

**TEACHING ABOUT ECONOMIC EFFICIENCY IN LAW
AND ECONOMICS COURSES:
CLARIFYING THE CONCEPTUAL PROBLEMS, EMPIRICAL
DIFFICULTIES AND NORMATIVE BIASES OF THE
EFFICIENCY CRITERION**

by

Gregory Crespi¹

April 9, 2016

Preliminary Draft

ABSTRACT

Law and Economics courses taught in law schools are sometimes criticized for inadequately explaining the normative criterion of “economic efficiency” and then applying this criterion throughout the course in a superficial and biased manner that pejoratively labels most governmental market interventions and wealth redistribution measures as inefficient. These criticisms have merit, and in this brief article I point out a number of conceptual problems, empirical difficulties and normative shortcomings of the economic efficiency criterion that students need to understand in order to later in their careers be able to effectively counter policy arguments that rest upon efficiency assessments.

The eight specific shortcomings of the economic efficiency criterion that I address in this brief essay are the pervasiveness of severe data limitations that render efficiency assessments unreliable, the lack of clarity as to whether willingness to pay should be measured by offer prices or instead by asking prices, the difficulty of obtaining honest and accurate responses as to willingness to pay from the persons surveyed, uncertainty as to the appropriate discount rate that should be used for discounting future policy consequences, the problem posed by endogenous preferences, the problem posed by the often-overlooked “problem of person-altering consequences,” the problematic nature of using willingness to pay

as a measure of social value, and finally, the problematic nature of using a normative criterion that does not give special primacy to rights.

TABLE OF CONTENTS

- I. INTRODUCTION
- II. THE ECONOMIC EFFICIENCY CRITERION DEFINED
- III. PROBLEMS INVOLVED IN MEASURING WILLINGNESS TO PAY
 - A. Sampling Limitations.
 - B. Determining Whether to Use Offer Prices or Instead Asking Prices to Measure Willingness to Pay.
 - C. Obtaining Honest and Accurate Responses.
 - D. Determining the Appropriate Discount Rate.
 - E. Valuing Policy Impacts When Preferences are Exogenous.
 - F. Valuing Policy Impacts in Light of the Person-Altering Consequences of Policies.
- IV. NORMATIVE PROBLEMS WITH THE ECONOMIC EFFICIENCY CRITERION
- V. CONCLUSION

I. INTRODUCTION

Law and Economics² courses taught in law schools are sometimes criticized for inadequately explaining and then summarily endorsing the use of the normative criterion of “economic efficiency”³ to evaluate policies, and then applying this criterion in a superficial and unreflective manner to pejoratively label most governmental market interventions and wealth redistribution measures as inefficient. Those critics argue that the efficiency criterion is not well defined in significant regards and thus allows for result-oriented manipulation by analysts, and that it is also often difficult or impossible to obtain sufficiently reliable data to be able to meaningfully apply the criterion, and that in addition the criterion has a pronounced bias in favor of the interests of the wealthiest members of society and does not accord any respect to rights. They also argue that these severe shortcomings of the criterion are not made sufficiently clear to students to enable

them in their later studies and careers to effectively counter policy

recommendations that are based upon efficiency assessments.

There is considerable merit to these arguments, and it is also troubling that much of the early advocacy and financial support provided for the introduction of Law and Economics courses into law school curricula, as well as the related advocacy for the broader use of the efficiency criterion in legislative, regulatory and judicial decision-making, has been shown to be due directly or indirectly to the efforts of certain wealthy individuals who embraced a pronounced right-wing libertarian political orientation, and whose large financial interests would be furthered by a more widespread and uncritical embrace of the efficiency criterion.⁴

The fact that several ideologically-oriented foundations established by these individuals provided much of the financial support for many law schools' Law and Economics courses and for judicial workshops that promoted efficiency as an

important social interest⁵ suggests that this insufficiently critical embrace of the efficiency criterion that often characterizes these courses and workshops may have been intentional, at least in some instances.

These powerful criticisms call conventional Law and Economics courses into some question. However, in my opinion these criticisms can be adequately addressed through relatively modest pedagogical changes and do not justify the more radical measure of eliminating these otherwise useful classes from law school curricula. My experience in teaching these courses for over 25 years is that it is possible to provide law students with a succinct yet comprehensive introduction to the concept of economic efficiency that does not only explain the concept and alert students to its current widespread application in policy analysis, but that also helps them to understand the numerous and severe conceptual, empirical and normative shortcomings of that criterion.

It is critically important for teachers of Law and Economics courses to explain the important concept of economic efficiency in some detail at the outset so that the students will then have an understanding as to why efficiency assessments of policies are in general useless or even positively misleading, so that in their later careers they will be better able to respond effectively to policy arguments based on such efficiency assessments. In this brief essay I will attempt to clarify the conceptual and empirical difficulties that are encountered when attempting to assess of the efficiency or inefficiency of a particular policy, and what is normatively problematic about using this criterion as an evaluative standard even if these measurement difficulties can somehow be overcome.

II. THE ECONOMIC EFFICIENCY CRITERION DEFINED

“Efficiency” is a protean term that means different things in different contexts. The particular concept of efficiency that is generally applied in assessing the economic consequences of policies is what economists more precisely call “Kaldor-Hicks efficiency.”⁶ The Kaldor-Hicks efficiency criterion has a straightforward and intuitively appealing definition. If the aggregate benefits of a policy to those persons favorably impacted by it, as measured by their willingness to pay for those benefits if they were required to do so, exceeds the aggregate costs of that policy imposed upon those persons unfavorably impacted by it, again as measured by their willingness to pay if required to do so to avoid those costs, the policy is then regarded as a “Kaldor-Hicks improvement,” and is often more colloquially described as simply being “efficient” or as “resulting in an increase in social wealth.” There is no additional requirement that compensation be paid by the policy’s beneficiaries to those persons burdened in order for that policy to be a Kaldor-Hicks improvement, it is sufficient if the aggregate benefits exceed the

aggregate costs. If a resource allocation is reached where no further Kaldor-Hicks improvements are possible then that allocation is described as being “Kaldor-Hicks efficient,” or, equivalently, as being “efficient” or “wealth maximizing” relative to any other possible allocation of resources.

The most straightforward application of this criterion, one that is usually presented early on in Law and Economics courses to illustrate the underlying concepts, is to conduct a geometric analysis of the efficiency of various forms of governmental intervention into a hypothetical free market that satisfies the behavioral and institutional assumptions for perfect competition. Using a simple supply-and-demand graph depicting such a market in competitive equilibrium, the aggregate benefits to the buyers of the good or service of the existence of such a market can be shown to be the size of the triangle⁷ under the demand curve and above the equilibrium price line. This area represents the aggregate willingness to

pay of those persons who buy in this market over and above the lower market equilibrium price they are required to pay for the good or service. Each of these buyers will obtain a benefit from their purchase equal to the difference between their personal “reservation price” for the good or service⁸ and its equilibrium price. The aggregate benefit to all of these buyers is referred to as the “consumer surplus” generated by this market.

In similar fashion, the aggregate benefits to the sellers of the existence of this market can be shown by the size of the triangle⁹ above the marginal cost-based supply curve and below the equilibrium price line, with each of these sellers obtaining a benefit from their sale equal to the difference between their personal reservation price for the good or service¹⁰ and its equilibrium price. The aggregate benefit to all of these sellers is referred to as the “producer surplus” generated by this market. It is then geometrically demonstrated to students that various forms of

governmental intervention into the operation of this market that go beyond the simple background enforcement of property rights and contractual obligations and restrictions upon collusion – for example, measures such as the imposition of price ceilings, price floors, sales taxes, income or payroll taxes, subsidies to buyers and/or sellers, or the allowance of collusion among buyers or sellers – will generally reduce the combined consumer and producer surplus that the market generates. Therefore such measures are usually judged to be inefficient.

In similar fashion, it is easily demonstrated that under the willingness to pay valuation framework any transfer of wealth from one group of persons to another will have an aggregate cost to those persons burdened of that exact amount of wealth, an amount that is equal to its aggregate benefits to those persons benefitted. Since any transfer of wealth will in practice necessarily involve some transaction

costs in arranging and carrying out those wealth transfers, when those transaction costs are taken into account the redistribution of wealth will be inefficient.

Now I do recognize that most if not all teachers of Law and Economics courses at about this point in the class will attempt to make clear to the students that not all markets satisfy the full range of perfectly competitive market assumptions that underlie these conclusions as to the general inefficiency of governmental interventions altering market outcomes. They can and usually do demonstrate through use of the same basic supply-and-demand framework that there will be market failures when there are positive or negative externalities of production or consumption that are not captured by the supply or demand curves, or where sellers or buyers significantly depart from rational actor assumptions. Under such circumstances markets may well fail to reach efficient results, and appropriate governmental interventions may to some extent rectify these

inefficiencies. But these important caveats as to the possible inefficiencies resulting from market failures are then often downplayed or even overlooked entirely later in the course when various laws are examined for their efficiency or inefficiency. In addition, the points that I will discuss below regarding why the use of the efficiency criterion is problematic even in those contexts where all of the restrictive assumptions of perfect competition are satisfied and when all costs and benefits are internalized are generally covered inadequately, if at all.

This simple framework of analysis demonstrating the general inefficiency of governmental interventions into perfectly competitive markets, absent market failures, and the inefficiency of governmental income or wealth redistribution measures, if not then developed in more detail so that the more subtle yet severe conceptual, empirical and normative limitations of the efficiency criterion are made clear, will tend to implicitly impart to students the general message that a

limited, laissez-faire government that provides a basic framework of law and order, but does that not otherwise interfere with market processes or redistribute market-generated incomes, will better facilitate the general welfare than will a more interventionist government. This simplistic, pro-free market orientation then often explicitly or implicitly suffuses the remainder of the course as various legal doctrines in property law, contract law, tort law, corporate governance, etc. are examined one after another for their efficiency or inefficiency. This repeated assessment of legal rules by an efficiency criterion whose conceptual and empirical shortcomings and problematic normative basis have not been first fully explained to students are what those persons who criticize these courses as biased in favor of the interests of the wealthiest members of society find most objectionable.

Law and Economics courses do not have to be taught this way. It is entirely possible for a teacher who is willing to devote a week or so of class coverage to the

important economic efficiency criterion to make clear to students its numerous shortcomings. The students hopefully will then in their later studies and careers have a much greater capability to effectively refute arguments regarding a policy's merits that rest upon efficiency or inefficiency premises.

Let me briefly discuss the serious problems that are inherent in the use of the efficiency criterion to evaluate the merits of policies. I will first address the substantial and arguably overwhelming conceptual and empirical difficulties that are involved in reaching a meaningful quantitative assessment of the efficiency or inefficiency of a policy. I will then address the problematic implicit assumptions involved in using efficiency as a normative criterion for assessing policies even if one is somehow able to meaningfully quantify those policy impacts.

III. PROBLEMS INVOLVED IN MEASURING WILLINGNESS TO PAY

The most severe problems that are commonly encountered in quantifying the willingness to pay-based impacts of a proposed policy measure are the following:

- 1) the practical limitation that it is usually not possible to survey more than a modest-sized sample of the likely affected current population when measuring a proposed policy's impacts, and that it is obviously impossible to survey those persons who will later be impacted by the policy but who are not yet born, 2) resolving the question of whether to measure the willingness to pay of the persons surveyed through their offer prices or instead through their asking prices, 3) the difficulty of obtaining honest and accurate measures of willingness to pay from those persons surveyed, 4) the problem of determining the appropriate discount rate to apply to future policy impacts before aggregating them with current impacts, 5) the problem posed by the possibility that the preferences of some or all of the persons impacted by a policy may be endogenous with respect to that policy, rather than stable and exogenously determined, and perhaps most intractably, 6)

the problem that one of the long-term impacts of any policy whatsoever is that it will eventually lead to a different genetic identity at the time of conception than would otherwise be created for all persons born thereafter through the end of eternity; i.e., it will eventually have universal “person-altering consequences” in that all of the persons who will later come into being after some period of time after the implementation of a policy will do so only because of the particular policy that was implemented, and these obviously momentous policy consequences for those persons will somehow have to be taken into account in valuing those impacts. Let me briefly address each of these difficulties in turn.

A. Sampling Limitations.

For the basic supply-and-demand market model commonly used to geometrically introduce the concept of economic efficiency in Law and Economics courses the reservation prices of all of the potential buyers and sellers are assumed

to be known and represented by the positions of the demand and supply curves, so that the problems of calculating the willingness to pay-based impacts of a governmental market intervention are reduced to simple matters of addition and subtraction. For real-world attempts to assess the efficiency of policies, however, the willingness of the impacted persons to pay to experience (or to avoid) the consequences of the policies are not so given and will have to be empirically determined.

In practice, unfortunately, whoever is directed to assess a policy's efficiency or inefficiency is usually not given anywhere near the resources that would be necessary to survey even a substantial fraction of the living individuals who would be impacted to some extent by the policy, who may easily number in the millions for significant governmental policies, and of course the analyst will be unable to survey at all those persons who will eventually be impacted by the policy once

they come into existence, but who have not yet been born at the time that the policy's impacts are being assessed. The analyst will typically be limited to surveying the responses of at most a relatively small sample of the current population and then will have to make projections based upon those findings in order to estimate the aggregate willingness to pay of the entire impacted current population. Based also on those sample survey findings the analyst will have to speculate as to how the affected members of future generations will regard the policy impacts, and then make willingness to pay estimates for those future persons accordingly. With all of this projection and speculation from an often very small survey data base even if relatively sophisticated stratified sampling techniques are utilized to enhance the representativeness of the selected sample this will limit the confidence that one can have in these estimates.

B. Determining Whether to Use Offer Prices or Instead Asking Prices to

Measure Willingness to Pay.

The use of the simple phrase “willingness to pay” masks an important yet unresolved (and perhaps unresolvable) issue in efficiency analysis. When attempting to ascertain the willingness to pay of a person who is benefitted by a policy, is the proper question to pose how much the person would be willing to pay to receive those benefits – commonly referred to as their “offer price” for those benefits– or is the proper question instead how much that person would demand to give up those benefits once they had been conferred upon them – commonly referred to as their “asking price?” Similarly, and more importantly, when attempting to ascertain the willingness to pay of a person who would be injured by a policy to avoid its costs, is the proper question to pose how much that person would be willing to pay to avoid the imposition of those costs – again their offer

price – or instead how much that person would demand for their consent to have those costs imposed upon them - their asking price?

Unfortunately, how the willingness to pay question is posed may significantly affect the answers obtained. Depending on the framing of the question these answers may differ dramatically in size, particularly when attempting to measure the impacts of a policy on those persons for which the policy is perceived as imposing substantial costs. The reason for this is that while a person's offer price to avoid a policy's costs is constrained by the amount of their wealth, and by what other objectives they have that they will want to use their wealth to satisfy, their asking price is not so constrained by their wealth and has no necessary upper limit. If a person is highly reluctant to consent to implementation of a particular policy as a matter of principle then that person's asking price for their consent could be very large, or even infinite in the sense that they might

refuse to give their consent to the policy's implementation for any amount of money.¹¹ An efficiency assessment of a policy will not provide meaningful guidance for to policymakers if there is an infinite cost included in the balance.

An analyst conducting an efficiency analysis will have to choose whether to use offer prices or instead asking prices for their measurements of willingness to pay, or whether instead to use some combination of the two forms of measurement. The efficiency criterion would be more useful if there was a clearly correct choice here as to the proper questions to pose, so that analysts would not be able to covertly manipulate the questions in order to reach a result that they favored for other reasons. But there does not appear to be a definitive resolution of this offer/asking price problem that is or can be incorporated into standard practice, given the vagueness of the general "willingness to pay" efficiency formulation.

One may choose to look to the operation of markets for some guidance here as to how to measure willingness to pay given that efficiency analysis is often regarded as an attempt to apply market-like assessment principles to evaluate governmental policies that primarily impact persons in non-market settings.

Buyers in markets will pay no more than their offer prices for the benefits of having a good or service, and this fact suggests that offer prices are perhaps the most appropriate measures of willingness to pay, at least for policy beneficiaries.

On the other hand, however, sellers do not have to part with their goods or services unless their asking prices are met, suggesting that asking prices may be the proper measure for valuing the costs imposed by policies. So if efficiency analysis is to be done in a manner that closely emulates the workings of markets perhaps offer prices should be used to measure benefits and asking prices to measure costs.

In practice, however, almost all efficiency analyses¹² use offer prices to measure both the benefits and costs of the policies being considered. But this approach seems to be based primarily upon the desire of analysts to avoid the possibility of having the assessment skewed heavily against the policy in question by very large or even infinite asking prices that perhaps only a very small minority of the impacted persons would declare, rather than upon some credible theoretical argument as to why the impacted persons' offer prices rather than their asking prices are the "correct" measure of a policy's negative impacts. This lack of clear direction as to how willingness to pay should be measured, and the possible sensitivity of the conclusions reached in an efficiency analysis with regard to how the willingness to pay question is posed, greatly undercuts the usefulness of the efficiency criterion as an objective measure of the merits of policies.

C. Obtaining Honest and Accurate Responses.

Even once an analyst one has determined the sample of the impacted population to survey, and has made their choice between using offer price or instead asking price measures of willingness to pay, there remains the question as to whether the responses obtained will honestly and accurately reflect true willingness to pay. There are several reasons why this may not be the case.

First of all, the questions posed to respondents with regard to the impacts of a proposed policy are usually not the sort of questions that people are asked to answer on a regular basis and therefore have some experience to draw upon in responding. Being asked how much one would be willing to pay, for example, for more stringent limits on the emission of certain toxic chemicals, or for a worldwide treaty banning blue whale hunting, or for a better freeway connection in a neighboring city that one occasionally visits by automobile are not like the normal routine shopping expenditure decisions that one makes every day. People may

give answers to such unfamiliar hypothetical survey questions that are not very well considered, particularly since they are not subject to the discipline of actually having to pay the declared amounts for the benefits of the policy, or having to accept the declared compensation for its costs. Economists generally favor on reliability grounds the use of “revealed preference” data that reflects actual spending behavior over merely hypothetical survey responses, but an efficiency analyst will unfortunately usually only have such hypothetical survey responses to work with.

A second problem here is that some of the persons surveyed may deliberately misstate their willingness to pay for strategic reasons. A person who would benefit from a policy has an incentive to overstate the size of this benefit in an attempt to make the policy’s implementation more likely, particularly since they will not be required to pay this stated sum. Similarly, a person who would be

burdened by a policy has an incentive to overstate the size of this burden in an attempt to discourage the implementation of the policy. On the other hand, the persons surveyed might be concerned that their responses might somehow provide a basis for later public assessments against (or compensation payments made to) them, which would encourage the beneficiaries to now understate the benefits, but would further encourage those persons burdened by the policy to further overstate those burdens. All of these strategic considerations that may be in play will undercut the confidence an analyst can have in the accuracy of their efficiency calculations.

D. Determining the Appropriate Discount Rate.

Some of the impacts of a policy will occur immediately, but other impacts will not occur for some period of time. This raises the difficult question of how to aggregate the different benefit and cost impacts of a policy that occur during

different time periods into a single bottom line figure. The standard approach taken in efficiency analysis is to discount all future impacts to a smaller present value before aggregation in a manner similar to what is done in conventional finance calculations of the present value of streams of future cash flow. Use of this procedure then presents the question as to what is the appropriate discount rate that should be used.

Ideally the benefits or costs borne by each person impacted by a policy in each future time period would be discounted by the impacted person's personal rate of time preference with regard to events occurring in that future period as compared to the current time period. Once again, sampling limitations will generally preclude such an individualized assessment and use of personal discount rates, and the usual convention in efficiency analysis is to use a single discount rate for discounting all future impacts of a policy to a commensurate present value.

The chosen discount rate usually reflects either the analyst's estimate of the average rate of time preference of the entire affected population over the relevant time periods, or instead is some bond market-based interest rate(s) paid in the market for long-run riskless investments over the relevant time period, such as the interest rate(s) on long-term US Treasury bonds of the appropriate maturity, or instead some estimate of the average annual rate of return on private invested capital over the relevant time period, or some combination of more than one of the above measures. In addition, there is controversy regarding whether such finance-theory based discount rates are appropriate for valuing very long-term consequences that even if quite large in magnitude when they occur will be reduced to insignificance by even a low discount rate, and what discount rate is appropriate when the cost impacts being valued are of the nature of loss of human life rather than merely financial consequences.

The lack of clear agreement as to what discount rates should be used in efficiency analysis presents a significant problem, because the assessment of policies that have substantial long-term impacts (which as I will discuss below includes all policies whatsoever because of their inevitable and universal genetic person-altering consequences) is very sensitive to the discount rates used to discount future impacts before aggregating them with current impacts.¹³ The sensitivity of the results reached in many analyses to an analyst's relatively unconstrained choice of a discount rate serves to undercut the objective significance of the conclusions.

E. Valuing Policy Impacts When Preferences are Endogenous.

A common assumption made in efficiency analysis is that people's preference structures are exogenous with regard to the impact of policies. In other words, it is assumed that they will value the impacts of a policy at the same amount

whether they express their willingness to pay either before or instead after they experience the policy's impacts. But this may not actually be the case, for a couple of reasons.

First of all, behavioral economists have established that in many instances people exhibit an "endowment effect" where they will value a good or service that they own more than they would value that same good or service if they did not yet own it, even if their underlying preference structures with regard to that good or service remain essentially unchanged after they take ownership. In such instances the aggregate willingness to pay for the benefits of a policy will vary depending on when that willingness to pay is expressed.

Secondly, and more troubling as a conceptual matter, it is possible that the implementation of a policy may significantly alter the preference structures of some of the people affected, perhaps dramatically so. Stated in a more technical

manner, preferences may be endogenous to a significant extent with regard to the policy under consideration. If this is the case it presents the question of whether an analyst attempting to ascertain the efficiency or inefficiency of the policy should determine the impacted persons' willingness to pay on the basis of their pre-policy implementation preference structures, or instead on the basis of their different post-policy implementation preference structures? In some instances the use of pre-policy implementation preference structures as the valuation baseline would seem to be more appropriate,¹⁴ but in other instances the post-policy implementation preference structures would appear to be more appropriate.¹⁵ It is unclear what is the proper way to address this problem, as a general matter, and this lack of clarity creates yet another opportunity for a result-oriented analyst to manipulate the results of the assessment.

G. Valuing Policy Impacts in Light of the Person-Altering Consequences

Policies.

The noted British philosopher Derek Parfit has elaborated upon what he refers to as the “non-identity problem.”¹⁶ This is a strikingly novel insight of great (and in my opinion under-appreciated) significance for policy assessment. I have relabeled this problem with what I regard as a more apt moniker as the “problem of person-altering consequences” in several articles that I have addressed in recent years to economists and lawyers and other policy analysts, rather than to the professional philosophers that were Parfit’s target audience.¹⁷

This problem results from the fact that even a policy with very minor social impacts will surely have enough effect on some human comings and goings to at least slightly alter the details of at least one act of human reproduction so that the female’s egg is fertilized by a different male sperm, out of the hundreds of millions

of sperm released in an ejaculation, than the sperm that would have succeeded in fertilizing the egg had the precise timing or other details of the act of conception not been slightly altered by the policy. This different sperm-egg fertilization will lead to a person being conceived and later born with a different genetic identity than would otherwise have been the case. In other words, now a different person will eventually be born. And as this genetically different person matures and goes through life they will surely affect other persons in ways that are different, perhaps radically so, than would have been the case had the policy not been implemented and the other “potential” person had been conceived and born instead. These effects from a genetically different person going through their life will also surely alter some later acts of conception from what they otherwise would have been in a similar fashion over time, leading to an exponentially growing cascade of genetic alterations stemming from the initial policy impact. After some period of time that is likely measured in mere decades rather than centuries the entire human

population conceived and born for the rest of eternity will all have different genetic identities than the future population that would have been born had the initial policy in question not been implemented. The policy will therefore eventually have universal person-altering consequences.

Another way to put this fascinating though somewhat troubling insight is that to note that one result of any policy whatsoever, even one with very minor initial impacts, will be to eventually create a future human population extending for eternity thereafter for whom that policy is a necessary condition of their existence. None of the members of that population would have been conceived and born had the policy not been implemented. This fact creates a real dilemma for attempts to estimate the willingness to pay for the policy's benefit and cost impacts of the many members of that eternal parade of future generations.

One would expect these future persons to each place a very high positive value on a policy whose effects were a necessary condition of their existence, certainly if asking price measures are utilized, but likely very high even if offer prices are aggregated, even if the policy in other ways had adverse effects on their lives.¹⁸ These untold billions or perhaps even trillions of persons in the aggregate would surely place an astronomically high value on the policy that resulted in their existence, even if some of the other consequences of that policy were adverse to their interests. Even if those future valuations are heavily discounted to a current present value the aggregate benefits are still likely to be massive (although obviously impossible of precise estimation), very likely large enough to greatly outweigh any negative policy impacts upon members of the current generation. In other words, any policy option whatsoever, including the null option of doing nothing, will result in massive net benefits of uncertain magnitude because of the particular very large future population that the policy will bring into existence,

regardless of any negative impacts upon the current population. This fact will render efficiency assessments rather useless for distinguishing among alternative policy options, and policymakers are of course likely to reject out of hand any evaluation method that essentially ignores policy impacts upon existing persons.

So what can be done here to salvage the usefulness of economic efficiency as an evaluative standard, given this seemingly insuperable difficulty? Analyses of the efficiency of policies are now invariably conducted on the basis of the demonstrably false implicit assumption that the same future persons will come into being whether or not the policy under consideration is implemented. With this facilitating assumption the policy impacts on the wealth of future persons are then assessed against the hypothetical baseline alternative of those same future persons coming into being but not experiencing the policy's impacts, which could then lead to either positive or negative estimated valuations of the policy impacts for each of

those many future persons, depending on the nature of those individual impacts.

Those impacts are then all discounted to present values for aggregation with the impacts upon current persons.

I want to emphasize that this working assumption that the same future persons will come into existence whether or not the policy in question is implemented is not merely a standard simplifying assumption that has been chosen in order to make otherwise difficult calculations somewhat more tractable without changing their basic character. This is instead the far more dubious use of a hypothetical counterfactual baseline for assessing policy impacts that cannot possibly occur, since it is beyond reasonable argument that any policy will have the dramatic person-altering consequences that I have described that will eventually create an entirely different future population. Those many efficiency analyses that simply assume away the problem of person-altering consequences by

pretending in this way that it does not exist are unfortunately not relevant to the real choices at hand between policy options all of which do have these person-altering consequences.

I see no way to incorporate person-altering consequences into efficiency analysis without leading to the counterintuitive and unhelpful but seemingly unavoidable result that massive but practically unquantifiable benefits will result for future generations that dwarf any impacts upon existing persons no matter what policies we pursue, since those future persons that will eventually come into being after any policy has been implemented will be created in part by that policy, and would therefore likely deem this particular policy impact to be of overriding beneficial significance. Even if the numerous other conceptual and empirical problems of the efficiency criterion that I have discussed above can somehow be

resolved I believe that the problem posed for efficiency analysis by person-altering consequences is likely fatal.

Given this problem it appears to be necessary to develop alternative evaluative criteria to economic efficiency that do not require valuing policy impacts upon specific future persons. Trying to think outside of the box here, one possible approach may be to simply assume that the massive benefits of uncertain magnitude for future persons that will result from any policy because of its person-altering consequences will in effect “cancel out” when comparing alternative policies, including the null option of not implementing any policy, and therefore only the impacts upon current persons need to be considered and compared in policy assessment. This approach would make policy assessment far more tractable. However, the approach appears to be unsatisfactory in that it would unduly favor those policies with current benefits and large long-term costs for

future persons,¹⁹ and unduly disfavor policies with current costs and large long-term benefits for future persons.²⁰ Effectively ignoring policy consequences for future persons in this manner does not appear to be any more credible than pretending that policies do not have person-altering consequences.

Another and more radical possible approach would be to try to develop evaluative criteria that focus instead on somehow valuing policy impacts upon the overall “human race” when it is viewed as a distinct entity that exists apart from the specific individuals that comprise it. It is unclear, however, what sorts of policy impacts upon this “human race” would be considered that are not simply aggregations of the impacts on individual persons, nor how those impacts would be quantified. Another even more radical tact might be to abandon consequentialist approaches and try to develop evaluative criteria that are based upon a non-consequentialist assessment of the merits of the goals being pursued by the policy

architects, rather than assessing the policy consequences for individual human beings. But should we really ignore consequences in evaluating policies? Or we could perhaps even embrace criteria that are more explicitly theological in nature and that do not rest upon either the intentions of the actors or the policy impacts upon human welfare, but only upon their congruence with God's grand plan? But is there any chance that we would be able to reach even a rough consensus on which theological premises are most accurate?

Determining how to take person-altering consequences into account in policy assessment is a very difficult question, and I concede that I do not have the answers at hand. But it is clear to me that simply ignoring the problem this poses for efficiency analyses by making a demonstrably false assumption as to the invariability of the genetic identity of future persons with regard to policy impacts is not a viable approach. This interesting and difficult conundrum should to be

made clear to students as part of their immersion in efficiency analysis. I have found that my students find the problem of person-altering consequences to be quite interesting, and they are quickly able to grasp its great significance for undercutting the usefulness of conventional efficiency analysis.

IV. NORMATIVE PROBLEMS WITH THE ECONOMIC EFFICIENCY

CRITERION

Even if one is somehow able to overcome (or chooses to ignore) all of the above-discussed measurement and conceptual difficulties in order to obtain a meaningful estimate of the aggregate willingness to pay-based impacts of a policy, there is still the fundamental question presented regarding whether economic efficiency is an appropriate normative criterion for assessing the merits of that policy. There are two major concerns here that deserve discussion in a Law and

Economics class, and that are fortunately relatively easy to convey to students in a succinct manner.

First of all, a person's willingness to pay to experience (or to avoid) the consequences of a policy are obviously not only a function of that person's preferences but also a function of that person's wealth. The efficiency criterion can be succinctly described as a "one dollar, one vote" aggregate decision criterion. It is definitely not the more democratic "one person, one vote" criterion, and the efficiency criterion systematically gives more weight to the preferences of more affluent persons in proportion to their relative wealth in policy assessment.

One can argue that the amount of an individual's wealth is positively correlated with productivity and with other socially desirable traits, at least to some substantial extent, and that this fact provides some social justification for giving greater weight in policy assessments to the preferences of more wealthy

individuals in order to provide incentives to people to further develop those desirable personal traits. But the modest degree of correlation between wealth and virtue obviously allows for many exceptions. Many individuals are wealthy through inheritance or good fortune without necessarily being productive or otherwise virtuous (anyone prominent come to mind here?), and other people are productive and/or otherwise virtuous yet poor, and so arguments can easily be made that would cut the other way against giving such overriding weight to relative wealth differentials in making social decisions. It is clearly a matter for ongoing and perhaps ultimately unresolvable debate as to whether a willingness to pay-based decision-making criterion is more appropriate, all things considered, than would be the more democratic alternative of giving equal weight to the preferences of each person, or than would be, as another example, the alternative of adopting a utilitarian-type criterion that focused more directly upon the relative

psychological impacts of a policy upon different persons apart from their relative wealth.

There is a sophisticated welfare economics literature that is not usually assigned in Law and Economics courses that makes clear that there are a large number of possible alternative “social welfare functions” that can be used to aggregate the impacts on different persons of a policy in order to evaluate its consequences.²¹ The willingness to pay-based algorithm is only one of many possible valuation yardsticks, and one with no special justification other than the fact that it closely parallels how the market prices for goods or services are set in competitive markets. The inherent arbitrariness of the choice of the economic efficiency criterion as an evaluative standard can be easily communicated to law students in a clear and non-technical manner.

Secondly, and perhaps even more importantly, there are no “rights” inherent in an efficiency calculation that need be respected. If a person’s willingness to pay-based assessment of the costs that a policy imposes upon them is exceeded by the net aggregate benefits conferred on the other impacted persons, then the efficiency criterion will still endorse that policy no matter what the size or nature of the costs that are imposed upon that unfortunate person. The use of an evaluative criterion that does not recognize any rights except the right to vote to the limit of one’s willingness to pay obviously co-exists uneasily with the fundamental values of our larger social and political system that is based on a Constitution that clearly embraces the idea of inalienable rights that are not to be infringed upon regardless of benefits thereby conferred upon other persons, especially if willingness to pay is measured by necessarily constrained offer prices rather than by asking prices.

V. CONCLUSION

Some commentators have criticized the Law and Economics courses offered in law schools as often being biased in favor of influencing students to uncritically apply the evaluative criterion of economic efficiency in policy analysis despite its conceptual and empirical shortcomings, and despite the fact that it privileges the interests of the wealthy and disregards rights. There is merit to these criticisms, and to some extent the introduction of these courses into law school curricula has been facilitated by funding provided by very wealthy persons who hoped to further institutionalize these normative biases within the legal community. But I do not believe that we need to throw out the beautiful baby of the valuable analytical tools and perspectives developed by economists in order to rid ourselves of the dirty bathwater of a poorly defined and ideologically biased efficiency criterion.

It is entirely possible for a teacher of this course to convey to their students an appreciation of the substantial conceptual and empirical difficulties involved in quantifying the efficiency consequences of a policy. It is certainly relatively easy to bring to the students' attention the rather obvious normative concerns raised by the use of a "one dollar, one vote" efficiency criterion that in addition does not recognize any limitations based upon rights. If these crippling shortcomings of efficiency analysis are effectively communicated then I believe that the bias of many Law and Economics courses in favor of the interests of the wealthy and against the consideration of rights that is often decried by critics will be eliminated. Students who are made aware of the limitations of the economic efficiency criterion that I have here described will hopefully be able to counter policy arguments that are based on an efficiency premise, and will understand the need to disregard or at least supplement efficiency assessments of policies with the application of other evaluative standards.

¹ Professor of Law, Dedman School of Law, Southern Methodist University. J.D., Yale Law School, Ph.D., University of Iowa.

² I will refer generically to those widely offered upper-level elective law school courses that are intended to apply basic microeconomic concepts to analyze the operation of various aspects of the legal system as “Law and Economics courses,” even though such courses sometimes are offered under different through similar titles such as the “Economic Analysis of Law” (the course title used at the Southern Methodist School of Law where I teach), or “Law and Economic Analysis,” etc.

³ When I use the terms “economic efficiency” or “efficiency” in this essay I am consistently referring to the Kaldor-Hicks efficiency criterion that is the evaluative criterion generally applied in practice, and not to the Pareto-efficiency criterion which because of its unanimity requirement has very little practical application.

⁴ See generally Jean Mayer, Dark Money: The Hidden History of the Billionaires Behind the Rise of the Radical Right (2016), at 92-119.

⁵ *Id.*

⁶ Richard Posner, Economic Analysis of Law (8th ed. 2-011), at 17-18.

⁷ This area will actually not be precisely triangular unless the demand curve is a straight line, but this fact does not affect the usefulness of this framework for illustrating this basic point regarding the determination of the size of the consumer surplus.

⁸ A buyer’s reservation price” for a good or service represents the highest price that they would be willing to pay for that good or service rather than to have to do without it.

⁹ This area will actually not be precisely triangular unless the supply curve is a straight line, but this fact does not affect the usefulness of this framework for illustrating this basic point regarding the determination of the size of the producer surplus.

¹⁰ A seller’s reservation price” for a good or service represents the lowest price that they would be willing to accept for that good or service rather than to keep it unsold.

¹¹ Consider, for example, the very large asking prices that might be expressed by some persons sincerely opposed to abortion as a matter of principle with regard to a policy whereby the federal government facilitated and fully financed all abortions.

¹² Many of which are done by governmental agencies and styled as “cost-benefit analysis,” a common and synonymous term for a Kaldor-Hicks efficiency assessment

¹³ For example, a policy whose costs occur largely in the near term and whose benefits are mostly long-term, such as many environmental protection policies, and whose benefits just barely exceed its costs using a 3% discount rate, will appear to be much more wealth-increasing if 1% discount rate is used, and may appear to be highly inefficient under a 7% discount rate.

¹⁴ Consider, as one extreme example, a policy that involved a forced brain lobotomy for one person who opposes that operation. It would appear to most people to be more appropriate to value the costs of that policy by that person’s willingness to pay to avoid that operation that results from the pre-lobotomy preference structure, than to value them by their willingness to pay based on their post-lobotomy preferences!

¹⁵ For example, policies that impose educational requirements of persons, or that provide them with environmental amenities that they have not previously experienced, have been shown to sometimes alter preference structures in a way that is more favorable to those impacts.

¹⁶ See, e.g., Derek Parfit, Reasons and Persons (1984), at 351-80.

¹⁷ *See, e.g.*, Gregory Scott Crespi, “The Endogeneity Problem in Cost-Benefit Analysis,” 8 Geo. J. of Law & Pub. Pol. 91 (2010), at 118-44; Gregory Scott Crespi, “The Fatal Flaw of Cost-Benefit Analysis,” 38 Envtl. L. Rep. News & Analysis 10703 (2008), Gregory Scott Crespi, “What’s Wrong With Dumping Radioactive Wastes in the Ocean? The Surprising Ethical and Policy Analysis Implications of Person-Altering Consequences,” 37 Envtl. L. Rep. News & Analysis 10873 (2007).

¹⁸ *See, e.g.*, Crespi (2007), *id.*

¹⁹ Consider the ocean radioactive waste dumping policies considered in Crespi (2007), *supra* n. 17.

²⁰ Such as an environmental protection measure that conferred very large benefits upon the members of future generations but that imposed small net costs on current persons.

²¹ [add cite to the welfare economics literature]