Prof. Bryan Caplan
bcaplan@gmu.edu
http://www.bcaplan.com
Econ 812

Week 1: Efficiency and Probability

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 usages.
   D. The two most common uses in economics are:
      1. Pareto efficiency
      2. Kaldor-Hicks (or cost-benefit) efficiency
   E. Since much of micro analyzes efficiency, it is important to understand these terms' precise meanings.

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. Slight variant - a situation is Pareto efficient if there is no way to make everyone better off. Note that in a perfectly continuous world, this is equivalent to the other definition. Why?
   E. 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.
   F. In a highly stylized theoretical setting, we will see that 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. More fruitful variant: Analyze the Pareto efficiency of ex ante *rules* instead of ex post results. But even then, someone somewhere is sure 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 that *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 your 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.

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. Now let's examine Landsburg's K-H analysis of drug legalization. Main insights:
        1. Taxes raised are a transfer, not a "benefit."
        2. Imprisonment and effort spent avoiding imprisonment is a deadweight cost.
        3. Theft is a transfer, but resources (time, tools, etc.) used to steal are a deadweight cost.
        4. Voluntary consumption is a benefit!
        5. Internalized losses (like loss of productivity) are already counted in consumption decisions.
    F. 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.
    G. Reducing deadweight costs is always a K-H improvement; if a situation is K-H efficient, deadweight costs are zero.

VI. Kaldor-Hicks Efficiency versus Utilitarianism
    A. Kaldor-Hicks efficiency is based on dollar valuations, not utility or happiness.
        1. You can know that I'm willing to pay $100 for something without having any idea about how much happiness it brings me.
        2. Similarly, you can know that something makes me very happy even if I have a low willingness to pay for it.
    B. Utilitarianism, in contrast, is precisely about maximizing happiness or pleasure. The main reason economists rarely officially use it is that it requires "interpersonal utility comparisons." Simply: How do you "add happiness"?
    C. People often say that utilitarianism just factors in the marginal utility of wealth, unlike K-H. There is a point here, though it is not necessarily true: People might be willing to pay for things other than happiness.
    D. Utilitarianism is often used to justify redistribution, but even on its own terms, this doesn't necessarily follow. The "utility monster" is the standard philosophers' counter-example.

VII. 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?!

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.

VIII. Moral Philosophy and Efficiency

A. Who cares about efficiency anyway? Does anyone seriously believe that the right action is always the one that does the most for K-H efficiency?

B. One popular reply: K-H efficiency combined with redistribution.

1. That still seems highly inadequate to me. What about desert and entitlement?

C. More moderate view: Efficiency is probably ONE of many consequences worth thinking. Why then should economists concentrate on it? Because they have special training for distinguishing transfers from DW costs, but no special training in moral philosophy. Economic analysis thus becomes a potentially useful input into the moral thinking of others.

IX. Probability, Objective and Subjective

A. Probability language allows us to **quantify uncertainty**. There is more to say in an uncertain world than "I don't know."

B. Least controversial interpretation: *objective* probability. Even when you do not know what will happen, you can still talk about *relative frequencies* of various observed events in the past.

C. But objective probability is problematic in many ways. Most notably, it implies that you cannot talk about probability of unique events. If you take this idea seriously, moreover, you will realize that every event is, strictly speaking, unique, so you could never apply probability to the real world!

- D. This leads us naturally to the broader but more contentious *subjective* interpretation of probability.
- E. A subjective probability is simply a *degree of belief* that a person assigns to a proposition. Simple axioms of probability:
  1. Beliefs range from impossible (p=0) to certain (p=1).
  2. Since *something* is certain to happen, the sum of all probabilities about an event must equal 1.
- F. Main objection to subjective probability: Realism. People rarely explicitly assign probabilities to events.
- G. My reply: Even so, people almost always have some probabilities in the back of their minds. Probabilities is like willingness to pay.
- H. Further objection: When people are asked difficult questions, they often say "I don't know."
- I. But what if they HAD to guess? In real life you must.
- J. Common sophism: "No one can 'know' X."
  1. If this means "No one can know X **with certainty**," then it's obvious but uninteresting.
  2. If this means "No one has any idea at all about X," then it is clearly false.
- K. Does probability theory rule out "surprise"? Not at all. The occurrence of the improbable, extreme events is inherently surprising.
- L. In practice, economists typically use the subjective interpretation of probability, but add assumptions that link subjective and objective probabilities. More on this later.

X. Conditional Probability and Bayes' Rule

- A. Subjective probability theory puts no constraints on pre-evidential beliefs, but it does restrict the way that people can *update* their beliefs when new evidence comes in.
- B. Conditional probability formula: P(A|B)=P(A&B)/P(B).
  1. Ex: P(2 heads|first flip is heads)=P(2 heads)/P(first flip is heads)=.25/.5=.5.
  2. Ex: P(child saw monster|says he saw monster)=P(child saw monster & monster)/P(says he saw monster). So if P(child saw a monster and monster)=10^{-9}, and P(says he saw monster)=.1, the conditional probability comes out to one-in-a-hundred-million.
  3. Note: Conjunction can never be more probable than either of the components!
- C. A more advanced formula, known as Bayes' Rule, lets us link the P(A|B) and the P(B|A). Bayes' Rule states that P(A|B)=P(B|A)*p(A)/[P(B|A)*P(A)+P(B|~A)P(~A)].
- D. Ex: P(child saw a monster|says he saw monster)=P(child says he saw monster|saw monster)*P(saw monster)/[P(child says he saw monster|saw monster)*P(saw monster)+P(child says he saw monster|did not see monster)*P(did not see monster). So if P(child

says he saw monster|saw monster)=1, P(child says he saw monster|did not see a monster)=.1, P(saw monster)=10^{-7}, and P(did not see monster)=$1-10^{-7}$, the conditional probability works out to $1*10^{-7}/[1*10^{-7}+.1*\{1-10^{-7}\}]=10^{-7}/[10^{-7}+.9999999]=9.999991*10^{-7}$.

E. Bayes' Rule provides a natural framework for scientists to relate hypotheses to evidence. Let A be your hypothesis and B be some evidence; then calculate P(A|B).

F. Ex: The P(minimum wage causes unemployment|Card/Krueger study's findings). Suppose P(CK findings|m.w. does cause unemployment)=.3, P(CK findings|m.w. does not cause unemployment)=.8, P(m.w. does cause unemployment)=.99, and P(m.w. does not cause unemployment)=.01. Then the conditional probability comes out to .3*.99/(.3*.99+.8*.01)=97.4%.

G. Do people update their beliefs "as if" they knew these formulae? Obviously, they do to some degree. We:
   1. ...run away when we appear to see a large fire
   2. ...meet reports of alien abduction with skepticism
   3. ...believe shocking disaster stories in the NYT, but not the Weekly World News.
   4. ...do not change our minds about the minimum wage when astronomers discover a new galaxy.

H. This is fortunate since game theory and information economics depend heavily on these formulae. After the midterm we will examine empirical evidence which points to some exceptions.

I. Application: What should you infer if you think you witness a 0-probability event?