. We showed above that if the discount factor is high enough, there is an equilibrium with cooperation. But there is at least one other equilibrium. Suppose each player defects in every period, regardless of what happened in the past. This will also be an equilibrium (no matter what the discount factor is). If the other player will always defect regardless of what you do, your best response is to defect also (this is your dominant strategy, and you may as well use it).

So we need to be careful how we talk about cooperation here. The condition $\delta \geq \delta_{min}$ is <u>necessary</u> for cooperation (if it does not hold, there is no hope of achieving cooperation, because even the most severe punishment will not prevent defection). However, it is not <u>sufficient</u> for cooperation. There is still a bad equilibrium where people always defect, so we need to think about whether the good equilibrium will actually occur. Cooperation is not automatic.

How does all of this relate to Ostrom? I mentioned earlier that she talks about

1. Discounting the future
2. Self-enforcing agreements
3. Mutual monitoring
4. Credible commitments

We obviously have discounting the future built into the model. This relates to things that Ostrom is concerned about. For example, in stable long lasting communities, people are likely to interact for many future periods, so δ is close to one, and thus cooperation is at least a possibility.

Furthermore, we have mutual monitoring. In order to use a trigger strategy, a player has to observe the actions of the other player and adjust their behavior accordingly. You may wonder what happens when people can make mistakes: for example, I might incorrectly think that you defected when really you cooperated. If I am using a trigger strategy this leads to a very bad outcome where I defect forever even though you were innocent, and you have no choice but to defend yourself by defecting too.

One way to deal with issues like this is for people to use strategies that are less extreme than trigger strategies, where they punish the other player for a while and then go back to cooperation. This avoids the problem of punishing someone forever by mistake, but it means that people who really are defecting will not be punished as severely. The result is that the discount factor δ must be closer to one in order to persuade people to cooperate. The discount factor may or may not be sufficiently high for this to work.

Let's now think about self-enforcing agreements. The idea is that the players might sit down and discuss their situation ahead of time, and agree on how they will behave. The agreement is 'self enforcing' if each person has an incentive to carry out their promises, given what they expect the other person to do.

This is a useful way to think about situations where there is more than one equilibrium. Suppose $\delta \geq \delta_{min}$ so that cooperation can be achieved through trigger strategies. Each player knows that this is better than having both people defect forever. Suppose they agree to use trigger strategies. Is this a self-enforcing agreement? Yes, because if you believe that the other person will use the trigger strategy, it is in your interest to do the same thing, so you will keep your promise to behave in this way.

What about credible commitments? The idea here is that a promise or a threat by some player is *credible* if it would be in the self-interest of that player to carry it out when the time comes. In the context of trigger strategies, questions about credibility arise in the following way.

Suppose I say that if you ever defect, I will defect forever. Is this threat credible? Would I really do it? One could argue that the answer is yes, because you are using a trigger strategy, so if someone ever defected in the past (it doesn't matter who), then you will defect forever. This applies even if you are the one who defected. I know this is your strategy, so I expect that your defection will be followed by further defections forever. This implies that it is in my interest to defect forever too. So my threat is credible.

But there is a counter argument. Suppose someone defected in the past and we are now defecting forever. Will we really continue to do this? Maybe not. What we are doing is not Pareto efficient, and we could get a Pareto improvement by returning to cooperation. So why wouldn't we just sit down together, agree to resume cooperation, and restart our trigger strategies from the beginning?

The problem is that if we would do this, and we understand from the beginning that we would do this, then no one should believe that a single defection will really be followed by an infinite series of defections in the future. But then the punishment is not credible. And if the punishment is not credible, why wouldn't people simply defect immediately? This logic undermines the whole idea behind trigger strategies.

These are complicated issues and I won't pursue them further. However, you should be aware that the credibility of punishments in a repeated PD game is not always obvious.

At this point I will shift gears and talk about the limitations of the repeated game model from Ostrom's point of view. There are several issues.

1. The model assumes that the players have a lot of background knowledge about the payoffs they face. For example they understand the physical characteristics of the common pool resource, the consequences of their actions, and so on. There is no room in the model for learning by doing. But in reality (for example, in the cases from chapter 4 about groundwater basins in the Los Angeles area), the process of learning about such things can be very important.

2. The model assumes costless information about the behavior of other players in the past. This is often unrealistic. Sometimes mutual monitoring is possible at a low cost, but in other situations it is necessary to hire specialists to monitor everyone's behavior, and this may not be cheap.

3. The model ignores the idea of graduated penalties. This means using small fines for small violations of the rules, larger penalties for larger violations (or repeated violations), and so on. Trigger strategies impose the maximum possible penalty for any type of violation (you will be punished forever regardless of whether you defected just once, or several times). Extreme punishments are often regarded as unfair, and may have bad consequences for everyone (maybe if we all defect, we destroy the resource). So in reality, some non-compliance may be tolerated, or small penalties may be used for small infractions.

4. The model ignores the need for specialized methods of group decision-making. The successful cases from chapter 3 have various bodies to solve these problems: associations, assemblies, committees, and so on. In general it is important that the members of these organizations are users of the resource, and have an opportunity to participate in making decisions about the resource.

5. The model ignores the problem of defining who is a player in the game (entitled to use the CPR) and who is not a player in the game (so not entitled to use it). It would be possible to generalize the game beyond just two players (we could have

N players). What would typically happen is that with a larger N, the lower bound δ_{min} would be closer to one, so it would be less likely that δ is high enough to get cooperation. This makes sense: it is usually harder to solve incentive problems in larger groups than in smaller groups. But even with this generalization, N would still be some fixed number, and Ostrom would still want to know how the set of players was determined.

There are other interesting issues about repeated games that will come up again when we get to the Ellickson book. At that point I will discuss other strategies (besides the trigger strategy) that people can use in an infinite prisoner's dilemma, what happens in laboratory situations when real people play PD games, and what happens in evolutionary situations where successful strategies tend to spread through the population.

For the moment, that's all I have on game theory. The next set of notes will discuss case studies from the Ostrom book starting with chapter 3.

Econ 354

Greg Dow

February 28, 2021

Notes on Elinor Ostrom, "Governing the Commons"

Chapters 3-6 (rent dissipation)

Before I discuss the case studies in chapters 3-6, I want to explain one further economic model. This one does not use game theory but it focuses on a crucial issue in the Ostrom book: if there is a CPR, who gets to use it? Why is it important to have rules that impose restrictions on the set of appropriators?

One point EO makes repeatedly is that in order to manage a CPR, there must be clearly defined boundaries, in two senses: the physical boundary of the resource itself must be known, and the set of people eligible to gain access to it must be clear. This means that some appropriators must be able to exclude others. These ideas are related to the concept of *rent dissipation*, which I will define below.

Suppose there is a lake and n is the number of people who go fishing in the lake. Think of the lake as a common pool resource. Let p be the price of fish (output), and let w be the wage that each fisher can get by doing something else, like working at McDonalds.

We will use the ideas of average product and marginal product introduced earlier for the Johnson and Earle book. See Figure 1 as you read these notes. We assume AP rises at first, hits a maximum, and then falls. Similarly MP first rises, hits a maximum, and then falls. For mathematical reasons, it must be true that when AP is rising, MP is above it; when AP hits a maximum, MP passes through it; and when AP is falling, MP is below it.

To convert the physical AP and MP curves (measured in fish) into curves measured in dollars, we multiply each of these by the price of fish (p). Therefore the curves in Figure 1 show p x AP and p x MP. Multiplying by a positive constant in this way doesn't affect anything important about the shape of the curves but it does make it possible to compare these two curves with the wage (w), which is measured in dollars.

To see how incomes are determined in the fishing industry, think about it as follows. If n people go fishing, the total catch is some output q(n) that depends on the number who are fishing; each fisher gets an equal share of the catch q(n) so each person has q(n)/n fish to sell; and each gets the price p per fish, so income per person is pq(n)/n or p x AP(n).

The local community has a total population of N. We assume that N is fixed throughout these notes. Each person can either go fishing (n people do this) or work at McDonalds

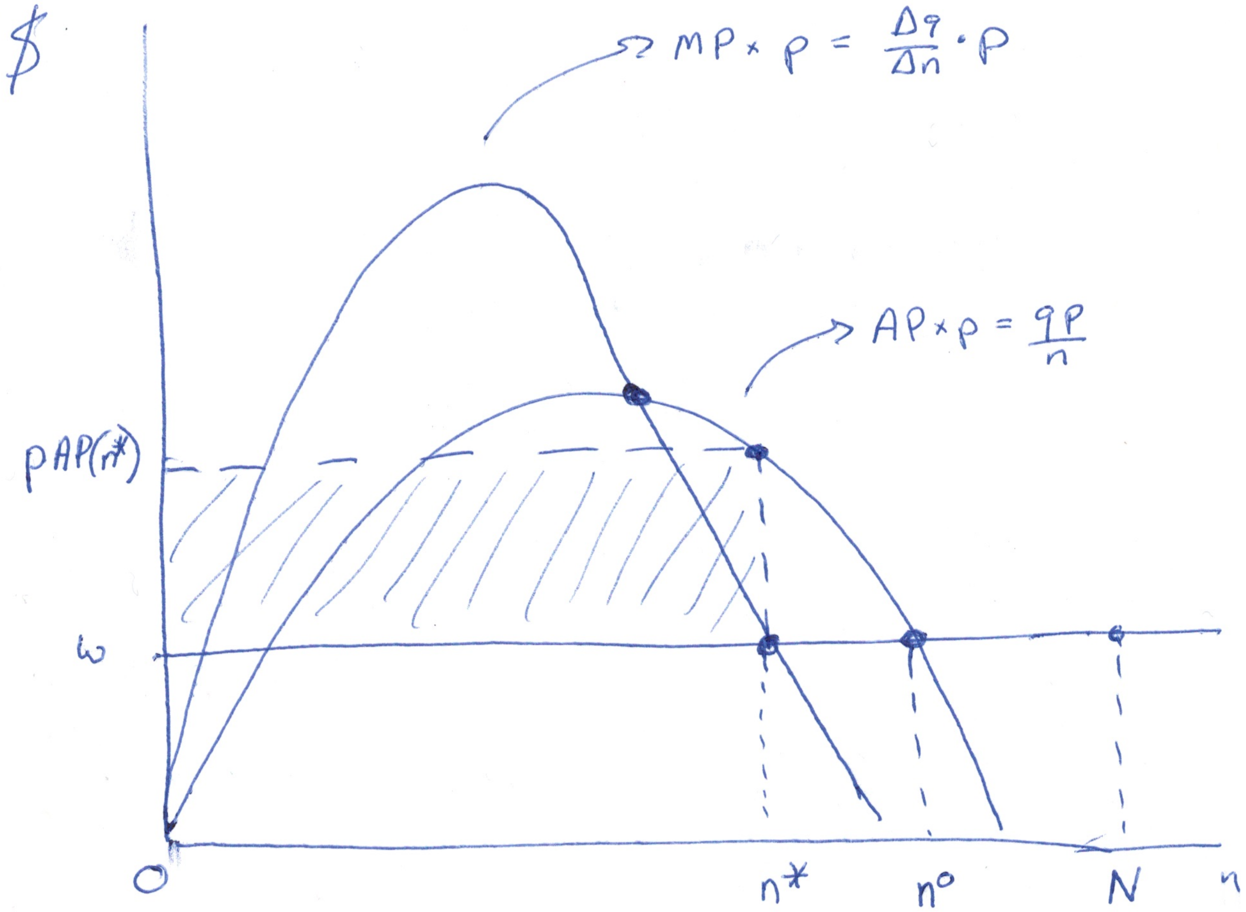

Figure 1

AP = average product = $\frac{q(n)}{n}$

MP = marginal product = $\frac{\Delta q}{\Delta n}$

p = price of fish

w = wage rate

N = total population of community

(N - n people do this). We are interested in what determines the level of n, in order to know how the total labor supply N is allocated between the two activities.

First, suppose there is open access: anyone is free to work in either activity. This means that in equilibrium an individual must be getting the same income from fishing and from McDonald's. If the incomes were unequal, some people would switch from the job with the lower income to the job with the higher income, so we don't have an equilibrium.

The incomes are equal when there are n^0 people going fishing and $N - n^0$ at McDonald's (see Figure 1), because someone who goes fishing receives the income $p \times AP(n^0)$ and someone who flips burgers gets w. At the point n^0 the income levels from these sources are equal.

Let's be clear about what is happening in the fishing industry: if n^0 people go fishing, the total number of fish caught is the output $q(n^0)$ determined by the number of fishers, each individual fisher has $q(n^0)/n^0$ fish to sell, and income per fisher is $pq(n^0)/n^0$ or $p \times AP(n^0)$. In an equilibrium with open access, this is equal to the wage w. Therefore in equilibrium everyone is receiving the income w, regardless of whether they catch fish or flip burgers, and everyone is indifferent about which of the two jobs they have.

From an economic point of view, this is not Pareto efficient. Suppose we define a Pareto improvement as a change in the situation that makes some people better off and does not make anyone worse off. This is slightly different from my earlier definition that a Pareto improvement has to make everyone better off simultaneously but it is a very similar idea.

Using my modified definition, it is easy to see that starting from n^0, we can get a Pareto improvement. All we need to do is reduce the number of people who go fishing to some level n' slightly below n^0 and have the $n^0 - n'$ people who no longer catch fish go to work at McDonald's instead. Everyone previously working at McDonald's is no worse off and the people who leave the fishing industry in order to flip burgers are no worse off either (they were getting incomes equal to w in equilibrium anyway, and they still get w in the burger industry). But the n' people who are still in the fishing industry now have a higher average product; that is, $AP(n') > AP(n^0)$. So these people now have higher incomes and are better off than they were before. This implies that we have a Pareto improvement, so n^0 was not an efficient way to allocate labor.

Now suppose the community hires you as an economic consultant. Their goal is to set a level of n (the number of fishers allowed on the lake) that will maximize the total income of the community as a whole. What would you recommend?

Total income for the N members of the community is $pq(n) + w(N - n)$ where $pq(n)$ is the total revenue generated by catching fish and $w(N - n)$ is the total income obtained by the people who work at McDonald's. We are assuming p and w are fixed: the town sells fish on an international market and can't affect the price p, and McDonald's sets the wage w at its world headquarters so the town can't influence that either. All it can do is choose n.

If we rearrange the total income algebraically, we get pq(n) - wn + wN. Notice that wN is a constant. The town has no control over w and no control over the total population N, which is fixed. Thus, choosing n to maximize total income is equivalent to choosing the level of n that maximizes

$$pq(n) - wn$$

An economist would call this expression rent, although you could also call it profit. The usual way of defining the *rent* from a resource is that it is the total revenue obtained from the resource minus the total opportunity cost of getting the revenue. Here the opportunity cost is wn because this is what people are giving up by not working at McDonald's.

We can do a little more algebra and express the rent in the form

$$pq(n) - wn = [pq(n)/n - w]n = [pAP(n) - w]n$$

This says that the rent is equal to the difference between the value of the average product and the wage, all multiplied by the number of people who go fishing.

It turns out that if you want to maximize the rent (and therefore total community income), you need to increase n as long as $p \times MP(n) > w$. The reason is that when this is true, the value of the extra fish p x MP caught by one more person on the lake exceeds the income w lost by having that person give up their job at McDonald's. This process should stop at the number of fishers n* where $p \times MP(n^*) = w$. If we continued to raise n beyond this point, we would get $p \times MP(n) < w$ and it would make sense to move some people out of the fishing industry and have them work for McDonald's instead (their additional wages would more than compensate for the lost fish).

Note: if you know calculus, this is easy to see. Maximize pq(n) - wn by differentiating with respect to n. The derivative of q(n) is the marginal product of labor in fishing. The first order condition says that p times MP must equal w. If you don't know any calculus, you can ignore this note.

Figure 1 shows n*, the point where the value of the marginal product from fishing equals the wage w. I assume that we are operating on the downward-sloping part of the MP and AP curves. In this case $n^* < n^0$ so when we choose n* there are fewer people fishing than the number n^0 we had with open access. The rent at n* is equal to the area of the shaded rectangle, which is

$$[pq(n^*)/n^* - w]n^* \qquad \text{or } [pAP(n^*) - w]n^*$$

From the graph you should be able to see that when w is below the maximum level of p x AP, so a positive level of rent can be achieved, n* will be <u>below</u> n^0 but <u>above</u> the level of n that maximizes average product. Therefore rent maximization is not the same thing as AP maximization. In this model, if you want to maximize total community income, you do not maximize AP. More people than that should go fishing.

On the other hand you definitely don't want to be at the open access equilibrium n^0. At that point, the rent is zero and total community income is wN. This outcome is exactly the same as what the total income would be if everyone worked at McDonald's and the town did not have a lake at all.

Thus, allowing open access destroys the economic value of the CPR. In equilibrium, no one is getting any benefit from the existence of the resource because it is being overused. In order to get some economic benefit from the lake, we have to prevent this overuse by the appropriators. It is exactly the same idea as the Tragedy of the Commons.

An economist would say that allowing more people than n* to use the lake leads to some degree of *rent dissipation*, which means that the town enjoys less rent than the maximum it could have had by choosing n*. When we are at n^0, we have complete rent dissipation, which means that the rent is zero.

The lesson Ostrom takes away is that it is vital to have clear rules about who is allowed to use a CPR. If there are no rules and anyone can use it, we dissipate all of the rent. So the first principle of good institutional design is to specify who is allowed to use the CPR.

At this point, you might be concerned about equity or fairness issues. You might say, yes it is bad to destroy the rent. But at the point n*, the people who are allowed to go fishing get an income of p x AP(n*) which exceeds w (this must be true if the rent is positive). Why should these people get more than the people who are excluded from the fishery?

This is a reasonable question. The issue is how the benefits from the CPR are distributed among the individual members of the community. Having n* fishers will maximize total income, but this doesn't say anything about how that income will be divided up.

There are at least two ways to bring about a fair distribution in this situation. First, you could randomize who gets to go fishing by having a lottery, where everyone has an equal chance of receiving a fishing license. In this system there will still be inequality after the licenses are handed out, but at least everyone had an equal chance to obtain one.

Second, the town government could sell fishing licenses. Assuming there are n* licenses available, the maximum someone would be willing to pay for a license is p x AP(n*) - w, because this is what they can gain from going fishing rather than working at McDonald's. Note that this amount is equal to the rent per person when n* people can go fishing.

The town can either set the price at this level (if it knows what it is), or alternatively hold an auction where the n* licenses are sold to the highest bidders. Either way, the revenue collected by the town will be equal to the total rent. After paying for their fishing license, each fisher will have the net (no pun intended) income level w so they don't get any more disposable income than the people working at McDonald's.

The government can then divide up the revenue it obtained from selling fishing licenses equally among all of the people in the town, or spend it on something that everyone likes (better schools, better roads, or whatever). This approach takes care of the equity issue by spreading the benefit of the CPR among everyone in the community, even if only a subset of people are actually allowed to go fishing.

Econ 354

Greg Dow

March 1, 2021

Notes on Elinor Ostrom, "Governing the Commons"

Chapters 3-6 (case studies)

Now that we have developed some background knowledge about game theory and rent dissipation, we are ready to talk about the case studies in EO's book and the lessons she derives from them.

Chapter 3 is about successful examples of CPR management. EO's criteria for success are simple: (a) the institutions lasted a long time and (b) the resource was not destroyed. These are crude criteria but there is no exact way to assess economic efficiency or equity, so (a) and (b) provide reasonable tests. As we will discuss below, there are other cases in the book that failed to pass these tests.

The success stories are as follows:

Torbel, Switzerland (pastures and forests)
Japanese villages (pastures and forests)
Spanish irrigation systems (water management)
Philippine irrigation systems (water management)

All of these systems are very old and very stable.

Chapter 4 is mostly about successes, but EO shifts the focus to 'institutional supply'. This means: how did people create new institutions in order to solve CPR problems?

Note that in chapter 3 we do not observe the actual process through which the institutions were created, but in chapter 4 we do.

The examples include the Raymond Basin, West Basin, and Central Basin. All of these involve groundwater and all are located in the Los Angeles area in southern California.

Chapter 5 is about what EO calls 'failures and fragilities'. This means that either the local people were unable to solve their CPR problem, or the institutions are in trouble. Cases:

Bodrum/Bay of Izmir, Turkey (fisheries)
San Bernardino County, California (groundwater)
Mawelle, Sri Lanka (fishery)
Kirinki Oya, Sri Lanka (irrigation)

Nova Scotia, Canada (fishery)

Chapter 6 is a general theoretical framework.

I will approach the issues in chapters 3-6 in two steps.

First, I want to discuss the *design principles* the successful cases have in common. The goal is to identify the elements or characteristics that good CPR institutions have. Or to put it more simply: what do good institutions look like?

Second, I want to discuss conditions in the physical and institutional environment that increase the *probability of success*. When is it likely that people will create good CPR institutions? When is it unlikely? Why?

Let's deal with the *design principles* first. My discussion will be based on pp. 88-102 (end of chapter 3); pp. 178-181 (end of chapter 5); and some of chapter 6.

EO summarizes eight design principles for good CPR institutions. She has a useful table on p. 180 where she shows that the successful cases satisfy all eight of the principles (for these cases, she says the institutional performance is 'robust'). For the remaining cases where institutional performance is labeled 'fragile' or 'failure', at least some of the eight design principles are violated.

I will run through each of the eight principles and discuss the reasoning behind them. Ostrom sometimes uses different labels or terminology when she is talking about her eight principles. This may be confusing, but the description of each principle below should make it clear what she is talking about.

1. Clearly defined boundaries. The boundary of the resource itself must be known, and the set of people eligible to use it must be clear. The reason is simple: if we don't have these features, it is impossible to stop people from using the resource. This leads to rent dissipation and can destroy the economic value of the CPR (see my earlier lecture notes on rent dissipation).

2. Appropriation and contribution rules. A successful CPR institution needs to have rules about the times when people can use it, the places where they can use it, the technology they use, and the quantity of the resource they can take. It also often requires rules about the responsibilities people have to maintain the resource.

   Even if we have clear boundaries and we have defined the set of people who can use the resource, we must still limit the number of units of the CPR each person can take (in order to avoid depletion problems), and we must impose obligations for people to maintain the resource (in order to avoid free rider problems on tasks like investment and maintenance).

The best way to set up these rules will depend on the type of resource (such as pastures, water supply, or fish), and also on local details: Is storage possible? How fast does the pasture grow back? At what time of year are fish available? And so on. It is good for the rules to be as simple and clear as possible.

3. Collective choice arrangements. It is sometimes necessary to change the day-to-day operating rules for appropriation or contribution. This means there has to be some institutional arrangement for doing this. Examples include:

   (a) Switzerland. There were 'alp associations' that included all local people who owned cattle, and the locals also voted on village statutes.
   (b) Japan. Village assemblies consisting of heads of households created rules.
   (c) Spain. There were autonomous 'irrigation communities' that elected chief executives. There were also 'water courts'.
   (d) Philippines. Members of the 'zanjera' elected leaders. There were also federations of local groups.
   (e) Los Angeles. Initially there were voluntary associations of water users. Eventually these evolved into 'special water districts' that were a kind of public enterprise.

   In all of the successful cases there was a high level of democratic participation by the users, and there were ways to adapt rules to changing conditions at relatively low cost.

4. Monitoring. There is a temptation for individuals in CPR situations to take more of the resource than what is allowed by the appropriation rules (a tragedy of the commons), or provide less maintenance effort than is required by the contribution rules (a free rider problem). The rules don't enforce themselves. As in a repeated game, the participants must collect information about the behavior of others and impose penalties when someone breaks the rules. We have seen earlier from the discussion of Type I and Type II errors that this must be done with a reasonably high level of precision, or the necessary incentives will break down.

   In all of the successful cases, there was either

   (a) Direct mutual monitoring as a by-product of other activities (I see how much water you take, or I may catch you in the forest out of season); or
   (b) The users made shared contributions to a monitoring system involving specialists. In the Swiss case this was a local official who controlled the number of cows; in the Japanese case there were detectives who patrolled the commons; in the Spanish irrigation case there were people called ditch riders who watched the canals; in the Philippine case there was a maestro who supervised maintenance labor; and in the Los Angeles water basins there were people called water masters. In each case, the monitors were not outside authorities like police or courts, but rather specialists who were hired or elected (and paid) by appropriators of the resource. The resulting

information was widely publicized and the monitors were accountable to the appropriators, so they could be replaced if they did not do a good job.

5. Graduated penalties. It is not only necessary to catch people who violate the rules, it is also necessary to punish them. If deterrence works well (people know they will be punished for breaking the rules, so most of them decide not to), then it won't be necessary to impose penalties very often.

   EO makes the point that in the cases she studied, some people became desperate enough to break the rules occasionally. But if they did it in a minor way, or not very often, they would usually get a slap on the wrist (a minor punishment). This helps maintain a sense of fairness, and it also sends a signal that violations will be noticed. This lets people know that the monitoring system works.

   If people violate the rules in a major way, or do it repeatedly, or grossly shirk in their contributions to joint projects, they are usually punished more severely. In extreme cases, they may be expelled and no longer allowed to use the CPR.

   Notice that graduated penalties differ from trigger strategies, which impose the maximum possible penalty for any violation. It might seem that this would give the most effective kind of deterrence, but it leads to problems:

   (a) This approach is not very flexible. Realistically, one should not expect or demand 100% compliance. Sometimes people have hardship situations or there are humanitarian reasons not to impose the maximum punishment.
   (b) Rigid rules like trigger strategies are highly vulnerable to mistakes. This takes us back to the Type I/Type II error problem. We should only impose big punishments when we are very sure that the person is guilty.
   (c) Extreme punishments may not be credible (people may not believe that they would actually be imposed, perhaps because they might require the people doing the punishing to bear high costs).

   I won't go into as much detail about the remaining three design principles:

6. Cheap conflict resolution systems (either among the individual appropriators, or between an individual user and officials or leaders).

7. Local autonomy (external authorities don't thwart attempts by the locals to solve their CPR problem).

8. Nested enterprises (decisions are often best made in a hierarchical way: there may be an operational level involving day-to-day decisions, a collective choice level where the operational rules can be changed when necessary, and a constitutional level where the collective choice procedures can be changed when necessary).

It is useful to compare EO's description of a good CPR institution with the idea of a local government. Governments normally have three main features:

(a) Legislative procedures to create rules
(b) Executive procedures to administer and monitor the rules
(c) Judicial procedures to enforce the rules and resolve disputes.

Fundamental to all of this is a definition of who is a member of the system and who is not (some concept of residence or citizenship).

The main difference between what Ostrom is describing and what a local government does is that a CPR institution manages a specific resource. A local government is not specialized in this way. Instead it is used for general purposes and it collects general revenue through taxation. But otherwise, they look quite similar to me.

At this point I want to switch topics. Now that we know what a good institution looks like, what can we say about *the probability that a good institution will be created*? My discussion of this topic will be based on pp. 136-142 (end of chapter 4) and chapter 6.

Before getting into this, I want to comment on the distinction EO makes between 'the first order dilemma' and the 'second order dilemma'. When she is talking about the first order dilemma, she is referring to the usual problems that arise in managing a CPR: free riding, shirking, taking too much of the resource, depleting it, and so on.

When she is talking about the second order dilemma, she means the following. If there is a CPR and people do not currently have good institutions, the economic value of the CPR will be low. If people could somehow create good institutions, the value of the resource would rise. But there is a free rider problem: why should an individual or a small group spend their own time, money, and effort to create a good institution when they bear most of the costs of doing this, and other people get most of the benefit? So the second order dilemma is that constructing a good institution is like supplying a public good, and each individual may contribute too little to activities of this kind.

EO thinks that under some conditions, it is highly unlikely that this second order problem can be overcome. For example, if people don't put much weight on the future, have low trust, choose actions independently, don't communicate with each other, can't make any binding agreements, and have no way to monitor or enforce compliance, then we are left with something like a one-shot prisoner's dilemma, and it is very unlikely that people will be able to create good institutions.

EO believes the following factors raise the probability that people will be able to create good CPR institutions.

1. <u>A small number of decision makers</u>. It is usually fairly easy for small groups to communicate and bargain with each other, and each person typically has a large

stake in solving the problem. But if there are many people involved, this raises the cost of negotiation and makes free rider problems more severe.

2. Relatively few participants are needed in order to obtain some collective benefit. This is a little different from the first point. Even if the total number of decision-makers is large, there may be a few 'big' users who can collectively make a major difference by reducing their own use of the resource. These people might be able to agree on an institutional structure that would be helpful, and they might face a less severe free rider problem among themselves. These people could be able to start the process of building an institution, and bring in the smaller users later.

3. People put a lot of weight on future payoffs. This could occur for various reasons such as having a relatively stable community of users, where individuals are likely to interact repeatedly in the future (low turnover in community membership). It may also be important whether a resource grows back rapidly or slowly because this affects the potential future reward from using less of the resource now.

4. Similarity, symmetry, or homogeneity of interests. If the users all have similar technologies, similar utility functions, and so on, they may have an easier time agreeing on what a good institution should look like and how the interests of the individual users should be represented. In such situations, it may be relatively easy to get a Pareto improvement by having everyone agree to cut back on their use of the resource in a parallel way. If there are several different categories of users, they have different technologies, different preferences, etc., it might be much harder to get a Pareto improvement (for example, it may be very costly for some people to reduce their use of the resource, and they won't cooperate unless they receive compensation from the other users).

5. The presence of people with leadership skills or other useful assets. This is fairly obvious: it is more likely that new institutions can be created if there are people involved who are good leaders, have relevant skills or training, are wealthy, are good at communication, are well connected politically, and so on.

You should take a look at the failure cases from chapter 5 in light of this list. You will probably find that all had problems with respect to one or more of these characteristics. For example, in the San Bernardino example there were many decision makers, it would have required cooperation from many people to get a noticeable collective benefit, there were different groups of users with different economic interests, and so on.

The last point I want to discuss involves strategies that might be useful in creating good CPR institutions.

Hierarchical organization. It is often possible to construct new institutions by building on existing institutions. You don't have to start from scratch. Moreover, it is often possible to organize large numbers of people in this way.

For example, consider the following levels of organization

Burnaby (the city where SFU's main campus is located)
Metro Vancouver Regional District (a set of local governments in the Vancouver area)
British Columbia
Canada
NATO, NAFTA, The World Trade Organization, the International Monetary Fund, the United Nations.

The MVRD solves a range of problems that go beyond what the city of Burnaby can do by itself. For example, you can think of the roads, bridges, and other transit facilities in the Vancouver region as a CPR. About 2.5 million people rely upon this transportation system. You might think these would be too many decision-makers and thus it would be impossible to organize them to manage the CPR. But in reality, it wasn't that difficult. There were already about 20 city governments in the area, so all that was necessary was to construct a set of Metro Vancouver institutions to coordinate things. The MVRD also handles other regional issues like economic development and environmental problems that go beyond what each individual town or city can do.

Similarly, British Columbia solves problems that go beyond the Vancouver area, Canada solves problems that go beyond the level of the provincial government. And so on.

As we learned from Johnson and Earle, when people encounter problems that cannot be solved by small-scale institutions, they frequently create larger-scale institutions. This is evident at the international level. We have a large number of organizations that include multiple countries (in some cases, almost all countries). These deal with security issues, economic development, and global environmental problems, among other things.

Hierarchy does not necessarily imply a lack of democracy. For example, people vote for the government of Burnaby, the government of British Columbia, and the government of Canada, even though these three levels of government stand in a hierarchical relationship with one another.

The point here is simply that if you already have some pre-existing institutions, and you need to solve a new problem, it is often a good strategy to construct new institutions that build on what you already have.

<u>Incrementalism</u>. The other strategy that is often useful in creating new institutions is called 'incrementalism'. The idea is that you don't have to do everything at once. In fact, this could be counter-productive. Instead, it makes sense to start with low-cost forms of collective action that lead to tangible short run benefits. One example might be setting up a voluntary organization to gather information about groundwater. Assuming this step is successful, now you have an organizational framework that could be useful for tackling more challenging collective action problems. In the process, people may develop more confidence in each other, higher levels of trust, and so on, even if they were previously suspicious of each other.

This technique has often been used in arms control negotiations (for example, between the U.S. and the Soviet Union during the cold war). Rather than trying for one enormous comprehensive agreement that solved every problem, it was often more productive to try for a series of modest steps where neither side felt completely vulnerable. Over time, it was possible to build up confidence in the monitoring procedures, obtain some small but tangible benefits, and establish relationships among the negotiators. After this had been done, people would sometimes try to solve harder problems.

Do strategies like hierarchical organization and incrementalism always work? No.

Think about the failure cases in chapter 5. What were some of the problems?

Turkish fisheries: some users had opportunities for quick economic gains, some groups had conflicting interests (for example based on the use of different technologies), there were no mechanisms for designing rules, etc.

San Bernadino County: the scale of the CPR was big, both in terms of geography and population; the resource was complex (15 different basins with different water sources); there was uncertainty about the nature of the resource and how severe the problem was; and there were conflicts of various kinds (small versus large pumpers, pro-development versus no-growth advocates, industry versus agriculture, local residents versus outside experts). The people who tried to solve the problem did not break it down into subparts and proceed incrementally, but it was a hard problem and it is not clear that they would have succeeded even if they had used this strategy.

Sri Lankan fishery (Mawelle): there was no way to prevent entry by users of the resource, which led to rent dissipation; there was a lack of local autonomy (and likewise for the Sri Lankan irrigation system, although this was a partial success); corruption enabled new entrants to bribe officials not to enforce the rules, and so on.

Nova Scotia fishery: the Canadian government tried to impose standardized rules from the outside, which led to a lack of local autonomy.

Conclusion: even if we know what good institutions look like (we have read chapter 3), there is no guarantee that it will be possible for users of a CPR to create good institutions. The probability of success depends on a variety of factors in the physical and institutional environment. There are some strategies that might boost the chances of success, but they do not always work. CPRs are not easy things to manage.