



GW Law Faculty Publications & Other Works

Faculty Scholarship

2013

The Role of Economics in Tax Scholarship

Neil H. Buchanan

George Washington University Law School, neilhbuchanan@gmail.com

Follow this and additional works at: http://scholarship.law.gwu.edu/faculty_publications

 Part of the [Law Commons](#)

Recommended Citation

Neil H. Buchanan, The Role of Economics in Tax Scholarship in BEYOND ECONOMIC EFFICIENCY (David A. Brennan, Karen B. Brown & Darryl Jones eds., forthcoming 2013)

This Article is brought to you for free and open access by the Faculty Scholarship at Scholarly Commons. It has been accepted for inclusion in GW Law Faculty Publications & Other Works by an authorized administrator of Scholarly Commons. For more information, please contact spagel@law.gwu.edu.

Chapter I:

(Neil H. Buchanan)

The Role of Economics in Tax Scholarship

Introduction

One of the fundamental tenets of modern tax policy analysis is that we should be concerned with so-called economic efficiency. Along with equity and administrability, efficiency is widely held to be a desirable and important goal. Indeed, to some analysts, efficiency is the most important of those goals, and perhaps the *only* appropriate goal of tax policy. Even for those who still take seriously non-efficiency concerns, however, efficiency is at least a central element of tax policy analysis, to be weighed against the other two goals (and, perhaps, some others).

All tax policy proposals are thus scrutinized to determine whether they enhance or reduce the economy's efficiency. If economic efficiency were a coherent concept, it would indeed be extremely useful in the analysis of tax policy, as well as in many other areas of the law. Unfortunately, there is no substance underneath the often-impressive superstructure of efficiency analysis. This makes it not just unwise, but affirmatively misleading, to base academic analysis of taxation – in whole or in part – on attempts to measure and maximize efficiency.

These assertions are likely to strike some readers as surprising, if not shocking. So much modern scholarly analysis takes for granted the usefulness of efficiency analysis that legal scholars might not stop to think about whether efficiency means anything at all. Efficiency is simply something that everyone thinks they understand. As I will argue in this chapter, however, there is “no there there,” when it comes to efficiency analysis. Efficiency is ultimately nothing

more than a seductive guise behind which unexamined assumptions and normative preferences hide.

This does not, however, mean that we must throw out everything that we think we know about using economics in analyzing tax policy. Knowing that efficiency is an empty concept does not mean that all bets are off – that there are no principled ways in which to analyze tax issues. It does mean that we can no longer pretend that there is a “clean answer” that is untouched by a scholar’s judgment, values, and choices. Therefore, what we gain by confronting the emptiness of the efficiency trope is the opportunity to sharpen our analysis of the real issues at stake in any assessment of tax policy.

Efficiency Analysis, and Its More Sophisticated Extensions

Efficiency is a word with multiple meanings. It often implies something as simple as turning out the lights when one leaves the room. In academic analysis, however, efficiency is a much more freighted concept. Also known as “economic efficiency” or “Pareto Efficiency” (in honor of Vilfredo Pareto, the 19th-Century economist who was central to developing the concept), “efficiency” as used by modern economic scholars purports to measure something much more subtle and important than whether an item is “inexpensive” or “cheap.”

Because of that important distinction, it is necessary to use a consistent label to describe economists’ non-intuitive notion of efficiency. Therefore, I will use Pareto Efficiency consistently throughout this chapter, to distinguish what economists mean by “efficiency” from what non-economists might imagine that the word implies.

This distinction is, moreover, not merely a matter of agreeing upon a common vocabulary. The word efficiency, despite its vagueness, carries with it a powerfully positive connotation. Describing something as inefficient is tantamount to condemnation, while being

efficient is much to be desired. It would be foolish, therefore, to allow scholars to take on the mantle of being “in favor of efficiency” when their analyses lead them to make policy recommendations that are anything but objectively desirable.

The definition of the word efficiency is also important because so many economists put great weight on the claim that their analyses are “value-neutral,” being merely a matter of cold calculation without taking into account one’s views on morality or philosophy. Describing something as “efficient,” therefore, ultimately relies on the familiar (but misleading) distinction between “positive” analysis and “normative” analysis, where positive analysis supposedly takes no position regarding matters of values and morality, while normative analysis allows an analyst to insert her own views about what is good and bad.

Pareto Efficiency is generally held out as a positive concept, with no normative content. If that were true, Pareto Efficiency would arguably be an improvement upon a world in which disagreements were mere matters of opinion. While one can readily see why such an approach is appealing, it is unfortunately not possible to remove normative concerns from supposedly positive analyses. Indeed, as I will discuss below, Pareto Efficiency is ultimately based on very strong – but usually unstated – normative assumptions.

The Standard Definition of Efficiency

But what does Pareto Efficiency mean? The standard definition is usually phrased as follows: “A situation is Pareto Efficient if it is not possible to make anyone better off without making someone else worse off.” The basic idea, therefore, is that Pareto Efficiency analysis involves comparing gains and losses caused by putting people in different situations, with a Pareto Efficient situation representing the maximum amount of “well-offness” in the society in question.

How do we measure how well-off people are? Doing so requires choosing a theory of value, that is, a basis to determine how much something is worth, and a metric by which to measure that value. Pareto Efficiency analysis, as currently understood, takes as its theory of value the so-called Willingness-to-Pay Principle. This means that value can be measured by asking how much someone would be willing to pay for something, and comparing it to how much people would be willing to pay for something else.

Once we can measure value, we can then compare the values of putting people in different situations. If a person possessed, say, five lamps and ten hats, then we could compare the value of that person's situation with having seven lamps and eight hats – or, for that matter, with having neither lamps nor hats, but having twenty bottles of wine.

This definition of value, therefore, purports to take individual preferences and personal sovereignty seriously. Rather than imposing an external measure of value on people, Willingness-to-Pay measures value by watching people put their money where their mouths are. Value is, therefore, determined subjectively but measured objectively.

Pareto Efficiency analysis involves applying this method of measuring value to alternative possible realities. From any starting point, an economist asks, is it possible to rearrange things so that no one is worse off than they are now, and at least one other person is better off than she is now? If so, then the current situation is not Pareto Efficient. We would only stop when we have reached the situation in which no further rearrangements are possible without leaving someone worse off.

Applying this concept to tax policy has become numbingly familiar in modern tax analysis. We measure the gains from any possible change in policy against the losses from such a change, and declare a situation as Pareto Efficient when all possible changes from that situation

involve greater losses than gains. If gains are possible, then achieving Pareto Efficiency would require us to capture those gains by changing the tax laws.

In the simple graphs used for standard tax analysis, the shortcut for determining whether a situation is or is not Pareto Efficient is to look for the existence of so-called “deadweight loss triangles.” In the typical analysis, there is one and only one point on the price/quantity graph in which there is no deadweight loss triangle, making that point the Pareto Efficient point.

Market solutions

One of the most important theoretical claims in modern economics is that “properly behaved” markets will reach a Pareto Efficient equilibrium on their own, without any need for planning or guidance on the part of the government (or any other extra-market entity). This claim is, therefore, simply a modernized version of Adam Smith’s famous “invisible hand,” in which a well-defined market allows self-interested individuals to act on their own behalf, seeking out advantage in transactions in a way that guarantees that every good and service ends up being consumed by the person who values it most highly.

If that is true, then it necessarily follows that any change in tax policy (which, in this simplified world, means the introduction of any tax or subsidy) would disturb a market’s Pareto Efficient outcome, creating deadweight loss where none had existed before. This means that the government’s decision to tax or subsidize *anything* must reduce the total well-offness generated by a market, compared to a market in which the government has not intervened.

In short, the standard economic analysis ends up being a brief for non-intervention by the government in markets, that is, for a *laissez-faire* approach to governing. If self-interested transactions end up maximizing value, then there is simply no need to undertake policies to try to

make people happier. Not only will such policies by definition fail (because aggregate happiness is already at its maximum), but they will inevitably make matters worse.¹

When Are Taxing and Spending Beneficial?

Of course, any competent economist would readily acknowledge that the analysis above is absurdly simplified, and that there are often situations in which it is possible for the government to engage in policies that will make people better off in the aggregate. Entire areas of scholarly analysis are based on exploring such possibilities. The key, of course, lies in the term “well-behaved,” as in the description above.

Using taxes to reduce pollution, for example, can (under very defensible assumptions) be shown to move from a non-Pareto Efficient situation to a Pareto Efficient one. If we recognize that there are costs and benefits that are not borne by the people who make decisions – that is, if there are externalities – then we can improve the way such a market behaves by adopting any of a number of approaches to forcing people to face the full consequences of their decisions.

Similarly, if we are concerned that not everyone acts in the cold, calculated, self-interested ways required by standard economic theory, then there might be a number of ways in which we can improve market outcomes by assisting people in overcoming various cognitive

¹ There is a serious problem with the very notion of aggregating preferences, that is, with being able to add together the amount of well-offness that different people feel. This problem is generally assumed away, with “well-behaved” economic models requiring that people’s individual happiness generally be unaffected by changes in other peoples’ happiness. For example, standard models generally assume away both altruism and malevolence. This set of problems is fundamentally important, but a discussion of them in any detail would require a technical diversion not appropriate here

biases. Or, when we acknowledge that not all markets allow for the kind of competition that is presumed by the standard theory – such as when a firm monopolizes a market and excludes new entrants – then we have opened up the possibility that the government can improve matters by creating or restoring the kind of competition required by standard economic theory.

Each of these possibilities is, as noted, not only well understood, but forms the basis for specialized study within economics. Behavioral economics, industrial organization, public finance, and other fields of specialization are dedicated largely to understanding when and how real-world markets diverge from the theory that supports *laissez-faire* economics. More sophisticated analyses confront the question of whether it might still be better to leave things alone, despite being in a non-Pareto Efficient situation, because of the likelihood that any intervention might (or might not) make matters still worse.

Many of the chapters in this volume dedicate serious efforts along these lines, taking into account real-world considerations that enrich the analysis beyond the basic anti-tax presumption that otherwise emerges from standard economic theory. Such analyses demonstrate that, even thinking within the confines of standard economic analysis based on Pareto Efficiency, there remains important uncertainty with regard to a wide range of possible policy interventions.

The Fundamental Problem with Pareto Efficiency Analysis

None of the discussion in the previous section, however, addresses the central claim of this chapter, that Pareto Efficiency is a fundamentally empty concept that ultimately fails to advance our understanding of tax (and other legal) analysis. Indeed, it might appear that the standard approach is supple and sophisticated, allowing scholars to extend a fundamental model beyond its elements, to include an impressively wide range of possible issues to be analyzed under a common rubric. This is, however, a mirage. The supposed expansiveness of the

economic approach is actually its downfall, making the approach either a veil for unstated assumptions and norms, or an undisciplined analytical method that can lead to any conclusion at all.

What is Wrong with Following Professional Norms?

As a strategy of advancing understanding and sustaining a productive argument, it is understandable that scholars would seek to extend rather than reject established analytical frameworks. Rather than disagreeing on the most fundamental bases of any question, it is often helpful to accept *arguendo* the majority of the assumptions of another's analysis, and then tweak those assumptions to demonstrate that a different outcome follows. This might be especially true for legal scholars, who are adept at pulling together insights from many disciplines, but who typically are not in the business of changing how any of those disciplines work.

When working within a discipline, of course, it is generally necessary to build upon the structure that has already been built. Many budding economists are tempted during their undergraduate or early graduate studies to try to tear the entire structure of economic theory down to its foundations, in an attempt to start over and reinvent modern economics. Their advisors sensibly tell them that this is not how knowledge accumulates, suggesting that the better path is to become conversant in the discussion that has been established in the profession, and then to add to that discussion at the margins.

Unfortunately, this combination of incentives facing various scholars – non-economists who are understandably hesitant to engage in deep critiques of established economic theory (if, for no other reason, because they do not feel competent in the nuances of a highly technical field), and economists who must accept the basic presumptions of their field – creates a dynamic in which the deep assumptions underlying economic analysis are generally ignored or

suppressed. Even so, a great deal of important work has been done along these lines, but it is nonetheless true that the vast bulk of economics-oriented analysis (both inside and outside of economics departments) continues to be based on a theoretical framework that has been shown to be incoherent.

The Baseline Problem

An essential assumption behind Pareto Efficiency analysis is that there is a neutral starting point that forms the baseline for measuring the consequences of people's economic interactions. That is, when we say that people compete fairly in markets, engaging in mutually-beneficial trades that reflect their own assessments of their desires and priorities, we act as if the constraints that shape people's decisions are somehow neutral.

In standard microeconomic analysis, for example, each person is assumed to be trying to maximize her well-being based on her preferences, her income and her assets, as well as on the prices that she faces as she considers making purchases. Each person is thus constrained by her circumstances, and she does the best that she can under those constraints. Nothing in Pareto Efficiency analysis, therefore, is understood to mean that people are in some objective sense happy or well-off, but only that people are in the aggregate as happy as they can be – *while operating under the constraints that the world imposes upon them.*

But what is the nature of those circumstances that constrain people, and how do those circumstances become the baseline used to measure Pareto Efficiency? The most obvious issue, of course, is the distribution of income and assets. It is widely known that, if one starts from a baseline in which one person possesses all of the productive assets in a society, there is a perfectly valid Pareto Efficient outcome in which that one wealthy person keeps everything and maximizes his happiness, while others starve.

One common response to this extreme possibility is to say that the aggregate happiness in society will go up when other people are made less miserable, because the marginal loss to the person who possessed everything must surely be smaller than the marginal gain to those who had nothing. While that seems likely to be true as an empirical matter, one cannot say that it is definitely true. If it is not, then Pareto Efficiency would still coexist with severe human misery.

We could, of course, simply assume that transferring resources from rich to poor will increase aggregate happiness. Doing so would allow us to label a redistributive tax system as “achieving Pareto Efficiency” rather than “creating deadweight losses,” but only because we are willing to make an empirical assumption in the absence of evidence.

One of the fundamental baselines of Pareto Efficiency analysis, therefore, is the initial distribution of resources that people possess when they enter the marketplace. If we change the initial distribution, then that new distribution of resources can just as easily be used as the baseline for analysis as any other.

In other words, there is nothing about the current distribution of resources that should give it analytical priority over any imaginable alternative. If, in Situation A, a person would start with one million dollars in assets, while in Situation B, that person would start with zero dollars in assets, then there is no consistent way to assess (from a Pareto Efficiency standpoint) a policy that takes one million dollars away from that person. Viewed from Situation A, this would amount to taking away his rightful endowment, leading to Pareto Inefficient results. Viewed from Situation B, however, this restores the person to his rightful initial endowment, guaranteeing Pareto Efficient results. Even something in between, such as a 50-50 split, does not save the analysis, because the question is what baseline to use when measuring Pareto Efficiency. There is no neutral way to elevate one endowment over any other.

The larger point, therefore, is that it is always possible to criticize as Pareto Inefficient the same policy that, from a different baseline, enhances Pareto Efficiency. The choice of baseline itself, however, is not based on any overarching principal.

The Legal Baseline Problem

While the baseline of income and assets is the most intuitively obvious area in which a theory based on “willingness to pay” (and, therefore, on ability to pay) falters, the question of defining a baseline ultimately implicates even more fundamental issues of ownership and rights.

Recall that the basic idea behind any measurement of Pareto Efficiency is that people will buy and sell goods and services in “free markets.” What is generally unacknowledged is the wide range of possible rules under which such “free” markets might operate. Even the freest of markets operates in the shadow of the law.

The basic presumptions behind buying and selling, after all, are that a person possesses the legal right of ownership for something, and that she is capable of transferring that right of ownership to someone else. The rules of property law, therefore, are fundamental in determining what people can and cannot do in the marketplace. The classic scam of selling the Brooklyn Bridge to unsuspecting tourists captures this notion in its extreme form, with a party purporting to be able to sell something to another person, when the seller does not even own what he is selling.

Even short of such outright fraud, however, the nature of ownership is deeply problematic and contingent. The possible rules of property law are so diverse as to defy the possibility of defining a presumptive baseline set of property rules (with all other possible property rules thus defined as Pareto Inefficient). Consider, for example, the range of rules regarding land ownership that have existed in England and the United States over their histories.

In most jurisdictions, even a legally recognized owner can be deemed to have lost ownership in the property by “adverse possession,” a rule which allows the state to sanction acquisition by a second person who has occupied the land without adequate contest by the nominal owner. Each jurisdiction, however, must set the rules determining what a current owner must do to adequately contest another’s adverse possession: the number of years can vary, the nature of “open and notorious” possession is contestable, and so on.

One might argue that even having a law recognizing adverse possession under any circumstances is a violation of the proper baseline. If a person owns property, then we could say that she will continue to own property unless she sells it, gives it away, or dies. Even to take that position, however, at least requires recognizing that an underlying property rule that currently exists nearly everywhere is “not the proper baseline.” That would necessarily imply that measuring Pareto Efficiency against the baseline of current legal rules is invalid. We would, instead, need to know what people would be willing to pay for all of their various goods and services if there were no adverse possession laws.

Arguing that such a minor law could not affect a large number of transactions is an insufficient response, especially because so many different current laws can be described as violating a proper baseline. Even if none individually appears to be significant, the totality of the arguable violations of the proper baseline cannot be so easily dismissed.

Moreover, even if it were possible to agree that ownership is only extinguished by sale, gift, or death, each of those issues raises its own set of baseline questions. Tax scholars know as well as anyone how difficult it is even to define when something is a gift. When is an arrangement merely a loan? Allowing a person to use property does not necessarily amount to giving a gift, nor necessarily a sale, either.

Sales, of course, move us from the realm of property law into the realm of contract law, which presents its own set of arbitrary choices of what counts as a valid agreement and what does not. Some laws prevent people from selling what they own (including those that prevent parents from selling their children – or themselves – into slavery), making the concept of ownership much more complicated than it initially appears.

In short, legal rules define all of the boundaries that define the limits of what can be bought and sold in open markets. There is no “natural” baseline for any of those rules. It might be tempting, however, to imagine that most of these property and contract rules are so well established, and are so unlikely to change, that we can confidently take them in their current form as the baseline. Making such an assumption, one might argue, would not allow us to measure Pareto Efficiency perfectly, but we could at least closely approximate a proper baseline.

Even here, however, the reality is that these basic legal rules are changing all the time. Intellectual property laws define entire categories of non-physical items that can be bought, sold, and owned. A law that would lengthen or shorten the number of years of copyright or patent protections would most definitely change a person’s willingness to pay for any particular piece of intellectual property – and, as a direct consequence, would also change their ability to buy other goods and services with the money that they earn from their intellectual property. Congress regularly legislates in these areas, and regulatory agencies must fill in gaps in those laws. Each such decision changes the baseline against which one would measure Pareto Efficiency, with owners of property viewing the lengthening or strengthening of property rights merely as an affirmation of their natural ownership rights, whereas others would be equally justified in viewing legal changes to shorten or weaken such rights as merely a move toward a different (and arguably better) baseline.

Finally, consider the complications involved in the rules governing property at death. Even setting aside the tax consequences of estate planning, the rules of succession must include line-drawing decisions affecting – among a wide variety of other issues – those who die intestate, the rules available to challenge the validity of wills, and any prohibitions against disinheriting spouses or minor children. Even the most notorious doctrine in all of property law, the Rule Against Perpetuities, has recently been abandoned by some jurisdictions, while being retained in modified form in others.

With all of those moving parts, therefore, it is simply impossible to define a neutral baseline against which we could say that a Pareto Efficiency analysis of all market decisions must be based. A person's willingness to pay for goods and services under one set of laws will necessarily differ – and could certainly differ substantially – from one's willingness to pay under another set of laws, even if the differences in the two sets of laws appear to be minor on the surface.

Taxes and the Government

One important implication of the baseline problem was recently developed in a notable book, *The Myth of Ownership*, by the legal philosophers Liam Murphy and Thomas Nagel. They primarily focused not on the range of possible laws that a government might pass to define the baseline of market transactions, but rather on the variety of ways in which that government might finance itself. Levying taxes on one or a combination of tax bases – income, estates at death, property more generally, sales, wages, and so on – allows a government to determine after-tax incomes in a way that cannot be evaluated meaningfully from the standpoint of Pareto Efficiency. Decisions by governments to levy taxes are thus as consequential for the baseline question as are decisions about property law, contract laws, criminal laws, and so on.

What is now commonly referred to as “the Murphy/Nagel” point is that it is not coherent even to refer to someone’s before-tax income, because the level and type of taxes collected will inevitably affect the government’s decisions and ability to pass and enforce its laws, which in turn will affect how much money a person can earn (and in many cases even the types of businesses in which she might work). Saying, “I would have had this much money *if there were no government,*” in other words, is to engage in a meaningless hypothetical exercise.

The concept on which the Murphy/Nagel point is built, however, is that the government itself defines and enforces the rules that make market transactions possible. Without guarantees of ownership, respect for the rules of transactions, and so on, there is no commerce. Commerce, therefore, logically requires government and taxes, because economic transactions presume the existence of a government that enacts and enforces the rules by which people transact.

Pareto Efficiency and Value Judgments

In the end, it might be tempting to dismiss all of the analysis above as needlessly complicating a simple situation. Sure, one might say (as I have actually heard people say), there are many different legal rules that we might adopt, but we are stuck with the ones that currently exist. We should thus use current reality as the baseline, and leave the possible alternative worlds to philosophers and science fiction writers.

As much as that argument might sound like an attractive form of pragmatism, however, it actually misses the entire point of all tax analysis – and, indeed, of all legal analysis. No one argues that the current set of laws is perfect. Those who use Pareto Efficiency as a way to evaluate policy make the claim that certain policy changes will move us toward a superior situation – that is, a situation that has smaller or nonexistent deadweight loss triangles – because

the preferred policies would “correct” some “imperfections” in the market that are currently preventing the market from achieving a Pareto Efficient outcome.

As the analysis above makes clear, however, one person’s imperfection is another person’s baseline law. Taking away something that benefits one person, based on the belief that it is an unnatural deviation from Pareto Efficiency, can only be justified if the thing that is being taken away is not part of the background of laws and regulations that allow us to determine Pareto Efficient outcomes in the first place. Because there is no hierarchy of laws that allows us to determine which are required and which are optional, however, there is no meaning to the statement: “This policy would improve Pareto Efficiency.” Even a more careful statement, such as “Under the baseline assumptions on which I am relying, this policy would improve Pareto Efficiency,” is nonetheless unhelpful, because the ready retort is: “Yes, and under other baseline assumptions, this policy would increase Pareto *In*efficiency.”

As a result, those who make evaluative statements about Pareto Efficiency are (usually unconsciously, as far as I can tell) generally making value judgments – normative assessments – but cloaking them in the language of efficiency and thus the supposed neutrality of positive analysis. By including within the baseline the values that one holds dearest (for example, views about what the employment laws should say about gender and race discrimination), one is then able to characterize the policies that one does not like as unnatural and economically wasteful.

An Example: When Is a New Law Efficient?

Given the relatively abstract nature of so much of the above analysis, a specific example might be especially helpful to clarify the point. An economist friend of mine used to constantly remind me that the notion of Pareto Efficiency is quite broad, because what counts as “well-offness” is in the eyes of all of the various people whose welfare we are trying to maximize in

the aggregate. Because there is no accounting for tastes, he argued, one must not fall into the trap of describing outcomes as Pareto Inefficient merely because of disagreements over the specific transactions that take place in a market. If enough people love auto racing, then profit-seeking media outlets will cover auto racing. If people love opera, then the government has no business banning or limiting the performance of operas.

The idea, therefore, is not just that markets act through the invisible hand to maximize well-offness, but that we can determine what people care about by what they buy and do not buy. People would then petition the government to make sure that the laws comported with their own preferences, to allow the laws to reflect their values and maximize their ability to use markets to achieve happiness.

This economist was, however, incensed one day when he learned that his state government had jailed a person for committing an extreme act of animal abuse. His argument was that this was “not Pareto Efficient.” I asked how one could possibly know that this did not maximize the aggregate of human happiness (by, for example, allowing people to know that animal cruelty would be somewhat reduced, offsetting the cost of incarcerating abusers). In the new legal regime that included stricter laws against animal abuse, one could argue that the regime was now accurately reflecting people’s preferences or, just as plausibly, that the new regime changed the baseline and thus moved us away from Pareto Efficiency. Without a clear baseline, one could not draw a conclusion either way.

This is not, moreover, merely a problem of not being able to measure people’s true preferences, because “true preferences” are themselves defined by the nature of the laws. People who lived in a world where bear-baiting was a regular activity could be described as having a “preference” for bear-baiting, but those who live in other places and times apparently do not.

Moreover, one cannot simply say that what the people decide through their government is presumptively the Pareto Efficient rule – at least, not without completely undermining the entire enterprise of making policy assessments. Many economists frequently inveigh against what the people desire, saying, for example, that minimum wage laws and rent controls (passed by democratically-elected bodies) are Pareto Inefficient. The nature of Pareto Efficiency analysis requires us to be able to say that some laws are unwise, even if they were passed by a legitimate government, because they do not maximize aggregate happiness. On the other hand, a presumption that all laws are Presumptively Inefficient is equally baseless, because some laws are necessary to create the baseline in the first place

My economist friend, therefore, presented what was really nothing more than his normative judgment that animal abuse is not serious enough to warrant jail time as a positive statement about Pareto Inefficiency – even though he was otherwise keenly aware that Pareto Efficient outcomes can be ones that particular individuals find quite unwise.

More generally, therefore, the “baseline problem” described here undermines claims that one can engage in amoral assessments of market outcomes, describing some outcomes as Pareto Efficient even if one personally might dislike the result of market transactions. Because Pareto Efficiency analysis requires knowing and agreeing on the proper baseline, which cannot be known, it is incoherent to describe something as Pareto Efficient, in either an absolute or comparative sense.

Economics Is Useful, If We Understand Its Limitations

Have offered such a counter-intuitive thesis, it is essential to be clear about its limits. Because so many scholars tend to think of “economic analysis” and “efficiency analysis” as synonymous, it would be tempting to imagine that this chapter represents an attack on “economic thinking.” It most definitely is not. Any competent analysis of taxation must certainly address

issues that are commonly thought of as “economic issues.” The point of this chapter, while important, is nonetheless quite limited: The Pareto Efficiency criterion is not a coherent or meaningful way to assess laws or policies.

That statement still leaves plenty of room for the use of “economic tools” to assess policies. The most important of these is the simple but powerful question of how people respond to incentives. This is especially potent in tax policy analysis, but it is useful in any assessment of public policies.

Consider, for example, the problem of health care costs. The primary economic question currently confronting the United States regarding health care is how to reduce the rising costs of such care. While one could imagine engaging in an analysis that attempts to determine whether any particular policy change enhances Pareto Efficiency, in an attempt to determine the optimal level of health care, the analysis above demonstrates that the outcome of any such analysis would not provide useful guidance.

On the other hand, we can make important strides in understanding the path of health care costs by looking at the nature of the incentives in the health care system. One of the most notable aspects of our current system, for example, is that health care providers are generally reimbursed on a per-treatment basis. This means that doctors and other providers make more money when they perform more procedures, no matter whether those procedures actually help the patients. A different reimbursement model would create different incentives, and thus lead to different behavior and potentially lower costs.

This suggests that thinking about the economics of the health care market is essential, even though Pareto Efficiency need not (indeed, should not) be part of that analysis. Simply analyzing – and, most importantly, trying to measure – the effects of various policies and rules

on people's decisions allows us to gather information that we can then use in forming policy judgments.

Similarly, analysis of tax issues is enriched when scholars try to predict how people will react to changes in the law. Will the amount of money donated to charities fall if the marginal tax rate on millionaires is decreased? Will employers lay off workers if the unemployment insurance tax is increased? Will businesses shut down if the estate tax is increased?

The core of good scholarly analysis of tax law, therefore, involves thinking through how changes in the law will change behavior, and assessing the consequences of those changes. While Pareto Efficiency analysis purports to do just that, it is ultimately nothing more than an elaborate superstructure that actually requires (but hides) value judgments, without offering any independent or positive insights that are not otherwise available.

As long as value judgments must be made, it is far better to own up to those judgments up front, rather than hiding them behind an incoherent notion of efficiency. The common strength of the remaining chapters in this volume is that the authors confront the real issues, including the moral choices facing policy makers, and draw careful conclusions based on thorough analysis.