Prof. Bryan Caplan
bcaplan@gmu.edu
http://www.bcaplan.com
Econ 854

Week 1: The Logic of Collective Action

I. The Many Meanings of Efficiency
   A. The *Merriam-Webster College Dictionary* defines "efficiency" as "effective operation as measured by a comparison of production with cost (as in energy, time, and money)."
   B. Economists occasionally do use "efficiency" in the dictionary sense - ratio of the value of output to input or something similar.
   C. But normally they use it in quite different ways, and unfortunately often equivocate between the various uses.
   D. The two most common uses in economics are:
      1. Pareto efficiency
      2. Kaldor-Hicks (or cost-benefit) efficiency

II. Pareto Efficiency, I
   A. Most of the famous theorems in welfare economics discuss Pareto efficiency.
   B. A situation is **Pareto efficient** iff the only way to make one person better off is to make another person worse off.
   C. Similarly, a **Pareto improvement** is any change that makes someone better off without making anyone else worse off.
   D. In theory, it is quite possible that people will voice objections to Pareto improvements for *strategic reasons*. So it is <u>not</u> equivalent to a demonstrated preference standard.
   E. In a highly stylized theoretical setting, Pareto improvements are conceivable. Ex: If everyone has identical preferences and endowments.

III. Pareto Efficiency, II
   A. Even so, there is a strong argument that, in the real world:
      1. Everything is Pareto efficient.
      2. Pareto improvements are impossible.
   B. Why? Almost any change hurts someone, and it is highly unlikely in practice that literally everyone can be compensated, that absolutely no one will be missed.
   C. Ex: I buy your watch. How will we compensate everyone who might have asked you the time?
   D. Rothbard's strange variant: Only count "demonstrated preferences." Then Pareto improvements happen all the time. But especially for an Austrian, this is bizarrely behavioristic.
   E. More fruitful variant: Analyze the Pareto efficiency of ex ante *rules* instead of ex post results. (This is the key intuition behind a lot of constitutional economics). But even then, someone is very likely to slip through the cracks.

IV. Kaldor-Hicks Efficiency, I

A. In practice, then, economists almost always switch to Kaldor-Hicks efficiency, aka "cost-benefit efficiency."

B. A situation is **Kaldor-Hicks efficient** iff the dollar value of social resources is maximized.

C. A **Kaldor-Hicks improvement** is any change than *raises* the dollar value of social resources.

D. Every Kaldor-Hicks efficient situation is Pareto efficient, but most Pareto efficient situations are NOT Kaldor-Hicks efficient.

E. Ex: You value a watch at $20, I value it at $30, the strangers you will encounter value my having the watch at $.10, the (different) strangers I will encounter value my having the watch at $.10.

1. If I have the watch, the situation is K-H and Pareto efficient.
2. If you have the watch, the situation is Pareto but not K-H efficient. Social value on the watch rises from $20.10 to $30.10, but your time-askers lose $.10.

F. Every Pareto improvement is a Kaldor-Hicks improvement, but most Kaldor-Hicks improvements are not Pareto improvements. (Return to above example).

G. K-H efficiency is often described as "potentially Pareto efficient" because if the value of social resources rises, then (assuming perfect continuity), you *could* compensate all of the losers by sharing the gain in surplus.

H. But what exactly does this "could" mean? Essentially, you could if transactions costs of arranging compensation were zero.

I. This bothers many people - why shouldn't the transactions costs count just as much as other costs? Ultimately, though, this is just another way of saying that Kaldor-Hicks improvements don't have to be Pareto improvements. No one said ever said they were.

1. When you judge whether something is a K-H improvement, you do count the transactions costs for the move itself.

V. Kaldor-Hicks Efficiency, II

A. K-H efficiency naturally gives rise to another concept: deadweight costs. If the value of social resources is not maximized, deadweight costs exist.

B. Everyone knows that you can *transfer* resources from one person to another. That's obvious.

C. Economists' marginal product: It is far less obvious that resources can be destroyed, leaving *no one* better off.

D. Ex: Piracy. It is obvious that pirates transfer treasure from victims to themselves. The deadweight costs of piracy are far less obvious. What are they? Treasure that gets lost in the fight, damage to ships, lost lives on both sides, etc.

1. The point is **not** that pirates make themselves worse off by piracy. At least ex ante, they don't. The point is that the pirates only gain a fraction of what the non-pirates lose.

2. This assumes, of course, that people don't *directly enjoy* fighting, watching gold sink to the ocean floor, etc.

E. Economists often criticize non-economists for thinking in terms of a "fixed pie" of wealth. In this sense, economists are more optimistic than the public. However, a corollary is that the pie can also *shrink*! In this sense, economists are more pessimistic than the public. With a fixed pie of resources, conflict at least has to benefit SOMEONE.

VI. The Comparative Institutions Approach and "Second Best"

A. Demsetz famously complained about the "Nirvana fallacy" - doing (K-H) efficiency comparisons while selectively relaxing important constraints.

B. His target was old-style welfare economics, where the solution to any market shortcoming was government involvement. The shortcomings of government - and even its basic overhead - were almost never factored in.

C. Classic example: P>MC.

1. Standard solution: Impose P=MC price control.
2. Secondary problem: With fixed costs, firms now lose money.
3. Standard solution: Subsidize them.
4. Tertiary problem: How can the subsidies be funded?
5. Standard solution: Taxes
6. But what about the DW cost of the taxes?!
7. And of course this still overlooks a wealth of problems. What *is* MC? Who awards subsidies, and what are their incentives? Etc.

D. Demsetz's lesson is that economists should use a "comparative institutions approach." Nothing in the real world is perfectly efficient. What fails least badly?

1. The Tale of the Emperor

E. When you add more constraints to a standard problem, the original optimum is usually no longer feasible. Economists frequently refer to the original optimum as a "first-best solution," and the new, worse optimum as a "second-best solution."

F. Example: Pricing subject to a P=AC constraint in a decreasing cost industry.

VII. Private Versus Social Benefits and Costs

A. Foundation of welfare economics: realization that private and social effects can differ.

B. Ex: A thief clearly enjoys private benefits of stealing. But looking only at the thief's benefits misses the big picture: The thief makes himself better off by making others worse off.

C. Ex: A person driving a polluting car is better off from driving, but that person isn't the only one who consumes the exhaust.

1. Contrast with: Worker safety trade-offs.

D. How to measure "social benefits"? The same way we always do: willingness to pay. If some people benefit and some people suffer from a policy, the net social benefits are the SUM of the private benefits (positive and negative).

VIII. Negative Externalities

A. The basic idea of the tragedy of the commons is that when no one owns a resource, it gets over-used.

B. Question: What exactly does "over-use" mean in economic terms?

C. Answer: It means that there are costly side effects, or "negative externalities," that selfish agents don't factor into their decisions.

D. How do you diagram negative externalities? In addition to the demand curve, draw a "social benefits curve." With negative externalities, the social benefits curve will lie <u>below</u> the demand curve.

E. Social optimum is at the intersection of the social benefits curve and the supply curve, but market equilibrium is at the intersection of the demand curve and the supply curve.

F. If the social optimum differs from the market equilibrium, it is typically called a "market failure."

G. Negative externalities are also often called "public bads," especially when the externalities are large relative to demand (so the socially optimal quantity is close to zero).

H. Ex: Pollution. People value better air, but polluters normally have no incentive to care.

I. The key: non-excludability.

1. There is no feasible way to exclude non-payers from the cleaner air.
2. Since you do not *have to* pay to use it, selfish people *will not* pay to use it.
3. And if no one will pay for it, why would selfish producers provide it?

IX. Positive Externalities

A. Positive externalities are the other side of the coin. Positive externalities are **beneficial** side effects that selfish agents don't factor into their decisions.

B. How to diagram? Draw a social benefits curve *above* the demand curve.

C. Positive externalities are also often called "public goods," especially when the externalities are large relative to demand (so the equilibrium quantity is close to zero).

D. Non-excludability is once again the key. If you can't exclude, there is no incentive to pay; if there is no incentive to pay, there is no incentive to produce.

E. Ex: Defense. People value defense, but how can suppliers be paid to provide it?

X. Understanding Externalities

A. David Friedman's two caveats:
    1. Must distinguish benefits from external benefits. (E.g. education).
    2. Must include both positive and negative externalities in your calculations. (Important case: "pecuniary externalities").
B. Further insight from Friedman: "It is easy to misinterpret problems of market failure as unfairness rather than inefficiency... The problem with public goods is not that one person pays for what someone else gets but that nobody pays and nobody gets, even though the good is worth more that it would cost to produce."

XI. Bad but Popular Examples; Good but Unpopular Examples
A. Some popular and plausible examples:
    1. Air pollution
    2. National defense
    3. Highways and roads (especially local roads)
    4. Law enforcement (especially victimless crimes)
B. Some popular but dubious examples:
    1. Education
    2. Health and safety
    3. Fire
    4. R&D
C. Some unpopular but plausible examples (depending on the society):
    1. Censorship
    2. Persecution of religious minorities...

XII. Fallacies of Group Action
A. Generalization of public goods theory: People often think in terms of **groups** acting to promote their *group interests*, just as individuals promote their self-interest.
    1. Workers/capitalists
    2. Women (and men?)
    3. Environment
B. But this is a fallacy of composition. Just because all members of group X would benefit if all members did something, it does **not** follow that it benefits any *individual* member to do so.
C. Ex: Suppose one worker decides to just stay home and watch TV while the other workers foment revolution.
    1. Case 1: Revolution succeeds, all workers (supposedly) enjoy a brave new world - including the couch potato.
    2. Case 2: Revolution fails, all workers continue to suffer under the capitalist system - but at least the couch potato got to watch some amusing television programming.
D. We do need to be careful before we assert that there is no selfish reason to contribute. Frequently there are "byproducts" and other "selective incentives" that make contribution selfishly optimal.
    1. Ex: Trotsky on military discipline

XIII. Individual Impact: Probability and Magnitude

A. Saying that "The same thing will happen whatever you do" is admittedly an overstatement. More precisely, "*About* the same thing will *probably* happen whatever you do."

B. In other words, you have to look at the **probability** you make a difference and **magnitude** of that difference, then weigh it against the cost of acting.

C. For example, it is possible that if you join the revolution, you will change the entire course of history. Possible, but not likely!

D. More relevant to public choice: the probability a vote matters and the magnitude of its impact.

E. Voting increases the probability that your favored candidates wins, but how much does it increase that probability?

F. And even if your candidate does win as a result of your vote, how much will policy change?

XIV. Calculating the Probability of Decisiveness, I: Mathematics

A. When does a vote matter? At least in most systems, it only matters if it "flips" the outcome of the election.

B. This can only happen if the winner wins by a single vote. In that case, each voter is "**decisive**"; if one person decided differently, the outcome would change.

C. In all *other* cases, the voter is **not** decisive; the outcome would not change if one person decided differently.

D. It is obvious that the probability of casting the decisive vote in a large electorate is extremely small. The 2000 election does not refute this. Losing by 100 or 1000 votes is a long way from losing by 1 vote!

1. You might however say that Bush did win by a single vote on the Supreme Court! But that is an electorate with only 9 voters.

E. There is a technical formula for "guesstimating" the probability of decisiveness using the *binomial formula*. (Brennan and Lomasky)

F. Suppose there are (2n+1) voters asked to vote for or against a policy.

1. Note: Assuming an odd number of voters avoids the picky problem of ties.

G. Then the probability that YOU are the decisive voter is the probability that exactly n voters out of the 2n voters other than yourself vote "for."

H. Now suppose that everyone but yourself votes "for" with probability p - and "against" with probability (1-p).

I. Then using the binomial theorem: $probability(tie) = \frac{1}{\sqrt{\pi n}}\left(4p - 4p^2\right)^n$

J. From this formula, we can see that the probability of a tie falls when the number of voters goes up. Why?

1. $\frac{1}{\sqrt{\pi n}}$ gets smaller as n gets larger
2. $\left(4p - 4p^2\right)$ is less or equal to 1. When you raise a number less than 1 to a larger power, it must get smaller.

K. This formula also says that as the probability of voter support goes above or below .5, the probability of a tie falls. Why?
1. When p=0, $\left(4p - 4p^2\right) = 0$; when p=1, $\left(4p - 4p^2\right) = 0$ too. In between p=0 and p=1, this term rises to a peak of $\left(4p - 4p^2\right) = 1$ when p=.5, then falls.

L. Intuitively, the more lop-sided opinion on a topic is, the less likely there is to be a tie. If everyone agrees, a tie is impossible.

XV. Calculating the Probability of Decisiveness, II: Examples

A. Let's work through some examples. Remember that the number of voters is (2n+1), not n.

B. Example #1: The close tenure vote. n=10, p=.5.

$$probability(tie) = \frac{1}{\sqrt{10\pi}}\left(4*.5 - 4*.25\right)^{10} = \frac{1}{5.60}\left(1\right)^{10} = \frac{1}{5.60} = .178$$, or 17.8%.

C. Example #2: The close county election. n=5,000, p=.51.

$$probability(tie) = \frac{1}{\sqrt{5000\pi}}\left(4*.51 - 4*.2601\right)^{5000} = \frac{1}{125}\left(.9996\right)^{5000} = \frac{1}{125}.1353 = .0011,$$

or a little more than 1-in-1000.

D. Example #3: The moderately close county election. n=5000, p=.53.

$$probability(tie) = \frac{1}{\sqrt{5000\pi}}\left(4*.53 - 4*.2809\right)^{5000} = \frac{1}{125}\left(.9964\right)^{5000} = \frac{1}{125}1.47*10^{-8} = 1.18*10^{-10}$$

, a little less than 1-in-8 billion.

E. Example #4: The moderately close state election. n=2,000,000, p=.51.

$$probability(tie) = \frac{1}{\sqrt{2{,}000{,}000\pi}}\left(4*.51 - 4*.2601\right)^{2{,}000{,}000} = \frac{1}{2507}\left(.9996\right)^{2{,}000{,}000},$$

a chance smaller than 1 in 10^{-100}! (My calculator just says 0).

F. Upshot: For virtually any real-world election, the probability of casting the decisive vote is not just small; it is normally *infinitesimal*. The extreme observation that "You <u>will</u> not affect the outcome of an election by voting" is true for all practical purposes.

XVI. Empirical Evidence on Collective Action Problems

A. One way to get a feel for the logic of collective action is to see how little participation in politics there is. Survey of adult Americans from Dye and Zeigler:

Activity	%
Run for public office	<1
Active in parties and campaigns	4-5
Make campaign contribution	10
Wear button or bumper sticker	15
Write or call a public official	17-20
Belong to organization	30-33
Talk politics to others	30-35
Vote	30-55

B. Many experiments have been run to help improve our understanding of collective action problems.
   1. Part of the design: Rule out "selective incentives" accounts of apparently unselfish behavior.

C. Standard design:
   1. I hand out a roll of 100 pennies to each person in the class.
   2. Then, people are allowed to secretly put any number of their pennies into a jar.
   3. You personally get to keep the pennies you *don't* put in the jar.
   4. I count the number of pennies in the jar; then I distribute **twice** that many pennies to the class, with each person getting the same share.

D. What maximizes the total income of the class? 100% donation by everyone!

E. What maximizes your private income (given 3 or more players)? 0% donation!

F. The first couple of times you do an experiment like this, you typically get moderate to high levels of donation - 50-80%.

G. Donation levels usually fall as you repeat the experiment with the same group. After a while, donation levels often bottom out at around 20%.
   1. For practical reasons, experiments usually only last a day or less. So we can still speculate about what would happen if people played this game 10 times a day for a year.

H. Donation levels usually decline as the number of participants rises.

I. The less secrecy there is, the higher the level of donation.

J. Conclusion: The "logic of collective action" appears to exaggerate the degree of human selfishness, but cooperation in these experiments is still far below the group-income-maximizing level.