NEW YORK UNIVERSITY SCHOOL OF LAW

FALL 2016
COLLOQUIUM ON HIGH-END INEQUALITY

The Mapmaker's Dilemma in Evaluating High-End Inequality

Daniel N. Shaviro
New York University School of Law
Guest commentator: Liam Murphy
New York University School of Law

November 21, 2016
Vanderbilt Hall-202
4:00-5:50 p.m.
Week Number 5

SPEAKER SCHEDULE FOR INEQUALITY SEMINAR
(Mondays 4:00 to 5:50 pm)

October 24 – Robert Frank, Cornell University
5 short pieces: (1) *Why Has Inequality Been Growing?*, (2)*Why Luck Matters More Than You Might Think*, (3) *Does Inequality Matter?*, (4) *Why have weddings and houses gotten so ridiculously expensive? Blame Inequality*, and (5) *The Progressive Consumption Tax*.
Guest commentator: K. Anthony Appiah, NYU Philosophy Department.

October 31: Kate Pickett, Department of Health Sciences, University of York
(1) *Income Inequality and Health: A Causal Review*;
(2) *The Enemy Between Us: The Psychological and Social Costs of Inequality*
(both co-authored by Richard Wilkinson).

November 7 – Ilyana Kuziemko, Princeton University Economics Department
Support for Redistribution in an Age of Rising Inequality: New Stylized Facts and Some TentativeExplanations (coauthored by Vivekinan Ashok and Ebonya Washington).

November 14 – Alan Viard, American Enterprise Institute
Progressive Consumption Taxation: The X-Tax Revisited (chapters 1-3)
(co-authored by Robert Carroll)

November 21 – Daniel Shaviro, New York University, School of Law
The Mapmaker's Dilemma in Evaluating High-End Inequality
Guest commentator: Liam Murphy, New York University School of Law

November 28 – Adair Morse, Haas School of Business, University of California at Berkeley
Trickle-Down Consumption (co-authored by Marianne Bertrand)

December 5 – Daniel Markovits, Yale Law School. *Meritocracy and Its Discontents*

THE MAPMAKER'S DILEMMA IN EVALUATING HIGH-END INEQUALITY

Daniel Shaviro*

Preliminary Draft

Please don't quote or cite!

August 19, 2016

NOTE TO READERS: This piece mainly consists of chapter 2 of a book-in-progress, entitled *Enviers, Rentiers, Arrivistes, and the Point-One Percent: What Literature Can Tell Us About High-End Inequality*. However, the first five pages are adapted from a law review article, based on the chapter, that will be appearing in the University of Miami Law Review.

High-end Inequality Versus Low-End Inequality

According to an old joke, a statistician whose head was on fire, while his feet were encased in a block of ice, reported that, on average, he was very comfortable. This mythical individual brings to mind the Italian statistician Corrado Gini, who devised the famous Gini coefficient, measuring statistical divergence from a perfectly equal distribution of, say, wealth or income. Gini not only created the measure that bears his name, but urged that it be used to express numerically the extent of a given society's material inequality.

The big difference between the mythical statistician and the real one is that the problem Gini missed relates to interpretation, rather than measurement. Under his coefficient, extreme inequality at both the top and the bottom of the social scale will not statistically offset each other, yielding a false reading of zero aggregate inequality, along the lines of the fire-and-ice example. Instead, each will raise the quantum of inequality that the measure detects. Yet the coefficient still has the defect of amalgamating two normatively distinct phenomena in a single numerical

* Wayne Perry Professor of Taxation, NYU Law School. I am grateful to Mary Louise Fellows and Edward Kleinbard for helpful comments on an earlier draft, and to the Gerald L. Wallace Matching Gift Fund for financial support.

expression. Low-end inequality and high-end inequality raise fundamentally different types of issues. They do not pose identical, mirror-image, or symmetric problems that should generally be analyzed jointly.

Consider first low-end inequality, or poverty as measured relative to the median in a given society. It matters because, if some people are worse-off than the rest of us, basic human beneficence supports trying to help them. Now, *how* we can best do this is controversial. And, to those of a sufficiently libertarian bent – who emphasize personal responsibility, in terms of moral desert, even when the poor could be aided without significant adverse behavioral effects – the question of *whether* we should help them may seem less obvious than it does to me. But the fact that addressing low-end inequality, if it can be done right, involves helping people who are worse-off than the rest of us, makes it a relatively easy project to embrace.

High-end inequality, or that pertaining to the super-rich, is different. Would the idea, in addressing it, be just to make very rich people worse-off, even if no one else gained thereby? From the standpoint of beneficence, why would we want to do that? Thus, the view that we should seek to reduce high-end inequality is easiest to embrace if it has harmful effects on people below the highest material level.

In evaluating whether and why this might be so, we should keep in mind what high-end inequality, in the United States and peer societies, actually looks like at present. Over the last thirty years, wealth and income concentration at the very top have been dramatically increasing, and already have reached levels unknown for a century. However, this has been almost entirely due to the rise of the top 0.1 percent in the wealth distribution, who held only 7 percent of the U.S. national total in 1979, but by 2012 held 22 percent (Saez and Zucman 2014)

Given this degree of high-end wealth concentration, the common practice of speaking about the "one percent" versus the "ninety-nine percent" actually misses the point. The 0.1 percent have been pulling away even from the one percent – and, for that matter, the 0.01 percent have been pulling away from the 0.1 percent, and the 0.001 percent from the 0.01 percent, in a process that economists call "fractal inequality" (Lowry 2014). Just as, in a fractal such as snowflake, one finds the "same amount of 'jaggedness' or 'unevenness' at every scale" (Easterly 2010), so, in data covering the last few decades, "one sees the pattern of growing inequality among the population as a whole replicated within any subgroup of that population" (Krugman 1994, 133).

There are a number of different grounds on which high-end inequality, when so sharply concentrated at the very top of the distribution, may be bad for everyone else. For example, it may lead to plutocratic capture of the political system by the super-rich, enabling them to extract rents and greatly reducing the system's responsiveness to all others' interests (Bartels 2010; Gilens 2012). Extreme high-end income and wealth concentration may also reduce economic stability, output, and growth (Stiglitz 2013). Claims to this effect call for conventional "hard" social science research,[1] which has indeed been ongoing, albeit well short of reaching consensus.

However, the rise of the 0.1 percent also raises a set of subtler, more intangible issues that require different modes of assessment. We are an intensely social species, and often a rivalrous one, prone to measuring ourselves in terms of others, and often directly against others. People thus "have deep-seated psychological responses to inequality and social hierarchy," creating the potential for extreme wealth differences to "invoke[] feelings of superiority and

[1] "Hard" social science research, such as that in economics and related disciplines, can be defined as that which relies on "theory, mathematics, rigorous methods, falsifiability, and replicability" in emulation of the physical sciences (Graham and Kantor 2007, 1).

inferiority, dominance and subordination" that powerfully "affect[] the ways we relate to each other" (Wilkinson and Pickett 2010, _).

In one view, this causes extreme inequality to be akin to pollution (Subramanian and Kawachi 2006, 149). According to recent research by the British social scientists Richard Wilkinson and Kate Pickett (2010, 19), high-end wealth concentration does not just reduce happiness for all groups – the rich as well as the poor – but even has measurable adverse effects on social trust, economic mobility, life expectancy, infant mortality, children's educational performance, teenage births, homicides and other violence, imprisonment rates, mental illness, drug and alcohol addiction, and obesity. While these claims likewise fall within the hard social science realm, at present they remain fiercely disputed (as in Snowden 2010).

However that debate proceeds, it cannot entirely resolve the psychological and moral issues that inequality raises. How deeply and widely felt are the sentiments of superiority and inferiority, or dominance and subordination? How unhappy do they make people, and is the pain at the bottom greater than the pleasure (if such it is) at the top? Are unequal power relationships morally objectionable for their own sake, even if people grow accustomed to and even comfortable with them? And if people in the 99.9 percent feel diminished by the economic gulf between themselves and those at the top, is this just a matter of socially destructive "bitterness" and "begrudg[ing] others their prosperity" (Brooks 2014) which they really ought to get over, and which policymakers ought to ignore? Or does it offer legitimate and important grounds for seeking to reduce high-end inequality?

In order to evaluate such issues, one needs to go outside the boundaries of conventional hard social science research – and in particular those of public economics. The problem with much economic analysis of high-end inequality is not, in the main, one of ideological bias in any

particular direction. Public economics methodologies can be, and have been, deployed on both sides of the debate regarding whether we should energetically address high-end inequality.[2] The problem, rather, is that a type of methodological tunnel vision that economists have adopted for good reasons, leading to high intellectual payoffs to our accumulated knowledge in many settings, has negative payoffs, unless duly supplemented, in this one.

This chapter therefore seeks to advance understanding of the following questions:

--What do we and don't we learn from the public economics literature regarding the issues associated with high-end income and wealth inequality?

--Why can't even such seemingly technical issues as the income tax rate structure at the top depend purely on standard economic analysis?

--What features of public economics as a discipline have produced both its triumphs and limits to its usefulness?

The rest of this chapter proceeds as follows. First I set forth what I call the "mapmaker's dilemma," which helps explain why modern economic analysis has achieved such wide-ranging intellectual triumphs, yet will prove inadequate as an all-in-one touchstone for analyzing high-end inequality. Then I discuss the crucial role of a very simple and indeed simplistic notion of "utility" in standard economic analysis. Finally, I discuss optimal income taxation, the main tool offered by modern public economics for the analysis of high-end (as well as low-end) inequality, and then offer a brief conclusion.

The Mapmaker's Dilemma

Economists and other social scientists, like mapmakers, aim to provide models of some part or aspect of the world. These models must combine being reasonably accurate with being sufficiently usable and useful. Unfortunately, these two objectives are often in direct conflict.

[2] See, for example, Diamond and Saez (2011); and Mankiw, Weinzierl, and Yagan (2009).

Hence, economists who are studying real world social or economic phenomena, such as high-end inequality, face a version of what I call the Mapmaker's Dilemma. That is, they must choose between how much accuracy, as opposed to how much usability, they are willing to sacrifice.

Leave it to Lewis Carroll to have identified one very clear and clean response to the Mapmaker's Dilemma. In *Sylvie and Bruno Concluded* – the second volume of a kind of follow-up to the *Alice* books that strews gleaming, beautiful diamonds of Carroll's delightfully hyper-logical nonsense amid gobs of gooey, indigestible sentiment – a mysterious visitor from a foreign land or world, known only as Mein Herr, asks the narrator:

"What do you consider the largest map that would be really useful?"

"About six inches to the mile."

"Only six inches!" exclaimed Mein Herr. "We very soon got to six yards to the mile. Then we tried a hundred yards to the mile. And then came the grandest idea of all! We actually made a map of the country, on the scale of a mile to the mile!"

"Have you used it much?" I enquired.

"It has never been spread out, yet," said Mein Herr: "the farmers objected: they said it would cover the whole country, and shut out the sunlight! So we now use the country itself, as its own map, and I assure you it does nearly as well."

Jorge Luis Borges apparently liked this passage enough to use it as the inspiration for a one-paragraph short story, fittingly named "The Exactitude of Science." Here he carries the narrative a step further. In a great empire somewhere, "the College of Cartographers evolved a Map of the Empire that was of the same scale as the Empire, and that coincided with it point for point." Succeeding generations, however, found this map so "cumbersome" that, "not without irreverence," they "abandoned it to the Rigours of sun and Rain. In the western deserts, tattered

fragments of the Map are still to be found, sheltering an occasional Beast or beggar; in the whole Nation, no other relic is left of the Discipline of Geography."

The Mapmaker's Dilemma has two distinct elements. First, miniaturization inevitably means loss of local detail. Second, usable maps must generally be flat, but the Earth is spheroid. While this hardly matters when the scale is sufficiently small, for maps of the entire world it leads to significant distortion. Different regions' shapes or relative sizes must be misrepresented, for the same reason that one cannot simply flatten out the skin of an orange. This offers considerable scope to choose the distortions that one finds most personally amenable, and then perhaps to forget that they are distortions. Perhaps it is not entirely coincidental that North Americans and Europeans still commonly use the Mercator projection method, dating back to 1569, which (while offering accurate shapes for the world's large landmasses) greatly exaggerates the northern continents' sizes relative to those of Africa and South America.

Economists, like real world mapmakers and unlike Mein Herr's countrymen or Borges' College of Cartographers, have leaned towards usability, albeit inevitably at the expense of perfect accuracy. This has served the field well. The rise of modern economics to the top of the academic pecking order (see, e.g., Schumacher 2014) reflects its many great triumphs in showing just how much one can explain by using very simple behavioral models that employ crudely reductive assumptions regarding human motivation. (More on these assumptions shortly.) Just as with maps, however, this comes at the dual cost of losing detail and flattening the underlying reality. Just as with maps, the flattening – in the sense of actually distorting important motivational inputs to behavior, not just simplifying them – matters more for a large-scale issue, such as the social evaluation of high-end inequality, than it does for a small one, such as understanding how equilibrium emerges in the market for vanilla beans or canola oil.

The intellectual progress that economists have made by pursuing very simple models speaks for itself as validation of their choice in responding to the Mapmaker's Dilemma. The danger, however, is that immersion in such models can lead one to forget the distortions and inaccuracy in cases where these are highly relevant. Still, the upshot is not that conventional economic analysis of high-end inequality should be abandoned, like geography in Borges' fictional empire by reason of its more extreme and opposite response to the Dilemma. Rather, the point is that we should remember to supplement the standard economic model, addressing its most important omissions and rounding it out as needed.

The Role in Economic Models of "Utility"

In evaluating how policymakers might respond to high-end inequality, the most pertinent economic literature is that in optimal income taxation, straddling public economics and welfare economics. However, before explaining the basics of how this literature approaches the issue, I will set the stage a bit, by delineating and critiquing some of its key underlying methodological and normative assumptions.

Economists like micro-foundations. For example, since a society consists of individuals, they generally prefer to start their analyses of market and other social interactions by building up from a model of individual behavior. This involves making assumptions regarding not just how, but to a certain extent why, a given individual makes the choices that shape her actions.

In the basic model that underlies, not just public economics and welfare economics but also price theory (exploring how markets operate), each individual has a utility function that we do not try to explain. It is just there. A utility function is a "mathematical function representing an individual's set of preferences, which translates her well-being from different consumption bundles into units that can be compared in order to determine choice" (Gruber 2013, G-11).

Just what and how much is being assumed or claimed, via the use of utility functions, is both contested and slippery. Moreover, just how much one needs to claim varies with the context. For example, we will see that, when using optimal income taxation to evaluate high-end inequality, one needs to make more capacious claims about utility than when using price theory to model how the price of canola oil is set in a perfectly competitive market.

The core ambiguity can be seen in the definition that I quoted above, which mentions both "well-being" and "choice." Of these two terms, choice is much easier to observe. Other people may see what I do, but they can only try to infer how I feel. But psychic wellbeing – however one ends up defining it – has more obvious normative significance. If I care about my own wellbeing and extend this concern to others (either from beneficence, or intellectual acceptance of their similarity and moral equivalence to me), then I will also care about their wellbeing, but not necessarily about their choices, other than as instrumental to their achieving wellbeing.[3]

The dichotomy between wellbeing and choice has been well understood in economics for a long time. For example, Alfred Marshall (1920) noted that, while "[u]tility is taken to be correlative to Desire or Want … desires cannot be measured directly but only indirectly, by the outward phenomena to which they give rise." Thus, he took comfort in the fact that, "in those cases with which economics is chiefly concerned, the measure is found in the price which a person is willing to pay for the fulfillment or satisfaction of his desire."

Since Marshall's time, however, economics has been on a wide-ranging imperialist binge. It now looks far beyond the study of commodity markets and international trade, to explore, for example, racial discrimination, drug addiction, marriage markets, dating strategies, and the right

[3] One can, of course, adopt if one likes an ethical framework in which we care about choice for its own sake, and not about wellbeing. This, however, is not the predominant methodological approach in economics (nor does it jibe with my own personal views).

to privacy. Even by Marshall's time, welfare economics had emerged, purporting to offer tools for the evaluation of aggregate social welfare. This commonly involves defining social welfare as a positive function of the psychic welfare that the members of a society would experience under different circumstances (such as the adoption of alternative government policies).

Marshall's way out of the maze, which was to focus on "the price which a person is willing to pay for the fulfillment or satisfaction of his desire," was later expanded and formalized by Paul Samuelson (1938), via the theory of revealed preference. Samuelson argued that consumer models could be "freed from any vestigial traces of the utility concept" (71) by relying on "the amounts of *n* economic goods which will be purchased per unit time by an individual faced with the prices of these goods and with a given total expenditure" (62). In other words, consumer choices, which at least in principle could be directly observed, were generally sufficient for economic analysis, without there being any need to worry about psychic underpinnings that could neither be directly observed nor compared to each other.

What perils were Marshall and Samuelson so understandably eager to avoid? One way to show this is by setting forth a maximally capacious version of the concepts of utility and utility functions. A "mapmaker" in economics whose inclinations were opposite to those of Mein Herr's countrymen and Borges' College of Cartographers – opting for maximum usability, rather than accuracy – might be tempted to posit the following: Given your underlying preferences, your mental state under any particular circumstances will always have a quantifiable hedonic utility score, in terms of the sensations that you experience of happiness, contentment, pleasure, absence of distress, and so forth. (Obviously, the difficulty of saying just what this utility *is* testifies eloquently to the underlying problem.) The higher your score – that is, the more "utiles," or units of utility, you feel – the happier or better-off you are.

In short, rather than assuming a can-opener, as per the old joke about the economist on a desert island who wants to open a can of food with no implements, the maximally capacious approach involves assuming a utilometer: a determinate quantitative gauge controlling the individual's behavior, and perhaps that she can even read. (While the utilometer also might control behavior automatically, in the manner of a thermostat,[4] even an analyst who was prepared to posit this scenario might feel compelled to recognize the fact that people often report consciously considering their choices, and having powers of introspection.)

A utilometer would certainly be a most convenient thing to have in practice, at least if it was sufficiently easy and cheap to operate. But despite its absurdity when described this bluntly, the state of affairs that it posits is not completely ridiculous (just partially so). If you had an internal utilometer that you could read, along with enough information about the world, it would empower you to rank all of your choices – concerning, say, how much to work, where to live, and what consumer items to buy. What is more, you would be able to rank them not just ordinally, or from best to worst, but also cardinally, or in terms of your degrees of relative preference. In actual fact, we can often do at least a rough small bit of this. For example, you may know, not just that you prefer pizza to pork chops, and pork chops to going to the dentist, but also that you regard the first of these two choices as presenting a much closer call than the second one.

While the internal or introspective objections to positing utilometers are bad enough, economists have tended to worry more about the external or evidentiary set of problems. This worry reflects the fact that, while we all have at least some direct experiential access to our own mental states, other people's feelings (and utility, insofar as there is such a thing) can only be indirectly inferred. Hence, even if one does not worry too much about the "zombie problem" in

[4] I ignore here the possibility, discussed in Chalmers 1997, that thermostats might have consciousness.

moral philosophy – involving the possibility that, while I (the observer) know that I actually feel things such as pleasure and pain, other seemingly sentient beings might merely look as if they do[5] – we face the apparent impossibility of making interpersonal utility comparisons. Who can say, for example, how great my subjective enjoyment of pizza or pork chops actually is, or my distress from going to the dentist, as compared to that of my neighbor, even if, in the same circumstances, we make exactly the same choices and even express ourselves identically?

This is where revealed preferences were thought to ride to the rescue. If both you and I would pay up to $20 for a pizza, and up to $15 for a pork chop, why not treat that as effectively the measure of the utility we each would derive from each item? "Desire or Want," then, if invoked at all, might simply be placeholders for the unknown and irrelevant underlying processes (conscious or not) that presumably generated the visible exercise of choice.[6]

Unfortunately for the use of revealed preferences, much evidence now shows that people often do not make consistent choices as expressed in terms of price (see, e.g., Bernheim and Rangel 2005). But even if people invariably expressed consistent valuations, the effort to substitute choice for utility would fare better with respect to some types of economic inquiry than others.

Thus, suppose one is observing a computer simulation of a marketplace with buyers and sellers, featuring interactions between "characters" that follow complicated algorithms but are no more sensate than a thermostat. For price theory, this would be good enough to generate testable

[5] See, e.g., Kirk 2009 ("Zombies are exactly like us in all physical respects but have no conscious experiences: by definition there is 'nothing it is like' to be a zombie. Yet zombies behave like us, and some even spend a lot of time discussing consciousness. This disconcerting fantasy helps to make the problem of phenomenal consciousness vivid, especially as a problem for physicalism.").

[6] Under such a view, it might simply be linguistically convenient to say, as a shorthand, that the pizza's utility to each of us equaled $20. Analogously, when discussing biological evolution, it may be convenient to use teleological language as a shorthand, without one's meaning to suggest that anything beyond blind processes is at work. An example would be saying that the "reason" our ancestors became bipeds is that it freed up their hands for other uses. One who said this might simply mean that positive natural selection for bipedalism was driven predominantly by the advantages associated with having free hands.

empirical propositions. One could even use it, in this setting, to detect "inefficiency" in the simulated market's operations. Thus, suppose that one of the characters was "willing" to pay up to $20 for an item that another was "willing" to sell for as little as $18, but that the transaction did not take place, because the simulation required payment of a $5 "tax." This is a classic example of tax-induced deadweight loss. Defined in terms of the characters' price points, the transaction would have generated $2 of surplus, if only it could have taken place. That is, but for the tax, the buyer would have gotten something it "valued" at $20, and that the seller "disvalued" providing at only $18. What is more, there would not, in this instance, have been any "tax revenue" generated.

Yet there would be no reason for us to care whether the characters in the computer simulation were getting the things that they (acted as if they) "wanted." If we operate from a principle of beneficence – generalizing from our own feelings to ascribe feelings that have similar moral importance to other, apparently comparably sensate beings – then actually caring about the characters' "frustration" would seem to rest on bringing "Want or Desire" back into the picture. Hence, actually objecting to the inefficiencies detected by price theory, no less than basing policy judgments on the use of welfare economics with its explicitly utility-based framework, requires defining utility in terms of internal sentiments that are actually felt by someone, albeit unobservably (at least as a matter of direct experience) by everyone else.

Is the main problem raised by using utility to connote psychic wellbeing the fact that we cannot make interpersonal utility comparisons, at least without engaging in unverifiable speculation? This was indeed, for many decades, though not as much today, the main concern that made economists eager to stick with revealed preferences insofar as they could. My own view, however, is that, just as one can safely ignore the "zombie problem" in all of the daily

social interactions that fill one's life – simply assuming that others' capacity to feel things is generally comparable to one's own – so we can generally ignore this problem in making social welfare judgments. Operating under the assumption that people are basically the same, in terms of the relationship between their revealed preferences and the true intensity of underlying mental states, seems not only polite and respectful, but also the best way of minimizing the potential size of one's errors in social welfare judgment. Suppose that I really cannot know who is the relative "utility monster" (Nozick 1974, 41) – that is, the person with stronger felt pleasures and pains, as between you and me. Even if we do in fact differ in this regard, albeit unknowably, a random guess would make the expected social cost of the error, given the 50 percent chance that I would get it backwards, greater than it would have been had I assumed psychic equality.[7]

Accordingly, in my view the more serious problem with a utility framework that frankly avows its reliance on "Desire or Want," conceptualized as if we had internal utilometers, lies on the introspective / commensurability side, rather than on the interpersonal comparability side. This will turn out to matter a lot, with respect to the use of welfare economics (including optimal income taxation) to assess high-end inequality, when such use fails to reflect appreciation of the Mapmaker's Dilemma, and thus of the need to address over-simplification and distortion. But before turning to what that framework both captures and misses, it is worth turning to how it typically models people's utility functions.

Main Characteristics of the Commonly Posited Utility Function

The standard economic model of people's utility functions follows two main principles. The first is non-satiation. That is, more of any item is always preferable to less of it, all else equal. In effect, there's always room for Jell-O (as a rather revolting advertising campaign once

[7] Abba Lerner (1944) showed that, under complete ignorance regarding who has which utility function, the optimal distribution of income (ignoring incentive effects) is completely equal. Among his key assumptions is that each individual's utility function features declining marginal utility for income, as I discuss below.

put it), and indeed for all other goods as well.[8] Under the second assumption, the extra utility that one derives from each extra unit of a given item (including the enjoyment of leisure) is always less than that produced by the preceding unit. The first slice of pizza yields more utility than the second, which yields more utility than the third, and so on ad infinitum even though one is presumed never quite to reach zero marginal utility given the principle of non-satiation.

Suppose Adam and Beth are each choosing how much pizza and beer to buy with the $20 that each has brought to the neighborhood Joe's. Adam might happen to prefer more pizza and less beer, while Beth might prefer less pizza and more peer. However, if pizza and beer provision is perfectly continuous – that is, if one can fine-tune how much of each one buys by fractions of an ounce or less – then each individual's choices will equalize the marginal utility that he or she derives from (a) the last unit of pizza consumed, and (b) the last unit of beer. Otherwise – say, if Adam faced the prospect of deriving slightly more marginal utility from his last pizza unit than beer unit – he would be able to increase his total utility, while still spending the same $20 overall, by instead purchasing slightly less pizza, and slightly more beer.

An assumption that frequently attracts adverse comment is that of consistent rational choice. Adam and Beth each seek to maximize utility, and thus are presumed to make the choices that, so far as they can tell from the information that is available to them, will have this effect. Moreover, how one formally presents the choices – for example, whether one starts with more pizza or more beer, where the two can readily be swapped with each other – is presumed to have no effect.[9]

---

[8] If one can save current resources for future use, the principle of non-satiation becomes more intuitively plausible than it is in a one-period world featuring just food.

[9] Robin West (1988, 868) has memorably expressed the absurdity of this view, if taken as a literal representation of reality. "[E]conomic man invariably knows what is best for himself, and he inevitably is motivated to seek it. He knows his own subjective welfare perfectly and pursues it relentlessly. He is the infallible judge, for example, of whether he 'would prefer' pushpin to poetry, alcohol to nutrition, or heroin to shelter. He knows best not only whether a Coke or a Pepsi would yield him greater pleasure, but also whether a liberal education or an

By now, however, it is not widely disputed, even within economics, that as Alfred North Whitehead (1978 ed., 79) put it, the assumption of perfect rationality is "palpably false – [people] are only intermittently rational – merely liable to rationality." Indeed, the burgeoning field of behavioral economics explores how people's choices may systematically depart from those one would expect from rational utility-maximizers. Thus, hyperbolic discounters (like the Grasshopper, in the parable of the Grasshopper and the Ant) fail to make adequate provision for the future, such as by saving for retirement. And an addiction to heroin or cigarettes need not, as the economists Gary Becker and Kevin Murphy (1988) posited, be rational, with any horrifying playout, or continual failed efforts to quit, merely reflecting that the benefits were front-loaded.

Even those among us who can rationally rein in their own irrational proclivities, in the manner of Odysseus having himself tied to the mast before the Sirens were within earshot, may be subject to manipulation via "choice architecture" (Thaler and Sunstein 2008). For example, suppose that gasoline costs $2.50 per gallon if you pay with cash, and $2.60 per gallon if you use a credit card. Rational consumers who were operating in accordance with the standard model would decide how to pay based simply on whether the convenience of using a credit card was worth the extra cost. But real world people tend to hate "penalties" more than they like "bonuses." Hence, it has been shown empirically that they will tend to use cash more, and credit cards less, if $2.50 is the posted price but there is a 10 cent per gallon penalty for using a credit card, than if $2.60 is the posted price but there is a 10 cent per gallon bonus for using cash. This violates consistent rational choice, unless one makes the model uselessly tautological by positing exactly the degrees of utility from receiving bonuses, and disutility from incurring penalties, that would serve ex post to "explain" (i.e., be consistent with) the behavior.

---

apprenticeship would better prepare him for life. His preferences perfectly mirror his subjective welfare, and his choices perfectly mirror his preferences. Thus, he relentlessly chooses what he prefers, prefers what he wants, wants what he desires, and desires what will maximize his subjective well-being."

One question that economists are still wrestling with is to what degree such rational choice problems can be domesticated – that is, treated as merely special exceptions to the standard model, to be dealt with on a targeted or ad hoc basis, without requiring fundamental rethinking. An example of ad hoc correction would be using behavioral "nudges" to increase retirement saving, if one believes that many people would otherwise save too little, as judged from the standpoint of their "true" preferences or welfare (see Shaviro 2015).

Going down this path raises the question of exactly where to stop. It also suggests more modestly reformulating continued reliance on a revealed preferences framework on merely institutional grounds. Even granting the pervasiveness of rationality problems, an individual may generally have both the strongest incentive of anyone to act in her own self-interest, and the best particularized "local" knowledge of anyone (other than, perhaps, intimates) regarding her actual preferences and circumstances. However, even if these considerations strongly support a social or political rule of thumb favoring reliance on revealed preferences as evidence of underlying utility, that is not the same as fully resurrecting the rational choice framework.

As it happens, for purposes of analyzing high-end inequality, problems with the rationality assumption are less consequential than those pertaining to the basic concept of utility. Again, utilometers would make it easy to judge people's subjective welfare under varying circumstances, and thus to decide when they are best off. Nature has not, however, so equipped us. In the absence of utilometers, the concept of utility has both descriptive and normative problems that, at least in some settings, can challenge its usefulness as an analytic framework. While the literature on these issues is vast, certain aspects of particular relevance to assessing high-end inequality bear noting here.

Descriptive Problems With "Utility"

Given utilometers' potential convenience, why don't we have them? The fact that evolution evidently did not see fit so to equip us presumably reflects that the brain is a very calorically costly organ to operate. Our experiencing positive and negative mental states is probably best explained as instrumental towards our making choices that will tend to favor survival and gene transmission.[10] A genuinely useful utilometer would have to do more than just tote up current sensations, which might alone be costly. It also would have to project the future sensations that one's model of the world (and of oneself) predicted would follow from making one choice or another. Making adequate decisions promptly and cheaply, using rough rules of thumb, may often be better than aiming to make really good decisions slowly and at high cost. Hence, evolution should not have been expected to select for such egregious over-design (and indeed maldesign) as equipping us with utilometers, even assuming that our primordial ape brains could have gotten there through gradual modification, one step at a time.

Our lacking utilometers is nonetheless bad news – or, one might alternatively say, liberating – from the standpoint both of describing crisply how people decide, and of evaluating what actually makes them better-off, rather than worse-off. This provides crucial background for how both economists and those most critical of conventional economics have approached the challenges of explaining behavior and evaluating people's welfare in practice.

Compression of the inputs to utility in economic models – Given how empirically messy (at best) the utility concept becomes as an attempted description of reality once we acknowledge that people lack utiles and utilometers, economists have unsurprisingly chosen to use models that employ radical simplification. Again, recall the basic price theory model in which utility, under fixed preferences, results solely from consuming market goods plus leisure, subject to non-

[10] I do not address here why people have consciousness, rather than operating zombie-style without it – a longstanding riddle that might be explained either in terms of its evolutionary benefits or as a naturally emergent property or byproduct of the requisite brain complexity.

satiation and declining marginal utility, in the face of a budget constraint. Nothing else matters in the basic model.

Taken as a literal representation of reality, this model is so reductionist as to be absurd. Obviously, there is so much more that affects how we feel about ourselves and about our lives. For example, we are intensely social creatures who care about status and relative position. What is more, our preferences clearly are changeable, whether it be habituating to a different-sized house, developing a taste for expensive wines, or gaining / losing tolerance for harsh winters.

Yet there are three main reasons why the basic model may often be useful in practice – going beyond its being simple and tractable, which would not alone justify using it if it bore no discernible relationship to our actual behavior and internal experiences. First, it is plausible regarding the inputs that it does consider. Even non-satiation, while clearly false about pizza slices at dinner, stands on strong ground in a cash economy where saving is feasible. How many of us would turn down a higher salary, all else being equal?

Second, in such classic settings for economic analysis as predicting how price and quantity will equilibrate in commodity markets, the basic model may offer all that one really needs. For example, if one wants to understand how taxing coal would affect coal use and overall carbon emission, the model offers a powerful tool both for framing the theoretical analysis, and for setting an agenda for concrete empirical research. The trick, of course, is not to forget that, in answering some other types of questions, a fuller and more realistic model of behavior and welfare might be needed.

Third, once one opens the door to a broader analysis, the entire framework becomes substantially more contestable. Thus, consider the evidence cited by economist Robert Frank (1985, 1999, 2007, Frank and Cook 1995) to the effect that, because people care enormously

about relative status, their wellbeing depends not just on their own absolute consumption levels, but also on relative consumption levels for "positional goods," such as housing and cars. This leads to arm's races in relative consumption, such as house size, funded by all of the participants having to work harder just to stay in (relative) place. Frank argues that concern about the resulting negative "positional externalities" supports imposing pollution-style taxation on high-end (or all) consumption. He further argues that high-end inequality greatly worsens these negative externalities, by triggering costly "expenditure cascades" (Frank, Levine, and Dijk 2010) as people at the top set ever higher standards for the assessment of relative deprivation, making it ever harder and costlier for those at lower wealth levels to keep up.

Frank's arguments are clearly important to the assessment of high-end inequality. However, once one adds them to the basic model, one may also need to consider such ripostes as the following:

--How many people actually care, and how much do they care, about the relative status effects of competitive consumption? In the words of a Frank-skeptic at the Cato Institute, "I don't doubt that some people are that way. My own solution is not to have such people as friends. But …. [w]ouldn't the proper thing be to persuade people not to care about others' income?" (Henderson 2007).

--Should other-regarding preferences of this sort be rejected, even if not ameliorable? Suppose one explains the phenomenon that Frank describes as mainly reflecting envy – although he views it largely in terms of the importance of social "context," leading to unconscious adjustment of one's consumption norms based on what one sees other people doing (Frank 2013 ed., ix-xi, 29-42). Many would agree that, if racists enjoy causing the members of disfavored

groups to suffer, we should disregard this ugly and hateful preference. But "[t]o say that racism should not count, but envy should," seemingly requires further support (Leonard 2008).

--What about positive externalities that might result from high-end inequality and luxury spending? Suppose we believe that "spending on today's luxuries lead[s] to innovation and higher standards of living for all income groups …. The wealthy pay extra to enjoy the benefits of new goods, which … will later become inexpensive and widely adopted" (Kashdan and Klein 2006, 422-423). How can we tell whether this is more important, or less so, than the negative externalities that may result from high-end inequality and resulting expenditure cascades?

The difficulties of the broader issues thereby raised, once we begin considering expansion of the basic economic model to reflect that humans are a highly social species, can properly support either, and indeed both, of two opposite responses. For the practicing economist, a natural response is to say: The uncertainty and controversiality of these issues suggests that I can reasonably ignore them in my work. At least, by sticking to the basic model, I can hope to illuminate *its* implications for understanding broader policy issues. These are of interest not just analytically (in an art-for-art's-sake way), but also because they genuinely are relevant inputs to an overall assessment. For example, if one is considering proposals to tax the rich, doesn't it matter who would bear the economic burden of the tax, and to what extent it would affect economic output? These are clearly issues which the basic model can help to illuminate.

A seemingly opposite response, but equally correct in its place, applies to the actual or hypothetical policymaker – that is, anyone who is trying to determine her overall bottom-line views regarding high-end inequality (whether or not she is actually empowered to implement them). Here, these broader issues *can't* be ignored, just because they are difficult and uncertain.

After all, they are potentially very important – perhaps, as we will see, dramatically changing the ultimate conclusions that one would reach via the standard economic analysis standing alone.

If the end result is to reduce somewhat the relative policy significance of conventional economic analysis with respect to analyzing high-end wealth inequality, as compared to the input of sociologists, psychologists, and happiness researchers, then so be it. There is still plenty of need for economic analysis of these issues. The point is simply that economists must share the spotlight, more than they have perhaps grown accustomed to doing in this era of their preeminence (Wolfers 2015).

Incommensurability and its broader significance – A further, and indeed deeper, set of issues raised by our lacking utiles and utilometers pertains to what is sometimes called the incommensurability problem. Even though we do not entirely lack ordinal and cardinal insight into our own preferences, our mental experiences often cannot be placed on a single common metric that runs continuously from best to worst. For example, would I rather have a good bottle of wine, feel I did my job well today, find that a mild skin irritation has eased, or hear from an old friend? Or, if I am thinking more macroscopically and down the road, should I prefer the type of life I would have in twenty years, and the type of person I would have become, if I undertook a career as a psychologist, a popular novelist, a lawyer, or an investment banker?

In a pure revealed-preferences framework, incommensurability is not a problem. If I must choose, presumably I will, tautologically establishing (within that framework) what apparently was best. Suppose, however, one agrees that subjective mental experience is what matters, and that choice has merely instrumental value – which is, however, imperfect – towards optimizing that. Then the fact that different experiences can feel so thoroughly incommensurate further widens the already open door for arguing that some types of experiences are actually

better than others for the individual herself. What is more, especially in cases where alternative experiences are not simultaneously accessible – for example, because they would require that one change or develop who one is over time – there is extra room for the assertion that some types of experiences are in some sense objectively "better" than others.

A thorough subjectivist might distrust such assertions, unless supported by hard research evidence of some kind, or at least by introspection that one believed was generalizable rather than idiosyncratic. In practice, such claims are often intermingled with normative claims that take the standpoint of an outside observer who would consider some sources of subjective wellbeing morally preferable to others even if all of them could be measured in terms of interchangeable utiles. Indeed, often the subjective claim seems clearly to be offered as backup for the normative claim (unless it is the other way around).

A good example from fiction is Aldous Huxley's *Brave New World*. Huxley plainly agrees on ethical grounds with the Savage, who rejects the highly medicated ease, comfort, and life of superficial pleasures that a future society purports to offer, saying: "I don't want comfort. I want God, I want poetry, I want real danger, I want freedom, I want goodness. I want sin…. I'm claiming the right to be unhappy." But Huxley also suggests that the society's endlessly repeated mantra, "Everybody's happy now," is not really true, other than at a very superficial level, as many of the elite Alphas, at least, are desperately thirsting for something more. *Brave New World* would have been more interestingly ambiguous had Huxley been willing to contemplate the scenario where soma and simulated thrills actually could "work" for everyone.

Even if we had utilometers that gave each mental experience a comprehensive hedonic utility score, there would still be possible grounds for normatively preferring some types of experiences for others. But the subjective claim that some types of mental experiences are

inherently better than others, for the individual herself, would be harder to support in that scenario. Hence, incommensurability plays an important role in creating space for debate about the relative subjective value of different types of experiences and lives.

Surely the most famous example of argumentation drawing on the intuitions made plausible by incommensurability is John Stuart Mill's assertion, in his classic work *Utilitarianism*, of two closely related points. The first is that beings with "higher faculties" are subjectively better-off than those without such faculties, even though they "require[] more to … [be] happy" and are "capable probably of more acute suffering." Yet, despite these concerns, "[i]t is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied. And if the fool, or the pig, is of a different opinion, it is because they only know their own side of the question. The other party to the comparison knows both sides" (Mill 2003 ed. 188).

The proof Mill offers of this proposition is self-evidently false. People don't actually know how pigs feel, nor is there any reason to think that Socrates truly knows what it is like to be a particular "fool." Indeed, the unabashedly extreme elitism behind Mill's implicitly giving himself the Socrates hat, and people he considers inferior to himself the fool or pig hat, is stomach-turning today in a manner that he surely did not anticipate. Yet the intuition that one would rather be wise than foolish (even absent a hedonic payoff) is a powerful one, even insofar as it reflects mere "pride" (as Mill admits) and egoistic self-identification.

Second, and relatedly, Mill distinguishes between the "higher" or intellectual pleasures, and the "lower" or animal ones (187). Unlike Jeremy Bentham, whom he elsewhere quotes as saying that, the "quantity of pleasure being equal, push-pin is as good as poetry" (85), Mill emphatically asserts that the higher pleasures are qualitatively better. Once again, he relies on

the ostensibly close to universal verdict of "those who are qualified by knowledge of both" (189).

Here, his proof might initially seem to be on stronger ground, as people capable of experiencing the "higher" pleasures will surely often know the other kind as well. Yet he once again undermines the persuasiveness of his argument by denying that counter-examples are relevant. Those who know both and prefer the "lower" pleasures may suffer from "infirmity of character," or have lost their "[c]apacity for the nobler feelings," or may "addict themselves to inferior pleasures" despite not preferring them. Hence, "[i]t may be questioned whether any one who has remained equally susceptible to both classes of pleasures, ever knowingly and calmly preferred the lower" (188-189).

This last sentence is, of course, tautologically correct if no possible counter-example would count. But a proof so circular and automatically self-validating is no proof at all. Still, the fact that different types of experiences may feel so incommensurate can make this type of argumentation, not just non-falsifiable, but even (albeit as a matter of personal taste) intuitively plausible. Then again, some people today might reverse Mill's hierarchy and – taking a dim view of the psychological process that Freud called sublimation – insist that the strongest animal pleasures, whether or not extending to pushpin, are actually more authentic, more important to welfare, and/or more intensely satisfying than the intellectual ones.

Mill's distinction, if accepted, appears likely to weigh in favor of supporting greater high-end inequality. After all, under his view, why not deny the "pigs" among us some of their shallow, animal-like, lower pleasures, if the tradeoff is that society can offer more "higher" pleasures to the elite who are capable of appreciating them? Thus, consider a poor society in

which only substantial high-end inequality would permit the existence of a wealthy patron class that could support the arts.

Incommensurability can also, however, be deployed in opposition to high-end inequality. Thus, consider the view that personal dignity, autonomy, and self-respect are indispensable to leading a good life (as in Dworkin 2011), and hence are more subjectively (not just morally) valuable than merely satisfying as many as possible of one's consumption preferences. Dignity and autonomy may be viewed as undermined not just by low-end inequality, which may compromise poor people's ability to achieve the good life, but also by high-end inequality, which may promote objectionable (as well as unpleasant) hierarchy and subordination.

Similarly, consider John Rawls' (1971, 62) normative focus on "primary goods," which he defines as "things that every rational man is presumed to want," reflecting their importance to pursuing a "rational plan of life." Rawls includes basic health as a natural primary good, and affords self-respect a "central place." Given the importance Rawls ascribes to primary goods' universal availability, they might weigh against tolerating high-end inequality even without regard to his "difference principle," which requires all real world social inequalities to work to the advantage of the worst-off individual (75-78).

Habituation and other sources of changes in preference or utility function – With readable utilometers and enough research opportunities, there would be no need for the standard assumption in price theory that people's preferences are fixed. One would be able to measure whether an individual's utility was (or could be) greater once her preferences had changed, even if she had no opportunity to choose her preferences at any time.

Things are more ambiguous, however, in a world without utilometers. Suppose there are two alternative states of the world in which people have distinct utility functions, and in each of

which they would rationally optimize given their preferences and budget constraints. The standard model offers no way of determining whether their true happiness or satisfaction or welfare (however conceptualized) is greater in one state or in the other. That determination would require their getting to choose between the two states of the world, based on knowledge about the utility they would experience in each.

Once we allow for the possibility that people's preferences will vary depending on the state of the world, conventional economic tools cannot be deployed to analyze the welfare effects of significant societal changes, such as in the degree of high-end inequality. Suppose, for example, that Robert Frank is correct in asserting that high-end inequality yields substantial negative positional externalities by generating heightened expenditure cascades. Then tax or other policy measures that aimed to address this by reducing wealth concentration at the top might improve welfare even if they looked highly inefficient within the standard framework. Likewise, if people are happier, all else equal, when they live in more egalitarian (or alternatively, in more hierarchical) societies, the standard model would not reflect that.

Habituation, to both good and bad circumstances, raises further complications. Surveys of self-reported happiness, in a range of countries and at different times, consistently reveal that there is a large positive psychic payoff when economic advancement reduces preexisting dangers of starvation, lacking shelter or medical care, exposure to violence, and so forth. However, the increased satisfaction from, say, a bigger house and higher-end consumption choices appears to be considerably more fleeting. People may initially report great satisfaction from improved

circumstances. However, once they have habituated to the "new normal," they end up reporting about the same levels of happiness as before.[11]

Then there is habituation to the horrible. For example, people who have been subordinated may learn to function better by lowering their expectations and finding a way to get along. If subordinated from birth, they may be worse-equipped to function in a more benign environment than they would have been if exposed to it earlier. A proper comparison of steady states may require looking down the road, and trying to evaluate how much better or worse-off people might be once their expectations had changed.

The costs of transition to a society that is either more equal or less so than the current one, while relevant to the overall merits of a change, should not be confused with the steady state issues. In addition, when gauging how bad (or not) high-end inequality is, one should distinguish between transitional and steady-state causes of discontent. A number of the literary works that I examine in this book suggest that changes in inequality – in either direction – may themselves cause anxiety and social strain. Once again, evaluating the significance of this may place us outside the realm where hard social science, and in particular economic, approaches can most straightforwardly be used.

Normative Problems With "Utility"

Should all utility and disutility, without regard to its source, cause, or character, count in the same way? Such questions are common fare in debates concerning utilitarianism and other variants of welfare economics. Nozick's "utility monster" is one example of a hypothetical that may evoke intuitive unease about counting all utility the same, and thereby treating people, in effect, as merely utility generators whose separate identity lacks first-order moral significance.

[11] One could rationalize such habituation, in evolutionary terms, as setting them up for further striving that might yield further payoffs, and as avoiding the "waste" of self-administered psychic rewards on the payoff to choices that lie in the past.

Other hypotheticals may evoke intuitions that relate to incommensurability, rather than to interpersonal utility issues.

A classic example involves the torturer and the victim. Suppose a torturer's sadistic impulses are so well-developed that he actually enjoys inflicting pain and humiliation more intensely than the victim dislikes having them inflicted. This may seem unlikely, suggesting that, if positive and negative utiles existed, we might be confident that the sum total would be negative. Suppose, however, that enough people with sadistic or vindictive preferences were watching the torture to ensure that the overall utile count would be positive. One still might be reluctant to conclude that this meant the act of torture was good on balance – even if one would reach this conclusion in a "ticking bomb" scenario where it was indispensable to saving lives.

A common intuitive response to the conundrum rests on arguing that the enjoyment of torture is an illegitimate or unworthy pleasure, hence not to be counted as if it were no different than savoring ice cream. In addition, if preferences are malleable, one might argue that the sum total of human happiness can be greater when people enjoy ice cream than when they enjoy watching torture. By discouraging links between one person's utility and another's disutility, we can hope to create circumstances where the sum total of utility can be higher.

To further illustrate the intuition, while lowering the thermostat just a hair from its level in torture debates, consider racism. If the members of one group enjoy subordinating and oppressing the members of another group, we might find it psychologically plausible that, at least on a per-person basis, the pain imposed exceeds the pleasure derived. In addition, however, one may have a moral intuition classifying the dominant racists' enjoyment of subordination as illegitimate and unworthy. We also know that racist sentiments are not entirely universal and irremediable, and it may be clear that, in the absence of such sentiments, society as a whole can

get to a better place. Hence, there are multiple reasons for declining to value the dominant racists' enjoyment of subjugating others, even in the mythical scenario where we are otherwise toting up all the utiles.

In the debate over high-end inequality, each side can try to invoke the racism analogy in its favor. For those who are anti-inequality, extreme wealth concentration at the top may be viewed as yielding relationships of dominance and subordination that are not wholly unlike racial injustice. And even if the poisonous sentiments on both sides would not yield to exhortation, counseling, or medication, they presumably can be mitigated by reducing the extent of the inequality that gives rise to these relationships.

For those who are anti-anti-inequality, whether or not affirmatively pro-inequality, the issue is one of envy – that is, gratuitous and invidious resentment of the rich by the poor. Suppose that someone who was below the top 0.1 percent would feel good about bringing the richest individuals down a few notches, even if it did nothing for her directly. Even though her animus comes from "below," rather than from a higher social position as in the case of a racist subjugator, she similarly attaches positive utility to someone else's disutility. This might potentially support a similar critique of the sentiment's unworthiness, remediability, and socially destructive character.

Before more closely examining these rival claims about the normative assessment of responses to inequality, it's worth noting one thing that both claims clearly get right. Each involves position-related utility, or the dependence of one's subjective welfare on one's vertical placement relative to other people. Only a true naïf – and perhaps one with Asperger's-level inability to grasp how most people actually think about social interactions – could believe that it is literally true that utility comes only from own consumption, rather than also being affected by

people's relative positions. *Of course* we have position-related utility, whatever the empirical and normative weight (or non-weight) that one ends up assigning to it. So the question, other than such sentiments' degree of importance in the big picture, is whether any grounds would justify differential weighting of particular types of position-related utility.

Battle of the Frameworks, Part I: Envy, Dominance Versus Subordination, or Something Else Entirely?: The anti-anti-inequality position is best-known as the credo of Mitt Romney during his 2012 presidential campaign, when he asserted that all public political debate concerning high-end inequality should be viewed as inappropriate, as it was just "about envy …. [and] class warfare." Asked whether *any* fair questions can be asked about wealth distribution, without its being just envy, Romney conceded that it was "fine to talk about those things in quiet rooms and discussions about tax policy and the like." Not, however, in a presidential campaign.

How might one who wanted to justify, not just woodenly assert, the case for dismissing envy's normative relevance go about doing so? Arthur Brooks (2014), president of the American Enterprise Institute, offers as illustration a comment once made by the pop singer, Bono, explaining a difference he had observed between the United States and his native Ireland. "In the United States … you look at the guy that lives in the mansion on the hill, and you think, you know, one day, if I work really hard, I could live in that mansion. In Ireland, people look up at the guy in the mansion on the hill and go, one day, I'm going to get that bastard."

Brooks then adds: "[P]sychologists have found that envy pushes down life satisfaction and depresses well-being. [It] is positively correlated with depression and neuroticism, and the hostility it breeds may actually make us sick." The solution, he argues, is twofold. First, increasing mobility towards the bottom of the income scale will induce people to think like Bono's American, rather than like his Irishman. (Brooks does not address whether high-end

inequality might inherently tend to reduce upward mobility – although recent research suggests that it does (Corak 2013).) Second, everyone should agree to avoid "fomenting bitterness over income differences[, which] may be powerful politics, but … injures our nation." So Romney ostensibly was right about the "quiet rooms" after all.

In his envy diagnosis, Brooks offers all three of the crucial elements for devaluing the sentiment that he describes. It is morally unworthy, remediable, and socially destructive – not to mention psychically self-destructive. Yet Brooks' focus is strikingly one-sided. All we hear about is the person who is staring up from below. We don't hear anything about the perspective or the actions of "the guy that lives in the mansion on the hill."

What if the Irish grandee does things that earn his downhill neighbors' hostility? Even in the absence of conflictual political or economic interactions, suppose he likes to impress them with his own social superiority and their inferiority. This would undermine all three elements of Brooks' case. We now may feel that the neighbors' hostility is more justified, even if we don't want them actually to "get" the grandee. Their side of the dispute may now seem less remediable, other than by addressing high-end inequality itself. And one could argue that what destroys social concord is the high-end wealth gap itself, not just one side's supposedly gratuitous reaction to a two-sided fight over status and power.

In the United States today, there can be little question that "class war" sentiments, if one wants to call them that, emanate from both sides of the divide between the top 0.1 percent and everyone else. Consider the ludicrous comments made by Silicon Valley venture capitalist and billionaire Tom Perkins (2014), who infamously wrote to the Wall Street Journal so he could ungrammatically "call attention to the parallels of fascist Nazi Germany to [sic] its war on its 'one percent,' namely its Jews, to the progressive war on the American one percent, namely the

'rich.'" To Perkins, apparently, expressing even mild criticism of our society's most powerful group is closely comparable to one of the greatest campaigns of organized mass murder in human history.

When one considers that, in the contemporary United States, the "extremely wealthy are objectively far wealthier, far more politically powerful and find a far more indulgent political class than at any time in almost a century," the claim "manages simultaneously to be so logically ridiculous and morally hideous that Perkins deserves every bit of abuse" that he got (Marshall 2014). Yet Perkins does not stand alone, even in having "his self-censor and/or editor fail[] him so miserably." For example, not long before, billionaire investor Stephen Schwarzman called proposals to tax hedge fund managers at the ordinary income rate faced by millions of Americans, rather than at special capital gains rates, an act of "war" that was "like when Hitler invaded Poland in 1939" (see Brooks and McQuaig 2012). The ranks of billionaires comparing even mild criticism of the super-rich to the rise of Hitler also includes Home Depot founder Ken Langone.[12]

While even three such anecdotes do not by themselves prove the existence of a broader trend, the contemporaneous rise of similar (if not always so extreme) anger and fear among members of the top 0.1 percent has been widely noted. This has led to a wave of recent articles asking such questions as why billionaires are so angry at the rest of us (Surowiecki 2014), why they are so "whiny" (Leopold 2010), and why they feel so "victimized" by political criticism that, as a historical matter, is par for the course or even relatively mild (Freeland 2012).

[12] In fairness to Langone, while he is evidently made uneasy by challenges to high-end inequality, he has recently emphasized the importance of addressing low-end inequality, arguing that, "if we don't do something about helping these people on the lower end of the pay scale, I think we're setting ourselves up for serious problems" (Ocasio 2015).

Josh Marshall (2014) argues that political "insecurity, a sense of the brittleness of one's hold on wealth, power, privileges, combined with the reality of great wealth and power … breeds a mix of aggressiveness and perceived embattlement." He thus views Tom Perkins' feelings, if not his gross lack of tact and proportion in expressing them, as "pervasive" among the super-rich. Accustomed to extreme deference in their daily business and consumer lives, they find it humiliating and intolerable that they might need to "run to the political class hat in hand" (albeit waving large checkbooks) in quest of protection and reassurance.

In Marshall's view, the "sheer scale of the difference" in people's daily experiences and circumstances means that the super-rich "live what is simply a qualitatively different kind of existence. That gulf creates estrangement and alienation, and one of a particular sort in a democracy where such a minuscule sliver of the population can't hope to protect itself alone at the ballot box …. The disconnect between perception and reality, among such a powerful segment of the population, is in itself dangerous. And it's led to … a significant radicalization of the politics of extreme wealth" (Marshall 2014).

This critique of high-end wealth inequality relies on concern about lost social capital and cohesion – a concern that goes back, in the social science literature, at least to the work of Emile Durkheim (1893) – rather than on the narrative of dominance and subordination that may be favored by commentators to Marshall's political left. However, both critiques support viewing the "envy" frame as unduly one-sided, in the sense of applying armchair psychology just to the 99.9 percent, not to the 0.1 percent, and as if in a social vacuum. Wherever one eventually comes out in deciding what normative weight (if any) to give position-related utility, it surely requires looking in depth at the entire picture.

Battles of the Frameworks, Part 2: Envy or Context? – The envy debate also matters with respect to the Robert Frank point regarding positional goods. The critique that the solution is to "persuade people not to care about others' income" (Henderson 2007) is part and parcel of the anti-anti-egalitarian dismissal of concern about position-related utility.

Frank himself, recognizing that "[t]here are good reasons to limit envy and other corrosive emotions" (2013, 42), responds that the ill effects of costly expenditure cascades rest on the universal relevance of "context" to people's utility and broader assessments (29-42). Even our evaluations of temperature rest on the applicable frame of reference. A sixty-degree day seems cold in Miami in November, but warm in Montreal in February (31-32). So a person who sees other people's large houses may come to want a larger one for herself, not just out of rivalry but due to her evolving perception of surrounding norms. "If you respect people's preferences and they experience psychological costs from relative disadvantage, why shouldn't those costs be taken into account in a welfare analysis?" (41).

Whether or not one accepts Frank's analysis of expenditure cascades and their link to high-end inequality, it helps to indicate a need for broadening both the descriptive and the normative analysis beyond the range of standard economic models. After all, at a minimum relative consumption *might* be genuinely subjectively important. And Frank is hardly the first to bring it to broad public attention. His most famous predecessor is Thorstein Veblen, the economist and sociologist whose classic work *The Theory of the Leisure Class*, first published in 1899, sets forth the theory of "conspicuous consumption" and – though less remembered today – "conspicuous leisure."

Veblen defines conspicuous consumption in light of competition for status, as distinct from the Arthur Brooks model of envy on one side and supposed obliviousness on the other.

Veblen (1912 ed., 36-37) argues, for example, that "to gain and to hold the esteem of men it is not sufficient merely to possess wealth or power. The wealth or power must be put in evidence, for esteem is awarded only on evidence. And not only does the evidence of wealth serve to impress one's importance on others and to keep their sense of his importance alive and alert, but it is scarcely less use in building up and preserving one's self-complacency."

The key prerequisite for achieving the desired reputational effects, according to Veblen, is "conspicuous waste" (100), since the main point is to show one can afford it. There is good reason to think that this rationale still applies. For example, another anti-anti-egalitarian from the Cato Institute, Will Wilkinson (2009, 6), offers a classic Veblenesque example (without realizing it) in the course of explaining how U.S. consumption inequality has changed since Veblen's era: "At the turn of the 20th century, only the mega-rich had refrigerators or cars. But refrigerators are now all but universal in the United States, even while refrigerator inequality continues to grow. The Sub-Zero PRO 48, which the manufacturer calls 'a monument to food preservation,' costs about $11,000, compared with a paltry $350 for the Ikea Energisk 318 W. The lived difference, however, is rather smaller than that between having fresh meat and milk and having none. The IKEA model will keep your beer just as cold as the Sub-Zero model."

As Timothy Noah (2012, 170) notes in response, "if getting rich is only a matter of spending more money on the same stuff you'd buy if you were poor, why bother to climb the greasy pole at all?" Is the owner of the Sub-Zero being stupid, or rather "[d]oes he know something Wilkinson doesn't?" Surely Veblen's explanation is partly responsible, even if the Sub-Zero also has practical or aesthetic advantages over the Ikea model that a super-rich consumer would appreciate even if he or she were the last person on Earth.

One important thing that *does* seem to have changed since Veblen's era is the role of what he calls conspicuous leisure, involving not just exotic vacations that yield showy souvenirs, but also conspicuous wasting of time to show that one need not work. Veblen wrote *The Theory of the Leisure Class* in an era when "wealth acquired passively by transmission from ancestors ... [was] more honorific than wealth acquired by the possessor's own effort" (1912 ed., 29). Hence his use of the label "leisure class," adopting an identifying feature that would not figure as prominently in similar analysis today.

The Theory of the Leisure Class predated the era of putatively heroic "job-creators" and superman CEOs, who ostensibly do "tremendous" things "advancing the public good" (Mankiw 2014). And it likewise predates today's rentrepeneurs, such as Tagg Romney and Chelsea Clinton, who like to pretend that the generous paydays that fall into their laps reflect their own talents and efforts, rather than the inherited benefit of having famous parents. Today, whatever remains of conspicuous leisure is closely intertwined – perhaps even more so than in Veblen's day – with high-end market consumption, as in the case of a St. Moritz ski vacation.

In short, it is "no longer apt … to talk of a leisure class … conceived of … [as] in hasty flight from anything tainted by work," when "our moguls of the moment are workaholics" (Fraser 2015, 295) who would rather "think of themselves basically as working stiffs" (Andrews 2016) than admit their increasing resemblance to a hereditary aristocracy. Yet conspicuous leisure's replacement by conspicuous economic success in no way rebuts Veblen's model of peacock's tail-style social competition through one's interactions with the market economy. Instead, it further exemplifies the model's continued relevance, even (or perhaps especially) in a post-rentier era.

The Veblenesque process of competitive display may help to explain why the social impact of the top 0.1 percent has been so great – as evidenced by recent years' intense focus and debate on questions of high-end inequality – even though the super-rich often cloister themselves far out of view of the rest of us. Consider the experience of airline travel, which I have personally seen grow ever more hierarchical during my several decades on the academic conference circuit. Airlines now regularly offer, not just first class and business class seating, but also premium coach seating categories, such as United's Economy Plus, along with as many as five distinct boarding groups, and multiple categories for security clearance.

Obviously, part of the motivation for all this, on the demand side, is to purchase tangible benefits. To the seasoned airline traveler, it is good for its own sake to have more leg room on a long-haul flight, and to get first crack at the overhead bins. Yet extra status seems so clearly to be a part of what travelers with deeper pockets (their own or someone else's) are purchasing that the rise of multiple gradations – and of spatial inequality in the skies (Berman 2014) – seems closely related to contemporaneous trends in high-end inequality.

But is this inference refuted by the fact that people at the very top of the income and wealth hierarchies commonly fly their own private and company jets, rather than needing to line up even in Group 1? Not necessarily, if those in Group 1 know about these people, and if that, in turn, is enough to transmit social messages all the way down to Group 5, and beyond Group 5 to people who never go to the airport or fly. This is the process that Robert Frank has in mind when he discusses expenditure cascades. Of course, it does not rebut the possibility that the general social impact of high-end inequality would be vastly magnified beyond current experience if the super-rich were not as cloistered and sequestered as so often they are.

The Existence of Bimodal Social Norms – One last point worth discussing, before we turn in detail to how leading economic models address high-end inequality, concerns the question of what empirical and normative priors one should bring to a broader analysis. In a world without utilometers, one cannot simply observe and record the "facts" regarding high-end inequality's welfare effects. One's underlying assumptions about people and society inevitably will play a role, and one should at least try to be conscious about this.

Since we live in what is still a formally egalitarian age, we may all too complacently nod our heads at the famous words in the U.S. Declaration of Independence, holding the "truth to be self-evident, that all men are created equal."[13] No less might we take to be obvious the words of the French Revolution's Declaration of the Rights of Man that "[m]en are born and remain free and equal in rights,"[14] although many in the United States might reject the Declaration's further claim that "[s]ocial distinctions may be founded only upon the general good." Yet history tells us that, however fervently one may oneself accept the case for equality (in whatever one deems the relevant sense, given the question "equality of what?"), it has not, at all places and all times, been considered even true on the whole, much less "self-evident."

As Wilkinson and Pickett (2010, 26) note, "[h]uman beings have lived in every kind of society, from the most egalitarian prehistoric hunting and gathering societies, to the most plutocratic dictatorships." If one looks at the history of civilization, "[i]nequality appeared as soon as society was born" (Milanovic 2011, ix). Ever since that time, marked inequality has been common, often accompanied by social ideals that compared society to a family that is ruled

[13] Of course, this did not prevent the United States from retaining slavery for almost ninety years after the issue date of the Declaration of Independence. Yet Americans were often highly conscious of its irreconcilability with the Declaration's principles. On the one side, not just abolitionists, but much softer critics of slavery such as Abraham Lincoln, emphasized the broader implications of egalitarianism. On the other side, antebellum Southerners such as the writer George Fitzhugh responded by lauding slavery, on the grounds both that most people (of all races) were naturally slaves, and that it was better for workers than a market system, under which they were on their own.

[14] This reference to formal legal equality presumably reflects the fact that revolutionary France had only recently abolished aristocrats' legal privileges.

by the father, and that lauded the scenario where each individual accepted his or her proper station.

Yet suppose one looks even further back into the past, rather than stopping at the dawn of agriculture and civilization. Prior to recorded human history, and for an estimated 90 percent of the history of our particular human species, "people lived in groups in which equality was quasi absolute" (Milanovic 2011, ix, see also Binmore 2006, 6.). Go back even further, and for at least two million years our ancestors "lived in remarkably egalitarian hunting and gathering – or foraging – groups. Modern inequality rose and spread [only] with the development of agriculture" (Wilkinson and Pickett 2010, 208). Thus, "[d]espite the modern impression of the permanence and universality of inequality, in the time-scale of human history and prehistory, it is the current highly unequal societies which are exceptional" (207).

In short, while broadening one's gaze may contradict the two Declarations' seemingly serene confidence that equality is a universal and indisputable human value, it may in the end support a more limited claim. Substantial social and economic equality has been the prevailing condition for most of our evolutionary history, and thus presumably is a condition to which we became (and may remain) well-adapted. However, it also appears to be the case that we adapted quite readily to hierarchy, once the material conditions had arisen to make it feasible. Thus, Wilkinson and Pickett (2010, 206), despite their strong egalitarian sympathies, agree that "human beings have had to develop different mental tool-kits which equip them to operate both in dominance hierarchies and egalitarian societies." And the strategies that we deploy to function effectively in a hierarchical setting "are almost certainly pre-human in origin" (207).

Why might all this matter for present analytical purposes? One point is simply that we should be modest about the universality of our own particular social values – even if, in the end,

viewing the alternatives that have prevailed at other times and places should end up strengthening, rather than weakening, one's attachment to these values.

A second point is that, if we value an egalitarian vision of society despite recognizing its historical contingency even as an ideal, we should not complacently assume that its place is secure. Other visions could potentially supplant it, as they have before. In particular, rising high-end inequality sits ill with it, beyond even just endangering political democracy.

I myself don't expect the imminent return of medieval Europe's rationalization of pervasive hierarchy via the Great Chain of Being (progressing by degrees from God to angels to kings to nobles to commoners to animals). There may, however, be a more modern expression of the view that we are fundamentally unequal. In particular, market triumphalism, extreme meritocratic values, and revivified Social Darwinism already can be seen at times to promote the view that "winners" deserve everything, "losers" deserve nothing, and that the former should be celebrated while the latter are despised and mocked.

A third, more speculative point relates to the social science research that Wilkinson and Pickett (2010) argue supports viewing high levels of inequality as bad for everyone in all sorts of ways. As I noted in chapter 1, they argue that it has measurable adverse effects on social trust, economic mobility, life expectancy, infant mortality, children's educational performance, teenage births, homicides and other violence, imprisonment rates, mental illness, drug and alcohol addiction, and obesity (2010, 19).

This contested research needs to be confirmed, refuted, or modified on its own terms – a process that one hopes is underway. Speculation about how it might relate to our long evolutionary prehistory will not settle anything in this debate. Yet the fact that high levels of inequality emerged only with the relatively recent rise of agricultural civilization could provide a

plausible intuitive explanation for Wilkinson's and Pickett's results, if they end up being confirmed.

By analogy, consider people's generally keen taste for fats and sweets. This was a highly adaptive trait during the countless millenia when food was often scarce and famine a continual risk. It is considerably less adaptive today for people who can access unlimited fats and sweets entirely at will, and when market forces reward the effort to stimulate our liking for them. So we live today amid pervasive health problems that are in principle wholly avoidable, but that reflect our being evolutionarily maladapted, in some respects, for present circumstances.

Might the pervasive ill effects that Wilkinson and Pickett claim are associated with high levels of social inequality, if verified, reflect a similar adaptive mismatch between our internal wiring and our present circumstances? Here the claim would be, not that doing what you like leads directly to a bad health result, but rather that competitive social drives lead to greater psychic stress in a highly unequal society than in the type that prevailed during most of our evolutionary history, even if one has the tools to adapt and cope in either society. While any such view remains speculative, it offers a more credible evolutionary perspective than positing that so keenly a socially competitive species as our own evolved to derive utility solely, or even principally once the basics are met, from own consumption of market goods and leisure.

Optimal Income Taxation: The Dominant Economic Framework for Responding to High-End Inequality

With all this as background, we now can turn to the question of how contemporary economics assesses issues of inequality. In the economics literature assessing what, if anything, policymakers should do about it, there is broad agreement that taxes (along with transfers, to

address low-end inequality) should be at center stage.[15] The term "taxes," for this purpose, generally refers to such means-related instruments as income taxes, consumption taxes, wealth taxes, and estate or inheritance taxes, that are deliberately designed such that a rich individual would generally be expected to pay more than a poor individual, even if the rate structure is flat rather than progressive.

The leading economic framework for evaluating the use of taxes and transfers to address both high-end and low-end inequality comes from the literature on optimal income taxation (OIT), founded by James Mirrlees (1971), who later won the Nobel Prize in Economics for his work. Two main points regarding this literature are worth emphasizing here. First, it exemplifies the Mapmaker's Dilemma, by embracing a narrow framework where only own consumption and leisure matter – making its analysis more tractable, but less complete and satisfying, not to mention highly vulnerable to challenge on intuitive grounds. Second, while in some ways optimal income taxation is quite radical in conception, at least until recently it has generally been thought to support only a surprisingly limited policy response to high-end inequality.

(1) *Optimal income taxation's narrow framework, and its consequences* – Voltaire famously remarked that the Holy Roman Empire was neither holy, nor Roman, nor an empire. OIT does somewhat better than this at living up to its own name. While generally not about "income" as we most commonly use this term (indeed, only *other* branches of the optimal tax literature address actual income taxation as such), it is admittedly about "taxes." As for "optimal," OIT involves constrained optimization, or being as optimal as one can, all things

[15] Indeed, a vigorous debate in law and economics concerns whether distributional issues should be addressed *solely* through the tax and transfer system, based on measures such as income, or whether there is also a significant role to be played by "legal rules" (such as those concerning tort liability, enforcement of contracts, corporate governance, rent control, the minimum wage, and so forth). See, e.g., Atkinson and Stiglitz (1976); Kaplow and Shavell (1994); Jolls (1998); Sanchirico (2000); and Fennell and McAdams (2015).

considered, when the first-best solution that it identifies is unavailable. In this sense, OIT is actually about optimizing among suboptimal choices, given the true optimum's unavailability. It thus addresses the effort to achieve "optimality" in one semantic sense of the word, but not in another.

While the OIT literature has burgeoned over time in multiple directions, its starting point in Mirrlees' work goes something like this. Suppose that people derive utility just from market consumption and leisure, each of which has declining marginal utility. Since Mirrlees employs a one-period model in which there are no savings or wealth (other than fully formed human capital), market consumption is funded purely through earnings from work. Indeed, market consumption and earnings are necessarily equivalent, since there is no next period for which one might want to save.

In Mirrlees' model, all of the people in a given society have identical utility functions, and differ only in ability or wage rate, defined as the amount that one can earn per unit of time or effort. However, these inputs are not directly observable. Ability is innate and unalterable, but it, too, cannot directly be observed. By contrast, earnings, which are the joint product of ability and time or effort, can be observed. Each individual chooses the work level, and thus the earnings given her wage rate, that maximizes her utility from market consumption plus leisure.

We have, therefore, an almost maximally restrictive model, so far as the things that actually might matter from the standpoint of personal welfare in a complex and unequal society are involved. Nothing beyond own consumption in a vacuum, implicitly in a world with utilometers, makes its way into the model. This turns out to have genuine payoffs, in terms of offering non-obvious insights regarding crucial parameters for the design of a real-world tax system, but it also eventually proves to be a straitjacket.

The model's analytic purpose is to inform a social planner, who not only can observe people's earnings, but also can tax them to fund a uniform cash grant. Mirrlees' social planner is a welfarist, or one for whom "social welfare is posited to be an increasing function of individuals' wellbeing and to depend on no other factors" (Kaplow and Shavell 2002, 24). He might either be a pure utilitarian, who weighs each individual's welfare equally, or either of two varieties of a weighted welfarist – that is, one who assigns greater social weight to the welfare or utility of worse-off than of better-off individuals. At the limit, a weighted welfarist might embrace what is sometimes called a Rawlsian maximin, under which increasing the welfare of the worst-off individual in the society by just one utile would be worth any quantum of welfare loss whatsoever to better-off individuals (so long as they remained better-off in the aggregate). This is called "Rawlsian" because it resembles John Rawls' (1971) famous "difference principle," under which inequality is permitted only if it works to the advantage of the worst-off – although it is not in fact identical, since Rawls was not a welfarist.[16]

Even under the utilitarian approach to OIT-style social welfare – which values greater material equality due solely to its aggregate welfare effects given declining marginal utility, rather than as good in itself – Mirrlees' set-up could support imposing a 100 percent earnings tax, with all of the proceeds being distributed pro rata, but for the fact that this would have disastrous effects on labor supply. (Indeed, labor supply, and thus everyone's market consumption, would be zero under the model, given that the only reason for working is to fund market consumption at the expense of leisure.)[17] In light of that concern, one must trade off the

[16] One key difference between Rawls and a welfarist was his normative focus on primary goods, rather than on utility. For example, civil and political rights, as primary goods, could not rightly be sacrificed even if this were to make people hedonically happier.

[17] The nineteenth century utilitarian writer Francis Edgeworth (1897, 553) did indeed note that declining marginal utility supports expropriating all wealth and distributing it pro rata, if one ignores incentive issues and indeed all other considerations..

utility gain from redistributing resources from high-earners to low-earners, against the utility losses that result from the tax's inefficiently discouraging work.[18]

Returning to the question of what the word "optimal" actually means here, the use of "income" (i.e., earnings) in Mirrlees' model is concededly suboptimal, or more precisely part of a tradeoff, given its effect on work incentives. The ideal solution would be to base the tax directly on ability, if only it could be observed. And the only reason in the model for taxing earnings, other than that they can be observed, is that they are a signal or tag, indicating or generally correlated with high ability. Worse still, once we alter the model to allow for more variation between individuals, so that two people with the same ability might choose different earnings, and two with the same earnings might have different ability, earnings' value as a signal of ability is degraded somewhat.

(2) *Optimal income taxation's combination of conceptual radicalism with apparent support for only modestly addressing high-end inequality* – OIT's focus purely on the marginal utility derived from consumption and leisure, and its consequent complete dismissal of property rights and entitlement to the fruits of one's labor, other than on prudential grounds relating to incentive effects, is startlingly radical. It therefore stands at some distance, not just from various intuitions that its critics can identify, but also from other branches of contemporary economics, which sometimes may appear to emphasize incentives to the exclusion of everything else. Neoliberalism and the "Washington consensus" of the 1990s, blamed by many for encouraging the adoption of pro-market policies that exacerbated inequality based on the view that it just did

[18] In general, the OIT literature finds that optimal tax rates, and the resulting size of the redistributive cash grant, are negatively correlated with people's labor supply elasticity. The higher this is, the greater the efficiency costs of a given tax rate. By contrast, optimal tax rates (and the resulting size of the grant) are positively correlated with the slope of declining marginal utility as the representative consumer's budget line rises. Tax rates and grant levels are also positively correlated with the degree of dispersion in ability, since the payoff to redistribution is greater if people are further apart to begin with. And they are positively correlated with the degree (if any) of pro-egalitarian weighting in one's social welfare function.

not matter or else would naturally take care of itself, are logically associated by many people with economists' rising political and intellectual influence in recent decades.

OIT potentially leans the other way. Yet in practice, even insofar as it had any significant influence on either beliefs or political outcomes, this does not appear to have been its main effect, at least until recently. Instead, if anything OIT appears to have encouraged the trend in recent decades towards adopting lower and flatter tax rates that may have contributed to the rise of high-end inequality.

This apparent political effect (unless there was no significant effect) reflects a key finding of the OIT literature, accepted for many decades although more recently challenged, to the effect that tax rates should be relatively flat (Slemrod 1990, 165; Mankiw, Weinzierl, and Yagan 2009). The reasoning that supports flattish rates is akin to, but less obvious than, that for generally supporting lower rates by reason of taxation's adverse incentive effects.

Suppose that, in the absence of incentive effects, we would agree under OIT reasoning that income above the mean – or equivalently *all* income, if uniform cash grants were used to get to the same place – should face a 100 percent rate. Why might concern about incentive effects support, not just lower rates, but also relatively flat rates?

The reasoning goes as follows. Suppose we are asking what tax rate should apply at $30,000 of income, which is roughly the 50th percentile in the United States, as opposed to at $150,000, which is roughly the 95th percentile. If incentive effects have exactly the same import at both levels, then shouldn't the marginal rate at $150,000 be much higher than at $30,000, given the assumption of declining marginal utility? While this argument might seem to make sense on its face, the problem is that incentive effects may not have the same overall import at both levels.

To illustrate the possible difference, suppose that people who are earning at least $40,000, placing them in the 60th percentile or above, are in most cases certain to keep right on earning more than $30,000. That is, suppose they are not considering (and do not face the risk of) having their earnings decline by that much. If we were to raise the marginal tax rate on earnings in the range below where they sit – say, by 10 percent for earnings between $25,000 to $30,000 – this would raise $500 from each of those higher-income people, without having any effect on their marginal incentives.

It's not that they wouldn't notice. After all, each of them would be $500 poorer, disregarding what the government did with the money. But their marginal incentives, regarding how much they would get to keep out of the last dollar they earned or the next one they might earn, would be wholly unaffected by this tax increase. Accordingly, a marginal rate increase for earnings in the range from $25,000 to $30,000 would be "free money" in efficiency terms – though not, of course in distributional terms – so far as all of these people were concerned.[19]

Now suppose instead that we are considering raising the marginal tax rate at around $150,000 of income. Just as in the case where we do it at around $30,000, (1) some people are wholly unaffected because their earnings are certain to be lower, (2) others have their marginal incentives unaffected because their earnings are certain to be higher, and (3) still others are right at the range where it affects their marginal incentives. Only, because we are looking at the 90th percentile, rather than the 50th percentile, the members of Group 2 out of the above three, which provides "free money" in efficiency terms, is far smaller. Accordingly, the tax rate increase at $150,000 may be likely to provide far less "free money," relative to the marginal distortionary effects, than the tax rate increase at $30,000. This pushes against marginal rate graduation, just

[19] Indeed, the income effect of inducing them to work more, so that they could replace the lost $500, would have a positive spillover effect insofar as it generated additional tax revenues that they ignored when deciding how much to work.

as declining marginal utility pushes in favor of such graduation. The exact effect on a given OIT analysis depends on, among other things, exactly how people are distributed across the range of actual and anticipated potential earnings. However, the net result, in many OIT analyses, has been to suggest that the overall rate structure should be roughly, even if not precisely, flat.[20]

Recent years have witnessed a prominent challenge to this view Two prominent and indeed A-list economists – Nobelist Peter Diamond and likely future Nobelist Emmanuel Saez (2011) argue, purely on the basis of a methodologically standard OIT analysis, that U.S. marginal tax rates should be steeply graduated, and indeed should probably exceed 70 percent at the top of the income distribution. They base this on certain technical issues that I discuss at length in Shaviro 2016.[21] These include asserting that the adverse marginal utility effects at the top from paying even significantly more tax are for a long time so small that "as a first approximation [they] can be ignored" as effectively indistinguishable from zero.

Is this how people earning income at that level actually would feel if they ended up paying significantly more tax? One suspects not.[22] There also is something paradoxical about assuming that the reason we can take money from the super-rich is that they don't actually care more than minimally. If it were true that they didn't much care, it might be hard to explain the

---

[20] Lest this not already sound neoliberal enough, there is more. As Joel Slemrod (1990, 164) notes, "raising the marginal tax at the [very] top [to] above zero [percent] distorts the labor supply decision of the highest earner [to his or her detriment] but raises no revenue." Accordingly, it has long been an accepted consequence of OIT reasoning that, at least "precisely at the top," the marginal tax rate should actually be zero. Indeed, this conclusion is logically irrefutable if one is a strict welfarist who views utility as depending solely on own consumption, and who assumes that high-end inequality does not yield any negative externalities. After all, under these assumptions any positive rate at the very top of the distribution – assuming that the individual who sits there can fine-tune his work-versus-leisure tradeoff to the very last penny – would reduce the highest earner's utility without (since zero revenue is raised) having any positive effect on anyone else.

[21] **[NOTE TO READERS: Shaviro 2016 is the University of Miami Law Review version of this chapter.]**

[22] Forbes journalist Rich Karlgaard (2011) notes that even people with a net worth, say, of $100 million often travel in circles that may cause this wealth level to feel subjectively as if it is not quite enough. "$100 million is really a 'tweener number. You can be Richie Rich among your upper-middle-class friends or a hanger-on in the superrich crowd. Choose the latter and prepare to have your ego smashed. You'd surely have to give up the Aspen and Maui homes or replace the Citation X [all affordable, along with a New York home, at the $500 million level] with a fractional jet-ownership."

strength of existing political opposition to increasing high-end rates. Diamond and Saez ignore such evidence because, following standard practice, they define how much the super-rich would relevantly care about high taxes purely in terms of the marginal utility of the lost consumption.

Placing OIT in a broader context

Diamond and Saez are to be commended for enriching the economics literature by convincingly showing that OIT *might* support doing far more to address high-end inequality than had previously been widely assumed. From a broader perspective, however, their contribution offers a classic example of the Mapmaker's Dilemma at work. On the one hand, it made perfect sense for them to stay within the literature's standard parameters, where their expertise lies, rather than trying to identify everything under the sun that might be relevant – or even crucial – to assessing what policymakers should do with respect to high-end inequality.

On the other hand, if we as a society are trying to decide how to respond to rising high-end inequality, it would be foolish to feel bound by their self-imposed restrictions. A model in which *only* the marginal utility derived (by the super-rich and others) from own consumption is deemed relevant, and in which position-related utility is disregarded – even though, without it, the intensity of today's political wars over taxing the rich would make no sense – cannot reasonably be viewed as offering the final word. One needs to supplement it with other inputs, both empirical and normative.

The idea that own consumption is all that really matters, in assessing high-end inequality, would be a lot more compelling if each individual lived on a separate planet, consuming available resources but neither seeing nor interacting socially with anyone else. In that type of a science fiction scenario, a benevolent social planner might have little to think about, beyond the question of where the resources that she could fit onto her spaceship (assuming it could hold

cargo, but not passengers) would create the greatest amount of happiness. But for human beings living in densely packed and heterogeneous societies on the planet Earth, this is an amazingly blinkered, reductive, and incomplete way of defining the relevant considerations.

Suppose that adding position-related utility to the analysis would indeed support doing more to reduce high-end inequality than seems appropriate in the separate-planets scenario, where people only care about utility from own consumption. Then favoring weighted welfarism, in lieu of utilitarianism, might be viewed as a very rough proxy for all that has been left out.[23] After all, if one lucked out in deciding just how much extra weighting at the bottom to apply, one might succeed in approximating the end point that one would have reached by counting everything that matters, without differential weighting.

Yet this cannot fully compensate for directly taking proper account of everything that matters. For example, whereas weighted welfarism, in a framework based purely on utility from own consumption, would never counsel adopting high-end tax rates that lie above the peak of the Laffer curve,[24] high negative externalities might get one there. In addition, using differential weighting, in lieu of directly considering everything that matters, tends to focus analysis and debate on the wrong variable. If we care about position-related utility, surely we ought to think about it directly, rather than employing an arbitrary weighting convention instead.

OIT's incompleteness in addressing everything that matters about high-end inequality makes clear the need for other inputs to one's analysis. Some of these may lie within the reach

---

[23] Leaving any such considerations aside, there is something paradoxical about assigning extra weight, within a welfarist social welfare function, to utility enjoyed by the worst-off individuals. In effect, this treats utility as if it could itself be subject to declining marginal utility. Rawls (1971) famously argued that applying infinite risk aversion from behind the veil might support giving absolute priority to the relevant interests of the worst-off individual. But a rational choice framework would more logically support maximizing expected utility from behind the veil, and applying risk aversion only as an application of this metric (for items that have declining marginal utility). See Harsanyi (1953; 1955): Kaplow (2008, 370-372),

[24] One might, however, in a standard OIT framework where one has incomplete information about potential revenue yields, adopt high-end tax rates that have a chance of being too high.

of hard social science literatures outside public economics – for example, in political science studies regarding policymakers' responsiveness to the interests and concerns of the bottom 99.9 percent. Yet one may also need to consider "soft" information that sheds further light on how societies with greater or lesser degrees of high-end inequality might feel on the ground – as well as on one's own underlying moral sentiments. Hence, the case for studying literature, as one among many of the potential soft inputs.