Econ 354

Greg Dow

March 31, 2021

Notes on Louis Putterman, "The Good, the Bad and the Economy"

Preface and Chapters 1-2 (Overview)

I will start with a few personal comments. Louis Putterman and I have been friends for almost 40 years. We met in the early 1980s when I was at Yale University and Louis was at Brown University (where he still is). Our research interests have often overlapped, and we worked together on some journal articles and book chapters during 1999-2000.

I understand how Louis thinks about a wide set of issues. However, we don't agree about everything. More importantly, you don't have to agree with him about everything. Don't worry about criticizing his ideas just because he is a friend of mine. What I care about is whether you provide good reasons (involving logic or evidence) for your opinions.

Louis had early interests in development economics, especially in Africa and China. He then moved into the areas of firm organization and experimental economics. Recently he has been interested in global economic history, and the origins of global inequality. This book summarizes LP's research interests in a non-technical way. Some chapters are easy and have lots of examples. Others are longer and more difficult.

I will group the chapters into four main topics.

1. Overview: the preface and chapters 1-2.
2. Human nature: chapters 3-5.
3. Institutions: chapters 6-7.
4. Global inequality: chapters 8-10.

My notes will be in separate files for each of these four sections of the book.

These notes are about the first part.

Overview: Preface and Chapters 1-2

I will begin by discussing some big concepts that will be used throughout the book, and then comment on chapters 1 and 2.

One very basic distinction involves positive statements versus normative statements.

A *positive* statement says how the world is. It may be true or false. Examples are (a) "it was sunny in Vancouver today" (true) or (b) "it was snowing in Vancouver today" (false).

A *normative* statement says how the world should be, or should not be. It may be widely accepted or it may be the view of a small minority. Examples include (a) the government should spend more money on education, or (b) loud music in restaurants is annoying.

When LP is providing examples, discussing historical facts, or reporting experimental results, he is making positive statements. When he is expressing opinions about what is good or bad, he is making normative statements. Notice from the title of the book that one of the things he wants to do is discuss normative issues.

The next big concept: what does LP mean by the term "progress"?

Usually *progress* means movement in the direction of goals that LP supports.

Examples of economic goals that LP supports include higher real wages, more productive technology, less poverty, and less inequality.

Examples of political goals that LP supports include more democracy and less war.

Examples of social goals that LP supports include longer life expectancy, higher literacy rates, a cleaner environment, more scientific knowledge, less crime, and lower stress.

Most of these goals are not controversial but they are still normative. Keep in mind that statements about how some goal could be achieved, or how close it is to being achieved, are positive. The normative part is the claim that the goal itself is good or desirable.

A large question for LP is why we haven't made more progress toward achieving these goals, and other similar ones. LP is especially concerned about why we haven't made more progress in reducing global economic inequality.

Next big concept: what does LP mean by "human nature"? He doesn't define it explicitly, but here is the definition I will use.

Human nature is a set of empirical generalizations about human behavior that are true for almost all societies regardless of their historical or cultural differences.

In this definition, when I say *empirical*, I mean based upon facts, observations, data, or evidence (as opposed to theory).

The concept of human nature is controversial (both whether such a thing exists at all, and what it may include). Some people would argue that the set of empirical generalizations that apply to almost all societies is very small, while others would argue that it is large.

Notice that I included the weasel word "almost" in the definition. I have found that when I make any generalization about human societies, a smart anthropologist can always think of one or more exceptions to the rule. I want to leave a loophole for cases like that. An empirical generalization can be true in a statistical sense, and be very important, even if there are a few individual cases where it doesn't apply.

Why is this idea important for LP's book and for the course?

1. Economists have to make some assumptions about human behavior. It would be nice if these assumptions applied to most or all societies. Social scientists from other disciplines might say that this is impossible because historical and cultural factors are too variable and too important. But still, economists have to assume something (or else give up).

2. We would like to know whether human nature imposes constraints on the kind of institutions we can have, or on how well certain institutions will work. This is an important practical consideration in thinking about how to achieve a better world.

According to LP, we can investigate human nature in various ways. We can learn some obvious things about it from casual observation. If we want to be more systematic, we can look at history, look at anthropological studies of other societies, run experiments, and so on. If we want to explain human nature (treat it as an endogenous variable), we can look at the biological evolution of human beings.

Here is a short summary of what LP thinks about human nature (from the preface, p. x):

"We humans are overwhelmingly predisposed towards looking out for ourselves and our immediate kin . . . but we're also highly social animals with predispositions toward empathy, toward positively and negatively reciprocating the helpful and harmful actions of others, toward caring about what others think of us, and toward wanting to think well of ourselves . . . cooperating toward common goals is part of what we do."

We have talked a lot in this course about institutions, but I haven't given you any official definition of this concept yet. LP does not provide a definition. There are many ways to define this word. For example, it is often treated as a synonym for 'organization', as in "SFU is an institution". Sometimes it refers to governments, markets, private property, marriage, and so on.

The two most common uses of the word in economics are the following:

1. An *institution* is a set of "rules of the game". Institutions determine what people are allowed to do when they interact with each other, as well as the payoffs they get from their actions. In this view, the rules are imposed from outside the social group. This definition was proposed by Douglass North, an economic historian who won a Nobel Prize in 1993. It can be helpful when you think of institutions as exogenous.

2. An *institution* is a "stable pattern of social behavior". We saw from Johnson and Earle that some societies do not have any central authority that imposes rules on everyone else. We saw from Ostrom that sometimes people develop their own institutions. We saw from Ellickson that people may ignore legal rules and use social norms instead. So it may make sense to look at what people actually do, rather than assuming they have to follow rules imposed on them from outside. This definition can be helpful when you think of institutions as endogenous.

Either point of view about institutions could be useful, depending on the context. For example, chapter 6 of the book will discuss human behavior in lab experiments. In this case, the experimenter determines the rules of the game, and it would make sense to say that the experimenter chooses the institutions the players must use (the North definition).

Or: we might be interested in some small-scale community without a government, and find that there are stable patterns of social, economic, or political behavior that persist over time, even though individual people come and go (maybe due to birth and death or migration). In a case like this, we might say that these persistent patterns *are* the social institutions, although no one consciously chose them.

LP studies artificial institutions in lab experiments, and also real institutions through the evidence provided by archaeology, anthropology, and history.

He believes this approach gives us valuable information about how well various kinds of institutions have functioned in the real world.

He also believes that property rights, markets, and competition are essential ways of harnessing self-interest in order to achieve economic prosperity. However, this isn't the whole story -- there is more to human nature than self-interest, and there are problems that can't be solved by markets. So we need additional institutions.

In LP's view, one important obstacle to progress is the conflict between self-interest and cooperation. At this point you should be thinking about the prisoner's dilemma, free rider problems, the tragedy of the commons, and so on. Ostrom and Ellickson both had a lot to say about these issues. LP believes institutions play a vital role in managing this conflict.

Chapter 1.

According to most economic historians, the Industrial Revolution started in England in the period 1750-1780. Real wages did not start rising until about 1820-1880, so for as much as one century, most working class people continued to live in poverty. Early in the Industrial Revolution, various people responded to this situation by trying to create utopian communities that would improve conditions for workers.

LP discusses the example of Robert Owen, who founded two utopian communities, one called New Lanark in Scotland, and the other called New Harmony in Indiana (USA).

In Owen's view there was no permanent human nature. Instead he thought that education and the social environment could shape attitudes and behaviors. If one believes this, then one can design ideal social institutions and train people to behave in whatever way might be necessary. In practice Owen emphasized the value of altruism and the desire for social approval, rather than self-interest.

I will briefly summarize the case of New Harmony. It was initially popular and attracted many recruits. However, there was a series of problems. There was no criterion for who could be a member, so anyone could join. Often the people who did join lacked the skills needed by the community. There were organizational problems like: Who is in charge? How do people get paid? What are the incentives for high quality or quantity of output?

Eventually these problems led to financial deficits. There was also a problem that class distinctions existing before people arrived tended to carry over into NH. A number of factional conflicts arose. Most people today regard NH as a failed utopian experiment.

LP draws several lessons from this case. The key one is that economic incentives and the way in which organizations are designed really matter. You can't ignore these things. If you do, the institutions you create are unlikely to be robust and sustainable.

Chapter 2.

We don't have a perfect world. What are the obstacles to further progress?

1. Is it mainly scarce resources?
2. Is it mainly human nature?
3. Is it mainly the institutions we have?

An economist might tend to think first about the scarcity of resources. In fact, people often define economics as a science that studies the allocation of scarce resources.

Clearly the planet has finite resources. These include good agricultural land, fresh water, forests and other ecosystems, fossil fuels, mineral deposits, etc.

But such things are only scarce in relation to population and technology. If we had a very small global population, we would have a lot of resources per person. If we had a very productive technology, we could use our limited resources very efficiently, and get a lot of goods and services from the resources we have. Institutions may also play some role in determining how much social benefit we can obtain from our limited resources.

One way to think about this is to picture the standard of living per person in a society as a function of several variables: resources, population, technology, and institutions. We can think about varying each of these factors while holding the others constant. In this view, the standard of living is the endogenous variable, and the others are exogenous variables. Holding other things constant, people tend to be better off when resources are abundant, when population is low, when technology is productive, and when institutions work well.

LP argues that the growth of world production of food calories has outpaced the growth of world population over the last 30 years or so. So at least in the recent past, technology was winning the race against population and living standards were rising for most people. However, we can't be certain that this trend will continue. Technological progress could slow down, or the environment could deteriorate, or resources could be depleted.

[Note: LP published his book in 2012. I have seen data indicating that in more recent years, the growth rate of world food output has fallen behind the growth rate of world population, perhaps due to climate change.]

In any event, LP does not think that continuing poverty at a global level is explained by the scarcity of resources, because he believes that technology can outrun population, at least in principle. Instead, he thinks poverty is best explained by institutions.

Take an example: why do we produce anti-baldness medication rather than anti-malaria medication? LP says the answer to this question is not determined by scarcity but rather by institutions. Firms produce what is profitable. This depends on the demand for one kind of medication relative to the other. In turn, demand depends on the distribution of income, which is influenced by institutions.

So should we get rid of markets and the profit motive? LP says no. That was tried in large parts of the world during the 20th century (the Soviet Union and China). In fact, the ideological foundations of the Cold War involved arguments about central planning versus markets, and state ownership versus private ownership. This came down to the issue of how economic institutions should be designed.

As it turned out, central planning and state ownership didn't work very well in terms of economic performance, and they were associated with political dictatorship. The Soviet Union collapsed, and China moved in the direction of less central planning, more use of markets, more private ownership, and more use of the profit motive.

LP thinks we need something like markets due to the prominent role of self-interest in human nature. But he also thinks markets cannot solve all social problems. All of the developed countries use government to regulate markets and redistribute income.

At the same time, governmental institutions have their own set of problems. Free riding limits the extent to which people become informed about policy issues, and the extent to which they vote. Political outcomes are influenced by money and special interests. And it is hard to convince 51% of the population to transfer income to people they have never met. It is even harder to convince them to do such things on the international level.

That's it for today. Notes on chapters 3-5 are coming soon.

Econ 354

Greg Dow

April 1, 2021

Notes on Louis Putterman, "The Good, the Bad and the Economy"

Chapters 3-5 (Human Nature)

Chapter 3 (self-interest).

This chapter is easy to read so I will just give a quick overview.

It is not difficult to find evidence that humans pursue self-interest. What does LP mean by this idea? A concern for one's own physical safety, food, shelter, clothing, comfort, and so on. He would also include resources invested in one's children, which from a biological point of view is a type of self-interest (we'll come back to this in chapter 5).

Economics has a long tradition (going back at least to Adam Smith) of assuming that people pursue self-interest and largely ignore the effects of their actions on others.

Standard economic assumptions about preferences include

1. People care about their own consumption of goods and services, not the consumption of goods and services by others.

2. More is better; so higher income or wealth is preferred to lower, because this makes it possible to buy more goods rather than fewer.

3. Other common assumptions in economics include the idea that people treat effort as costly, leisure as a good; they prefer consumption today over equal consumption in the future; and they are often risk averse.

Examples of self-interest provided by LP in this chapter:

1. Firms pursue profit, and their shareholders want them to, even for products that may be harmful to consumers or the environment.

2. Immigration motivated by the desire for a higher standard of living, and any efforts by people in the receiving country to block immigration because they are afraid it will lower their standard of living.

3. Education: people go to school in the hope of obtaining higher incomes and better jobs [note: professors hope that students also want to learn things!]

4. Marriages often involve sorting by income (high income people tend to marry each other, which implies that lower income people also end up marrying each other).

5. Residential neighborhoods often involve sorting by income (there are neighborhoods where most people have high incomes, and neighborhoods where most people have low incomes).

6. There is often a lack of trust in markets with informational asymmetry (for example, the used car market). Buyers don't automatically believe what the sellers say about the quality of the goods in such markets, and their suspicions are frequently justified.

7. Deng Xiaoping is famous for introducing stronger material incentives in China, which led to faster economic growth.

These are all examples of what LP would call 'polite' self-interest (they do not involve the use of force).

Economists routinely assume that this kind of self-interest exists (including some lying about private information) but they usually also assume that people will respect property rights (or at least that property rights are easily enforced).

However, in the real world there are a lot of examples of self-interest involving the use of physical force at the expense of others. These include crime, human trafficking, slavery, serfdom, colonialism, civil wars, and dictatorships.

Under normal conditions, in well-functioning societies, these things are usually avoided or suppressed, at least to a large degree.

Good institutions protect us from the worst forms of self-interest, especially forms that involve the use of force in non-legitimate ways. However, with bad institutions, things can get really bad. So in this relatively obvious way, institutions are important.

Chapter 4 (better angels).

Again, lots of examples in this chapter, and I will just provide an overview.

The main idea is *compassion*. This means an ability to feel what others feel, and to care about it.

Economists have traditionally been reluctant to make interpersonal utility comparisons, which means a comparison of the welfare of person A with person B. The usual attitude in the profession has been that this is too difficult, or somehow unscientific. Economists have also traditionally been reluctant to assume that other people (non-economists) make judgments of this sort about each other.

But in the real world:

1. People who know each other well do this all the time. It is not unusual for someone to say that person A is happier than person B, and these are not meaningless statements.

2. People routinely assume that when an extra dollar is given to a poor person, this raises the welfare of the poor person by more than if the dollar had been given to a rich person.

Also notice that Ellickson's whole theory of social norms assumes that people can add up utilities, so he is assuming that people make interpersonal comparisons in everyday life.

Examples of situations where people feel compassion for others, and this influences their behavior (including their economic behavior):

1. People sometimes risk their lives to save a stranger (such as volunteer fire fighters).
2. People sometimes risk their lives to fight for abstract principles based upon religion, nationalism, freedom, or democracy, for example.
3. People are routinely helpful to strangers they will never meet again (such as tipping in restaurants when traveling, or spending time giving directions to lost tourists).
4. People sometimes make career decisions where they sacrifice higher incomes in order to help others or do things that are socially meaningful.
5. People contribute time and money to support political parties despite the existence of large free rider problems, and sometimes support policies that run contrary to their own self-interest (such as paying higher taxes in order to help poor people).
6. People often help sick or disabled family members, and sometimes adopt genetically unrelated children.
7. People donate time and money to charities.
8. People donate time and money to support religious or artistic activities.

So there is a lot of economic behavior that doesn't seem consistent with pure self-interest.

A couple of observations about happiness as measured by survey questionnaires:

1. Self-reported happiness seems to depend more strongly on the quality of one's social relationships than on one's income.

2. In rich countries, income seems to affect happiness more through relative position than through the absolute level of income (so comparisons with other people matter).

The bottom line for this chapter:

Life is not just about competition, fraud, and violence. Cooperation occurs in virtually every society. It is probably a hard-wired part of human nature.

Chapter 5 (evolution and human complexity).

If we want to understand where human nature comes from, we need to understand how humans evolved biologically.

But first, I will say a few more things about economics.

LP distinguishes between *homo economicus* (H.E.) and *homo sapiens* (H.S.).

H.S. refers to real human beings (it is the biological label for our species).

H.E. refers to a simplified, fictional set of assumptions about how humans behave. These assumptions are often useful for economic models but are not meant to be the full picture.

H.E. is sometimes called "a rational economic agent". This usually implies self-interest and internal consistency of preferences, the stability of preferences over time, updating of beliefs in response to new information using the rules of probability, and so on.

Note: I would express this a little differently from LP. I would say a person is rational if they have consistent and stable preferences, and use the rules of probability. However, I would not include self-interest in the definition of rationality. In my definition a rational person could care about the consumption of others (they could have compassion).

It is easy to ridicule H.E. by arguing that it is not a 'realistic' portrait of human beings. People from other social sciences often make such arguments. But no model can be fully 'realistic', because every model is a simplified version of reality. The question is whether the model is useful, given the purpose for which it was designed (the purpose of a model could be accurate prediction, causal explanation, policy analysis, or whatever). We don't want to include every possible complication in a model; this would make it useless.

So don't make the mistake of treating a model as a literal description of the world. I like the analogy of a road map and a road. A model is like a road map; reality is like a road. No one would want a road map to include every fact about reality; it should just include information that is useful in getting from point A to point B.

When an economist assumes people are rational or self-interested or both, a question you should have in mind is whether these assumptions are likely to be useful in light of what the economist is trying to do. If the goal is studying behavior in financial markets, then assumptions like this may be completely sensible. But if the goal is studying charitable contributions, maybe it will be necessary to drop the part about self-interest.

Over the last several decades economists have become a lot more flexible about the kinds of assumptions they are willing to make about human behavior. Experimental economics has been an important part of this trend. This is a large and growing field, and its results are often published in top journals. The idea is to see how people behave in laboratory

situations. Often the results seem surprising in relation to theoretical expectations. For example, people are often willing to give up some of their own money in order to punish others who behave 'unfairly' (either unfairly to the person doing the punishing or to some third person). This seems inconsistent with standard ideas about self-interest. We'll talk more about this in chapter 6.

Now let's move on to

Biological evolution.

For those of you who may not have come across these ideas before, I will explain some basic concepts. Think about a population of some kind (in general this could be a species of plant or animal, but I will focus here on humans). There is genetic variation within the population on certain characteristics. For example, in a human population there could be variations in height across individuals due to differences in their genes (i.e., their DNA). Some genetic traits might be a good fit with the current environment while others are less good. For example, in certain environments it might be helpful to be tall (you can reach things that are up high), while in other environments it might be helpful to be short (you can squeeze through small spaces). People who have characteristics that are well suited to the current environment tend to be successful at surviving, reproducing, and raising their kids. But people who have characteristics less well matched to their environment tend to be less successful in these ways. Over many generations, genetic characteristics that are well adapted to the current environment tend to spread through the population because the people with those characteristics have more surviving kids, while other not-so-well-adapted genetic characteristics become a smaller fraction of the population and may eventually disappear completely. This process is called *natural selection*. It is the central idea in the field of evolutionary biology.

People often ask whether a particular characteristic (like height) is determined by nature (genes) or nurture (the environment one grows up in). But often both factors play a role, and they interact in complicated ways. Your adult height depends both on your genes (your DNA) and how much food you received when you were a growing child. More generally, the environment can affect which genes are switched on or off.

LP makes the point that there are basic similarities of emotion and social interaction across many species: humans, chimpanzees, dogs, elephants, dolphins, whales, etc. In short, there are many 'social animals'. So some of what we call 'human nature' may not really be limited to humans.

The large size of human brains probably reflects the complexity of our social interactions, not just the fact that we use tools. There is a general tendency for social animals to have large brains relative to their body size.

Natural selection does not automatically imply that self-interest will prevail. You might think that only the most selfish would survive. However, all social species have bonding or caring behaviors among individuals. In reality it is often true that cooperation leads to

evolutionary success (recall the success of the tit for tat (TFT) strategy in the repeated prisoner's dilemma game, which we discussed in connection with the Ellickson book).

From the perspective of evolutionary biology, one of the things that makes humans very unusual compared to other animal species is that we have an exceptionally high amount of cooperation among genetically unrelated individuals.

It is not hard to explain why genetically related people would tend to cooperate. You share 1/2 of your genes with each of your siblings, 1/8 of your genes with each of your first cousins, and so on. So a gene that says you should help your relatives, even at some cost to you, can spread through a population over many generations. For example, if you have such a gene, there is a 50/50 chance that your sister has the same gene. If you save her life and she has kids, some of those kids will carry that gene, and so on. Biologists call this process *kin selection*.

A bigger question is why some people help other genetically unrelated people. This is where humans are really distinctive.

In general, you might expect natural selection to favor free riders. If everyone else in a group is cooperating, and you defect, your payoff is bigger than theirs. This should give you a higher probability of survival and reproductive success. So we would not expect natural selection to favor strategies in which people cooperate all the time, or maximize aggregate welfare. Genes that promote altruism toward unrelated people at the expense of the person carrying the gene should tend to disappear.

BUT: this does not rule out *conditional cooperation* (as in TFT). You could still use a strategy that says "I will help you if and only if you also help me". This is often called *reciprocity* or *reciprocal altruism*. People behaving this way are called *reciprocators*.

Next we consider the concept of *group selection*. Suppose groups of people compete with one another to get natural resources. This does not necessarily involve conflict or warfare. It could just be that some groups are more successful at getting food or other resources (think about the small foraging bands described by Johnson and Earle).

Now suppose groups whose members cooperate with each other are better at getting food than groups where every individual member is pursuing their own self-interest. Groups with lots of cooperators (or reciprocators) might survive and reproduce more successfully than groups with more defectors (or free riders). If this is true, we might expect natural selection to favor genes for reciprocation.

BUT STILL: there is a temptation for people in mostly cooperative groups to defect once in a while. Cooperative groups need to find ways of suppressing this kind of behavior. It can be done if the cooperators punish people who defect, shirk, free ride, etc. A problem though: what if punishing defectors has a cost to the punisher? Then punishing defectors is a public good. Someone who lets other group members do the punishing should get a higher payoff. So there is still a free rider problem.

Group selection can overcome this problem under certain conditions:

(a) There is a big advantage to being in a group with a lot of reciprocators, compared to a group with relatively few people of this type.

(b) The cost of punishing someone who defects or violates social norms is small (for instance, you might just not interact with them any more).

(c) Reciprocators tend to sort themselves into groups that have other reciprocators (e.g. you can identify defectors in advance, or fairly quickly after interactions begin, so you can screen them out or avoid them).

LP says that as an empirical matter, reciprocity is universal in all human societies (it is part of human nature). It almost certainly has genetic foundations. Furthermore, it is reinforced in practice by repeated game and reputation mechanisms that harness self-interest for the purposes of cooperation. [This is sometimes called *enlightened self-interest*, meaning that people can see how cooperating with others may be beneficial to themselves in the long run.] LP points out that many institutional mechanisms serve to screen out people with a record of bad past behavior, or ostracize them, or prevent them from being promoted to higher levels within an organization, etc.

Now let's go back to the idea that humans evolved in a context of small mobile foraging bands. We want to consider what characteristics might have developed through natural selection in societies of this kind. LP thinks that 'human nature' includes the following.

Reciprocation (as discussed above).

A desire for social approval. People often give up some income in order to get approval from others. We can imagine what other people think of us (we know what we think of them, and we can imagine that they look at us in a similar way). So we behave in ways that make other people think well of us, including being willing to help them if needed. Specific norms within a society are often backed up by this social approval mechanism. The specific norms may be culturally determined, but the general concern for approval probably has a genetic basis, because individuals who had such concerns tended to be more successful biologically (they got help from others when they needed it).

A desire for self-approval. Going one step further: people often take actions that they know others would approve of, even when no one else is watching. This involves ideas about self-control, or the ability to resist temptation (Ellickson talked about this). For example, you might not steal money from others even if you believe you can get away with it. You can think of having a conscience as a kind of short cut that enables you to do what others would approve of, without having to think much about the details of the situation (such as whether people happen to be watching at the moment).

Mental accounts. It is probably universal that people keep track of the past behavior by others in relationships, have ideas about whether a particular relationship is currently in or out of balance, and retain mental records of what they owe to others, and what others owe to them. This is related to the "Even Up" strategy in Ellickson.

Us versus them. It is probably universal that people can define themselves as members of groups and can feel a sense of identity, solidarity, or loyalty toward the other members of their own group. At the same time, they may behave in ways that are suspicious, hostile, or dehumanizing to outsiders. This likely comes from the fact that our distant ancestors evolved in small groups that competed with other small groups. But people are clearly flexible about the criteria they use for defining group membership, and there is nothing inevitable about racism, religious intolerance, etc. Many successful societies have been highly diverse in terms of race and religion. Although an ability to think in terms of 'us versus them' is probably built into human nature, the details about how such groups are defined clearly come from culture and history.

At this point you may be thinking: not everyone behaves in the ways described above, or at least it is a matter of degree. Some people care a lot about social approval while others care relatively little, and so on. LP agrees about the importance of individual variation.

In the next couple of chapters, an important theme will be that people vary in the degree to which they pursue self-interest. Some people actually do care a lot about reciprocity, fairness, and cooperation. Other people only obey cooperative norms when it is in their self-interest to do so (for example, if institutions punish people who violate the norms).

LP's view on this: if you take a large random sample of any national or ethnic group from anywhere in the world, you will obtain a similar distribution of attitudes on this spectrum. Different human populations have statistically similar mixes. Why? Because attitudes of this kind tend to reflect hard-wired genetic variation that is a product of human evolution. Note here that biological evolution doesn't make everyone identical; there is always some variation within a population. LP is just saying that the statistical distributions are similar for all human populations.

This view is controversial. Some people in other social sciences would argue that much of what LP regards as 'human nature' is actually driven by culture or history. But I tend to be sympathetic with LP's perspective.

A few closing comments on biological evolution:

In economics, we normally treat preferences as exogenous. One nice thing about the biological framework is that we can treat preferences as endogenous, at least to some degree, by explaining how they evolved.

I think some of our preferences are genetically hard-wired, others are culturally shaped, and others are due to random individual experiences such as eating strawberry ice cream when you were young. But biology gives us a way to think about certain important parts

of this: for instance, why most people have concerns about fairness. It would be difficult to explain such observations without relying on some kind of biological mechanism.

Another point is that biological evolution is usually a very slow process, at least for large mammals like us. Humans have evolved over millions of years. For many genetic traits, it might take 10,000 years or more to see any noticeable biological change. But we have only had agriculture for the last 10,000 years, and cities for the last 5000 years.

Before that, everyone was in small mobile foraging bands. So this was the environment that shaped our biological evolution, and there has probably not been very much change in human nature over the last 10,000 years. Cultural change is much faster.

Today we live in radically different societies: cities with millions of people, nations with tens or hundreds of millions of people (two with over a billion), markets that extend over the entire world, etc. It is amazing to me that we have been able to develop institutions to organize people on this scale, given the small foraging groups in which humans evolved. Can we continue to solve institutional problems arising on the scale of the modern world? We don't know. The modern world is a huge experiment and we have no way to predict what the results will be.

Econ 354

Greg Dow

April 2, 2021

Notes on Louis Putterman, "The Good, the Bad, and the Economy"

Chapters 6-7 (Institutions)

These two chapters are both about institutions. Chapter 6 discusses evidence about how experimental subjects respond to artificial institutions created in the laboratory. Chapter 7 explores the historical development of real economic institutions.

Chapter 6 (experimental economics).

There are two main issues in chapter 6:

1. How do people of different "types" interact? (cooperators versus free riders etc.)
2. How do these interactions affect the performance of institutions?

We will run through various games, with a focus on what theory predicts versus what is actually observed in the lab.

The Voluntary Contribution Game (this is a classic type of experiment on public goods).

People are put into groups of 4. We call the players i = 1, 2, 3, 4.
The players are anonymous and interact through a computer.
Each player starts with $10 received from the experimenter.
Each player i = 1, 2, 3, 4 chooses some amount x_i between 0 and 10 to contribute to a group fund. The total contribution is $x_1 + x_2 + x_3 + x_4$.
The experimenter adds 60% to the fund, so now the total is $(1.6)(x_1 + x_2 + x_3 + x_4)$.
This amount is split equally among the individual players (it is divided four ways).
The game ends.

In this game, the payoff to person 1 can be written as

$u_1 = 10 - x_1 + (1.6/4)(x_1 + x_2 + x_3 + x_4)$ or after some algebraic manipulation,

$u_1 = 10 - (0.6)x_1 + (0.4)(x_2 + x_3 + x_4)$

This generates the payoffs in the table shown on p. 123 of the book.

Does player 1 have a dominant strategy? Yes: set $x_1 = 0$. No matter what the other players do, this always maximizes the payoff of player 1.

This is a symmetric game so for the same reasons, each player has a dominant strategy. Therefore, there is a dominant strategy equilibrium (DSE) where $x_1 = x_2 = x_3 = x_4 = 0$.

In the DSE, each player's payoff is 10. Is this Pareto efficient? No. To see why:

Suppose each player put in the same amount x. Then player 1's payoff would be

$u_1 = 10 - x + (1.6/4)(4x) = 10 + (0.6)x$

The same is true for players 2, 3, and 4. This payoff is maximized by x = 10 (everyone makes the maximum possible contribution). In this case each player's payoff is 16. This would clearly be a Pareto improvement compared to the DSE where each player gets 10.

So if everyone is selfish, everyone gets 10. If everyone is cooperative, everyone gets 16.

According to my definition earlier in the course, this is a prisoner's dilemma, because (i) the game has a DSE and (ii) the DSE is not Pareto efficient.

Another way to put this: each player's contribution has positive externalities for the other players, but if the players are selfish, they ignore these positive externalities, so the group as a whole ends up worse off.

Game theory predicts that we should observe the DSE outcome, at least when the game is played a finite number of times. But what do people really do?

1. In the simplest version of the game, with no repetition, on average people put in about $5. This is neither the DSE nor is it the Pareto efficient contribution level. Furthermore there is a lot of variation across individuals in the contributions they make (LP would say that this reflects differences in the 'types' of the players).

2. If the game is repeated 10 times, and everyone knows this in advance, on average the contributions fall over time. This could be due to learning about the nature of the game (but see point 4 below). People contributed more when it was the same group of people for all 10 rounds rather than having group membership randomly reshuffled after each round.

3. If the experimenter raises the group reward from 1.6 to 2.0, average contributions rise. If the group reward is lowered from 1.6 to 1.2, the average contribution falls. So people do appear to be responding to economic incentives. On the other hand, in all of these cases, the DSE prediction is that people should contribute zero.

4. If the experimenter announces a surprise 10 more rounds, contributions jump back up to an average of $5. It is as if this announcement restarts the game, indicating that the earlier decline in contributions was not just about learning.

Hypothesis from LP: we can explain all of these results by heterogeneity among subjects. People are not all the same. Some people are more interested in cooperating than others. This is consistent with the individual differences observed in the one-shot game.

A more detailed version of the hypothesis:

(a) If someone is a cooperator, they first check to see whether others will cooperate [this explains why a lot of people start off by making contributions above zero].

(b) If someone is a cooperator, they punish non-cooperators [this explains why the average contribution tends to drop over time; the cooperators discover that some of the other players are not going to cooperate, and they punish these people by reducing their own contributions].

(c) Each time a new group forms through the random reshuffling of members, this pattern repeats.

(d) If the group remains the same, an unexpected restart after 10 rounds causes the cooperators to hope that the defectors have learned a lesson. Also, the defectors may pretend to be cooperative for a while in order to get a larger payoff near the end. So the restart gets contributions to go back up temporarily.

We will come back to this game later and think about how changes in the rules can lead to better outcomes.

<u>The Ultimatum Game</u> (this game is very famous and very simple).

There is a total of $10 available from the experimenter.
There are two players (A and B) who can't communicate in advance.
The players are anonymous (they communicate only through a computer).
A proposes a split (x_A, x_B) of the $10, where x_A is what player A will get and x_B is what player B will get. These must be whole numbers (no fractions) and must add up to 10.
B says yes or no to A's offer. If B says yes, each person gets what A proposed. If B says no, each person gets zero (the experimenter keeps the money).

Game theory says we can predict the result using a method called *backward induction*. This means starting at the end and working back to the beginning. Here is the logic.

If $x_B > 0$ then player B should say yes, because getting x_B is better than getting zero.

Player A correctly expects B to behave this way, so A will choose x_B to be the smallest possible positive number. Given that only whole numbers can be used, this is $x_B = 1$.

Player A then takes the rest of the money: $x_A = 10 - x_B = 9$.

This is a clear prediction, and it is clearly wrong!

This game has been played many times in many societies around the world, with different histories, cultures, levels of economic development, and so on.

The outcome is almost always 5/5 or 6/4, with B saying yes.

Sometimes player A proposes 7/3, 8/2, or 9/1. Player B often rejects such offers, and the rejections become more common as the inequality becomes more extreme.

One possible explanation: B wants to develop a reputation for toughness. But the design of the experiment does not give B any payoff from having such a reputation (no one will find out what B did). Of course, some players may not believe this.

A more widely accepted explanation is that the players tend to be thinking about fairness.

One interpretation the players might place on the game: A does not deserve more than B, because the roles in the game were assigned randomly (it was a matter of luck). If A tries to exploit this unfair advantage, B wants to punish A, although this means B will have to give up some money to do it. As A's proposal becomes more unequal, B is more inclined to punish A for several reasons: (i) the extent of the unfairness increases, which makes B more annoyed; (ii) the cost to B of punishing A falls (B is sacrificing less by saying no); and (iii) the size of the punishment to A is rising (B can force A to sacrifice more).

Furthermore: it is not just that player B has a sense of fairness. It is also important that player A expects this, and takes it into account when deciding on an offer. Of course, A might also have a sense of fairness, but even a totally self-interested A would be wise to consider the possibility that B might be concerned with fairness.

A suggestion from LP: maybe anger is how nature gets people to punish free riders, even when the process of punishing them costs something to the punisher.

<u>More on the Voluntary Contribution Game</u>.

Now that we have seen the results for the Ultimatum Game, we can think more about the role of punishment in the voluntary contribution game.

Suppose we have two stages in the voluntary contribution game. The first is the same as before. In the second stage, experimental subjects learn the individual contributions of others, and can punish other players if they want, but it will cost them something to do it.

Again, game theory suggests using backward induction. If you only go through the two stages once, no player should punish anyone at stage two. This costs you something and there is no benefit (except maybe some emotional satisfaction).

But if everyone anticipates that no one will punish at stage two, then we are back to the original game at stage one. Adding the punishment stage shouldn't change anything.

The same backward induction idea works if there are finitely many repetitions. There should be no punishment at any stage, and therefore no cooperation at any stage.

But in the laboratory (with a known finite number of repetitions):

1. Many people do punish.
2. This normally involves high contributors punishing low contributors.
3. Contributions tend to rise over time rather than dropping.
4. There is no drop in punishment in the last round (after many rounds).

Often these experiments involve university students (they are convenient experimental subjects from the standpoint of the professors doing the research). However, the results generalize to non-student populations, people from poor societies where $10 is a lot of money, and so on.

LP's conclusion: many people are reciprocators. They are nice to other people who are also nice, but they punish people who are not nice. However, not everyone behaves in this way. Some people are just free riders and don't even try to be nice.

This supports the idea that individuals are frequently heterogeneous in situations where cooperation is beneficial to a group.

Some institutional solutions that can be useful:

1. Let cooperators (or reciprocators) form their own groups and expel people who repeatedly free ride. This works well.

2. Use majority voting when deciding whom to punish, in order to avoid situations where the low contributors punish high contributors. This also works well.

3. Let people see the performance of other groups (give them information about the contributions, the punishments, etc.). Groups tend to imitate the behaviors used in more successful groups.

4. Allowing people to talk with each other is very helpful (even though game theory says it shouldn't matter). Verbal promises are important. Many people do not like to violate their promises. Also, if they do, they may face even more punishment.

<u>The Trust Game</u> (this is also a very famous and influential game).

Note: all steps in this game are private, and everyone knows the rules in advance.

To begin, the experimenter gives $10 to A and $10 to B.

1. A puts \$x in an envelope to be given to B, where $0 \leq x \leq 10$. In practice this involves the use of whole numbers.

2. The experimenter adds \$2 for each \$1 contributed by A. So now the envelope contains 3x.

3. B opens the envelope and sends \$y back to A, where $0 \leq y \leq 3x$. B keeps the rest of the money.

In terms of algebra, the payoffs are

$$u_A = 10 - x + y$$
$$u_B = 10 + 3x - y$$

The total payoff is always $20 + 2x$. This is aggregate welfare measured in dollars.

Again, use the backward induction method to get a game theory prediction. B has a dominant strategy: no matter what x is, B's payoff is maximized by choosing $y = 0$.

If A correctly anticipates that B will use her dominant strategy, A will send $x = 0$.

The outcome is that each player gets 10. This is often called the "no trust" equilibrium.

Is this Pareto efficient? No. To see why, we will construct a Pareto improvement.

Suppose $x = 10$, which is the maximum A can send. This implies the total payoff is $20 + 2(10) = 40$. B's choice of y divides up this amount. According to the rules, B can't send back more than $3x = 30$. If B sent back exactly 30, A would receive a payoff of 30 and B would receive a net payoff of 10. The result (30, 10) makes A better off and B no worse off, compared to the outcome (10, 10) in the no trust equilibrium.

More generally, A can send the maximum of $x = 10$, which ensures a total payoff of 40, and B can send back any amount that gives A something between 10 and 30. This makes both players better off than with (10, 10). To see what is going on, look at Figure 1. It is similar to the graphs we used for the Ellickson book.

What happens in experiments like this? Most A players send some positive amount $x > 0$. Many B players send back some positive amount $y > 0$.

The first observation may not be too surprising. If A expects B to send back any amount $y \geq x$ then it makes sense for A to send x. But why does B ever choose $y > 0$? And why does A expect B to do this?

As in earlier games, many B players may be concerned with fairness. Moreover, at the stage where B has to make a decision, A has already shown trust in B by sending some money. Many of the B players want to show that they are worthy of A's trust. Most of

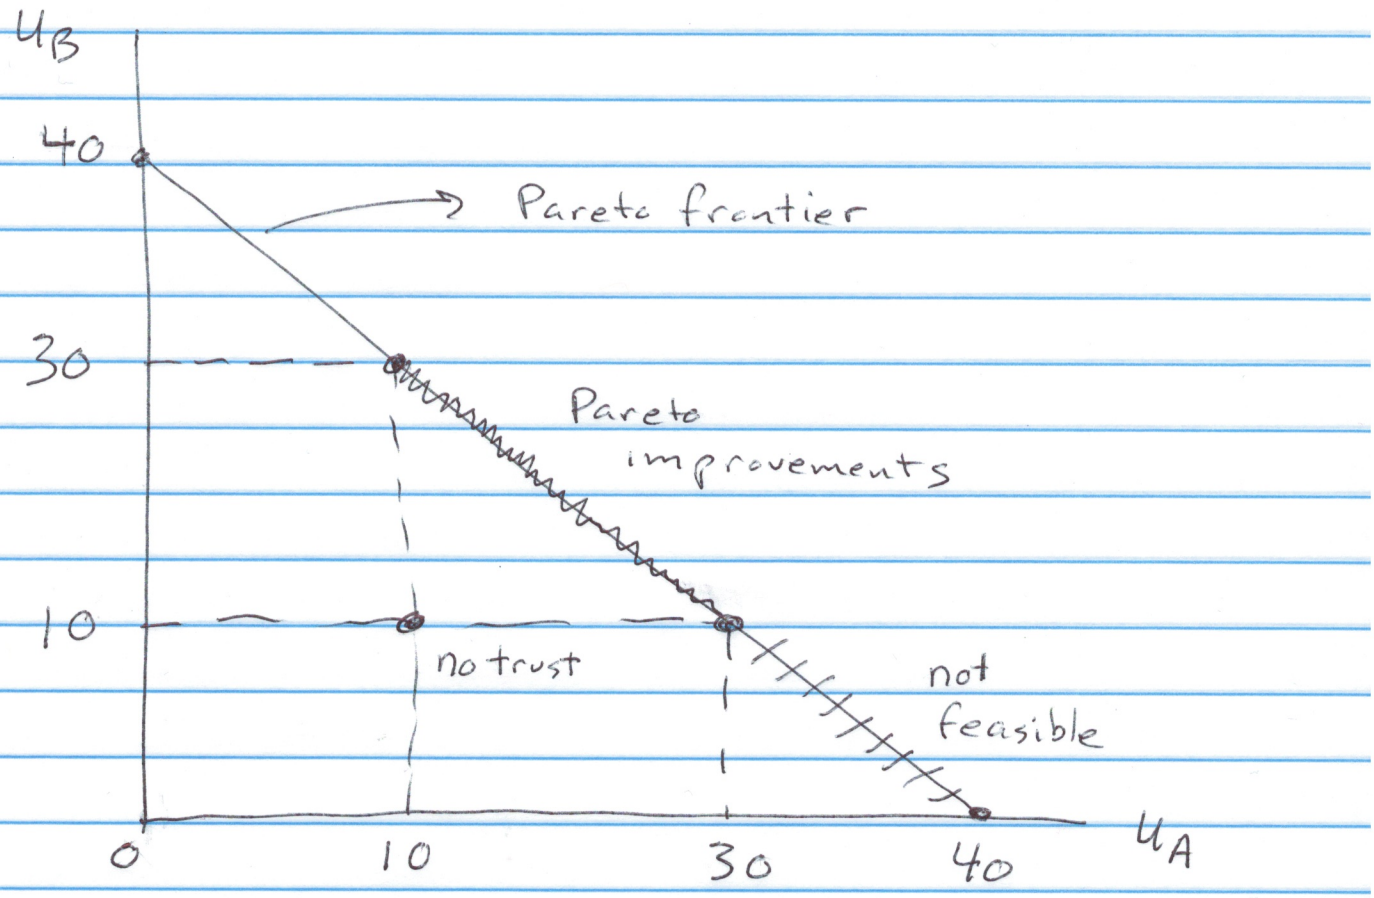

Figure 1

Pareto Improvements in The Trust Game

the A players anticipate (correctly) that many B players have motives of this kind, so A doesn't assume that B will be completely self-interested, and is willing to take a chance on B being trustworthy.

What about letting people discuss the situation in advance? Game theory says it should not matter. However, it does. In this case, there is more trust and more trustworthiness. When discussion is permitted, a majority of pairs agree to a (20, 20) split and B lives up to her promise.

LP draws several lessons from this.

1. In the real world, trust is cheaper than contracts and lawyers. Or to put it another way, trust reduces transaction costs. So it is good to encourage such attitudes or beliefs, if we can find ways to do it.

2. But we also know that people are heterogeneous. Not everyone can be trusted to cooperate. Some people are opportunists or free riders, and will be selfish.

3. So we need good institutions that either
   (a) screen out the opportunists in advance (so participants can be trusted); or
   (b) create incentive systems where the opportunists don't do very well.

4. In the real world, a good reputation is often valuable. When reputations exist, the opportunists face a choice: they can pretend to be nice, or they can take the short run gain while revealing that they are not nice and losing opportunities to interact with nice people in the future. LP says that in the laboratory when an opportunist lost their reputation, they ended up in the "dog group", where all the other players were also opportunists and everyone received a low payoff, because the players in that group did not cooperate with each other.

That's it for chapter 6.

Chapter 7 ("the economic rules of the game").

How did real economic institutions evolve? Why do we have the ones we have today?

Before getting into this, I will make a few preliminary comments. In chapters 8-10, LP will focus on global economic inequality. The idea will be that a constant human nature plus changing technology leads to changing institutions, causing changes in inequality.

A different way of explaining inequality (this is a popular story but LP rejects it) is to say that it results from differences in ability. E.g., smart people get rich, dumb people don't.

LP criticizes this idea because it ignores the demand side of the economy. What abilities are in demand in a given society? Hockey players? Electric guitarists? Brain surgeons? Traders in the stock market?

People having certain unusual skills may be well paid in our society, while in many other societies they would not be. The demand for a particular skill depends on a wide range of factors such as history, culture, technology, institutions, etc.

Now let's start from the beginning.

Hunter-gatherer societies were (and are) generally egalitarian. Why? Such societies are organized into small-scale social groups with limited specialization, not much storable wealth, and a high value placed on social cohesion. So there is very little inequality.

However, even HGs have a concept of possession. In fact, ownership or possession of resources is widely recognized in other animal species such as dogs.

[I will tell you a short personal story to illustrate this. I used to have a dog who liked to chew on beef-flavored sticks. He understood that the sticks belonged to me until I gave him one. After that, it belonged to him. If I tried to take away a stick, he would grip it tightly in his teeth, growl at me, and pretend to be angry until I let go. He clearly knew this was a game, and he liked to play it. He enjoyed pretending that he might bite me, although he never did. Of course, I always let him win, and he knew I would.]

LP wants to explain the origins of two economic institutions: property rights and trade.

To think about property rights, LP uses the Hawk-Dove game, which is famous among evolutionary biologists. There are two animals wanting the same food item or the same good location. Note that they are not literally hawks and doves (species of birds). The names are meant to suggest the difference between an aggressive strategy (hawk) and a peaceful strategy (dove). An animal using the hawk strategy is willing to fight over the resource, while an animal using the dove strategy is not willing to fight.

In my discussion I am going to use the payoff matrix in Figure 2. This is different from what LP does in the book, but I think my version makes things a bit clearer.

(A) \ (B)	hawk	dove
hawk	−1, −1	3, 0
dove	0, 3	1.5, 1.5

Figure 2

The Hawk - Dove Game

So now look at Figure 2. The general idea is that if both of the animals are hawks, they fight and could be injured, so they get (-1, -1). If one is a hawk while the other is a dove, the hawk gets the food for free, which is worth 3, while the dove gets nothing, which is worth 0. Finally, if both are doves, they share the food peacefully and get (1.5, 1.5). In my version of the game, the total payoff is always the value of the food (3) except in the case where fighting occurs.

You will see that if B is acting like a hawk, the best reply for A is to act like a dove (this gives 0 instead of -1). But if B is acting like a dove, the best reply for A is to act like a hawk (this gives 3 instead of 1.5).

Does A have a dominant strategy? No. The optimal strategy for A depends on what B does. Because the game is symmetric, B does not have a dominant strategy either. And therefore, there is no dominant strategy equilibrium (this is not a prisoner's dilemma).

We need to use a new equilibrium concept called *Nash equilibrium*. This is defined to be a combination of strategies, one for each player, such that neither of the players wants to change strategies, given what the other player is currently doing.

Or to put it another way: in a Nash equilibrium each player maximizes their own payoff, taking as given the strategy of the other player.

You can see from the payoff matrix that the Hawk-Dove game has two Nash equilibria. In the first, A is a hawk and B is a dove. This gives the payoffs (3, 0). Starting from this strategy combination, neither player wants to switch to a different strategy. In the second Nash equilibrium, A is a dove and B is a hawk. This gives payoffs (0, 3). Again, starting from this strategy combination, neither player wants to switch strategies. The question is: do the players have any way of coordinating on a particular Nash equilibrium?

A common solution to this problem is called the "bourgeois" strategy: if you arrived first, you act like a hawk, but if you arrived second, you act like a dove. Another name for this strategy is "finders, keepers". The player who already has possession is prepared to fight in order to keep the resource, while the other player leaves.

Biologists find that behavior of this kind is very common across many animal species.

Early humans lived in groups or bands (recall Part I of Johnson and Earle), so property rights over resources or territory would have involved ownership by groups, not single individuals. However, the general idea is the same.

Groups would often have had kinship ties with nearby groups (recall the discussion of kin selection from chapter 5 of LP). This would encourage sharing of food or territory. But two groups whose members were not genetically related would probably have had an "us versus them" attitude, leading to potential hostility in conflicts over scarce resources. In such cases, the Hawk-Dove game may be a reasonable description of how early property

rights were created. This idea links up with ideas about human biological evolution and human nature, plus anthropological observations of the HG societies that survive today.

New topic: trade. Once property rights exist, variations in natural resources can lead to variations in comparative advantage across different individuals or groups. In cases of this sort, trade can lead to Pareto improvements.

The easiest way to explain this is through the use of an *Edgeworth box*. You should look at Figure 3 as you read through the following explanation.

There are two individuals, A and B. There are two goods, 1 and 2. The total supply of good 1 is w_1 (the width of the box) and the total supply of good 2 is w_2 (the height of the box). E is called the endowment point, and it describes how the goods are distributed at the beginning of the game, based on initial property rights.

Person A is endowed with the quantity w_{A1} of good 1 and w_{A2} of good 2. To see person B's endowment, you have to treat the upper right corner of the box as the origin point for B. Using this origin, B is endowed with w_{B1} units of good 1 and w_{B2} units of good 2 (see the arrows, which show how quantities are measured for each person). Obviously $w_{A1} + w_{B1} = w_1$ and $w_{A2} + w_{B2} = w_2$ because the individual quantities for each good have to add up to the total supply.

Any arbitrary point x in the box describes some feasible allocation of the goods between the two people A and B, keeping the total supplies fixed. So we always have $x_{A1} + x_{B1} = w_1$ and $x_{A2} + x_{B2} = w_2$.

Each person has preferences over allocations of the goods, which can be described by indifference curves (recall that we introduced this concept in a model for Johnson and Earle). In Figure 4, I have drawn some indifference curves for each person so you can get the idea. The indifference curves for A bend in the usual way relative to A's origin point. The indifference curves for B bend in the same way, as long as you remember to treat the upper right corner of the box as the origin for B.

Why do the indifference curves look like this? First, remember that all points along the same indifference curve are equally desirable for that person (they give the same utility level). Because 1 and 2 are goods, keeping a person equally well off means that as we add more of one good, we must subtract some of the other. Thus, the ICs slope down. Second, we are assuming diminishing marginal utility for each good, which gives the convex curvature (again, taking into account the difference in the origin for A and B).

Figure 4 shows the indifference curves for A and B passing through the endowment point E. These are labeled u_A^0 and u_B^0 respectively. You can think about these labels as giving the initial utility levels for each person. I have also indicated a shaded area where both A and B would be better off than they are at point E (in the shaded area, each person would be on a higher indifference curve and have more utility).

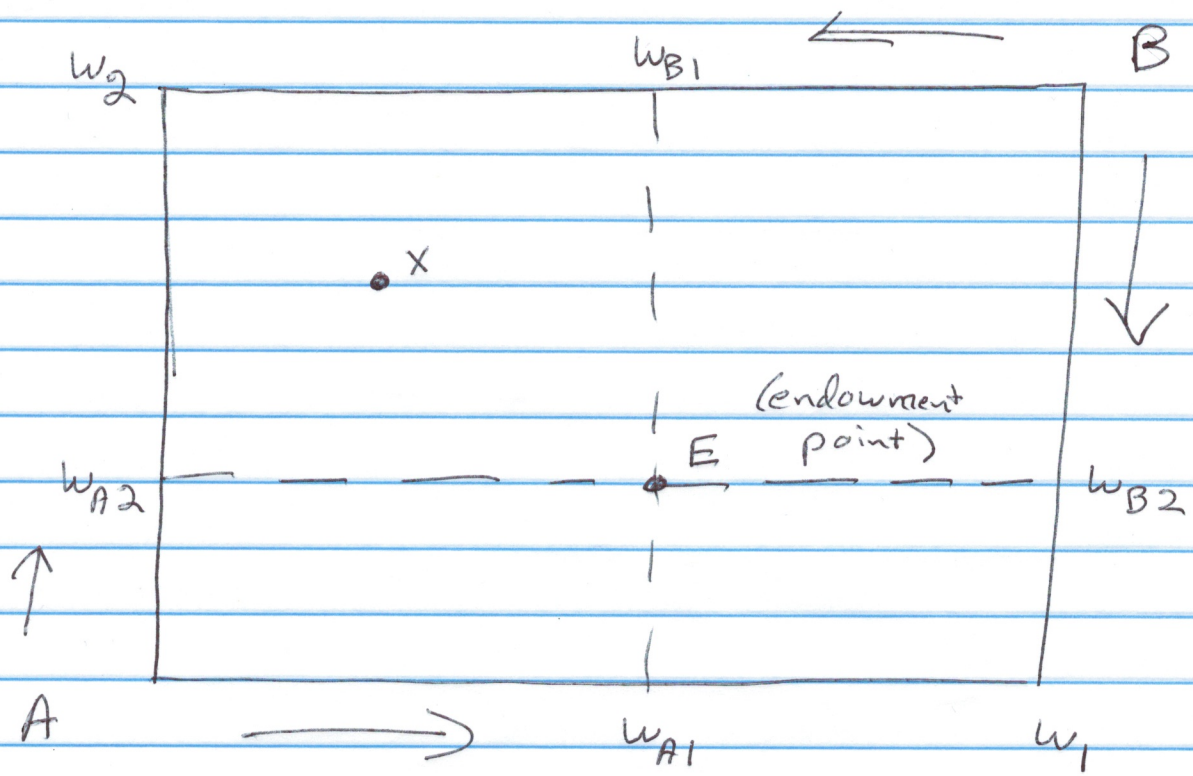

Figure 3

Edgeworth Box with Endowment Point

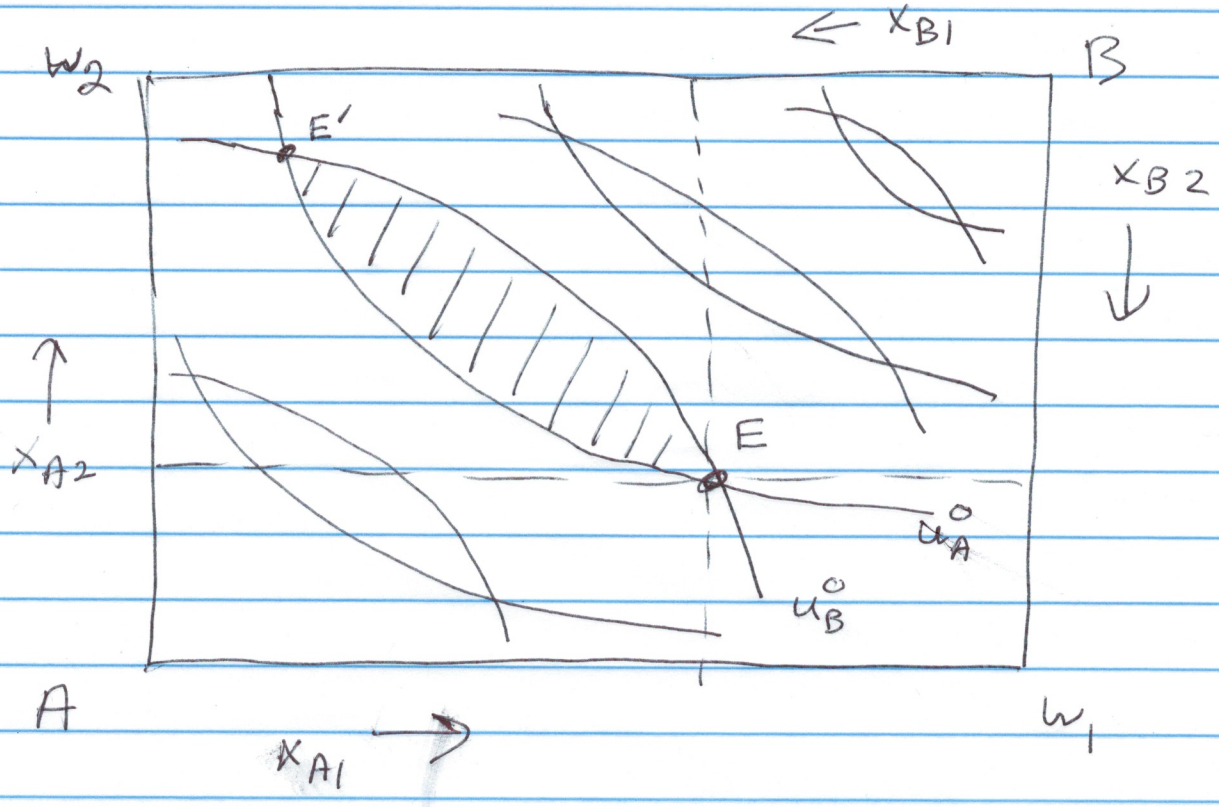

Figure 4

An Edgeworth Box with Indifference Curves

If the graph looks like this, can we get a Pareto improvement starting from E? Yes. Just go to any point in the shaded area, where both people are better off. Notice that this will work any time the IC for B through point E is steeper than the IC for A through point E.

A similar argument works when A's IC through the endowment point is steeper than B's IC through the endowment. To see why, imagine that the endowment is located instead at point E' in Figure 4, and consider the effect of moving into the shaded area.

Now we know two things. First, any time we can get a Pareto improvement, we are not yet being Pareto efficient. Second, we can always get a Pareto improvement if we are at a point in the box where the slopes of the indifference curves are not equal to each other.

Putting these together, a necessary condition for Pareto efficiency is that we must be at a point in the box where the slopes of the indifference curves through that point are equal. In general, there are many such points (see Figure 5). One example is point x*, where A receives the utility level u_A* and B receives the utility level u_B*. Starting from point x*, there is no way to obtain a Pareto improvement; any further change would make at least one person worse off.

The set of all points where the slopes of the indifference curves are equal is called the *contract curve*. This is shown in Figure 5. In general, we don't know much about the shape of the contract curve: it may be linear, it may be curved, it may pass through the origin points at the corners of the box or it may not, etc. These details depend upon the preferences A and B have, which will determine their indifference curves.

The idea of trade is that A and B start from some endowment point, which usually is not on the contract curve. As they trade, they move to new allocations in the box and obtain Pareto improvements. We know these movements result in Pareto improvements because each trade is voluntary so neither person has to accept any change that makes them worse off. Eventually they reach a point along the contract curve. Starting from this point there cannot be any further Pareto improvements, so trade stops, and we have a Pareto efficient allocation of the goods.

A few points. First, we are not saying much about which point on the contract curve is the final outcome. This generally depends on the bargaining abilities of the two people. All we know is that neither will be worse off than they were at the endowment point.

Second, A and B could be groups rather than individuals. The ideas would be the same as long as each group has a meaningful set of indifference curves.

Third, we are not relying here on Ellickson's idea of aggregate welfare. There is no need to add up the utilities of individuals A and B. All we require is for each of the individuals to have a consistent set of preferences about their own consumption of the two goods.

Fourth, we could add production activities to the model, although of course this makes everything more complicated.

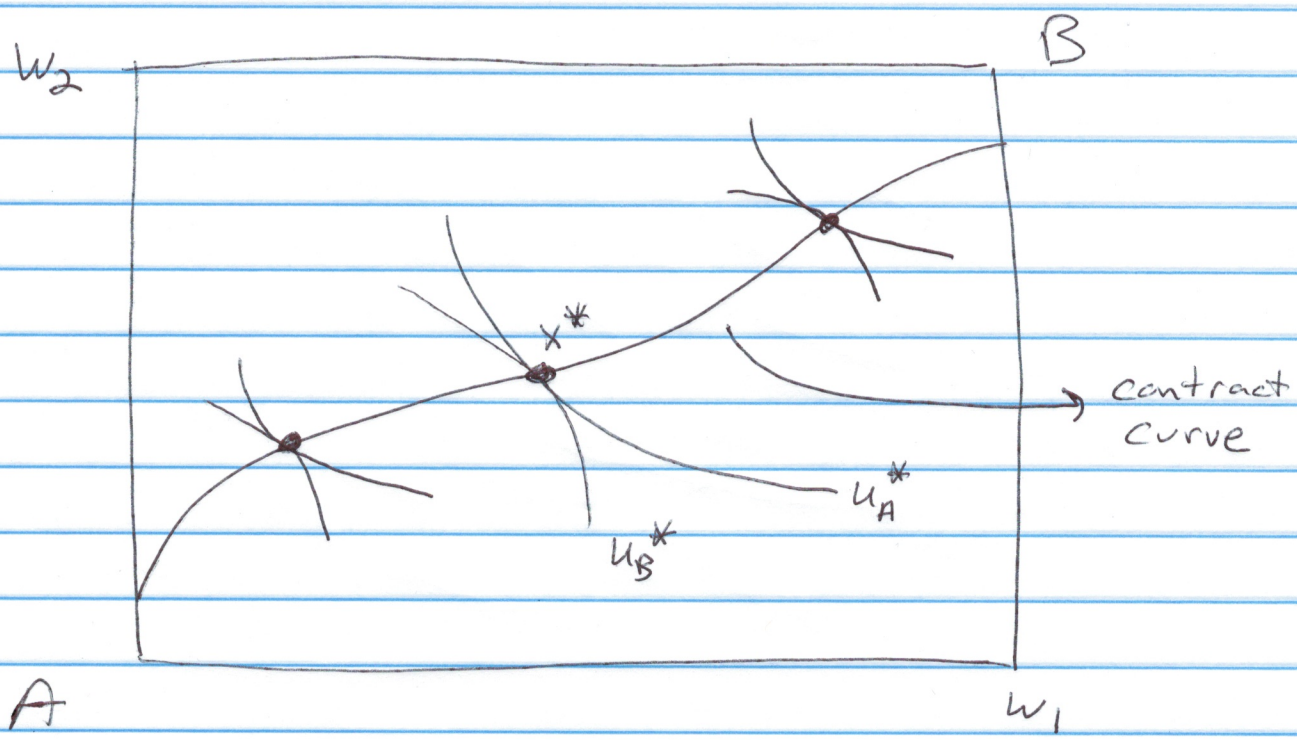

Figure 5

An Edgeworth Box Showing The Contract Curve

Fifth, if we had many individuals of type A and many individuals of type B, and all of these individuals were competing with each other in the trading process, we could talk about supply and demand curves for the goods, equilibrium prices of the goods, etc.

Bottom line: the Hawk-Dove game gives us a way of thinking about how property rights could have been established and the Edgeworth box gives us a way of thinking about how early trade could have occurred. LP believes that property rights and trade are very old hunter-gatherer institutions.

Now I want to give two examples where LP argues that changes in technology led to changes in institutions, and both of these together led to changes in inequality.

The agricultural revolution. This was a change in technology. LP thinks it caused a shift in property rights from the group level to the household level. Why?

1. Technology: early farming did not involve significant scale economies. It was an activity that could be carried out on the household level. {Note: there are a couple of issues with this. First, Johnson and Earle might disagree, because they believe technological innovation was one reason for larger-scale groups to develop. Also, LP does not explain why hunter-gatherer societies would have had group property rights -- did HG societies have scale economies in their food collection activities? Maybe this had some connection to hunting in teams, or sharing of food?]

2. Incentives: there were smaller free rider problems within families. Also, maybe it was necessary to have household-level property rights in order to motivate people to carry out land clearance, tilling, weeding, and other agricultural activities.

The effect on inequality: material assets like land or animals could be accumulated and inherited within families. Powerful individuals or families accumulated greater wealth over time and defended it. This led to the use of military force, stratification, and so on.

LP skips very quickly over the Johnson and Earle story about the evolution of chiefdoms and the state. His story would probably be similar, except that he doesn't deal much with the fact that elites are social groups, not just wealthy individuals or households.

The industrial revolution. Again this is a change in technology, and again LP thinks it caused both a change in institutions and in the degree of inequality.

1. Industrialization led to much more specialization and much more use of markets. In such societies, almost everyone specializes in some form of labor, gets income from supplying labor, and buys most of the consumption goods they want.

2. Scale economies were often important for industrial production. As a result, the markets for products had to be large geographically and/or in terms of population.

Firms also had to develop new types of organization, which raised new incentive issues (production in large teams and so on).

3. Property rights became increasingly individualistic, rather than remaining at the household level. For example, individual people could own shares in firms.

Other effects included more pressure for political democracy, as well as more separation between work and other parts of life (consumption, residence, family, etc.).

LP's story about why inequality increased as a result of the industrial revolution:

1. Specialization of work roles tended to split up bundles of skills into individual skills, where each skill had a separate wage based on supply and demand.

2. Financial wealth became more important. There is no limit on the amount of financial wealth an individual person can accumulate.

LP believes that since the industrial revolution, private property and market exchange have led to faster technological change and increasing productivity. He regards this as a good thing (it channels self-interest in socially useful directions).

But in market economies self-interest has costs such as fraud (for example, the financial meltdown of 2008-9 had a lot to do with fraud in the financial markets). Also there may be a need for very detailed contracts and excessive use of lawyers (these are institutional safeguards against opportunism, but they have high transaction costs).

Trust and reputation are nice alternatives when possible, but they don't always work (for instance, people may not have detailed enough information about the previous behavior of other people, or may not put enough weight on future interactions).

AND: markets don't solve problems involving public goods. These can include national defense, crime prevention, law and order (including protection of physical safety as well as property rights), environmental protection, quality of food and drugs, basic scientific research, etc.

Self-interest leads to large free rider problems in supplying public goods of this kind. In order to overcome these free rider problems, we need governments with the power to tax and regulate economic activities.

Governments also play major roles in education, health care, and infrastructure. These sectors involve large positive externalities, and they also contribute to productivity and economic growth within the private sector.

Econ 354

Greg Dow

April 3, 2021

Notes on Louis Putterman, "The Good, the Bad, and the Economy"

Chapters 8-10 (Global Economic Inequality)

This is the last set of lecture notes on the book by Louis Putterman.

Chapter 8 ("The Origins of Global Inequality").

In this chapter, LP is primarily concerned with the differences in per capita income across nations (why some nations are rich while others are poor).

Another important set of issues involves inequality among individuals within each nation, but LP has less to say about that.

There are a number of common hypotheses about the sources of global inequality.

1. Exploitation (rich nations exploit poor ones, or did in the past).
2. Luck (rich nations happened to have important technological innovations).
3. Geography and climate (rich nations happen to be located in better places).
4. Biology (racist ideas about the inherent superiority of various human groups).

LP takes a historical approach to the causes of modern economic inequality. He wants to know how it developed from roots in the distant past. He makes the point that important technological and institutional innovations did not occur everywhere simultaneously, and he argues that the differences in timing help to explain global inequality today. He thinks geography and climate played a role, but history also mattered.

Especially, we need to look at two changes in technology (the agricultural and industrial revolutions) and early changes in institutions (such as the emergence of cities and states).

Around 20,000 - 13,000 years ago, everyone on Earth was a hunter-gatherer. There were no significant differences in technology (only minor differences in techniques, depending on the local natural resources); minimal differences in the sizes of the foraging bands and in population densities (although again with some minor differences in population density according to the richness of local natural resources); and minimal differences in standards of living across different regions of the world.

Around 10,000 years ago, early transitions to agriculture began in the Middle East, India, and China. These developments later spread into Europe and North Africa. Independent forms of agriculture developed in Mexico, South America, and a few other places.

[Why did this happen? We would need another course to explain the locations and the timing of the agricultural revolution. In my view, this was mainly about geography and climate change. The key point is that some regions shifted to farming at an early stage, others did later, and some never did.]

Around 5000 years ago, we started to see the development of cities, states, writing, metal working, professional armies, and greater long-distance trade. An important point here is that there is a link from technology to population to institutions. Regions that had early farming also had substantial population growth and developed state institutions sooner.

Meanwhile, many other regions of the world did not have agriculture at all, or it arrived relatively late from elsewhere. Some regions did have agriculture but did not have states. Some regions had both but were relatively isolated from events in Asia and Europe (such as sub-Saharan Africa, the Americas, and Australia).

By about 1500 AD, there was a set of societies extending from western Europe through the Mediterranean, the Middle East, India, China, and Japan that had been exchanging ideas about technology, institutions, and culture for thousands of years. This occurred along trade routes, through migration of people, through military conquests, and so on.

These societies tended to have early agriculture, early states, high populations, towns and cities, and advanced military technology. LP believes these early advantages from before the industrial revolution have had persistent effects in the modern world.

At this point I want to spend some time discussing the Malthusian model of population dynamics. This comes from an English economist, Thomas Malthus, who published a famous book about population in 1798 (note that this date is around the same time that the industrial revolution was getting underway in England, although this was probably not obvious to Malthus at the time).

The key point: before the industrial revolution, improvements in technology tended to raise population in the long run, but not the standard of living. The Malthusian model can explain why this was true.

Suppose that when an adult has more food, s/he has more kids who survive to become adults. This occurs for several reasons: women with better nutrition have higher fertility rates, both men and women with better nutrition have lower mortality rates, and kids with better nutrition have lower infant and childhood mortality rates.

Let's work out the economic implications of these ideas using a series of graphs. Figure 1 shows food per adult (y) on the horizontal axis and the number of surviving kids per adult (r) on the vertical axis. For reasons described above, we have an increasing function r(y), so as y goes up, r also goes up. Note that if food per adult is equal to y*, we have r(y*) = 1, so on average each adult will be replaced by one surviving adult in the next generation. In this case, the population stays constant from one generation to the next. If $y > y^*$ then

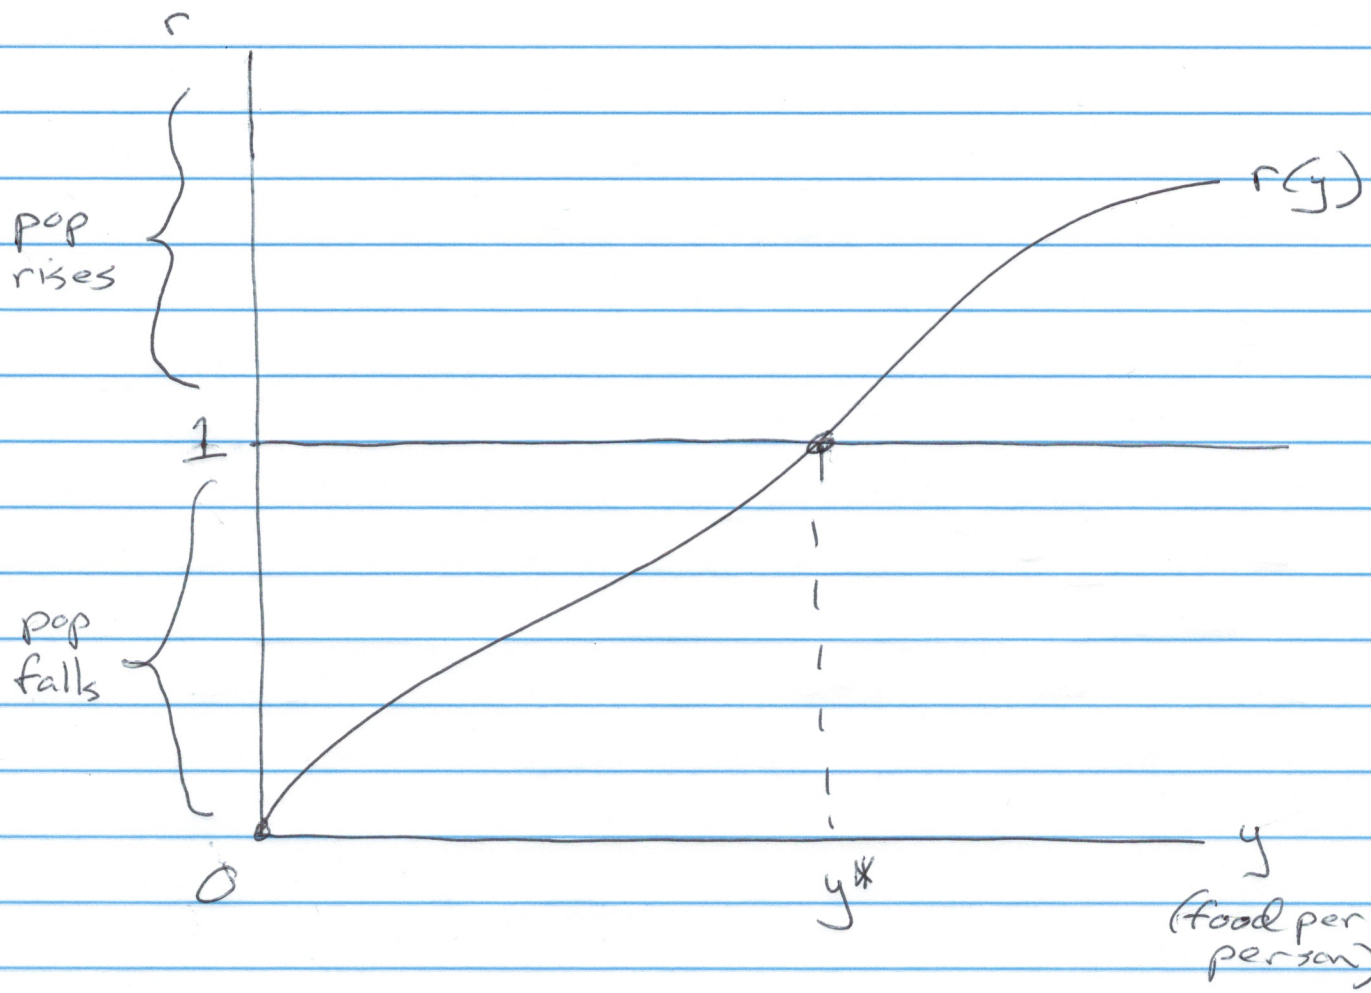

Figure 1

Food Per Person and Surviving Kids Per Person

$r(y) > 1$ and the population increases over time. If $y < y^*$ then $r(y) < 1$ and the population decreases over time.

The next step in the argument involves the relationship of total population (N) to food per person (y). Figure 2 has total population N on the horizontal axis. We think of this as the number of adults, and also as total labor supply (assume each adult has one unit of labor time). The vertical axis shows total food output (Y). There is an increasing relationship between N and Y. With more labor, the society can produce more total food. However, this curve is concave due to diminishing returns. Because inputs of land and other natural resources are fixed, output increases when labor inputs increase, but at a decreasing rate.

Recall that the average product of labor (APL) is defined to be Y/N (total output of food divided by total input of labor). In Figure 2, APL is indicated by the slope of a ray from the origin to a point on the Y(N) curve. For example, when labor input is N_0 the APL is Y_0/N_0, which is the slope of the dashed line from the origin to the corresponding point on the curve. As you can see, diminishing returns implies that as N goes up, the slope of the ray from the origin goes down, so APL must decrease.

Food per person is $y = Y/N = APL$. So Figure 2 says that if we have diminishing returns, food per person (y) will decline as the population (N) rises.

Finally, go to Figure 3 where the horizontal axis shows the total population (N) and the vertical axis shows food per person (y). Ignore the dashed curve $y_1(N)$ for the moment and focus on the solid curve $y_0(N)$. Suppose $y_0(N)$ is obtained from a fixed production curve as in Figure 2. Now consider how population will be determined in the long run.

To have an equilibrium population in the long run, from Figure 1 food per capita must equal y^*. Otherwise, population would either be rising or falling. From Figure 3, this implies a population equal to N^*. At N^* we have $y = y^*$ and thus $r = 1$, so population is constant. If $N < N^*$ then $y(N) > y^*$, so $r(y) > 1$ in Figure 1 and population rises. If $N > N^*$ then $y(N) < y^*$, so $r(y) < 1$ in Figure 1 and population falls. In either of these cases, we move toward N^*, which is therefore a stable equilibrium population level.

Now suppose technology improves, so it becomes possible to produce more food from any given level of labor input (we continue to assume land and other natural resources are fixed). This shifts up the entire production curve Y(N) in Figure 2 and also shifts up the entire y(N) curve in Figure 3 (at any given level of N, we now have more Y, so the ratio $Y/N = APL = y$ must rise). This causes a shift from the solid curve $y_0(N)$ to the dashed curve $y_1(N)$ in Figure 3.

In the short run, before population has time to change, we stay at N^* but move vertically from y^* to y' as shown in Figure 3. So the short run effect of technological innovation is to raise the standard of living (food per person). But in the long run, we are now on the new curve $y_1(N)$ with $y' > y^*$. This leads to population growth for the reasons in Figure 1. As population rises, we move down the $y_1(N)$ curve and the standard of living y falls, as shown by the arrows in Figure 3. This continues until we reach the new population

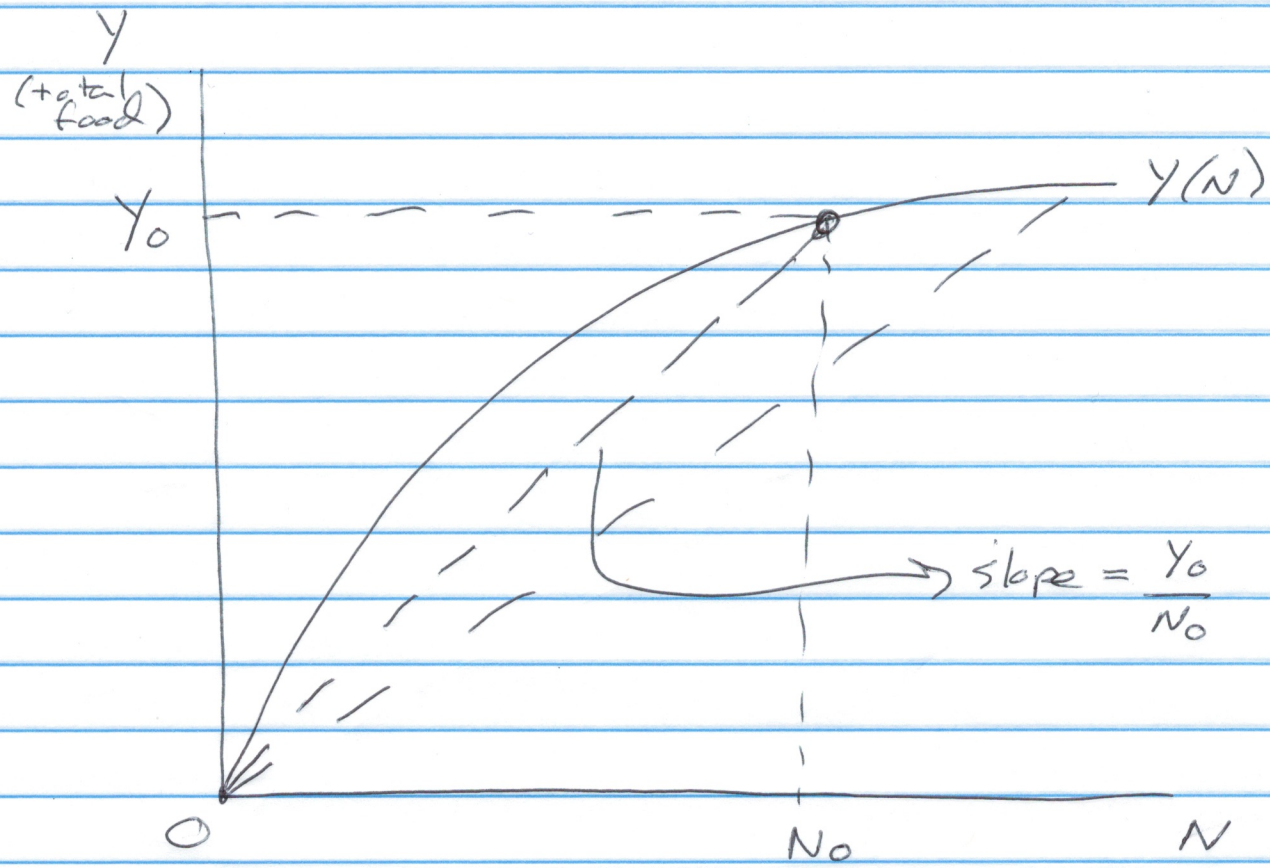

Figure 2

Relationship Between Population and Food Output

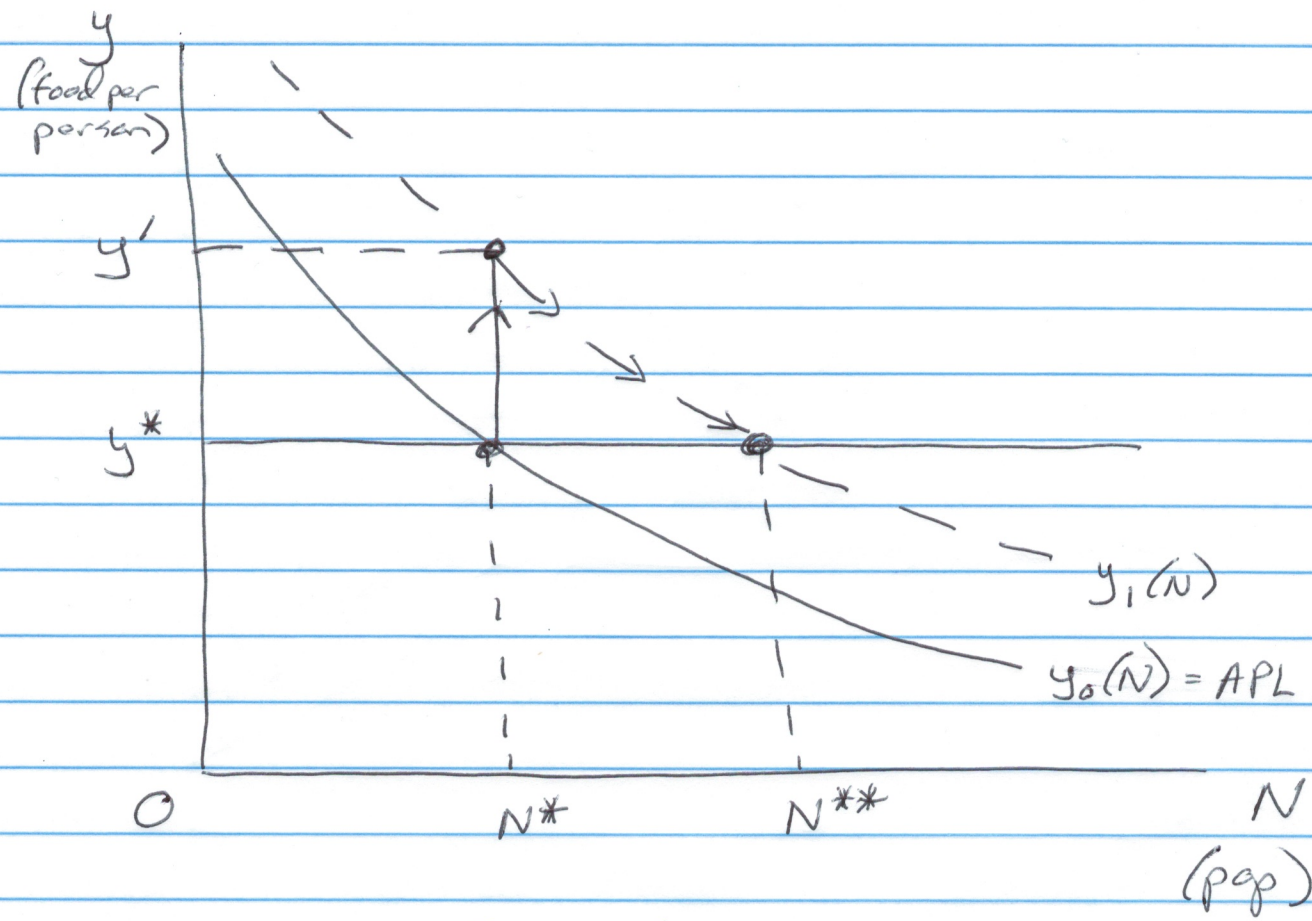

Figure 3

Technological Progress, Population, and Food per Person (Malthusian Model)

N**, where we get back to y*. At this point, population reaches a new equilibrium and growth stops. We have now returned to the original standard of living y*.

Thus in the Malthusian model the long run effect of technological progress is to increase population, while leaving the standard of living unchanged. The same would be true if a change in some other exogenous variable shifted up output per unit of labor (for example, a better climate).

The whole process would operate in reverse if technology went backward, or the climate became worse. In such cases, the short run effect would be for people to become poorer. The long run effect would be a decrease in population, with the standard of living rising until it returned to where it was in the original equilibrium.

I have ignored inequality in the Malthusian model, but it could be included. In the early states, there were generally a few very rich elite people and a lot of poor commoners. In this situation, technological progress could increase total population while making elites richer and commoners poorer, with long run food per person staying at y* on average.

There is a lot of evidence that the Malthusian model is a good description of the world before the industrial revolution. Almost all economists who do research in this area are willing to accept some version of this model (perhaps a more complex version than the one I have described here).

Many people (both archaeologists and economists interested in prehistory) think that the agricultural revolution made people worse off relative to their hunter-gatherer ancestors (poorer diet, poorer health, less leisure, shorter life expectancy). How could this happen when the Malthusian model says y* should have remained constant? The answer is that early agriculture could have shifted up the r(y) curve in Figure 1, which implies a lower level of y* (keeping in mind that in the long run, an equilibrium population must have r = 1). I won't go into the reasons for the shifting r(y) curve in detail, but it is related to the more sedentary agricultural lifestyle and probably an increase in the economic value of children in an agricultural society.

The key point is that since the industrial revolution, population dynamics have changed, and the Malthusian model no longer applies today.

This has involved a *demographic transition*: in the modern world, when people get richer they tend to have fewer kids, not more kids as the Malthusian model assumes.

Why? There are several theories about how the demographic transition happened. These theories are not mutually exclusive; several of them could be true simultaneously.

1. Kids are a direct economic asset for their parents in an agricultural society, but this is less true in an industrial society. Thus there is less incentive to have large families. One result is that societies develop laws against child labor and require kids to be educated.

2. Parents want to invest more in the education of each kid, so they want a smaller total number of kids.

3. Women join the industrial labor force, which increases the opportunity cost of having kids (women have to give up some wage income in order to have them).

4. Contraceptive technology has improved a lot.

5. Government programs now provide support to old people, so it is less necessary for parents to have numerous kids in order to ensure economic support when old.

For some combination of these reasons, what starts to happen is that in rich countries the rate of growth in productivity due to technological progress becomes greater than the rate of growth in population. The difference in growth rates leads to rising income per person and therefore a rising standard of living.

At a global level, the rate of world population growth is now slowing down. Depending on which experts you believe, world population might stabilize at an equilibrium level (a zero growth rate) between the years 2050-2100. Beyond that point, any further technical progress would go entirely to raising incomes per person, not population. This assumes that technological innovations are not canceled out by the depletion of natural resources or a deteriorating environment.

Now go back to before the industrial revolution (by the way, most economic historians say the IR started around 1750-1800 in England, so it is still a relatively recent thing).

With Malthusian dynamics, even though incomes didn't rise much, there was a positive feedback loop where higher population led to faster technological progress (more people meant more innovators), and better technology led to higher population (due to Malthus).

This feedback loop was especially important in Eurasia (the land mass consisting of Asia and Europe together). This was the largest land area on the planet, it had the most total people, and it had the greatest opportunities for the spread of ideas.

Other regions like sub-Saharan Africa, the Americas, and Australia were more isolated, had less advanced technology, and had lower population densities. The technology that did exist was generally invented locally rather than borrowed from somewhere else.

When Europeans started colonizing the world around 1500 AD, such regions were often easy targets. According to LP, what else explains who got colonized?

1. Places with especially low population densities (e.g., hunter-gatherer societies).
2. Disease patterns (in some places, the local people had essentially no immunity to European diseases, resulting in massive loss in local populations; in other places, settlement was limited because Europeans were vulnerable to the local diseases).

3. Climate (some parts of the world had similar environments to those of Europe, so domesticated plants and animals from Europe were successful there).
4. Geography (some places had valuable agricultural land, other valuable resources, and were easy to find or travel to).

[I would add one more factor. It was easier to colonize areas where people did not have powerful political institutions like chiefdoms or states that could organize resistance.]

Much early European colonization occurred before the industrial revolution, which as I mentioned above, did not begin until around 1750-1800. There has been a lot of debate about whether colonization helped stimulate the industrial revolution, whether colonies were a net economic gain or loss for the colonizing powers, and so on. I will not pursue those issues here.

One interesting point: Europeans had more trouble colonizing Asia than most other parts of the world, and didn't make significant inroads there until after the industrial revolution was underway. They never did colonize Japan and had limited success in China. This is consistent with LP's general view that across Eurasia, a number of distinct regions all had relatively high population density, relatively advanced technology and relatively complex state institutions. Thus, it was difficult for one part of Eurasia to colonize the other parts.

That's all for chapter 8.

Chapter 9 ("Winners and losers in the race to industrialize").

Before the IR, inequality in per capita incomes across regions of the world was small, despite large differences in technology. Malthus explains why: good technology was offset by high population.

After the IR, there were two main effects.

1. Regions that had early technological and institutional advantages before the IR tended to industrialize sooner.

2. The demographic transition meant that places with early industrialization started to have fewer kids sooner, and as a result their incomes started to rise sooner.

This is reflected in ratios of incomes between the richest and poorest countries. In 1500, this ratio is estimated at 3:1. In 1870, it was 9:1. In 1960, it was 38:1, and in 1990 it was 45:1. According to LP, this growth in inequality reflects a combination of the two factors mentioned above: differences in the timing of industrialization, and the differences in the timing of the demographic transition.

However, the demographic transition is ultimately a good thing for two reasons:

(a) The demographic transition makes rising incomes possible. Probably for the first time since foraging societies, a majority of people can have better nutrition, better health, more comfort, more leisure, and longer life expectancy. The IR would not have led to these outcomes if the population had continued to increase as rapidly as productivity.

(b) The demographic transition alleviates human pressure on the natural environment because it implies a lower human population relative to finite natural resources, as compared with what would have happened without the transition.

Nevertheless, this transition has led to more inequality across nations at least temporarily, because different nations have gone through the transition at different points in time.

There is a lot of debate about why the IR happened first in England, rather than China or other places. Was it due to a difference in resources? A difference in institutions? This is similar to asking why agriculture began in certain specific places. It is a hard question.

Historically speaking, industrialization spread first to nations that were similar to England in technology, institutions, and culture, and to those that were geographically close. This may have been due in part to competitive pressures among nations within Europe. It then spread to non-European countries with a long history of agricultural technology and state institutions, such as Japan and other areas within East Asia. Industrialization did not tend to "leap frog" to areas without these favorable preconditions.

Summing up LP's view:

Places with early agriculture tended to have early states. And places having both early agriculture and early states tended to have several advantages around 1500 AD: more urbanization, higher population density, and higher incomes (at least for the elites).

Countries that were ahead in 1500 AD tend to have either the highest incomes today, or the highest rates of economic growth (they are catching up to the leaders).

When looking at incomes today, these correlations become stronger if we adjust for the migration patterns since 1500 AD. LP concludes from this that skills, norms, and culture are important for economic growth; it is not just a matter of natural resources and climate.

LP then goes into a discussion of various development strategies that have been tried over the last century or so. These include

(a) Import substitution (protection to domestic industries facing foreign competition; the problem is that the protectionism tends to be permanent, resulting in industries that can't compete with the rest of the world).
(b) Heavy industry first (popular with central planners in the former Soviet Union).
(c) Export-oriented development (this strategy has generally worked best).

LP emphasizes that export-oriented development needs to be combined with government investments in infrastructure and education, plus political stability and the rule of law. In other words, institutions matter a lot.

This is a common view today among development economists. You may want to read a book by Acemoglu and Robinson called "Why Nations Fail" (2012). The authors are a leading economist and a leading political scientist. They provide many case studies to show that good institutions are essential for sustained economic growth.

Although we usually think of markets and governments as doing different things, and in political debates they are often treated as substitutes, there is considerable evidence that in the real world they are complements. If you want to encourage private sector growth, you should also support growth-enhancing investments by the public sector.

On days when I am feeling optimistic, I think there may be a feedback loop where good institutions tend to encourage economic development, and economic development tends to encourage good institutions. The second part of this loop could arise through more education, the growth of cities, access to better communication technology, and so on. These things tend to put pressure on elites to reform and open up the political system.

That's it for chapter 9.

Chapter 10 ("Hope for a Better World?").

What do we need for a better world, or "progress" in a broad sense?

LP and I have a similar view: we need both better technology and better institutions (each of these is necessary and neither alone is sufficient).

Historically, institutional innovation has been just as important for the quality of human life as technological innovation.

And institutions don't improve all by themselves. Groups of people have to organize, experiment, and work hard, and they often have to fight against oppression by elites.

Let's go back to human nature. LP would say that key economic institutions like property rights and markets tap into self-interest, and can channel it into socially useful directions, like investments in physical capital, investments in human capital, and also technological innovation. This is good. Without these things, there would be little hope of increasing living standards for the majority of the world's population.

Another essential thing is the demographic transition (the end of Malthusian population dynamics). There cannot be significant progress for the majority of the population if the gains from technology and investment are canceled out by population growth.

BUT: many people are still poor, with large inequalities both within and across nations.

There are some grounds for optimism. As LP stressed in earlier chapters, there are non-selfish aspects of human nature. He thinks that democracy, civil liberties, and equality under the law tap into deep feelings about fairness, so these institutions are consistent with human nature and can be expected to spread, at least under favorable conditions.

Here is a list of some ways in which the world has made institutional progress over the last century or two, at least in many countries:

1. Slavery has become illegal almost everywhere.
2. Former colonies have become independent almost everywhere.
3. More nations have at least some degree of political democracy.
4. Women are increasingly gaining equal rights with men.
5. Most legal forms of racial discrimination have been abolished.
6. Movements for environmental protection have become stronger.

In addition we have created global institutions to deal with a number of specific problems like security, trade, development, health, and the environment.

LP asks: can institutional changes reduce the poverty and negative externalities resulting from markets? He thinks so, and points out that governments in rich countries tend to do

a lot for their own citizens. For example, they usually provide social insurance, health care, education, sanitation, environmental regulation, and so on.

Furthermore, institutions can significantly diminish poverty by tapping into feelings of fairness, empathy, and reciprocity. However, "us versus them" attitudes are a problem. They tend to limit willingness to help the poor, especially when the poor are defined as members of "other groups", either within nations or internationally.

On the problem of global economic inequality:

LP thinks it is unlikely that people in rich countries will simply donate enough money to lift people out of poverty in poor countries.

But he argues that people are often willing to supply assistance aimed at specific goals:

(a) Raising productivity (better health and education, plus technology transfers)
(b) Lowering trade barriers, so poor countries can export more to rich countries
(c) Improving political institutions
(d) Providing loans to entrepreneurs

[Question: if we accept LP's view of human nature, why are some people willing to do these things? Is it reciprocity? That seems unlikely. But maybe it is compassion.]

The bottom line for LP: Institutions need to be designed and improved in light of realism about both the good parts and bad parts of human nature.

I would add a couple of other things. First, modern institutions are vastly more complex than anything our hunter-gatherer ancestors had. As I have mentioned previously, I find it amazing that our species evolved in small foraging bands but can somehow manage to organize enormous modern societies. However, there are reasons to believe institutional innovation will continue, and that's an encouraging thought.

Second, human nature certainly has flaws but it has gotten us this far. We don't want to overestimate what humans can accomplish, but we don't want to underestimate it either.

I will close with a few insights I hope you have gained from this course.

1. Human society has not always looked the way it does today.
2. Institutions are important, complicated, and can be improved.
3. Economics is useful for understanding how institutions work, and maybe also for understanding how they can be improved.