

Joel Slemrod

What We Know About Tax Systems, How We Know It, and What We Still Need to Know



What We Know About Tax Systems, How We Know It, and What We Still Need to Know¹

Joel Slemrod
University of Michigan

Finnish Centre of Excellence in Tax Systems, Tampere
September 1, 2022

1. Introduction

This is my third visit to Finland and it is my second to Tampere. It is great to be back and a pleasure to be part of the first event of FIT, the Finnish Center of Excellence on Tax Systems. I have long known that Finnish economists are excellent in the research of tax systems long before the beginning of this center (and its naming), as several Finnish economists have visited Michigan, given talks at Michigan, and I have learned much from their research.

I also know that tax systems research is not the only thing Finland is known to be excellent at. This morning in my hotel room, I googled what Finland is best at, and learned that Statistics Finland has published a long list of such things. I know you all probably know these things, but you have the best governance in the world, you have the highest trust in other people--at least in Europe, and maybe in the world--you drink more coffee than anyone else in the world, and finally, which has more relevance than you might guess to me, you have more heavy metal bands per inhabitant than any other country in the world. Some, but not all, of these points of excellence will figure in my remarks today, to which I now turn.

2. Tax Systems

A lot of my research has been about what I call *tax systems*. It is in the title of this Finnish center of excellence, and it is in the title of my talk, so let me define what I have meant by tax systems. A tax system is a set of rules and regulations and procedures that includes three basic components:

The first component is what tax economists have been studying mostly for the last many decades: it defines what events or states of the world trigger tax liability, what rates are applied to these bases, and what the effects of these things are. This is the bread and butter of tax economists, what we call tax *bases*, such as income, or consumption, or payroll, or wealth, or property. What bases should be used for taxation and what *rates* should be applied to these bases, are central issues.

¹Keynote address given by Professor Joel Slemrod (University of Michigan) at the inaugural seminar of The Finnish Center of Excellence on Tax Systems (FIT), University of Tampere, Finland, on September 1, 2022. A Finnish translation of the speech will be published in *Kansantaloudellinen Aikakauskirja* (The Finnish Economic Journal), Vol. 118, No. 4 (2022).

So that is the bread and butter of tax analysis. Tax systems analysis adds two important additional components.

The second component is that it specifies who or what entity must remit the tax liability, and by remit, I mean who actually writes the check, who hands the money over to the tax authority. Now if you had only done up to intermediate economics, you probably were taught that this doesn't matter. For example, it doesn't matter for a consumption tax whether the consumer remits the tax, or the retailer remits the tax or, as in most countries other than the U.S., whether all businesses remit the tax through a value-added tax system. Although that's what an intermediate textbook will tell you, I believe that actually it does matter quite a bit in many situations.

It is no coincidence that, of all the countries in the world that have consumption taxes, almost all use a value-added tax rather than a retail sales tax, which is different predominantly in who remits the tax. In retail sales tax, like most U.S. states have, it's all remitted by the retail business. In a value-added tax, tax is remitted by businesses throughout the production and distribution chain. If it didn't matter, why would it be that 160 out of 161 countries use one method and one country, the U.S., uses the other method? It matters because it is much more administratively efficient to levy a value-added tax.

It also turns out that if you look across the world, about 85 per cent of taxes are sent in (i.e., *remitted*) by businesses, either because they are taxes nominally levied on businesses or because they are taxes levied on labor income that employers remit to the government. I owe tax on my salary from the University of Michigan, but most of the money is actually sent in by the University of Michigan to the tax authority, and that's how it works in almost every country. About 85 per cent of taxes in Europe and the United States are remitted by businesses, and in the one study I know of concerning a developing country, India, (I know of it because I was a co-author), about 85 per cent of taxes in India are remitted by businesses. Collecting taxes from businesses is more efficient than monitoring people in their role as employees, interest receivers, and so on. So who remits matters for administrative efficiency, and it matters to minimize tax evasion.

The third important aspect of a tax system is that it details the procedures for ensuring compliance with the tax law. A country cannot just announce a tax system, sit back and let the money roll in, because as soon as people – most people, at least – realize that there is no consequence for not remitting the money, it will stop rolling in. So, every country has rules for ensuring compliance, including information-reporting requirements, specifying the consequences of not remitting one's tax liability, what the penalties are, what the audit rules are, and so on.

These three components primarily compose the tax system. I hope and expect that the center for excellence here in Finland will be focusing on them.

3. Building Blocks of Tax Systems

Some of the building blocks of the study of tax system are as follows:

- Multiple sources of the cost of taxation: excess burden, administrative and compliance costs

- Multiple types of behavioral responses: real, avoidance, and evasion
- Attention to non-standard (i.e., non-rate, non-base) tax policy instruments, such as audit coverage, disclosure policy, and taxpayer education policy
- The tax base elasticity as a summary measure of the marginal social cost of raising revenue through higher tax rates
- The tax evasion elasticity as a summary measure crucial to assessing the optimal tax gap

Let me elaborate on these building blocks. It is crucial to recognize the multiple social costs of taxation. Of course, there is the administrative cost of running the tax authority. It turns out that the cost that the taxpayers incur other than the money they remit--the time to figure out their tax liability, to maybe plan about minimizing their tax liability--tends to be about ten times higher than the administrative cost, in the countries that have been studied.

The *excess burden* is what we economists mean when we say that one cost of a tax system is that it induces taxpayers to change their behavior to reduce their tax liability. They do things they wouldn't otherwise do because of taxation. Sometimes that is a good thing if we're talking about, say, environmental taxes, but mostly it is a bad thing. For example, we work less or invest less.

We economists spend a lot of time thinking about how businesses and people respond to taxes. There are three basic kinds of behavioral response. There is evasion--illegal things people do to reduce their tax liability. There is avoidance, which is how people change their behavior in ways that are not "real", meaning they don't work less, they don't invest less, they don't change whether they smoke cigarettes or not, but they do things that are legal that reduce their tax liability. For example, consider the introduction of R&D credits. We know that the accountants at some firms suddenly realize that things they have been doing all along they never "realized" were R&D, but now that they are subsidized, they realize it's R&D! That's a kind of avoidance activity, a re-characterization, or relabeling, of what they have been doing all along.

A tax systems approach also pays attention to policies other than the rates and the bases. How does a tax agency decide whom to audit, the audit coverage? How much information about taxation, both aggregate and taxpayer-specific, is disclosed? You all no doubt know that, in many countries of the world, you can learn details of the tax filings of other citizens. I have studied this in Norway, and I know Norway is not the only Nordic country where this is true. In Pakistan, you can go on a website and see in a pdf file what other people in Pakistan say their taxable income is and how much tax they remit. Who needs to transmit information to the tax agency--employers, banks, importers? These can be critical aspects of tax systems.

And finally, a tax-systems approach focuses attention on two summary measures of how people and businesses respond to a tax system. The first is the tax base elasticity. Think about the income tax, where reported income is the base. How does reported income change when the tax rate goes up? Answer: it almost always goes down. How much does it go down? Under some very weak assumptions, the answer to the question of how much the tax base goes down when tax rates go up, which we call the tax base elasticity, is a summary measure of the marginal social cost of raising those tax rates.

Similarly for evasion, the tax evasion elasticity, which is a quantitative measure of how responsive evasion is to things like the fraction of returns that are audited, is a summary measure crucial to assessing the *optimal tax gap*. Perhaps the term optimal tax gap is not familiar or natural to some of you. Many countries now try to measure the tax gap: what fraction of the tax liability that should be remitted is not. In the U.S., our best estimate is that it's about 15 per cent. I'm going to guess that in Finland it is less than that, and in many countries it is certainly more than that. But what is the optimal tax gap? Certainly, the optimal tax gap isn't zero. It would be way too expensive to eliminate every dollar of tax evasion, in the same way it would be too expensive to eliminate every last bit of street crime. To evaluate what is optimal, one needs to know the responsiveness of evasion to the various policy measures one can use to reduce it.

4. What Do We Know about Tax Systems?

I have three questions in my title, the first being what do we know about tax systems. There are hundreds, maybe thousands, of articles written about this, and I'm going to summarize what we know in one slide, as that's all I have time for. Here's some things we know.

I mentioned that we know that compliance costs, the time and money taxpayers spend in complying with the tax system, dwarf the administrative costs. In the U.S., if you take the IRS budget, you have to multiply it by about ten to learn what the compliance costs are, and they are both social costs.

We know that tax bases not covered by information reports, for example self-employment income, tend to have very high evasion rates. In the U.S., our best guess is that for income covered by information reports, like wages and salaries, the evasion rate is about 1 per cent, while for income not covered by information reports – largely self-employment income – it's about 56 per cent. Huge difference. So, a natural policy question is how wide the coverage of the information reports should be? It is not costless to ask citizens and businesses to report information to the IRS. What is the optimal tradeoff?

We also know that if a country is trying to achieve significant tax progressivity, it cannot avoid taxing capital gains, and taxing capital gains is very problematic for a lot of reasons. I'm going to come back and talk about that.

And also, as I have suggested, we know that taxpayer response to tax policy is pivotal to evaluating the policy. The magnitude and nature of the response affects the ultimate incidence, the revenue yield, and the efficiency cost of taxation. So, for example, FIT is studying and will be studying how emigration responds to tax policy. We need to know the answer to this question in designing certain aspects of tax policy. If it turns out that some policy is going to cause all the great data scientists in Finland to move out of the country, that is a social cost that needs to be considered.

5. How Do We Know It?

I turn now to *how* we know what we know about tax systems and, in so doing, give you a sense of how the FIT researchers will go about their business.

5.1. Learning from the Past

Traditionally, the way we economists have learned is to examine the past. We look at data from around times when the tax system has changed and see if we can get measures of how businesses have responded to these changes, how individuals have responded to these changes, how prices have responded, and so on. We sometimes call these policy changes *natural experiments*. Of course, for the most part, the governments are not changing their policy in order to provide knowledge about it, but when they do change the tax system, we can learn.

As Kaisa Kotakorpi mentioned earlier today, the central challenge in trying to learn from the past in this way is identifying what we call the counterfactual. As a concrete example, say the U.S. were to increase the top personal tax rate from 35 per cent to 40 per cent and we have data and look how high-income people reacted. What we need to know, though, is how people would have changed their behavior over this time period if it weren't for this increase in the tax rate. One can naively assume that people would have gone about doing whatever they did exactly before if it weren't for the tax change, but that just may be very wrong. And if it is wrong, we will ascribe too much, or too little, impact to the tax rate change. So, identifying what would have happened without these tax changes, is crucial. In some settings, we need a control group, a group of businesses or individuals who weren't affected by the tax change so we can compare the behavior, the response of the people who were affected.

In some cases, it is important to know *why* the tax policy changed. To see this, what if it turns out that tax cuts always happen when the economy is doing badly. If that were true, then we have a problem, because when the economy is doing badly, looking forward, the economies tend to recover. If we didn't take that into account, we might conclude, oh look, whenever there is a tax cut, the economy gets better. Well, it turns out that most economies would have gotten better then anyway, so you are going to vastly overstate the positive impact of the tax cut.

So, we learn from the past and we will continue to learn from the past, and we have developed all sorts of clever econometric methods to deal with both of these problems. Let me now focus on some new ways to learn about tax systems.

5.2. New Developments

In my view, the most promising development in the economic analysis of taxation in the past fifteen years or so, is the greatly increased willingness of tax agencies to share their data with academics, under strict protocols. In every country I know of, when the tax agency is sharing data, they are very careful about how that data is used and who gets to look at it.

A couple decades ago, I was on the advisory board for the Internal Revenue Service (IRS) back in the U.S.A. I would go to its meetings twice a year, and the IRS folks would tell us about all the great data they had. I would plead with them, can you let us look at it, and try to learn from it. "No, no, the lawyers won't let us, it's too complicated to manage." Well, I am thrilled to say that has changed, and it has changed in part because of the Nordic countries' tax authorities'

willingness to share data. This has taken off all over the world, in developed countries and developing countries.

Now that we can look at these data, under strict protocols, what that enables us to do is really impressive. We can conduct much more granular analyses. In the old days, a tax agency might release a random sample of data, but now we often get to look at the universe of tax returns, enabling analyses of demographic groups and, as I will discuss in a minute, learning from “bunching” in the distribution of tax reports.

The partnership of tax agencies with academics has really tremendously expanded our knowledge of the impact of enforcement policies on tax evasion. As you can no doubt guess, tax evasion is by its nature hard to measure. If it were easy to measure, the tax agency could measure it, they could detect and punish evasion, and there wouldn't be much tax evasion. So, by its nature, it is hard to measure. But, with the tax agencies much more willing to share with academics what they know, what they suspect about tax evasion and the impact of alternative policies on measures of evasion, we learn a tremendous amount.

5.3. Randomized Controlled Trials in Tax Policy

Kaisa also mentioned randomized controlled trials. About twenty years ago or so, there was a revolution in the empirical analysis in economics. It was a way to get around the problems I mentioned earlier with counterfactuals and the difficulty of finding a control group by using randomized controlled trials (RCT's). In labor economics and development economics it became all the rage, and twenty years ago or so I am sitting in my office thinking, well that's it for tax research, we're done. If development economists are going to be able to do these great RCT's, and labor economists are going to be able to do these great RCT's, how are we in tax ever going to get to do that? What country is ever going to randomize the tax system? What country would ever randomly assign a tax rate of 35 per cent to half the country and 20 per cent to the other half so we can learn how people respond? The answer, I thought, was no country. it's never going to happen. I thought we were doomed to irrelevance.

Well, I was kind of wrong. It is true that there haven't been many RCT's about tax rates. But there've been now a score of RCT's about tax *enforcement*. Here's how it works. Start with the taxpayer population, and randomly divide it into two groups, say according to where in the auditorium you are sitting. Assuming that who sits on the right and who sits on the left here are randomly different, I'm going to divide you into two groups and those of you on the right are going to get a letter from the tax authority saying “Hey, by the way, we noticed last year that you only reported 80 per cent of the income you reported year before, and we're a little worried and so we're watching you.” Those of you on my left don't get that letter. Because you were randomly chosen, on average the folks on my right are exactly the same as the folks on my left, on average. So, I can look forward and see if, after the letter goes out, the two groups behave differently, and we can learn pretty well what the impact of that letter is, at least in the short term.

Of course, RCT's are not perfect. One problem is that if we notice that one group changes their behavior compared to the other group, we don't know why exactly, and when we don't know why,

it is a little bit problematic to generalize the causal mechanisms. And, if we are thinking of doing the same policy in another country, it might be important to know why. But RCT's are hot now in taxation, and tremendously valuable.

5.4. An RCT Example

One RCT I worked on with the IRS concerned a group of businesses that they thought were at risk for a particular kind of non-compliance, but they didn't have the resources to go after them. In case they obtained more resources, they maintained a list of about 30 000 businesses that they thought were at risk. In the RCT, we split the 30 000 into three groups: 10 000 got no special treatment, 10 000 had a revenue officer going to the business, knock on the door, and deliver the message in person – we suspect noncompliance, and here is why we suspect it – and the other 10 000 got a letter from the IRS conveying the same information. We then look forward and see how they changed their behavior. On average, these three groups were identical before the enforcement treatments, and so any statistically significant difference in behavior going forward had to be because of the policy. We learned that knocking on the door has an immediate and large (positive) effect on tax remittances, lasting at least a year. The letters affected compliance substantially less, and less persistently. Does that mean that if the IRS wants to get more compliance from these kinds of firms, they should send revenue officers out? No. Why not? Because although personal visits are more effective, they are also much more expensive to implement. So, one has to carefully compare the impact, in terms of more revenue, against the cost of achieving it.

Another interesting thing we learned from this study is that there are network effects. We can look to see whether, after a knock on your door, other firms in your network changed. We investigated geographical networks. Did businesses in your 9-digit zip code (small postal area) change their behavior? Did firms you own, or firms that own you, change their behavior? What about firms that use the same auditor or accountant you use? The answer to that last question was yes. So, what probably happened is you got a knock on the door, and on a Friday night at the bar you mentioned it to your accountant, or the next time you spoke to your accountant on the phone, you say, hey by the way, guess what happened to us, to our business. And the accountant then tells her other clients, and they change their behavior.

Given the impressive success of RCT's, I now tell tax agencies everywhere the following story. If you are thinking of changing some policy (maybe not tax rates, but other tax system policies), before you do it, you should pilot it in an RCT, so randomly do the policy for some fraction of the taxpayer population, which is going to be cheaper for sure, and then you can learn how big an impact, if any, to expect. Sometimes this story works. One thing I have learned is to never, in conversations with tax agencies, use the word "experiment". Understandably, government agencies don't want to be seen as experimenting on the population. Randomized controlled trial is the ticket, not experimentation.

5.5. Bunching

I mentioned earlier the term bunching, which is another hot way to learn about the impact of taxation. The analysis of bunching could hardly be done without access to administrative data, because you need a lot of data to do it convincingly.

The idea is to look at how taxpayers respond to particular characteristics of many tax systems. These characteristics have the wonderful names of *kinks* and *notches*, kinks being where a small change in your behavior changes your marginal tax rate. In Norway and the U.S., and in most countries, in the income tax system there are brackets so that if your income jumps from one bracket to the next, which could happen if you have one dollar more in taxable income, your marginal tax rate, how much you owe on the next dollar, changes discretely. That's an example of a kink.

If taxpayers respond to their marginal tax rate, as we know they often do, we should be able to look at the difference in behavior of the people who are just below the kink and just above the kink; they are quite similar in income, but face distinctly different marginal tax rates. We can also look to see how many people end up right at, or very near, the kink. These people were happy to work when their marginal tax rate was 20 per cent, but once it jumped to 28 per cent, they said that's it, it's not worth working a single hour more, and they pretty much stopped right there, at that amount of income when the marginal tax rate abruptly increases. How many people stop right there, at or very near the kink, is a measure of the elasticity of response I have been talking about.

Even more fun than kinks, although the name isn't as exciting, are notches. In notches, when the taxpayer changes her behavior slightly, not only does the *marginal* tax rate change, but the tax liability changes. So, for example, if there is some credit you can get, say for retirement savings, but the credit is limited to people whose income is below say 20 000 euros a year, you get this big credit if you earn 19 999, but if you report one more dollar, you lose a thousand-dollar credit. That's a (big) notch.

You would expect lots of people who have retirement savings to report to the IRS that they have 19 999 euros of income, either because they make sure that they don't earn more and thus lose their credit, or they do earn more, but they only report 19 999 euros. Either way, the number of people who are reporting just below that notch is an indicator of how flexible people's behavior is, their elasticity. There have been some great studies using the magnitude and nature of bunching--how many people operate right at the notch, or the kink--to learn about elasticities.

There are many kinds of notches, some based on the characteristics of the taxed activity. An importunate example is the tax treatment of corporate debt and equity. In most countries the tax treatment of corporate debt and corporate equity are very different, with debt payment being considered a deductible expense while payments to equity holders are not deductible. But what is a debt instrument and what is an equity instrument? Well, in the U.S, there is a twenty-factor test, and as I speak there are many very smart people with MBA's trying to figure out how the next security issuance of their company is going to be just debt-like enough that the IRS will consider it to be debt. Thus, there will be bunching with respect to the characteristics of securities.

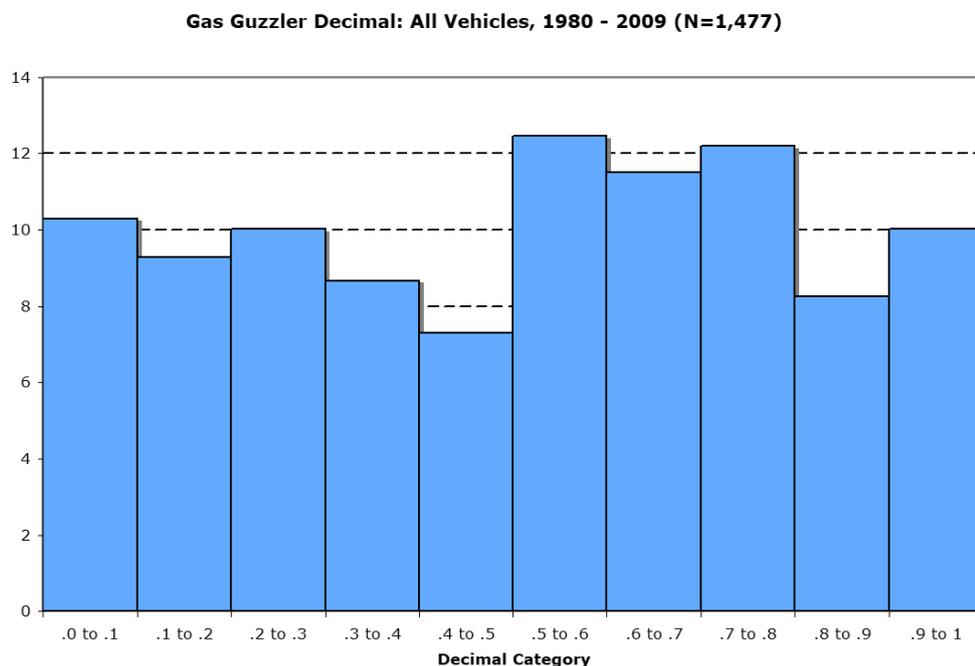
5.6. A Bunching Example

Here's another fascinating example of tax-system bunching from a paper I did about a decade ago. It's about the U.S. Gas Guzzler Tax, which is a tax on cars that's higher, the less fuel efficient a car is. Fuel efficiency is measured by miles per gallon (mpg), so the lower the miles per gallon, the bigger the tax. The motivation was to incentivize people to buy more fuel-efficient cars, and companies to produce more fuel-efficient cars.

The Gas Guzzler Tax was designed with many notches. So, there is a certain tax within a one mile per gallon wide band, and if the vehicle's mpg is just below the threshold, the tax goes up by a lot. For example, going from 14.5 to 14.4 miles per gallon increases the tax by \$900. How flexible (i.e., elastic) is car design with respect to this tax? In particular, do car manufacturers re-engineer their cars to achieve the higher fuel efficiency and reduce the tax by a lot? It's got to be good for the consumers and probably the car manufacturers to do that.

Well, in economics it turns out that, for some questions, one diagram gives the answer, and this is one of those questions. And here it is:

Graph 1: Bunching of Car Model MPG's in the U.S.



The graph shows the distribution of the mpg number after the decimal for almost 1 500 car models in the U.S. over a thirty-year period. The bars are based on the decimal of miles per gallon, so the first on the far-left side are the number of models whose miles per gallon is between 14.0 and 14.1, or 17.0 and 17.1. The crucial fact is that the threshold (that is, the notch) is always at

0.5. For example, redesigning a car from one whose miles per gallon is 17.47 to one whose mpg is 17.52 reduces the tax due substantially. Moving from 17.37 to 17.42 would not. Thus, if car manufacturers are responding to this, you should see a lot more models in the .5 to .6 range, attracting a lower tax, compared to the .5 to .4 range. The graph shows exactly that. The University of Michigan is about 40 miles from Detroit, Michigan, where a lot of cars are made, or used to be made, and a lot of car research still happens. My co-author, James Sallee, and I talked to people who were in the room when they were designing their car models, and we said, does this come up? The answer was “you bet it does!”

In sum, there is significant bunching at the miles per gallon just over those thresholds so that a lower tax applies. And here’s another proof that this isn’t just a random occurrence. The Gas Guzzler Tax doesn’t apply to SUV’s or trucks. (Why not is a different issue.) When we construct the same diagram for SUV’s and trucks, we don’t see any bunching.

One last thing. Those of you who have ever been to a car dealer in the U.S. will know that, when you are looking to buy a car, there’s a sticker in the window that says “miles per gallon in the city” and “miles per gallon on the highway”. It turns out that those calculations of miles per gallon are different than the calculations used for the Gas Guzzler Tax. They are both measuring miles per gallon, but they are different. There’s also bunching in the distribution of mpgs shown on the stickers. There’s a lot more models just at 20.0, compared to 19.9, with the explanation being that the manufacturers understand that people have left-digit bias and 20 sounds a lot more attractive than 19.9.

6. Things We Still Need to Know

Let me turn now to the third and final segment of my talk, the things we still need to know. Luckily, there are a lot of things we still need to know, or else it would be really bad for the graduate students in the audience, and it would be a little bad for me.

As all self-centered people would do, I’m going to focus on things I’m thinking about now. They are not necessarily the most important things in the world, but I think they are important. I’ll talk about three, and then I’ll get to just a list of few other things that I think you young folks should start to think about.

6.1. A Few Hot Topics in Taxation

So here are a few hot topics.

- How best can one apply the findings based on an analysis of country A, to country B?
- Are gender-blind and race-blind policies gender-neutral and race-neutral?
- How best can a country increase taxes on the rich in the face of difficulties with taxing capital gains?

6.2. Some Thoughts on Extrapolating across Countries

In the last 20 years, empirical tax research has flourished all over the world. There is now fascinating research in scores of countries. That raises the classic question of external validity: can a finding based on data in, say, Finland, be applied to inform policy in, say, the United States? How externally valid are these results? What can the U.S. learn from FIT, and vice versa?

6.2.1. The External Validity of Estimated Tax Base Elasticities

One thing I have been thinking about recently is how experiences from a wealth-tax country, inform the discussion of wealth taxation in the U.S. In the last couple of years, the U.S. has had a fairly lively debate about whether we should have an annual wealth tax. Where does one look for evidence about what its impact would be? Well, you've got to look at countries that have a wealth tax, or once had it. But can you confidently extrapolate from this evidence?

This question is related to the tax base elasticity, because it turns out that there is no reason to think that a tax base elasticity is the same in all countries. In fact, there are reasons to think there are factors that vary systematically across countries that would affect the tax base elasticity, such as the country's attitude towards inequality, its attitude towards privacy, the efficiency of their administration and the degree of globalization. If you have a country where exit and entry, emigration and immigration, are really easy, you would expect the tax base to be more elastic, because you could have brain drain – or the opposite.

I have been working on a systematic way to estimate what factors across countries affect these elasticities, so you can get a quantitative measure of how generalizable an estimate of this elasticity based on data from, say, in Finland is to the U.S.--do we have to double it for the U.S. or do we have to halve it?

A particular interesting example of this question is whether the Norwegian wealth tax experience generalizes to the U.S. As background, as of 1990 there were twelve European countries that had an annual wealth tax, and now there are three, and Norway is one of them. There has been a lot of good research about the impact of a wealth tax on things like reported wealth and inequality in Norway. How should we in the U.S. learn from this? I would argue that, to answer this question, one has to keep in mind that Norway is not a random country. There are reasons that Norway has a wealth tax, and the U.S. doesn't. Are those reasons relevant to how we can extrapolate what we know about the Norwegian experience to the U.S. experience? I believe so.

Take one example. Norway has public disclosure of tax information, so in Norway, you can go on a website and find out how much your neighbor reported in taxable income. In the U.S., that has a 0.0000001 per cent chance of happening in my lifetime. (I admit that is not that long a horizon.) This is relevant, because to have an annual wealth tax, the government needs reports of your wealth every year. That is a matter of privacy to a lot of people in America, and obviously in Norway people feel different about privacy because they have had this disclosure for many years. You can't just extrapolate from Norway to the U.S. about the cost of implementing a wealth tax.

6.2.2. How are Finland and Norway Different?

In Table 1, I present some data about how Finland and Norway are different, from each other and from the U.S. I already mentioned a couple of things in my introduction: you have the best governance, you have a lot of trust in other people, you drink a lot of coffee, and you have a lot of heavy metal bands. The first two might actually matter for tax policy, but in the table, I look at some other characteristics that might matter more. In Finland and Norway, people on average believe that the government should be more active in redistributing income, they have more confidence in government, they think tax cheating is less acceptable, and their country is more globalized. All of these things should, in principle, affect the attractiveness of different kinds of taxes, and should also affect how well you can extrapolate from the experiences in Finland and Norway to the U.S. Just because they are different, doesn't mean you can't learn anything, but you've got to be careful.

Table 1: How are Finland and Norway Different?

	Finland	Norway	U.S.	OECD Ave. (Std Err)	Std Dev Diff. Finland-U.S	Std Dev Diff. Norway-U.S.
1. Egalitarianism	6.99	6.85	5.71	6.29 (0.70)	1.83	1.63
2. Confidence in Government	2.32	2.61	2.13	2.21 (0.28)	0.68	1.71
3. Acceptability of Tax Cheating	1.80	1.70	2.05	2.01 (0.42)	0.60	0.83
4. Globalization	83.9	77.7	67.4	76.1 (10.1)	1.63	1.03

1. Is "governments tax the rich and subsidize the poor" an essential characteristic of democracy? 1-10, 10 = Absolutely yes
2. How much confidence in government? 1-4, 4 = a great deal (scale reversed)
3. Is cheating on taxes acceptable? 1-10, 10 = always justifiable
4. Economic Globalization Index, 0-100, 100 = max

Data sources: World Values Survey (1-3) and KOF Globalization Index (4)

6.3. Are Gender-blind and Race-blind Policies Gender-neutral and Race-neutral?

Are gender-blind and race-blind policies gender-neutral and race-neutral? In the U.S., the tax system does not mention gender anymore, although it used to, and it doesn't mention race. Does this mean that these tax systems are in fact neutral to race and gender? The answer is no, in general, it could be. For example, it could be that Black households make choices and have characteristics that are correlated with how generous or ungenerous the tax system is, in which case the tax system isn't race-neutral, even though race isn't in the tax code.

In fact, it looks like in the U.S., the tax system is biased against Black households of given income and age, even though race is not measured. It also looks like, based on some preliminary work of my own, that the U.S. income tax system is biased in favor of women, even though gender isn't mentioned in the tax code. Why? Because women get more child-related tax credits such as the earned-income credit and child tax credit. Whether that gender disparity is a good idea or not depends on the social externalities from the activities that women do more of that qualify for the credits – just because there is this disparity, it doesn't mean it's a bad thing. The point is that, in both the gender and race examples, the tax code is blind, but the policies may not be.

6.4. Taxing the Rich without Taxing Capital Gains

The third and final hot topic is the taxation of capital gains. In the U.S., and probably everywhere, they are received disproportionately, in fact incredibly disproportionately, by rich people. But capital gains are often exempt when they are taxed, or tax preferred, and they are taxed upon sale rather upon the accrual of the gain. So how can a country have a truly progressive tax system, if it can't get at capital gains? It's a very tricky tax design problem, because in a realization-based system increasing rates will trigger lock-in, and taxing gains upon accrual or at death raises concerns of liquidity and the valuation of private businesses.

I learned recently that, in Finland, capital gains are basically fully taxable upon realization, at a rate up to 34% (the highest in Europe), along with a modest inheritance tax regime. What can the U.S. learn from the Finnish experience, keeping in mind the concerns about the limits of cross-country extrapolation I raised earlier?

7. I Could Go On...

There are a lot more than three important and intriguing questions we need to know more about, but I don't have the time to address any more in detail. Instead, I'll close by raising without much comment a few more.

How does the presence of behavioral anomalies affect conventional wisdom on optimal tax design? There's been a mountain of evidence showing that many people, when making tax-related decisions, do not behave as our workhorse rational-actor model predicts. But does this substantially change our view of the appropriate tax system?

The U.S. has just enacted a provision that subjects some very large corporations to a tax based not on taxable income, but rather on the income they report on their financial statement. Will taxing financial statement income degrade the value of accounting information to investors?

The U.S. has now had the Foreign Account Tax Compliance Act (FATCA) in place for several years, and scores of countries have committed to the Common Reporting Standard (CRS) for automatic exchange of financial information to other countries. These measures are designed to combat tax evasion through holding wealth in tax havens. We do not, though, have convincing causal evidence of the effectiveness of these regimes yet. Have FATCA and CRS reduced tax evasion via foreign accounts?

The bottom line is that this is a wonderful time to embark on an ambitious program of tax systems research, as the Centre of Excellence in Tax Systems Research promises to do. New data and new analytical methods promise to shed light both on new tax issues as well as perennial but difficult policy choices. I very much look forward to learning from the research they will generate in the years to come.