

University of Chicago Law School

Chicago Unbound

Journal Articles

Faculty Scholarship

2003

Should Legal Rules Be Used to Redistribute Income?

David A. Weisbach

Follow this and additional works at: https://chicagounbound.uchicago.edu/journal_articles



Part of the [Law Commons](#)

Recommended Citation

David A. Weisbach, "Should Legal Rules Be Used to Redistribute Income?," 70 University of Chicago Law Review 439 (2003).

This Article is brought to you for free and open access by the Faculty Scholarship at Chicago Unbound. It has been accepted for inclusion in Journal Articles by an authorized administrator of Chicago Unbound. For more information, please contact unbound@law.uchicago.edu.

Should Legal Rules Be Used to Redistribute Income?

David A. Weisbach†

The University of Chicago Law School has been central to the development of law and economics over the past half century. Central to law and economics is the norm by which one should judge the consequences of legal rules. The overwhelming majority of law and economics scholarship looks solely to efficiency to evaluate legal rules. The question I address in this Essay is whether or why this is appropriate.

More specifically, in thinking about legal rules, we must ask whether they should be designed to redistribute or whether they should merely be efficient. For example, we might want to know whether to have pro-tenant landlord/tenant rules on the theory that tenants are systematically poorer than landlords. Should we have pro-consumer warranty rules on the theory that consumers are systematically poorer than producers? And should we have pro-plaintiff nuisance rules?

A common thought might be that legal rules that redistribute income are a good idea because they help the poor and disadvantaged. Laws of this sort can be viewed as a means of doing social justice. Law is not some sterile instrument that comes out of a vacuum. Instead, law is a means of improving our society.

A common counterclaim in law and economics is that legal rules should not be used to redistribute to the poor. Instead, legal rules should be efficient. While we see the efficiency claim often, it is not always clear on what it is based. Do law and economics scholars simply not like the idea of redistribution? Are they uninterested in social justice?

I argue that legal rules should not be designed to redistribute to the poor—but this is not out of a lack of concern for distribution or equality. My argument is that there is a better method of addressing these concerns. In particular, the tax system is a better tool for redistribution of income than legal rules. We should, therefore, use the tax system rather than legal rules to address income inequality, and, correspondingly, legal rules should not systematically favor the poor. With-

† Professor, The University of Chicago Law School. This Essay was delivered as the Coase Lecture at the University of Chicago Law School on February 19, 2002.

out the tax system, this conclusion would not necessarily hold and legal rules might optimally be set to redistribute. Thus, I claim a central place for the tax system in thinking about legal rules.¹

I start by discussing redistribution more generally. Then I turn to a comparison of legal rules and the tax system as methods of redistribution. Finally, I consider some possible extensions of the argument and some counterarguments.

I. REDISTRIBUTION

Why might we want to redistribute income? There are a variety of reasons. For example, the rich might be altruistic toward the poor so that redistribution helps both. Alternatively, redistribution might act as a form of social insurance. Any of us could easily be down and out someday due to misfortune. Redistribution hedges this risk.

The most commonly cited reason for redistribution, and the one I will focus on, is that individuals have declining marginal utility of income. All this means is that as individuals get wealthier, each dollar is less important. One reason that individuals might have declining marginal utility of income is simply that individuals tend to satisfy their most important needs first. Once they have enough income to satisfy these, they move on to lesser needs. This type of planning or maximization means that the first dollar one receives is more important than the millionth. Another reason we might believe in declining marginal utility is just introspection—the extra brass and teak fittings on a gazillionaire's yacht are likely to be less important to him than food or housing to a pauper.

If individuals have declining marginal utility of income, redistribution can be a good idea because the loss in utility to the wealthy from whom we take a dollar is not as great as the gain to the poor who receives the dollar. If we ceased the analysis here, we might conclude that significant redistribution is appropriate. In fact, at the extreme, we would redistribute until everyone had equal incomes. So long as there is declining marginal utility, it is always a good idea to take from the wealthier and give to the poorer until all income is equal.

The problem is that redistribution is costly. Not only are there the costs of the administrative apparatus of redistribution—measuring income, collecting the money, etc.—there are costs in terms of incentives. In particular, if we take income from people, then each hour of work or unit of savings brings less to them. For example, if you previously earned \$10 per hour and we start taking 40 percent of your income, your take home pay is now \$6 per hour. This means that the

¹ When I use the term "tax system," I mean it broadly to include the transfer system as well. That is, the tax system includes government payments to individuals as well as collections.

tradeoff between working and doing other things is changed. Work is relatively less attractive than leisure. As a result, we can expect people to work less. For example, why spend those extra hours moonlighting at a second job for not much reward?

This distortion, which is generally called a distortion in labor/leisure choice, results in efficiency losses. The overall size of the pie, including how much there is to redistribute, goes down because people work less. A metaphor that is frequently used is that redistribution is a leaky bucket. When we take money away from one person to give to another, we lose some during the transfer.

What is the size of the distortion? There are, as one might guess, many studies, and, as one might also guess, they produce conflicting results.² It turns out to be very difficult to measure. One story is that the distortion is small in part because people tend to work relatively fixed hours. If your job demands forty hours per week, you don't have the choice to reduce it to thirty-five hours when tax rates go up.

Another story is that the distortion is large, particularly for various subclasses of individuals. The moonlighter has the choice of whether to work the extra hours. Second earners in a family are also thought to be very responsive to changes in tax rates. And even primary earners putting in their forty hours can be responsive to tax rates in a variety of subtle ways, such as how much they invest in education or training, how hard they work for advancement, and similar non-hourly elements of work. This empirical debate will continue for the indefinite future. For our purposes, unless you think the distortion is zero, there will be a cost to redistribution.

In addition, regardless of the absolute size of the cost initially, it is generally thought to go up with taxes; so the more we wish to redistribute, the higher the cost. Worse, it is thought to go up very fast—with the square of the tax rate—so even small distortions can become very large if tax rates are high enough.

The million-dollar question is how we balance these concerns. Redistribution seems like a good idea, but it is costly. How do we decide how much to do? To make this decision, we need a theory of what is good for society—of how much we should be willing pay to achieve redistribution. Such an expression of beliefs about this tradeoff is called a social welfare function. It tells us how well off a society is based on the welfare of the individuals in that society. This means that it tells us when more redistribution at a given cost is a good idea. I

² See, for example, the papers contained in Joel B. Slemrod, ed, *Does Atlas Shrug? The Economic Consequences of Taxing the Rich* 141–461 (Harvard 2000) (examining the impact of taxing the rich on a wide range of behavior).

will, for purposes of this discussion, put two restrictions on the social welfare function.

First, I will restrict the social welfare function to be a function of the welfare of individuals. Society is better off when individuals in society are better off. Concepts like justice can affect social welfare, but they must do so by making individuals better off rather than by satisfying an abstract ideal. Second, I want the social welfare function to be consistent with Pareto efficiency. This means that if the welfare of one individual goes up and the welfare of nobody else goes down, we should think of society as better off.

What are some prominent examples of social welfare functions that meet these criteria? One of the most prominent is the utilitarian social welfare function. Utilitarians hold that we can measure how well society is doing by adding up the utilities of the individuals in society. Each person is weighted equally under this social welfare function.

The philosopher John Rawls proposed a different measure of social welfare known as the maximin. Although he was concerned with values other than individual welfare, the maximin, roughly described, measures the welfare of society according to the welfare of the worst off individual in society. Rawls's formulation is generally thought to be more egalitarian than the utilitarian social welfare function.

We can also come up with intermediate versions that are more egalitarian than the utilitarian social welfare function but less so than the maximin. One example is the product of individuals' welfare. While we can debate the merits of various social welfare functions, I am not going to concern myself with this choice except to ask that it be explicit and meet the few minimal conditions I have mentioned. I should note before moving on that there are many very difficult problems with this formulation of what makes a good society. For example, we must decide whether to count hateful preferences in social welfare—if someone is happy when another is injured, do we count that perverse happiness? We also have to decide who is included. Do we count foreigners? What about embryos or animals? There is a whole literature devoted to studying these problems, but for purposes of this Essay, I will assume that we have a social welfare function and get on with the analysis.

Once we have had our ethical debates and chosen a social welfare function, we can think fairly clearly about redistribution. It actually becomes an optimization problem. We want to maximize social welfare given constraints about the costs of redistribution.

The original work along these lines was done by James Mirrlees in 1971, and he was later awarded the Nobel Prize in part for this

work.³ The actual mathematics are very complex and of little interest to the nonspecialist. It might be helpful, however, to demonstrate the principles with a simple example because the discussion so far has been somewhat abstract. An article by Joe Bankman and Tom Griffith has an example that serves well for our purposes.⁴

Suppose we have three individuals in a hypothetical society. (Obviously in a real society, there are more than three individuals. We can think of the three as three representative classes of individuals in a more realistic society.) Their wage rates are as indicated in Table 1.

TABLE 1

<i>Individual</i>	<i>Wage Rate</i>
Alice	\$10/hour
Betty	\$20/hour
Cindy	\$40/hour

Assume that each has a utility function that combines both their consumption and their time off from work, or leisure. In particular, assume that their utility is the log of the product of consumption and leisure, or:

$$\text{Utility} = \text{Log} [\text{consumption} \times \text{leisure}]$$

The reason for using the log function is that it goes up, but at a declining rate, which gives the individuals declining marginal utility. I use the product of consumption and leisure so that the individuals will want to balance the two—both count, and if either is too low utility will decline. For example, if the individual is a beachcomber and has only leisure and no consumption, utility will be low. Similarly, if the individual is a workaholic and has consumption but no leisure, utility will decline. Our hypothetical individuals want a mix of free time and wealth, like most individuals in the real world. Note also that the utility function gives us implicitly the labor/leisure distortion because it determines how much individuals will reduce work when they are taxed. It tells us their tradeoff between work and leisure.

³ See, for example, J.A. Mirrlees, *An Exploration in the Theory of Optimum Income Taxation*, 38 Rev Econ Stud 175, 175 (1971) (exploring questions of “what principles should govern an optimum income tax; what such a tax schedule would look like; and what degree of inequality would remain once it was established”).

⁴ See Joseph Bankman and Thomas Griffith, *Social Welfare and the Rate Structure: A New Look at Progressive Taxation*, 75 Cal L Rev 1905, 1945–65 (1987) (summarizing the welfarist and nonwelfarist theories of redistribution and the data as of 1987 concerning the efficiency effects of redistribution).

Given this society, we can set a tax rate or rates to raise revenue. I will assume that we set a single rate on everyone and use the money raised to pay for a public good that benefits all equally. We can think of this as simply giving them back money—the public good counts as consumption in their utility functions. This system is progressive because how much they pay in is based on a percent of income, so Cindy pays more than the others, but they all get back the same amount. We can do the calculation with more complex tax structures, but it just makes it more complicated without changing the basic ideas.

We can calculate how much each individual will work and their utility for any given tax rate under this tax and spending system. We then plug the result into our social welfare function and determine overall social welfare. The goal is to find the tax rate that maximizes social welfare.

Below I've calculated social welfare for the utilitarian and Rawlsian social welfare functions.

Figure 1 is the utilitarian case. The graph shows how total social welfare, the sum of utilities, varies with the tax rate. From examination of the graph, we can see that social welfare is maximized when the tax rate is approximately 31 percent.

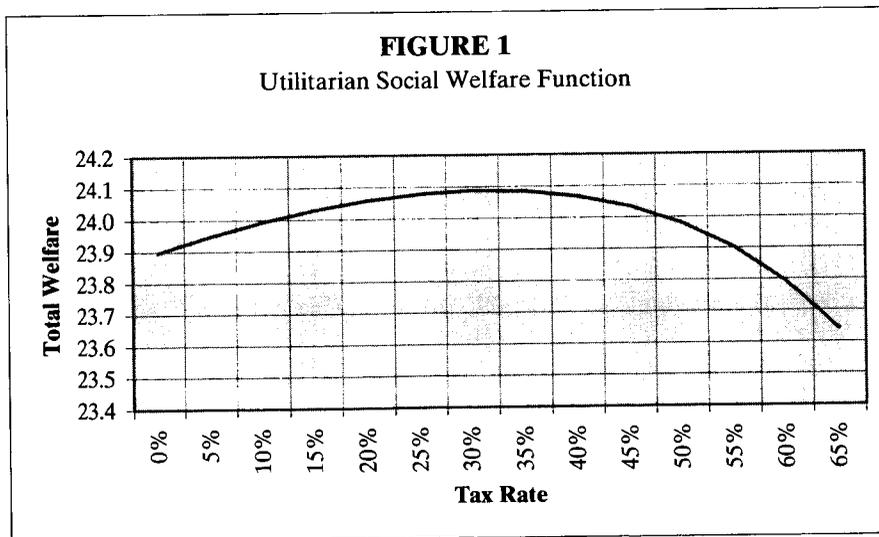


Table 2 shows the income and welfare numbers at the optimum. Even at the optimum, our sample society is not even close to completely equal. The wealthiest individual, Cindy, has more than three times as much income as the poorest, and is better off in terms of utility as well. Notwithstanding this inequality, further redistribution would be a mistake because it would reduce total utility.

TABLE 2
Optimum with Utilitarian Social Welfare Function

	<i>Income</i>	<i>Utility</i>
Alice	\$118	7.61
Betty	\$201	7.98
Cindy	\$366	8.49
Total	\$686	24.08

Figure 2 represents the same society, but using a Rawlsian social welfare function. Note the change to the graph. Previously we had total utility on the y-axis. Now the y-axis has the utility of the worst off individual. The optimal tax rate has gone up to somewhere around 60 percent.

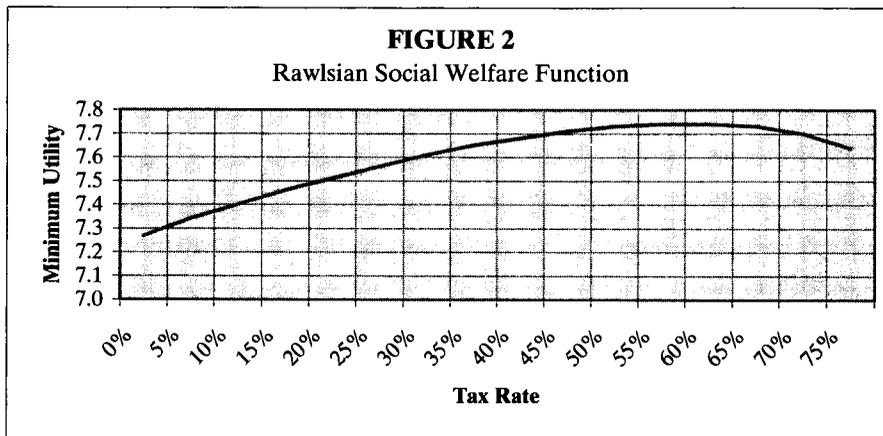


Table 3 shows the income and welfare numbers at this optimum. The numbers are more equal. The best off now has only two and one half times the wealth of the worst off. But there is still substantial inequality in both income and utility. Cindy is still much better off than Alice. Note also that total wealth and total utility went down. Total wealth went down from \$686 under utilitarianism to \$480 under the minimax. We have traded off total wealth for more equality.

TABLE 3
Optimum with Rawlsian Social Welfare Function

	<i>Income</i>	<i>Utility</i>
Alice	\$96	7.74
Betty	\$144	7.86
Cindy	\$240	8.19
Total	\$480	23.79

Why would a Rawlsian social welfare function leave us with this sort of inequality when it cares only about the welfare of the poorest? Why wouldn't it tell us to have more redistribution? The reason is that we cannot help the worst off any more. Raising the tax rate to transfer more to her actually makes her worse off because of the extra taxes she has to pay and because the efficiency losses from taxation reduce the amount available to redistribute to her. The pie shrinks so much, the efficiency losses are so high, that we cannot do any better.

Is this an artifact of our limited tax structure? We are limiting ourselves to a single tax rate, so raising Cindy's taxes also raises Alice's taxes. It is hard in this structure to take from Cindy without also hurting Alice. The answer, it turns out, is not very much dependent on our limited model. Even under quite general tax structures there is still substantial inequality with a Rawlsian social welfare function. The leaky bucket remains very important.

That's the background. The core idea is that redistribution of income is probably a good thing. It most likely increases social welfare. Unfortunately, redistribution creates inefficiencies, such as the distortion in the choice between labor and leisure. Redistribution is a leaky bucket. Therefore, we have to trade the benefits of redistribution with its efficiency costs, and we end up with some desirable amount of redistribution that may leave us far from equal. We are now finally ready to discuss the main topic, which is whether legal rules should be used to redistribute income. I want to compare legal rules and the tax system as methods of redistributing income.

II. THE LEGAL SYSTEM COMPARED TO THE TAX SYSTEM AS A METHOD OF REDISTRIBUTING INCOME

Suppose we have a legal rule and a tax system that have similar redistributive effects. They both, for example, take similar amounts from those with high incomes and give to those with low incomes. Which one do we prefer?⁵

⁵ The argument made here is primarily due to Louis Kaplow and Steven Shavell, *Why the Legal System is Less Efficient than the Income Tax in Redistributing Income*, 23 J Legal Stud 667, 667–68 (1994) (arguing that employing legal rules to redistribute income “distorts work incentives fully as much as the income tax system” and creates other inefficiencies); and Steven Shavell, *A Note on Efficiency vs. Distributional Equity in Legal Rulemaking: Should Distributional Equity Matter Given Optimal Income Taxation?*, 71 Am Econ Rev 414, 414–17 (1981) (arguing that “despite imperfect ability to redistribute income through taxation, everyone would strictly prefer that legal rules be chosen only on the basis of their efficiency”). For additional discussion, see Chris Sanchirico, *Deconstructing the New Efficiency Rationale*, 86 Cornell L Rev 1003, 1069 (2001) (arguing that law and economics is unjustified in relying exclusively on the principle of efficiency); Chris Sanchirico, *Taxes versus Legal Rules as Instruments for Equity: A More Equitable View*, 29 J Legal Stud 797, 820 (2000) (arguing that “the ability to redistribute by income tax does

Consider a legal rule that redistributes income by varying from the efficient result. The legal rule redistributes income, which is good. It also creates inefficiencies. Because it redistributes, it reduces the return from working, and therefore causes labor/leisure distortions. But, and this is the key fact, it also deviates from the efficient legal rule, so it causes inefficiencies along that boundary as well. It will cause individuals to take too much or too little care, breach contracts inappropriately, under- or over-invest in property, and so on. The legal rule affects two margins: the decision to work and the decisions relevant to the legal rule.

Compare that to an income tax. It too causes a labor/leisure distortion. We have the same disincentives to work. But the income tax does not cause inefficiencies along other margins. It will not change incentives to take care, incentives to breach, and so forth. This means we can get the same redistribution at a lower cost. Or we can get more redistribution at the same cost. But either way, the tax system is a more efficient method of redistribution than the legal rule.

That, in a nutshell, is the main argument. We can develop the intuition further by stating the argument slightly differently. Suppose we are considering a legal rule that favors the poor. Say we deviate from efficient tort damages so as to redistribute toward the poor. The wealthy in this case will expect to pay more or receive less in tort, and there is a corresponding benefit to the poor. Suppose instead we make the legal rule efficient and adjust the tax system to redistribute the same amount by increasing its progressivity.

By making this adjustment, we have identical distributional consequences. The wealthy pay more and the poor receive more. There are, therefore, identical benefits from redistribution and also identical distortions of the labor/leisure choice. But we have increased the efficiency of the tort system because now we resort back to efficient tort rules. So we are better off.

In fact, strong versions of this argument claim that we have made a Pareto improvement—we have redistributed in exactly the same way by using the tax system but increased total resources because the tort system is more efficient. The argument is consistent with *any* of the social welfare functions that we might pick, regardless of how egalitarian we might be. Therefore, we get our conclusion: The income tax is the better instrument for redistribution and legal rules should, therefore, not be so used.

not make it appropriate to ignore redistributive goals in evaluating legal rules”); and Louis Kaplow and Steven Shavell, *Should Legal Rules Favor the Poor? Clarifying the Role of Legal Rules and the Income Tax in Redistributing Income*, 29 J Legal Stud 821, 834 (2000) (responding to the criticism of Sanchirico and maintaining that the “use of legal rules to redistribute income is generally less effective than relying exclusively on the income tax system”).

The argument just made assumed that legal rules were actually able to redistribute income fairly well and merely argued that the tax system would do better. But there are separate reasons to believe that the legal system is often an ineffective instrument for redistributing income at all. These arguments go under the rubrics of “contracting around” and “haphazardness.”

The contracting-around argument states that attempts to redistribute through legal rules will often be defeated because prices will adjust to offset the effect of the legal rule. For example, if we require landlords to make their apartments habitable through an implied warranty of habitability, landlords will simply raise the rent to cover their costs.

The initial story seems quite intuitive here. We can examine the extent to which prices change in various markets and therefore when the effects of legal rules get passed on. Sometimes prices might not rise, so legal rules might seem to work.

This story, however, may be less persuasive than it appears. The core problem is that legal rules change the very product under consideration. For example, a mandated warranty changes the good being sold from simply a good to a good plus a warranty. Buyers will value the benefit of the warranty, perhaps more or perhaps less than its cost. The extent to which prices change reflects how much buyers value the warranty. If they value it highly, we might expect them to agree to pay much more for the product with the warranty. And if they do not value it at all, they may not be willing to pay more. Therefore, the extent of the price change is misleading. A big price increase (relative to the cost of the legal rule) might indicate that the legal rule has made buyers better off, while a small price increase might indicate that the legal rule has made buyers worse off. The intuition should be completely flipped from what we might have first thought.⁶

To make it messier, individuals will value the warranty differently. Some will value the benefit of the legal rule more than its cost and others less. When prices adjust, there will be redistribution among consumers. If our goal was, say, to redistribute from sellers to purchasers or landlord to tenants, we might instead have mostly redistributed among purchasers or among tenants. It is not even clear what it means to have pro-tenant or pro-consumer rules in this case. Sorting out the

⁶ For a more complete analysis, see generally Richard Craswell, *Passing on the Costs of Legal Rules: Efficiency and Distribution in Buyer-Seller Relationships*, 43 *Stan L Rev* 361 (1991). See also Henry E. Smith, *Ambiguous Quality Changes from Taxes and Legal Rules*, 67 *U Chi L Rev* 647, 649 (2000) (arguing that analogizing legal rules to taxes is more insightful when one assumes “that commodities and services are bundles of attributes that can vary from specimen to specimen,” and “that measurement of such attributes is costly”).

actual effects becomes very difficult. The theoretical questions are hard, and the empirical questions are probably even more difficult.

For example, we are not entirely sure whether labor laws help workers as a whole. They might, for example, help workers who are employed but make finding work more difficult for those who are unemployed. Similarly, it is not clear that consumer protection laws help consumers. I don't want to take a position on this. There is a large literature on these various rules. Suffice it to say that some probably help their intended beneficiaries, and some probably do not. The only point is that attempting to redistribute income through legal rules is a precarious exercise.

The other standard argument against using legal rules is known as the haphazardness problem. Consider a pro-plaintiff tort rule. We might want such a rule because, say, plaintiffs are systematically poorer than defendants. The problem with such an approach is that it is likely to be both under- and over-inclusive. It is under-inclusive because many poor individuals are never part of a tort suit, so they would not benefit from the rule. Similarly, the rule is over-inclusive because there will inevitably be some rich plaintiffs who benefit.

Once again, as in the contracting around case, the argument gets much more complex once we get into the details. For example, even if people are never subject to a lawsuit, they might bargain in the shadow of the law and, therefore, receive benefits or detriments indirectly. Similarly, in the tort case, insurance prices will be based on expected damages, so to the extent individuals buy insurance, they are all affected even if they are never party to a lawsuit.

Before drawing any conclusions about the strength of the contracting around and haphazardness arguments, we must compare the legal system to the tax system. The argument must always be a comparative argument—it doesn't help to criticize one method of redistribution without comparing it to other methods. The tax system might very well be subject to the same criticisms as the legal system. Prices might adjust to undo the effect of tax redistribution, and the tax system might be over- and under-broad, taxing the wrong people sometimes and missing the right people other times.

While these criticisms of the tax system are undoubtedly true, I tend to think that they are likely to be less severe in the tax system than in the legal system. The reason is that the tax system is a dedicated system that can be tailored to minimize these effects. But there is no a priori or slam-dunk argument that tells us this will happen. We have to do a detailed institutional comparison. Still, my initial view is that it is extremely unlikely that the legal system would compare well.

Taking all of these arguments together, the double-distortion argument and the problems with contracting around and haphazardness,

I believe the case against using legal rules to redistribute to the poor becomes almost overwhelming. The tax system is designed to measure the variables relevant to redistribution and act only on those margins. It is hard to imagine that legal rules are likely to do a better job.

One key point to note is that I have not argued that legal rules should be efficient. All the argument has shown so far is that legal rules should not be used to redistribute income. But income is not the only source of inequality. Race, gender, disability, or health all might be sources of inequality in our society. If we value equality of all sorts, not just income equality, we might want to redistribute on these bases.

Maybe legal rules should be used to redistribute on other grounds and, therefore, should not be efficient, even if they should not be used to redistribute income. Here is an example. Suppose the severely disabled are less well off than others of equal income. Their costs may be higher, their needs greater, and so on. An income tax won't help much here. Income won't necessarily correspond to disability. We can even imagine the case where everyone has the same income but differs along this dimension. Under these circumstances, an income tax would have nothing to do even though there is inequality in our society.

Legal rules, however, might help. We could adjust the legal rules to favor the disabled and thereby redistribute away from everyone else and to the disabled. For example, we could, as we actually do in real life, impose a mandate on employers that they provide accommodations for the disabled to allow them to work.

Nothing in the argument so far tells us not to do this. The income-tax-only argument is only about income redistribution. It does not tell us that legal rules should not be adjusted to favor non-income sources of inequality. That is, we need to distinguish between an argument for efficient legal rules (which I have *not* made) and an argument that legal rules should not be used to redistribute income (which I have).

Note that in the disabilities example, legal rules might seem to favor the poor. For instance, this would be the case if the disabled are systematically poorer than others and legal rules are designed to favor the disabled. But this could easily have gone the other way. If the less well off class happens to be wealthier than others (the idea might be that they are less well off at any *given* level of wealth even though they tend to have more wealth) then the adjustment would favor the rich. Another way to phrase this point is that if we think of the income tax as the best tool for handling distribution between income classes, we should think of other types of redistribution as working within an income class. Accommodation mandates for the disabled can be thought of as redistributing between individuals with the same income but with and without disabilities.

While the arguments I have made so far have concerned only income and leave room for this type of redistribution, we should not jump to the conclusion that legal rules are a good method of redistributing on this basis. The arguments about income redistribution might very well translate to redistribution on other bases.

For example, much has been written about whether rules that appear to help a given class, such as the disabled, really do—this is the contracting around argument. There is literature indicating that the disabilities law reduced overall employment of the disabled while perhaps increasing wages (for those who kept their jobs).⁷ This means that it might primarily redistribute within the class of disabled rather than from the healthy to the disabled.

Haphazardness is also a serious problem—the disabilities law seems to be both over-broad and too narrow. It might apply to individuals who are not really worse off because of their disabilities and it does not apply to the seriously disabled who are completely unable to work.

Finally, a version of the double distortion argument also might translate to this case. Suppose we compare a direct redistribution of money to the disabled with a modification to legal rules to help the disabled. Direct redistribution would have the same redistributive benefits and costs but would not skew employment relationships. Because legal rules skew employment relationships, direct redistribution may be superior to legal rules.

I don't want to condemn the various disabilities laws without serious and focused analysis of those rules. There may be much to be said for them. There are really only two points I am trying to make. First, the analysis about whether legal rules should favor the poor leaves room for this sort of redistribution. The primary argument I have made is not an argument for efficiency; it is an argument against using legal rules to redistribute income. Second, if we believe that using legal rules to redistribute in favor of the poor is inherently flawed, then we should also be leery of using legal rules to redistribute on other bases. There are good reasons to think that legal rules should be efficient and not redistribute at all.

That is the basic argument. Let me mention three potential counterarguments. First, a common criticism of the arguments I have made is that adjustments to the tax system are not feasible. Just be-

⁷ See Christine Jolls and J.J. Prescott, *Disability Discrimination Law, Disabled Employment, and the Relative Role of "Reasonable Accommodations" Requirements and Firing Costs* *33–34 (draft on file with author) (examining the various state legal regimes prior to the enactment of the ADA and concluding that the "legal requirements imposed by the ADA were causally linked to the declines in disabled employment").

cause we identify a social problem, say inadequate housing, doesn't mean that the tax system will be adjusted to resolve it.

This argument is correct to some extent. If the tax system is unavailable, we would want to resort to the next best alternative. If our only alternative is legal rules, we may want to use them. But there is no particular reason to believe that legal rules will be more available than taxes. The tax system is adjusted all the time—it is famous for constantly changing. Moreover, the extent of redistribution is a constant topic of discussion, so much so that many observers complain that it dominates discussion to the exclusion of other considerations. Moreover, if some political coalition would block additional redistribution through the tax system, there is no reason to believe it would not block the same redistribution when done indirectly through legal rules. It might fight even harder against legal rules because of the additional costs imposed by the inefficiencies.

Maybe using legal rules gets around these coalition problems because courts can make decisions about redistribution using legal rules, and courts are largely outside the political process (although, in fact, the vast majority of legal rules are made by legislatures). But one needs a complex view of the appropriate roles of courts and legislatures to argue that if redistribution cannot be passed through the democratic process, it should occur in a more expensive manner through courts.

The feasibility intuition is very strong among lawyers. I'm not sure exactly why—perhaps because lawyers are focused on legal rules, so they consider them more malleable than the tax system. Perhaps it is because lawyers tend to see what looks like an immediate injustice in a given case without thinking systematically about the problem. On close examination, it has little force.

The second common criticism of the argument is that the tax system is a mess. While using legal rules has problems, the tax system is riven with loopholes, it is costly to administer, and it is byzantine in its complexity. It is not the efficient and sleek method of redistribution portrayed so far. How can I defend such a monstrosity?

I don't intend to defend the income tax from these attacks. I make them myself all the time. But there is no reason to believe that other broad-based redistributive systems would be any better along these lines and they could easily be much worse. All the problems that lead to loopholes and complexity in the income tax would pervade any alternative attempt to measure and redistribute income.

The third counterargument to my thesis is that people might respond differently to redistribution through the tax system than

through legal rules.⁸ For example, people might notice the redistribution through the legal rules less, and, therefore, redistribution through legal rules might distort work incentives less. To make this slightly more formal, suppose that legal rules apply only infrequently, so that there is a low probability of them applying. A common claim is that people underestimate the impact of infrequent events. If true, people may systematically underestimate the redistribution through legal rules and, therefore, not adjust their labor/leisure choice.

We can't know the truth of this claim without more investigation. I have been assuming that people respond to real incentives, that there are no systematic, longstanding illusions. If there are such illusions, there might be many opportunities to take advantage of them, not just with respect to redistribution. I would think, however, that to the extent there is any significant redistribution, it would be difficult to retain illusions of this sort.

In addition, if insurance is available so that the legal rule applies through insurance prices, it is much more difficult to tell a story that involves systematic bias. And if we are relying on lack of insurance, so that legal rules apply only probabilistically, haphazardness becomes a serious problem. Therefore, I do not think arguments of this sort get us very far.

CONCLUSION

My conclusion is quite simple: Legal rules should not be used to redistribute income. Let me give a slightly softer conclusion as well. Even if one rejects some of the particulars of the argument—and there is literature that does⁹—there is no simple case for redistribution using legal rules. The existence of the tax or other direct redistributive systems means that one should not think of legal rules as necessarily addressing injustices posed by the facts of any given case. Instead, we must think systematically and comparatively about using legal rules and other systems to redistribute.

⁸ See, for example, Christine Jolls, *Behavioral Economics Analysis of Redistributive Legal Rules*, 51 Vand L Rev 1653, 1657 (1998) (“[T]he costs of redistributive legal rules are more likely to be *underestimated* than to be *overestimated*, and, thus, [] work incentives may be distorted less by redistributive legal rules than by taxes.”).

⁹ See, for example, Sanchirico, 86 Cornell L Rev at 1069 (cited in note 5) (“Both the verbal and the formal versions of new-rationale reasoning fall far short of a coherent argument for efficiency.”).

