What Public Policy Can Be: An Interview with Matthew Adler

MATTHEW ADLER (Philadelphia, 1962) is the Richard A. Horvitz Professor of Law and Professor of Economics, Philosophy, and Public Policy at Duke University since 2012. Before coming to Duke, he was the Leon Meltzer Professor of Law (2006) at University of Pennsylvania Law School, where he had taught since 1995. Adler received a B.A. in History from Yale University in 1984, was awarded a M.Litt. in Modern History from Oxford University in 1987, and received a J.D. from Yale Law School in 1991. Adler had stints as a management consultant at Oliver, Wyman & Co, in New York (1986–1988), and as an associate at the litigation department of Paul, Weiss, Rifkind, Wharton & Garrison, also in New York (1994). Adler served as law clerk for Judge Harry Edwards, U.S. Court of Appeals for the D.C. Circuit (1991–1992), and for Justice Sandra Day O'Connor, U.S. Supreme Court (1992–1993).

Adler's research traverses welfare economics, normative ethics, and legal theory. He is the author of *New Foundations of Cost-Benefit Analysis* (2006; co-authored with Eric Posner), *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis* (2012), *Measuring Social Welfare: An Introduction* (2019), and is currently working on the manuscript *Risk, Death and Well-Being: The Ethical Foundations of Fatality Risk Regulation*. Adler is the founding director of the Duke Center for Law, Economics, and Public Policy, and co-founder (along with Ole Norheim) of the Prioritarianism in Practice Research Network.¹ He served as co-editor of *Legal Theory* between 2008–2017 and *Economics and Philosophy* between 2017–2022.

The *Erasmus Journal for Philosophy and Economics* (*EJPE*) interviewed Adler about his formative years (section I); his work on the theoretical foundations of public policy, zooming in on welfare-consequentialism and social welfare functions (section II), welfarism and interpersonal comparisons (section III), the ethical deliberator and the role of the philosopher (section IV); and, finally, his views and visions for interdisciplinary

¹ An edited volume based on the work of this network has just appeared. See Adler and Norheim (2022).

EJPE's Note: This interview was conducted by Måns Abrahamson and Akshath Jitendranath, editors at the *EJPE*.

Footnotes throughout the interview are editorial annotations directing readers to the relevant literature discussed in the text.

work in law, economics, and philosophy, as well as his advice for graduate students in the field (section V).

I. FORMATIVE YEARS

EJPE: Professor Adler, perhaps you can start by saying something about the trajectory of your career as a scholar of Law and Economics and the events, people, or writings that have had a particular influence on the development of your interests.

MATTHEW ADLER: I was actually a history major undergrad and then ended up at Yale Law School, where I graduated in 1991. Yale Law is a great place in terms of Law and Economics. For example, Guido Calabresi taught there for a very long time. I ended up taking more classes in constitutional theory than in Law and Economics—Yale Law is also very important in constitutional theory. So, although I had some exposure to Law and Economics at Yale Law School, my serious interest didn't really blossom until after I started law teaching at University of Pennsylvania Law School. And I was lucky to have Eric Posner as a colleague there at the time, who had started a year before me. We started thinking about costbenefit analysis, had conversations, and then wrote an article that came out in 1999 and that was the basis for the 2006 book, *New Foundations of Cost-Benefit Analysis*, where we challenged the traditional Kaldor-Hicks view of the justification for cost-benefit analysis.² So that was a formative series of events.

Then I shifted from cost-benefit analysis to social welfare functions because it seemed to me that we could do even better than cost-benefit analysis in terms of tracking social welfare. And at that point I was lucky to have another colleague, Chris Sanchirico, who taught me a lot. He's a PhD economist and had been on the economics faculty at Columbia University. He had a big influence in terms of my understanding of Law and Economics. And since then I have been lucky to collaborate with a lot of people. For example, I encountered Marc Fleurbaey, I think, maybe in 2008 or 2009, when he came to a conference at UPenn and I have since worked with him on various projects.³

-

² See Adler and Posner (1999, 2006).

³ Adler and Fleurbaey co-edited *The Oxford Handbook of Well-Being and Public Policy* (2016); other joint work includes Adler and Fleurbaey (2018a, 2018b), Adler et al. (2020), and Adler et al. (2023).

The other thing I should say is that I started teaching at the University of Pennsylvania Law School in '95, which at that time had (and still does have) a very strong legal philosophy group. Not a philosophy of economics group, but a legal philosophy group. And a lot of my interest in philosophy—and knowledge, really—started with conversations with those colleagues.

As a follow-up question for clarification: Was the challenge to the Kaldor-Hicks criterion that it is insufficient as a moral criterion to justify cost-benefit analysis?

Exactly. In the United States there's a weird disjunction between the view of cost-benefit analysis in applied and in theoretical welfare economics. Kaldor and Hicks proposed their notions in the 1930s and the 1940s and there is this immediate reaction where a lot of problems were discovered. For example, problems of intransitivity—Scitovsky talks about that.⁵ Indeed, one of Arrow's motivations in his book was thinking about Kaldor-Hicks efficiency. But at the level of applied economics and in terms of governmental practice, Kaldor-Hicks efficiency is seen as the touchstone. And also, I think—at least in the U.S.—in Law and Economics, where it is associated with the intuitive notion of making the pie bigger. One often hears the expression, 'just make the pie bigger, why isn't that a good thing?'. Eric Posner and I were driven to respond to that. Also, part of what motivates Kaldor-Hicks efficiency is the notion that one cannot make interpersonal comparisons. So, Eric and I did a double move. One was to say, 'of course, we can make interpersonal comparisons'; and the second was to think of cost-benefit as a rough proxy for overall well-being and to argue that justifying cost-benefit analysis by way of overall wellbeing is much more persuasive—at least if overall well-being has some kind of intrinsic normative weight. Kaldor-Hicks efficiency doesn't.

-

⁴ A Kaldor-Hicks improvement, named for Nicholas Kaldor and John Hicks, is an economic reallocation of resources such that those that are made better off could hypothetically compensate those that are made worse off, so that a Pareto improving outcome *could* be achieved. A situation is said to be Kaldor-Hicks efficient, or equivalently, is said to satisfy the Kaldor-Hicks criterion, if no potential Kaldor-Hicks improvement from that situation exists. See Kaldor (1939) and Hicks (1939).

⁵ What Tibor Scitovsky demonstrated was this. It is possible that if an allocation x is deemed to be superior to another allocation y by the Kaldor-Hicks criterion, then by a repeated application of the same criteria, we can also show that y is superior to x. See Scitovsky (1941).

⁶ See Arrow (1951).

Thanks, we'll come back to this later. But let's start by zooming in on some of the parts of your trajectory that you described before. Starting from the early days, what made you choose history?

I was lucky to go to Yale college, which is a wonderful place. And one of the great things about the undergraduate major there is that you can major in a subject without having that be your exclusive focus. I actually went in thinking that I'd be pre-med and afterwards go to medical school. Then I was in this wonderful, intensive freshman humanities program called Directed Studies, which is one of the famous things at Yale—in six classes over the course of the freshman year, you read great books in literature, political philosophy, and philosophy. And after that I realized that I really wanted to spend a lot of time taking classes in the humanities. Although I also took a lot of economics, I took a fair bit of math, I took some science. But Yale was very, very strong in the humanities, including history—they had a lot of great professors in history.

Also, in terms of my background, I'm Jewish. My family, grandparents, great-grandparents, left Europe because of antisemitism in the early 20th century; that led me to be interested in studying German history. So, I ended up majoring in history. I then received this scholarship to go to Oxford to get a Master's in history, which I enjoyed very much. But I realized that I didn't really want to go on to do a PhD in history—that I had more theoretical interests. I feel lucky that I did history, but also that I had all these classes in math in secondary school, and then in college—along with history—economics, math, philosophy, and so forth, that I could then use when I pivoted to more theoretical pursuits.

Do you remember any specific book or set of books that left a deep impression on you from your undergraduate days?

That's interesting. In Directed Studies we read a lot of different things. We read *The Iliad*, for example, and *The Odyssey*, which are incredible. We read Kant (in translation, of course) and we read Hume, and the funny thing is that at the time I was more of a Kantian. I had a good friend and we had debates: he was a Humean, and now I've ended up being much more of a Humean than a Kantian. But I was exposed to both Kant and Hume.

After you did your history studies, you went on to work as a management consultant for a while. How did that come about?

I finished my Master's studies at Oxford in 1986. It was a great time to be there, wonderful place, and the dollar was very strong so one could live well as a student. There were a lot of Americans at Oxford who were there via various scholarships and what have you, but then, when that was coming to an end, we were trying to decide what to do. In 1986 Wall Street was booming and, I mean, that sort of seemed interesting—maybe I should make a little bit of money. So, I worked at Oliver & Wyman—which is now a huge firm, but when I worked there it was quite small—and I enjoyed that, but I realized after two years that I wanted to do something more academic. I was fortunate at that point to be admitted to Yale Law School. I went to law school thinking that if things worked out, I could become a law professor and if not, law seemed pretty interesting and there are obviously a lot of things you can do with a law degree. So, it was a short stint in the private sector. (And, by the way, some of the skills that I developed in the private sector still prove useful. Chapter five in *Meas*uring Social Welfare, where I have this sort of simulation model of risk regulation, is based on an Excel spreadsheet. I learned to do Excel as a management consultant.) In a nutshell, I went to law school with an eye to getting into the legal academy and thinking that law would engage my theoretical interests more than a history PhD would do.

You also served as a law clerk in the U.S. Court of Appeals for Judge Harry Edwards and later in the U.S. Supreme Court for Justice Sandra Day O'Connor. What can you tell us about these experiences? Have they shaped you in any particular way?

They were really wonderful experiences. I was very fortunate to get those jobs. What's interesting is that, as a law clerk you are really working for the judge. For example, each of the nine Supreme Court justices has four clerks. You are working for the justice. You are hired by the justice. The four clerks work in chambers. And it's a very small unit: there is the justice, the four clerks, and maybe a couple of assistants, but that's chambers. You also interact with other clerks. The same is true on the Court of Appeals. At that point, maybe there were 10 or 11 judges, each had three clerks. So, it's very intimate, and intense—huge amounts of work.

The main job of the clerks was twofold. First, before the case was decided, you as a clerk would write a memo for the judge analyzing the case. And that was supposed to be your impartial take on it. Neither the judge nor the justice ever said to me, 'this is the way you should analyze it'. It was rather, 'here are the briefs, here is the opinion, write a 10- or 15-page

memo analyzing the issues here'. You'd then discuss that with the judge. Then (and this was the second main part of the job) the judge might well end up writing an opinion in the case and you were told to prepare a first draft of the opinion. And the opinion would then be circulated to other chambers. It was a very intellectual experience. Which in a way was kind of inspiring. Of course, judges have political views, and there might well have been some politicking behind the scenes among the judges, but fundamentally that was not what the clerk's job was about. The job was doing legal analysis.

There were many issues on which the analysis was helpful. That is, after drilling down on the questions in the case you'd realize that one answer was the legally right answer, that one answer was the more persuasive answer. Now, and this is really at the Supreme Court level, there are cases where the court is simply not bound by prior law, where the court in effect is making law—although the court doesn't say that. But even at the Supreme Court level, and certainly at the Court of Appeals level, there are lots of cases where, again, this exercise of legal analysis could point you to a right answer, a better answer. It's a good question, which I haven't really thought about, but I think that the experience as a law clerk of finding right answers through legal research has in turn helped to motivate and inspire me with respect to my scholarship.

Would you say that was a way of going to philosophy graduate school? Because you are reading a lot of texts and there is a pressure to read these texts quickly and analyze them in a very impartial way. That mirrors a lot of graduate training in analytic philosophy at least.

That's interesting. It's a good question. Being a law student and then law clerk are highly analytic. And actually, the thing I like most about teaching is the analytic part. People say, 'oh, you teach law school, you teach legal doctrine, that's boring'. But for me the best thing about the classroom experience of law teaching is that it's analysis. It's not memorizing the rules. What it is, is: 'Okay, so this is the case, this is the issue, should the court decide the issue this way or that way? How will we analyze that using these legal materials?'. That's analytic reasoning. And the same is true, obviously, as a clerk. The difference between judging and philosophizing is that judges are constrained by the legally authoritative texts. That's it. You're never going to say: 'Well, why do we have this constitution? Maybe we should change the constitution'. You cannot. There's a limit to how deep you can go. But philosophers are bound by nothing.

I've already mentioned this, but some of my research development is just the serendipity of my academic career. I was fortunate to get a teaching job in '95 after a brief time in practice. And I just happened to have colleagues who were really good philosophers, I mean, *really* good philosophers. And so, I experienced arguing with them, about bedrock stuff. For example: we all have these deontological intuitions, 'you shouldn't kill someone, you shouldn't kill someone even to save five lives'. But to really drill down on deontology, and to make sense of its content, scope, and justification, is not easy. The people I was working with at the University of Pennsylvania Law School at the time, people like Michael Moore, Heidi Hurd, and Steven Perry, were some of the greatest deontologists and legal theorists. I ended up not going in the deontological direction, but the depth of their involvement and the seriousness of the analysis were really inspiring to me.

Talking about moving from working as a clerk, to coming back to academia. In a more light-hearted way, in Measuring Social Welfare, you use the following example to drive home the point that preference satisfaction accounts of well-being can be problematic if people's preferences are based on poor information:

Felix, remembering all the courtroom dramas he has watched, prefers to be a litigator over alternative careers. If he were better informed, he would realize that much of a litigator's time is spent reading documents or writing briefs—which doesn't interest Felix—and not making dramatic closing arguments. Felix's actual preference for the litigator career diverges from what he would prefer, careerwise, if he knew more. (Adler 2019, 48)

You, yourself, worked as an associate at the litigation department of Paul, Weiss, Rifkind, Wharton & Garrison (New York) in 1994. Is this example by any chance autobiographical?

That's funny. Maybe a little bit. Paul Weiss is a great firm, a great litigation firm. Not just commercial practice, but also pro bono practice. For example, much of the litigation against the Trump 'Muslim ban' was undertaken by Paul Weiss. So, it's really a fantastic firm. But yes, it became clear to me pretty quickly that I didn't want to be a litigator. After my law clerkships, we had moved back to New York from Washington D.C. because my wife was finishing her PhD in history at Columbia. I said to Paul Weiss, you know, I'm likely to go on the law teaching market and try to get an

academic job, I'm probably not going to stay here. And their response was, 'that's fine'.

But the story I like to tell about these large law firms—the term that is often used in the U.S. is 'Biglaw'—is this. In the Biglaw firms, typically the time it takes to become a partner is maybe eight years or ten years. So, the image I had, before I started at Paul Weiss, was: you have this group of associates, and everyone sits around for eight years or ten years, and at the end of that they pick a couple to be partner. That's not true. What happens is that you have people who are associates who, at these big firms, are typically very talented people. And most of them decide that they don't want to become Biglaw partners. That is, people end up going into the academy, going to work for government, going to work for a smaller firm, going to work in-house at a firm, getting out of law entirely. Because what it takes to become a partner at that kind of firm is not just intelligence, but also a tolerance for an incredible amount of hard work. And a willingness to have the work be driven by clients. Because at the end of the day, at that kind of firm, you have a client: you are litigating a matter because the client comes in and says, 'here is this problem I need you to solve'. So, it's those particular skills, motivations, and interests that it takes to become a partner. And so, what in fact happens is that most people say, 'that's great, but I'll stay here for a couple of years to get the training and then I'll leave'. And at the end of the eight or ten years there are only few associates remaining who still want to become partner, and one of them is chosen. Again, I was very happy during my year at Paul Weiss. I think, however, that I would not have been a good litigator. In part because I really value intellectual freedom. I mean, one of the great things about being an academic is control of your own time. I just value the ability to decide that today I'm just going to read this book and think about this. I was very fortunate to get the job at UPenn Law School, I was there for 18 years. That worked out very well and then I moved to Duke.

You already mentioned that there was this idea to pursue a career in legal academia, but the courses you started out teaching were in Administrative Law, Food and Drug Law, and Theories of the Administrative State. This, at least for us, sitting in Europe or India, or wherever else, seems like a world apart from the work you're known for these days, at least among the readers of this journal. What led you to start working on cost-benefit analysis, in particular, and the foundations of

public policy, more generally? You've already spoken about the influence of Eric Posner in this regard. Is there anything more that you want to say about this?

One of the things I should say is that the U.S. Court of Appeals for the D.C. Circuit is the federal appellate court that deals with a huge volume of administrative law issues. Administrative law in the U.S. is the law governing administrative agencies—for example, the Environmental Protection Agency, the Food and Drug Administration, the Department of Transportation, and so forth. A lot of the cases involving challenges to rules or decisions by those agencies go to the D.C. Circuit. As a law clerk on the D.C. Circuit, you will have a year during which a lot of the work will be on administrative law. You will come out of that either loving administrative law or hating administrative law. For whatever reason, I came out of that clerkship just loving administrative law. The issues are often somewhat technical, they're quite detailed, but I found that really interesting.

I actually thought that what I wanted to do as a law professor, was to write a treatise on administrative law. I started law teaching with that in mind because, back in the 1950s, '60s, and '70s, there were these great books about administrative law—I mean really, really good books about the doctrine. And then, of course, I got involved in other things. But that's why I started teaching administrative law, because of my experience as a clerk on the D.C. Circuit. And as a new professor your Dean says, 'you're going to teach this'. So I was like, 'yeah, sure, I'll teach administrative law, teach food and drug law, theories of the administrative state, this is all about administrative law, this is where I start'.

Cost-benefit analysis is very important to U.S. administrative law. Since the Reagan administration, major administrative agencies, by presidential order, have been required to engage in cost-benefit analysis. Basically, what the executive order says is that all regulations, unless the stature requires a different approach, should satisfy a cost-benefit test. And, in addition, for large regulations, agencies have to produce extensive cost-benefit documents that are then reviewed by OIRA (the Office of Information and Regulatory Affairs), a cost-benefit oversight office. So, it's really because I was teaching administrative law that I got interested in cost-benefit, which in turn led to writing with Eric about that, which in turn led to... It was a sort of winding path. But every step made sense at the time. It's just that when you look back you may think, 'how on earth did you end up moving from administrative law and food and drug law, to writing about measuring social welfare?'.

This leads to our next question. Like you said, you have written extensively about cost-benefit analysis and also about the foundations of public policy. So, what field would you call your academic home? Would you say you're doing Law and Economics? Or Philosophy of Economics? Or maybe PPE? Or something else entirely?

In a way I'd say Law, Economics, and Philosophy. Something I have been persuaded by throughout is that there are huge returns to engaging philosophically with the foundations of economics. And one of the things that traditional Law and Economics doesn't do enough, and certainly economists don't do enough, is to engage with the philosophy. So, I think philosophy is going to have to be in there, which is why I wouldn't just say 'Law and Economics'. I wouldn't say philosophy of economics, because that in a way takes economics to be primary. I mean, I think philosophy is equally important. Let me say that my interests are normative. At the end of the day, my aim is to try to figure out what to do, to provide normative guidance. And I think Philosophy, Economics, and Law are on a par with respect to that aim. Philosophy helps us in drilling down on these basic questions, such as, 'should we be consequentialists? Non-consequentialists? What's welfare?' And so forth. Economics—welfare economics—is hugely significant in providing all of these formal results and concepts that are very important for my purposes. One could also say that it's Philosophy, Economics, and Policy. Concerning PPE, I haven't done much on political theory, so I wouldn't use that label. So, let's say Law, Economics, and Philosophy.

That's interesting. You said, your interest is in practical reasoning, 'what ought we to do?', either at an individual level or socially. That's the motivating question. And you go to whichever discipline is relevant to answer that question. The thing is, some of these disciplines require a kind of training in formal reasoning, not just to contribute to that literature, but even to understand that literature. But your background was in history and law. Did you always have some 'formal chops' as part of your training, or is it something that you had to learn 'on the job' to do the kind of things you wanted to be doing?

At the end of the day, ethics is about what to do. And there are different relevant tools—there's conceptual reasoning, there's the process of arriving at a reflective equilibrium, which is philosophy. I was lucky to have a really, really good secondary school education. I went to a school in New York called the Fieldston School, which is an excellent school. And then I

had a really good undergraduate education at Yale. So, for example, in college, I took multivariable calculus, I took real analysis, I took mathematical economics, I took game theory. I also, in Directed Studies, did enough philosophy to have some background in that field. So when I decided that I wanted to be more engaged in philosophy, be more engaged in economics, I had enough tools.

One of the things I say to students is, 'take as much math as you can, take as much philosophy as you can, so then you have the tools'. You can read these literatures and then you can decide what to research. And, by the way, if you're interested in something and you feel like it's too hard for you to make a contribution, you can find a co-author.

This leads nicely to our next question. Throughout your career, you have co-authored both books and articles with scholars both within your discipline (law) but also other disciplines. Can you tell us a bit more about these experiences and the value you see in these types of projects? And, further, what might be some of the barriers to collaborate in and across disciplines?

I've been lucky to collaborate with interdisciplinary people. So, again, my first collaborator was Eric Posner, who is probably at this point one of the most famous law professors. He's been incredibly influential and is incredibly interdisciplinary. We were both law professors who were doing interdisciplinary work. There's Chris Sanchirico, who's a law professor but with a strong background in economics. I've worked extensively with Marc Fleurbaey. We co-edited the *Oxford Handbook of Well-being and Public Policy*, and Marc is a serious economist but one who also has many, many philosophical interests. And similarly, I've done a lot of work on applying social welfare functions to risk regulations with a group that includes Jim Hammitt, Nicolas Treich, and Maddalena Ferranna.

You have to find a co-author that you get along with. That you enjoy working with. Who brings something to the table that you don't have (and the converse also has to be true). But I also think it cannot just be strategic: you cannot say, 'I don't know this, I need this, so I'm collaborating with this person'. The collaboration has to really, really work. The most important thing I have found is, 'do the work'. A collaboration where you

⁷ See Adler and Posner (1999, 2000b, 2006, 2008), as well as the co-edited volume *Cost-Benefit Analysis: Legal, Economic, and Philosophical Perspectives* (2000a).

⁸ See Adler and Sanchirico (2006).

⁹ See Adler, Hammitt, and Treich (2014), Adler et al. (2020), Adler et al. (2021), Adler et al. (2023), and Ferranna, Hammitt, and Adler (2023).

expect that the other collaborators are going to do more than you is not going to work. All my collaborators have been equal partners, they have done huge amounts of work. I've seen these multi-author collaborations where one author thinks another author is going to do something, and vice versa, and it falls through. Everyone is going to have to be fully engaged. And there is going to have to be flexibility. You cannot insist that the article is going to have to be written in this way, or that this result is important as opposed to that result.

This recent book, *Prioritarianism in Practice*, which appeared in 2022, involved a whole team of people writing about the application of prioritarianism in different areas. 10 There is absolutely no way I could have done that myself. One of the things I would say about edited volumes more specifically—I have done a bunch of these edited volumes—is that the way to push along an edited volume is to have a conference and ask people to present the first draft of their chapter at the conference. 11 Contributors can delay writing a first draft indefinitely over email, but no one likes to show up at a conference empty-handed. I'm grateful to Duke University for funding the Center for Law, Economics, and Philosophy, which I've used to organize conferences.

Is there any way in which collaborating has influenced you?

Yes. For example, Eric is incredibly successful because whenever he writes, he wants to know what the bottom line is. He's had a huge influence across many, many different fields. So, for example, now there is this movement in the U.S. about anti-trust with respect to labor relations, breaking down the way in which firms have used monopoly power to limit the opportunities of workers. And Eric has written a seminal book about that.¹² The point is, he's very focused on, 'OK, so, what are we saying? What's the recommendation? Can we say this clearly without losing the forest for the trees?'. So that was very helpful to me. If I had written *New* Foundations of Cost-Benefit Analysis myself, the book would have been three times as long and much less influential.

Maybe you could say something more about your personal convictions. You said before that at UPenn you had colleagues who had more of a

¹⁰ See Adler and Norheim (2022).

¹¹ Edited volumes include, Adler and Posner (2000a), Adler and Himma (2009), Adler and Fleurbaey (2016), and Adler and Norheim (2022).

¹² See Posner (2021).

deontological bent of mind. But you stuck to a more consequentialist viewpoint. Is that a conviction you always held?

Jules Coleman—a great legal philosopher, whose legal theory class I took at Yale—said something to me once that really stuck in my head, which is: 'As scholars, we cannot answer the ultimate questions, but we *can* answer the penultimate questions'. If I write a book taking on the debate between consequentialism and deontology, I'm not going to make much of a contribution. It's not like I'm going to suddenly persuade deontologist to become consequentialists. That's just not how philosophy works. But what I can do is make progress within consequentialism. I find welfare-consequentialism persuasive. Is my credence one hundred percent? No. I find it to be persuasive, plausible, and worth pursuing.

Maybe this is something I learned as a management consultant, which is: as a management consultant, we always advise companies to pursue market opportunities. As a scholar, of course, if you write about something nobody cares about, nobody is going to read your work. But if you write about the thing that everybody else is writing about, you're also not going to make much of a contribution. So, it struck me that there was a lot of work to be done fleshing out welfare-consequentialism. Both at the level of theory, and at the level of application. And I find welfare-consequentialism to be quite plausible. So, that's what I've done.

II. WELFARE-CONSEQUENTIALISM AND SOCIAL WELFARE FUNCTIONS

Let us then turn to welfare-consequentialism, and more specifically the work you have engaged in, in recent years, since the publication of New Foundations of Cost-Benefit Analysis. This work is helpfully summarized in the content and title of your last book, Measuring Social Welfare. Is it right to say that this has been motivated by the need to find a better alternative to cost-benefit analysis as a basis for public policy evaluation? If this is correct, which you also hinted at in your previous answers, why do you take the view that cost-benefit analysis is an insufficient basis for public policy evaluation? And how have your views on this evolved over time?

Indeed, a lot of my work, starting in 2006, has been motivated by the need to improve on cost-benefit analysis through this notion of the social welfare function (SWF). I see myself as proceeding on three tracks. The first is filling in the philosophical details—for example, the nature of well-being, or the functional form of the social welfare function. The second is

applying the social welfare function in different policy domains, showing that this framework is actually useable. There is, of course, this literature in economics where SWFs are used: optimal tax theory. James Mirrlees wins the Nobel Prize for using social welfare functions to think about taxation.13 And, also, climate economics.14 But there are other fields where it's not used much. My view is that the social welfare function framework is an all-purpose policy tool. It's important to demonstrate that—to persuade people that it's actually useable, that it's not just a theoretical device. So, the second aspect of my work has been to do that. I've written about social welfare functions and both risk regulation and, to a lesser extent, climate economics. 15 And this Prioritarianism in Practice project was an attempt to have lots of people in different policy areas doing that. And then the third track has been to make the apparatus more accessible. Measuring Social Welfare was meant to provide an overview of social welfare functions that was sufficiently rigorous but more broadly accessible than Well-Being and Fair Distribution. As I already said, all three tracks are intended to get us to move beyond cost-benefit analysis—with cost-benefit analysis as the target because this is the dominant policy analysis approach in the U.S., where I live and teach, and in applied economics.

It seems to me that there are really three justifications for cost-benefit. There is the so-called neoclassical justification, which points to Kaldor-Hicks efficiency. That's problematic for the reasons I have alluded to, which go back to Amartya Sen. ¹⁶ Either the policy is actually Pareto superior, in which case just say that; or, merely potentially Pareto superior, in which case there is a potential—an unrealized potential—to convert the policy into an actual Pareto improvement, which doesn't seem to be much of a justification. The second defense of cost-benefit is the defense that Eric and I mounted in our work: cost-benefit analysis is a rough proxy for overall well-being. But why have a rough proxy? Can we not do better? Cost-benefit analysis uses money as the measuring rod for well-being. But money is an imperfect measuring rod given diminishing marginal utility.

.

¹³ James Mirrlees was, together with William Vickrey, awarded the 1996 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel (see Royal Swedish Academy of Sciences, n.d.). Mirrlees' work—building on the work of Frank Ramsey (1927)—laid the foundations of optimal tax theory (see Mirrlees 1971).

¹⁴ See Botzen and van den Bergh (2014) for an overview of the use of SWFs in climate economics.

¹⁵ On the application of the SWF approach to risk regulation, see Adler, Hammitt, and Treich (2014), Adler (2015), Adler et al. (2021), and Adler (2023); on the application of the SWF approach to climate change, see Adler and Treich (2015), Adler et al. (2017), and Adler (2018b).

¹⁶ See Sen (1979, 24-25).

The third justification for cost-benefit, which I think is probably the best, is the Kaplow notion. Kaplow actually accepts a social welfare function as the foundation for policy analysis but says that we should limit the direct use of social welfare functions to the tax system, and we should use costbenefit analysis for non-tax policies.¹⁷ I think that's not fully persuasive for various reasons, but again note that Kaplow and I agree that the foundation in terms of designing government and in terms of figuring out what policy techniques different institutions should use, is a social welfare function and not cost-benefit analysis.

So, to label stages as well. The social welfare function approach that you propose has three key components. First, there is a well-being measure, which translates each of the possible outcomes (a complete description of a world) of a policy choice into a list ('vector') of interpersonally comparable well-being numbers, which quantify how welloff each person in the fixed population would be in that outcome. Second, there is a rule for ranking alternative vectors, or lists of well-being numbers. Third, there is an uncertainty module, which orders policies understood as probability distributions across outcomes. So, a few questions of clarification about the framework, before we get to more substantive issues about this framework.

First, there are further sub-components to this approach which consist in the restrictions imposed on the three basic components (and we will get to some of these restrictions below). This has the implication of ruling out a broader class of social welfare functions. Those that do not allow interpersonal comparisons of well-being, for example; or, nonwelfarist social welfare functions, for example. What was the motivation for introducing and scrutinizing only those social welfare functions that conform to the constraints imposed on the distinct components? We ask because the social welfare function approach is more general than the class of social welfare functions discussed in the book. The social welfare function approach, as I see it, is based on the social choice, or the theoretical welfare economics, literature on so-called social welfare functionals—the SWFL approach—which originates with Amartya Sen. 18 There are various good review chapters, but, for example, Bossert and Weymark have a fantastic chapter. 19 The SWFL approach imagines

¹⁷ See Kaplow (1996, 2004, 2008).

¹⁸ See Sen (1970a).

¹⁹ See Bossert and Weymark (2004).

that we have a set of outcomes and a profile of utility functions for N individuals, and then we have a mapping from a given profile of utility functions to a ranking of the outcome set. The so-called welfarism theorem shows that if you adopt certain assumptions—in particular, the independence-of-irrelevant-utilities axiom—this mapping takes the form of a single ranking of utility vectors. The SWFL approach is somewhat narrower than the approach I adopt, because of the assumption that we should understand people's well-being in terms of utility. Each person has a utility function mapping the outcomes to utility numbers—a utility function that tracks her preferences.

What I do is that, instead of having a profile of utility functions, I have a well-being measure. This well-being measure converts an outcome into a vector of well-being numbers. In that sense it's more general than the SWFL approach. I *am* adopting the welfarism assumption. The most general formulation of my approach would be that we have a well-being measure, a set of outcomes, and a function that takes any given well-being measure to a ranking of a set of outcomes. But I'm assuming, as per the welfarism theorem, that that mapping will take the form of a single ranking of well-being vectors.

Now, one of the things that the SWFL literature does is to look at different possibilities for this ranking, including possibilities that don't make assumptions about interpersonal comparisons. For example, various kinds of dictatorships. A result that I cite in chapter two of *Measuring* Social Welfare is that, if well-being is not interpersonally comparable, we will end up with a dictatorship. In the book, I have this notion of 'admissible rescalings of the well-being measure', which is based upon the SWFL literature. Suppose that we have a theory of well-being, which could be a preference-based theory, a hedonic theory, an objective-good theory, or some other kind of well-being theory. That theory is going to tell us what the admissible well-being comparisons are. In particular, it's going to tell us if we can make, first, intrapersonal comparisons of well-being levels, differences, and ratios, and, second, whether we can make not just intrapersonal comparisons of that sort but also interpersonal comparisons. Corresponding to this specification of admissible well-being comparisons is going to be a set of admissible transformations of the well-being measure. So, in particular then, if we believe that the only admissible well-being comparisons are *intra*personal comparisons of levels and differences, then any individual-specific positive affine transformation of the well-being measure will be admissible. An individual-specific positive affine transformation is going to disrupt the interpersonal information; but if you think that such interpersonal information doesn't exist, this transformation is admissible because it will preserve all the *intra*personal information about well-being levels and differences. But then there is this theorem, which says that, if we require the social welfare function to be invariant to individual-specific positive affine transformations, and also satisfy the Strong Pareto principle, we are going to end up with a serial dictatorship.²⁰ Non-dictatorship is a weaker requirement than anonymity, which in turn I take to be fundamental to ethical reasoning. In short, if I want a social welfare function to satisfy at the very minimum non-dictatorship and more broadly anonymity, I'm going to require some interpersonal comparability.

The second question of clarification is, again, why the identification of the social welfare function approach with well-being. The object of inquiry in the book is well-being, but the social welfare function approach can also be deployed in the evaluation of other distinct moral concepts, like justice. Justice might be related to well-being, but it's not the same thing. We can also use the social welfare function approach to measure poverty, to measure inequality. What was the motivation to identify the social welfare function approach as an approach to evaluate well-being? Richard Arneson, who is a great philosopher, has an article, "Welfare Should Be the Currency of Justice". 21 I don't view a focus on well-being as being inconsistent with a focus on justice. There is a literature, going back to work by Jerry Cohen, Sen, Dworkin, Arneson, that asks: 'To the extent we are concerned about distribution'—and I'm concerned about distribution—'what should the currency be?'.22 Rawls says that the currency for distributive justice is primary goods.²³ Rawls talks about the principles of justice applying to primary goods. My view is that, to a first approximation, the currency for justice is well-being. But we can then think about what a just allocation of well-being would be. For example, we can say that the correct view regarding the just allocation of well-being is not utilitarianism, it's, let's say, egalitarianism, or leximin, or prioritarianism.

My view is that a concern for justice, or a fair distribution, and consequentialism are not mutually inconsistent. And that a concern for justice, or a fair distribution, and a concern about well-being are not inconsistent.

_

²⁰ See Adler (2019, 45-46).

²¹ See Arneson (2000).

²² See Cohen (1989), Sen (1980), and Dworkin (1981a, 1981b).

²³ See Rawls (1971).

Indeed, my 2012 book is titled *Well-Being and Fair Distribution*. What I was trying to get at with that title was that you can be concerned about fair distribution *and* about well-being. And the same with both poverty and inequality. In practice, poverty metrics and inequality metrics are often applied to income. But if you view income there as just a proxy for well-being, the view is going to be that, if you care about poverty and inequality, it would be better to apply the metrics to well-being.

A last question of clarification. What was the motivation to apply the social welfare function approach only to public policy? The field of Law and Economics has a long tradition of applying tools from economics to a variety of domains, like contract law for example. (Like the Coase theorem.) Do you think an avenue for future research is to apply the social welfare function approach to other domains of law as well? Absolutely. Again, this relates to the path-dependence of an academic career. So, again, I graduate law school; I work as a clerk for a judge (Judge Edwards, who is, by the way, a very great judge, he's still on the bench) where a lot of the work is administrative law; I start teaching administrative law (administrative law is public law, it's not private law, I have never taught a private law class); I become interested in the use of cost-benefit analysis as part of public law (that is, thinking about big regulations and so forth); so that leads to my interest in social welfare functions and public law. But there is no question that we can think about legal doctrines for an area of law—such as the doctrines of contract law, the doctrines of tort law, or the doctrines of property law—as objects that can be evaluated using a social welfare function. Would it maximize, for example, prioritarian social welfare, if, for example, the doctrines of tort law looked this way as opposed to that way? So, yes, it's a fruitful question and something, I think, that should be pursued.

Thank you. Young scholars reading: PhDs waiting to be written here! Stephen Hawking is supposed to have said: 'Someone told me that each equation I included in the book would halve the sales. I therefore resolved not to have any equations at all'. This is relevant here because one needs a certain competence with formal reasoning to engage with the social welfare function approach. Not least because informal reasoning can be quite treacherous in this domain of inquiry. At the same time Hawking is not wrong to suggest that the use of notation can come

in the way of the broader uptake of arguments that use formal reasoning. What are your views on the trade-offs involved in using formal reasoning, especially within domains where it is relevant to public policy? I really think it's a question of audience. Obviously, if you're writing for a public, even an educated public, don't use equations. Because even educated people may not be comfortable with equations. But the same will be true of philosophical language. If you're talking about public policy, don't use the term 'agent-relative side-constraint.' Members of the public don't know what an agent-relative side-constraint is. Don't use the term 'supererogation'. So, let's be clear: if you're writing a popular book, and I think that's a very important thing, even a book for an educated public, you are going to have to be very careful about using formal notation or about using technical nomenclature.

On the other hand, if you're writing for scholars, then you should use whatever linguistics tools and approaches, mathematical tools and approaches, match the inquiry. And here I would say, it's a trade-off. There are people who have very advanced mathematics and are very sophisticated as a conceptual matter. That's really difficult to do. So, obviously, at the high-end, someone like Amartya Sen. But most of us are not like that. So, there is a trade-off. The level of mathematics that you would use in a paper in *Econometrica* is going to be higher than the level of mathematics that you're going to use in, say, Economics & Philosophy. Economics & Philosophy—which I've been privileged to edit for a number of years does allow math, but we are not going to have the level of math that you might have in the most mathematical journals. On the other hand, the level of conceptualizing that would be tolerated, or demanded even, in a philosophy journal such as *Ethics* is not going to be acceptable in an economics journal. So, there is a trade-off. But I would say that, if we are talking about welfare economics, normative economics, and its twin in philosophy, which is moral philosophy, scholars in those fields should have some degree of comfort and openness to both conceptualizing and mathematics. Because it seems to me that both of those tools have been incredibly fruitful. And scholars who say, 'I'm just not going to do it—if there is math there, I'm not going to read it', or, conversely, 'if there is conceptualizing, I'm not going to engage in it', I think, are just making a mistake.

Now, the third question is, what if you're writing not just for other scholars, or for the public, but for students? That's in a way an exercise in textbook writing. In writing *Measuring Social Welfare*, I felt like I was

writing for students, and also, by the way, for law professors. You may have colleagues in other fields who are not deeply engaged in either economics or philosophy, but might be interested in those topics, and so, at that point, you try to lower the barriers of entry. But at the end of the day, it seems to me that a lot of the power of the social welfare function approach *is* the mathematics. At two levels: one is that a lot of the results, the most important results, depend on the mathematics—so, Arrow's theorem, or Harsanyi's aggregation theorem, for example, depend upon the mathematics. The other thing is that the tool itself is a mathematical tool—this is the point. The tool allows for systematic and tractable policy analysis by converting outcomes into vectors of well-being numbers; by converting the problem of policy situations under uncertainty into the problem of maximizing the expected value of some ethical utility function; and so forth. So to say, 'let's just drop mathematics', is just to say that we are not going to use the tool as it stands.

Now, to go a little deeper into the weeds. The social welfare function approach that you present is consequentialist, in the following sense: a policy P ought to be pursued only if the state of affairs x resulting from P is 'optimal' or 'best' with respect to a given rule. By that we mean that it's at least as good as any other outcome or state of affairs that could have been pursued instead. Now, some deontologists might take issue with the whole set-up—and that's as you said earlier unproductive as a debate to engage with—but our question here is internal to consequentialism: Why think that consequentialism requires optimization? Alternatively, you might respond by weakening the consequentialist claim to the following: a policy P ought to be pursued only if the state of affairs x resulting from P is 'maximal', by which we mean it's not strictly worse than any other state of affairs y that could have been pursued instead (with respect to a given rule). Then, our question is why wouldn't consequentialism require optimization or going for the best? In short, what sort of action-guidance is required, you think, of consequentialism?

That's interesting. I view consequentialism as the view which says that ethical guidance is ultimately grounded in a single—that is, an agent-neutral—ranking of possible worlds. And specifically, I'm a *welfare*-consequentialist to a first approximation. So, ethical guidance is grounded in an agent-neutral ranking of possible worlds that satisfies certain basic

axioms, in particular Pareto-indifference, strong Pareto, and anonymity, that express the core of welfarism.

By the way, what does the ranking involve? I believe that it's a well-behaved ranking, so I believe in transitivity, but I believe that the ranking takes the form of a quasiordering, or a preorder, which allows for incompleteness—there could well be incompleteness in the ranking. Then we have the question of 'what to do?'. What action should we choose in light of the ranking? In terms of my work, that's what I call the 'uncertainty module'.

There are lots of further assumptions one might make to make consequentialism tractable, but that I don't think are essential to the approach. Number one is assuming that well-being is measurable. Actually, in Well-being and Fair Distribution, I allow for the fact that there are going to be pockets of incomparability in terms of well-being itself. Let me put it this way: we have an agent-neutral ranking of worlds that's grounded in well-being. Well-being itself, or the well-being account itself, is captured in a ranking of what I call 'histories', where histories are pairings of individuals and outcomes—such as, 'being Matt in this world'. And also a ranking of pairs of histories, so as to capture differences. Those rankings too—the rankings of histories and history pairs—are quasiorderings. So, there can be incompleteness at this foundational level with respect to well-being. The assumptions that the well-being rankings are complete and measurable by a well-being measure, and that the social welfare function takes the form of a complete ranking of well-being vectors, I mean, these are incredibly important for tractability, but I don't view them as being part of the bedrock of consequentialism.

Now, finally coming in a long-winded way to your question: there is this further question of 'what should we do?'. Do we always optimize? Some people believe we should satisfice. I actually think the most important thing is to have the rankings. We want to go from this ranking of outcomes to a ranking of choices. And if there is a best choice, it seems to me—I'm not a satisficer—we should choose the best choice. But the most important thing is to have the ranking. And, by the way, if the set of choices is infinite and not compact, there may not be a best choice—in that case it's possible to have an infinite sequence of choices, each better than the one before. It seems to me that fundamentally what consequentialism is about is arriving at a ranking of choices, which might have incomparabilities, from this welfare-based ranking of outcomes, which also might have incomparabilities.

So, the idea is that there is a ranking. The ranking is what constitutes the basis of the choice. Other than that you have no fixed views on action guidance: if you want to optimize (and if you can optimize), do it; if you want to maximize, maximize; if you want to satisfice, satisfice (and so on)?

This is very useful for me, because it summons me to be very precise as to what my views are. It seems to me, the most important thing is to develop this welfarist ranking of worlds and to use that to arrive at a ranking of choices, for any choice situation that an agent might confront. That's the guidance. The guidance is the ranking. And then on top of that, it seems to me that, if given the ranking and the choice situation there is a best choice, then you should choose it. But there might not be. And given incomparabilities it may be that the best you do is to maximize as opposed to optimize.

The previous question is of course whether or not consequentialism requires completeness, as you brought up.

I think not. Joseph Raz's work has pointed to allowing incompleteness.²⁴

A related question—that has implications for some of the avenues for future research that you specify in the last chapter of Measuring Social Welfare—is this: Does consequentialism require transitivity?

I have always worked with an assumption of transitivity. John Broome's view is that transitivity just falls out of the logic of betterness: if you really understand what it means to be better or worse, that means transitivity. In effect, a conceptual argument. I'm not sure I buy that. It seems to me that asymmetry is part of the logic of betterness. If I say, 'this cup of coffee is better than this glass of water, but this glass of water is also better than this cup of coffee', I don't understand betterness. And, similarly, it seems to me that symmetry is part of the logic of equal-goodness. If I say, 'this cup of coffee is equally good as this glass of water, but, by the way, this glass of water is not equally good as this cup of coffee', then I don't understand equal goodness. I'm not sure I'd say the same thing about transitivity.

-

²⁴ See Raz (1986, 325-335).

²⁵ Broome (2004) holds that "transitivity is a feature of the logic of 'more ... than'" (50). Indeed, "that comparative relations are transitive is self-evidence" and, consequently, "it does not itself need to be supported by argument, and not much argument is available to support it directly" (51; see also 1991, 11–12).

On the other hand—and this relates to the point that I see ethics as giving guidance—transitivity is just incredibly useful in terms of its inferential properties. Imagine that I have this cup of coffee, this glass of water, and this bowl of ice cream; if I know the cup of coffee is better than the glass of water, and the glass of water is better than the bowl of ice cream, I can then conclude given transitivity that the cup of coffee has got to be better than the bowl of ice cream. I don't need to make a direct comparison. If you give up transitivity, then for every pair of items, to figure out whether one is better-than or equally-good-as the other, you need to make a direct comparison of the two, which is just cognitively overwhelming.

I've have never adopted a view that ethical criteria are simply criteria of rightness. It seems to me that the main reason, or a major reason, we care about the betterness ranking of worlds is because we want to use it to give guidance to us. And, by the way, we are cognitively limited, so something that is very useful is to have the guidance take the form of a transitive ranking.

The other point is that, some of this is simply the economist's approach of looking under the streetlight because that's where the light is: there are just a huge number of results that follow from transitivity. We are just not able to say very much without transitivity. I guess if there were overwhelming intuitively persuasive examples that involve intransitivity, then I might be motivated to give up transitivity. But because I'm not aware of such examples, and I think the deliberational—that is, the tractability—benefits of transitivity are massive, I'm going to, in line with most researchers (Larry Temkin has this research agenda to think about consequentialism giving up transitivity), stick with transitivity.²⁶

III. WELFARISM AND INTERPERSONAL COMPARISONS

The social welfare approach that you present is not just consequentialist but also welfarist. This basically means that it considers only vectors (or lists) of utility numbers that persons realize under different policies, or institutions, or systems. A prominent concern articulated by authors like Amartya Sen and John Roemer is that how that vector was gener-

_

²⁶ See Temkin (2012).

ated is of no independent moral concern under welfarism, but it is (intuitively at least) of import in assessing well-being. How would you respond to that general criticism of welfarism?

All I would say here is that, in both *Well-Being and Fair Distribution* and in *Measuring Social Welfare*, the last chapters in those books have said, 'until this point in the book, I've assumed welfarism, but actually we should move beyond welfarism by taking account of responsibility, by doing what John Roemer has done'.²⁷ Richard Arneson has been doing this for a long time, in fact Roemer was inspired by Arneson's work. Arneson says that the currency of justice is strictly speaking not welfare, but opportunity for welfare.²⁸ I think that's right. I think the social welfare function approach should be refined to take account of opportunity for welfare in a way that Roemer has done. But there are also a lot of difficult questions that arise in trying to do this.

Would it be correct to say that you see, for example, Sen's work or Roemer's work as within the broad class of the approach you're advocating? You don't see it as antagonistic to the social welfare function approach—it's a friendly refinement.

On the question at hand, Sen has written about the possibility of a non-welfarist consequentialism, for example, a consequentialism of rights.²⁹ That I view as being pretty removed from what I'm doing. I think that the next refinement for the social welfare function approach is to do what Roemer has done, which is to focus not on well-being, but well-being adjusted for considerations of responsibility, opportunity, or desert. That's motivated by specific intuitions about justice, that there's a difference between people who end up badly off through brute luck as opposed to option luck. As for a consequentialism of rights, at that point I would object, 'why should I care about rights as such, I care about welfare'; moreover, the distinction between a consequentialism of rights and deontology, although still there, becomes much thinner.

While the social welfare function approach is compatible with many different views of well-being, like experientialist or objective-list accounts, the view that you focus on in your work is, what can be called,

_

 $^{^{\}it 27}$ See Roemer (1996, chap. 8). For more on Roemer, see his interview in $\it EJPE$ (Roemer 2020).

²⁸ See Arneson (1989).

²⁹ See Sen (1982).

a 'laundered' preference-based view of well-being.³⁰ Would it be fair to say that that is your preferred conception of well-being, either in terms of what constitutes well-being, or—more in line with what Hausman would call the 'evidential view'—that suitably laundered preferences are reliable indicators of people's well-being even though they do not necessarily constitute well-being?³¹

My view on this has actually evolved. First, let me say, there is an ambiguity in the literature on well-being as between preferences and desires. Desires are pro-attitudes towards state of affairs—'I desire something'. The desires are then either wholly satisfied or not. So, 'I want a vanilla ice cream cone at some point in time' and then, later, that desire is satisfied or not. As opposed to preferences, which are what Hausman calls 'total comparative evaluations'.³² It's a total evaluation of, in principle, whole possible worlds or whole lives. Of course, humans cannot think in terms of whole possible worlds or life histories, but at least there can be total comparative evaluations of models of those. That, of course, is the conception of preferences that economists use. So, economists think of preferences as being rankings of whole worlds, or lifetime bundles, or something like that. And then there are utility functions that representing preferences thus understood. So that's number one.

Number two is that I don't think of preference-satisfaction as either constituting or necessarily being a good indicator of well-being. There are some people, by the way, who think of desire-satisfaction as constituting well-being, which to me seems to be crazy. Obviously getting what you want is a good thing, but there are lots of other good things. Nor do I think of preferences or preference-satisfaction as constituting well-being. But I also don't agree with Hausman that they are just good evidence of well-being.³³ I would take a third approach, which is that, to the extent

³⁰ The notion of 'laundering' preferences was coined by Goodin (1986). In broad strokes, the core idea behind preference laundering is that only preferences that satisfy certain conditions should 'count' in (social) welfare analyses. For example, preferences that are formed in light of full information, are based on rational deliberation, and are self-interested. On reasons for laundering preferences and conditions on welfare-relevant preferences, see Adler and Posner (1999; 2000b; 2006, chaps. 2, 5) and Adler (2012, chap. 3; 2019, chap 2).

³¹ See Hausman (2012, chap. 8).

³² See Hausman (2012, chap 1). According to Hausman, preferences (as they are for the most part, and should be, understood in economics) are judgments (thereby 'evaluative') ranking alternatives (thereby 'comparative') in terms of everything the agent regards as relevant (thereby 'total').

³³ See chapter 7 in Hausman (2012) on why the satisfaction of (even suitably laundered) preferences does not constitute well-being, and chapter 8 on why (suitably laundered) preferences can yet be considered reliable indicators of well-being.

that individuals are autonomous persons, we should defer to their preferences with respect to their well-being. So, in deciding whether one life-history or another would be good for someone's well-being, if that person is an autonomous individual, then up to some limit we should defer to her in deciding that. That's my view—that's the view that I have come to.

I teach at a law school, and legal scholars think a lot about deference the legal system is full of deference. For example, officials defer to courts about what the law means. And I think there is a similar structure with respect to well-being. We defer to people who can have preferences with respect to their well-being. Now, the tension is of course that people's actual preferences might be quite uninformed, irrational, formally problematic. So, I would look to people's fully-informed—or at least well-informed—and rational—in the sense of meeting certain rationality conditions—preferences. But, of course, there is a tension here with regards to the laundering. For me at least, the rationale for looking to someone's preferences is that she is authoritative with respect to her well-being, so it's tied to her autonomy: if we launder them too much, then the preferences have nothing to do with her; if we don't launder them at all, then her autonomy is not fully engaged. Again, my views on this have evolved over time: we want a suitable degree of laundering to capture the preferences she would have were she thinking, or deliberating, autonomously.

That's interesting. Is it then really about autonomy and less so about well-being? So, would you classify your view as an account of well-being, or is it about respecting autonomy?

No, it's an account of well-being. For example, I as a moral decision-maker care about someone's well-being. That is, I care about what is good for her, what is good for her life. But I recognize that the person is an autonomous person. She's autonomous, so I defer to her about her well-being. It's complicated, it's not autonomy as an input to well-being, nor is it a concern for autonomy apart from well-being, it's rather that I, in light of the fact that she is autonomous, am deferring to her about her well-being.

This line of thinking also—and I haven't written about this—helps resolve the question, 'what is a self-interested preference?'. This is a huge issue in the literature.³⁴ Preferences have to be restricted in some way. Satisfying my purely ethical preferences doesn't advance my well-being. But it is not clear how we restrict the preferences. One way to do that is

³⁴ See Adler (2012, 174–181) for a discussion of the importance of, and the difficult in, specifying 'self-interested' preferences when making welfare judgements.

to say that the restricted preferences are preferences in light of the person's judgements about well-being. The moral evaluator is trying to decide whether one life-history or another advances someone's well-being, the subject there also has a conception of well-being, and so the evaluator is going to defer to the subject's preferences that are motivated, or produced, by her judgements about well-being.

So, autonomy is the guiding criterion for thinking about what conditions we should impose on the preferences we use when we do welfare analysis?

Yes.

In the past you have talked about conditions of non-remoteness and being fully informed, fully rational. 35 Are these conditions you still subscribe to?

Yes. Full information and rationality because we want to engage a person's autonomy, we want the deliberation to be autonomous, which means thinking with good information, thinking rationally. Non-remoteness because there needs to be some self-interest restriction or screen. Some of my preferences are for things that are remote to me in the sense that satisfying those preferences is not going to advance my well-being. This goes back to Derek Parfit's so-called *Stranger on the Train* example, which at this point is pretty well-known in the literature on well-being.³⁶

In the broader literature on laundering preferences, there is on the one hand a discussion of conditions imposed on laundered preferences for reasons of well-being—full information, non-remoteness, rationality, as we have talked about—but then there are also authors in the literature who discuss imposing additional constraints based on other, non-welfare-related, reasons. Harsanyi, for example, writes that we shouldn't count people's anti-social preferences when we do social welfare analvsis because that would clash with the reason we are doing social wel-

³⁶ The example goes as follows: "Suppose that I meet a stranger who has what is believed to be a fatal disease. My sympathy is aroused, and I strongly want this stranger to be cured. Much later, when I have forgotten our meeting, the stranger is cured" (Parfit 1984, 494). Many would agree with Parfit that it seems implausible that the satisfaction of this preference (for the stranger's betterment) would improve my well-being.

³⁵ See, for example, Adler and Posner (2006, chap. 2) and Adler (2012, chap. 3).

fare analysis in the first place—"general goodwill and human sympathy";³⁷ Dworkin is worried about double-counting—we shouldn't use people's external preferences because then people's fundamental right to be treated as an equal might be disregarded in the social welfare analysis.³⁸ Are these types of non-welfare conditions something that you would apply to a plausible account of laundered or autonomous preferences?

It's a good question and I do not necessarily have fully formulated views. First, let us take the example of sadistic or anti-social preferences. I think that, if someone is sufficiently depraved in their preferences, then at a certain point we do not treat them as autonomous. We don't defer to their preferences at all. We just fall back on whatever our view of well-being is. The picture here is: I as a deliberator have my own view of well-being, which can be a hedonic view, it can be an objective view that mixes various goods—that would be, for example, the view that I would use in thinking about the well-being of children. If an adult is sufficiently deprayed, then I am, in effect, going to treat him as a child and not look to his preferences at all. The question then is, 'what if we have someone who is an intermediate case?'. We are still deferring to their preferences, but some of their preferences are anti-social, or sadistic, or what have you. Presumably, or hopefully, those preferences will be laundered out. If she is not depraved, then with good information and thinking well, those will not be the preferences that she has after idealization. She likes to torture cats, but this is an isolated depravity, so it is to be laundered out.

On Dworkin's concern about external preferences and double-counting: I would view that as addressed by the non-remoteness condition. My purely ethical preference is not part of my well-being, so that should be screened out by non-remoteness. Now it is complicated, because it is true that, to the extent that I make the realization of certain ethical preferences my own project, those can advance my well-being. Say, I devote my life to protecting a certain eco-system: there is some eco-system, let us say, in Alaska that is endangered, I devote my life to protecting that, so I have not merely an ethical preference that the ecosystem be saved, but it has become my project and that advances my well-being. This again goes

_

³⁷ See Harsanyi (1977, 647). Anti-social preferences are, for example, preferences based on "sadism, envy, resentment, and malice" (647).

³⁸ See Dworkin (1978). 'External' preferences are preferences "for the assignment of goods and opportunities to others", which can be contrasted with 'personal' preferences that are preferences "for [the person's] own enjoyment of some goods or opportunities" (234).

to how we distinguish between remote and non-remote preferences. Once we have done that, there is not an additional double-counting problem.

There is also the question how you synthesize the preferences of people with different preferences, or formulate what you do with utility functions representing different preferences. Because utility is only unique up to a linear transformation, these different utility functions have got to be scaled in some way. One way to scale them is through a high-low rule. So, there is an option that, for example, everyone strongly dis-prefers, and an option that everyone strongly prefers, and you scale all the utility functions so that you give zero to the first option, and one to the second—one kind of high-low rule. If you think about things that way, there is no double-counting problem. Everyone gets the whole numerical range between zero and one to assign to all their possible histories, from the worst to the very best. The bottom line is that, by laundering preferences, screening for remote preferences, and appropriately scaling utility functions, I think you can handle both Harsanyi's and Dworkin's worries.

Following up a bit more. You wouldn't impose any substantive conditions on laundered preferences? We might have a worry about idiosyncratic preferences, such as Rawls' case of the Grass-counter.³⁹ You would bite the bullet in this case?

This case takes us back to the relation between objective-good and preference accounts of welfare. On my current view, it is not that the preference account is a competitor to an objective account as a substantive matter; it is rather that we are deferring to the individual with respect to what the balance of objective goods is. Let me also say that, in a way, a hedonic account, which says that well-being reduces to pain and pleasure, or happiness, is simply one type of objective-good account. The hedonic account is an account that says: the only well-being good is pain, pleasure, happiness. At the end of the day, substantively, I would say, any well-being account is going to be an objective-good account of some kind. And then the question is going to be, 'is the subject sufficiently autonomous that we are going to defer to her about the balance of good, or not?'. If she *really* just counts grass-counting as the only good, maybe that suggests that she is so depraved that we are not going to defer, but if we are deferring, I wouldn't build in any substantive conditions.

³⁹ The case goes as follows: "Imagine someone whose only pleasure is to count blades of grass in various geometrically shaped areas such as park squares and well-trimmed lawns. He is otherwise intelligent and actually possesses unusual skills, since he manages to survive by solving difficult mathematical problems for a fee" (Rawls 1971, 379).

A final question concerning conditions on laundered preferences: in your 2012 book, you talk about substantive conditions and procedural conditions, as we have talked about, but then also historical conditions.40 One historical condition could be that people's preferences are not problematically adapted to their circumstances—for example, they are not based on manipulation or on problematic kinds of social pressure. That ties in interestingly to the idea of autonomy. In light of this, how would you look at the problem of adaptive preferences?

That's a good question. I have certainly taken a stab at the nature of wellbeing, particularly in chapter three of the 2012 book, but most of my work has not been about that—most of my work has been either on the implications of a preference view for the *measurement* of well-being, for example, that is my work on extended preferences in the 2016 handbook; or, on the social welfare function framework more generally, which can incorporate a whole range of theories of well-being. 41 So, I cannot hold myself out as having well-developed views on this and how to deal with adaptive preferences.

That said, my current thinking would be that we are trying to screen out individuals who are below the threshold of autonomy, or if above the threshold, preferences that are not the result of engaging their autonomy. So that is how we should think about screening out adaptive preferences. And I think that is going to be a matter of historical conditions. It is also going to be information. This goes back to something that Richard Brandt suggests. He talks about cognitive psychotherapy as the process for laundering processes. 42 So I would give people information not just about what might happen if their preferences are satisfied but also information about their own history. Part of full information is information about the origin of your preferences. That is going to be part of the solution to adaptive preferences, but it may not be the full solution. Imagine that, through a terrible upbringing, you have been brainwashed into terrible preferences; you are now given information about that, but the brainwashing is such that the information does not dislodge the preferences. But, in general, we can make some progress on the problem of adaptive preferences by conceiving of information about the origin of your preferences—that is, about your own history—as part of what full information means.

⁴⁰ See Adler (2012, 171).

⁴¹ See Adler (2016b).

⁴² See Brandt (1979), who names the "process of confronting desires with relevant information, by repeatedly representing it, in an ideally vivid way, and at an appropriate time" 'cognitive psychotherapy' (113).

This is interesting if we look at moving from laundering preferences in theory—what the conditions would look like and how to characterize laundered preferences—to actually trying to elicit laundered preferences and use them in practice when we do welfare analysis. Presumably there are two ways to approach that: on the one hand, we try to construct some circumstance or a context, when eliciting the preferences, that roughly satisfy whatever conditions we apply, to make sure that people are autonomous—we have some methodological procedure to make sure that people are fully informed about their history and their future and reason rationally; or we make some guess about what those preferences might look like and then impose them on people in some sense. Have you thought about this stage of preference laundering? That's a good question, and I'm not an expert on preference elicitation, but I think you are right. For my project, I would be trying to elicit people's well-informed, rational, and restricted (that is, non-remote) preferences. How do I do that? The direct approach is to give them all the information I think they would need and to de-bias them so as to make sure that the preferences meet rationality conditions, such as being transitive and so forth. And then just elicit the preferences. This is a kind of intensive survey method for eliciting these idealized preferences.

Let's imagine that out of this survey, or based on people's behavior, I get some data and I have to figure out what the preferences are. That is going to be the estimation stage. Presumably, at the stage of estimation, I am going to have to make certain assumptions about what the preferences might plausibly be. Imagine someone that has a preference for eating a saucer of mud but only on Tuesdays—this is an example inspired by Anscombe and Parfit.⁴³ And imagine that this is a preference that is fully autonomous, that the individual holds with good information and thinking rationally. To me as the elicitor, as the experimenter, that seems so weird—one, a preference for a saucer of mud; two, the preference is only for eating mud on Tuesdays—that in estimating the preferences, I am just probably going to rule that out. At the estimation stage we have to have some model of what the preferences might be. There is an infinity of possible models and I have to rule out some, so my judgement as to which preferences are plausible is going to come into play there.

_

⁴³ See Anscombe (1957) on the example of the saucer of mud, and Parfit (1984), who discusses the case of a man who "always prefers the prospect of great suffering on a Tuesday to the mildest pain on any other day. [...] Simply because the agony will be on a Tuesday" (123).

This ties quite nicely into the next question, which is a point about alienation, which you talked a bit about before too. Tyler Cowen has this point, for example, that "with perfect information and cleansed preferences, I would prefer red wine to Coca-Cola, but this hypothetical taste has little bearing on my welfare in a world where my preferences are not cleansed" (1993, 262). Again, it would be this worry that laundered preferences would not be suitably connected to people—they would not have motivational force.

The broader concern related to this, also looking at eliciting or using laundered preferences in practice, would be that by invoking people's laundered preferences instead of their actual preferences, analysts are given quite a few degrees of freedom to decide which preferences should count, which should not count, and we might then have concerns about these analysts imposing their own views on well-being on people instead of deferring to people. Or that some perspectives are removed from the analysis. How would you look at eliciting laundered preferences in practice in terms of who's in charge of the process and how the process should look like?

On the latter question, obviously the analyst is going to have to sincerely want to elicit the laundered preferences. The analyst needs to have—which is all really a matter of theorizing about well-being—her own understanding of what good information is, what the rationality conditions are, what non-remoteness is, and so forth. She has to sincerely want to elicit that. If she doesn't, that is an implementation problem—not a problem in theory.

The first part, the alienation problem, is a problem in theory. I believe it is Peter Railton who has a solution.⁴⁴ Think about it this way. Imagine that either my full-informed self—fully informed, rational, developing non-remote preferences—or my ideal advisor is ranking worlds. Let's say it is myself fully informed. My fully informed self is ranking possible worlds such that in the possible worlds I may not have full information. So, with full information, I am considering the possibility that in the worlds being ranked, I don't have full information. Or, for that matter that, in the worlds being ranked, my actual preferences are not the same as my fully informed preferences. Imagine that if I had full information about the process of producing red wine, I would much prefer red wine

⁴⁴ Railton (1986) proposes an 'ideal advisor' account of well-being, which holds that "an individual's *intrinsic* good consists in attainment of what he would in idealized circumstances want to want for its own sake [...] were he to assume the place of his actual self" (16).

to cheap beer. But lacking such information, I have no interest in red wine. Now imagine we have two worlds, x and y, in which I lack information about the process of producing red wine. So in both those worlds I lack the information; in x I have the wine, in y I have the beer. My fully informed self would prefer y—that I have the beer. Because the fully informed self takes account of the actual information of the subject in the worlds. That is, in effect, Railton's solution and I think it resolves the alienation problem. We are not imagining a ranking of worlds, such that people are fully informed in the worlds, we are simply imaging that the evaluator is fully informed.

The next question takes a leap to interpersonal comparisons of wellbeing. That is something that is required in the social welfare function approach. One might, however, find your argument in favor of interpersonal comparisons of well-being surprising. We say this because you write in Measuring Social Welfare: "[...] So we shouldn't accept such an SWF and thus shouldn't adopt a view of well-being that makes only intrapersonal comparisons" (46); similarly, you write that, "a plausible theory of well-being should allow for intra- and interpersonal comparisons of both well-being levels and differences" (125). This argumentative strategy returns at a number of points throughout the book. Isn't this reasoning backwards? Surely the correct theory of well-being should be determined independently of any concerns about interpersonal comparability and should then be applied in the framework? My underlying epistemology here, which a lot of people have when it comes to moral epistemology, is coherentist—it's reflective equilibrium. Specifically, here, we refine our views about ethical evaluation taking account of our views about the nature of well-being, but we also refine our views about well-being in light of the nature of ethical evaluation. We go back and forth until we have equilibrium. I would say, as a matter of strong intuition, except for a certain corner of economics, both ordinary people (as far as I can tell) and philosophers of well-being accept interpersonal comparisons. There are hypothetical cases: Amy has a high income, good health, lots of friends, an engaged intellectual life; Bob has a low income, is in terrible health, no friends, is bored and miserable. It seems clear that Amy is better off than Bob. Only economists are going to doubt that.

And then, in terms of ethical evaluation, we start with the thought that, if we are welfarists, then the ranking of worlds should depend on

the pattern of well-being. We then find out, when we try to operationalize this for the purposes of a social welfare function, that there are going to be real problems in light of this axiom of 'invariance' (which I talk about in chapter two of *Measuring Social Welfare*), if well-being is not interpersonally comparable. So that is an additional data point in favor of welfare being interpersonally comparable.

Now, you can imagine starting the process of arriving at reflective equilibrium by thinking, 'I cannot make interpersonal comparisons'. (By the way, why would you think that? You probably think that because you have a preference view of well-being; people have different preferences, how do we make interpersonal comparisons?) You start with a theory of well-being which says I cannot make interpersonal comparisons. You then realize that, in light of this 'invariance' idea, your social welfare function is going to require an interpersonally comparable well-being measure. At that point, you are out of equilibrium: either you have to give up the notion of a social welfare function; or you have to give up your initial view about well-being, namely, that it is not interpersonally comparable. If that disequilibrium is resolved by changing the view of well-being, I would not view that as being illegitimate—that is simply that nature of reflectiveequilibrium reasoning.

The social welfare function approach that you present treats interpersonal comparability in a binary way—either we have such comparability, or we don't—the broader social choice approach, however, does not require us to treat interpersonal comparability in this way. For we can have partial comparability—a whole class of social welfare functions that admits what Sen calls partial comparability.45 What was your motivation to restrict the class of social welfare functions that you present to the class of social welfare functions that give us some sort of interval or ratio scale comparability, while excluding other social welfare functions? Is it just a matter of keeping it manageable, or was there some other motivation?

In chapter two of Measuring Social Welfare I have this two-by-two table with admissible well-being comparisons. 46 We can talk about either intraor interpersonal comparisons of well-being levels or differences. If a particular type of comparison is *inadmissible*—'admissible' or 'inadmissible' meaning being eligible or not, recognized or not, by our theory of well-

⁴⁵ See Sen (1970b).

⁴⁶ See Adler (2019, 42).

being—then one can never make well-being comparisons of that type. If it is *admissible*, there still might be pockets of incomparability. For example, assume that I think that well-being levels are comparable both intrapersonally and interpersonally—that is, I think that well-being level comparisons both intra- and interpersonal are admissible. Thus I have a well-being ranking of histories that is not limited to the histories of one person (intrapersonal level comparison) but includes the histories of different persons (interpersonal level comparisons). Now, I am going to say that this well-being ranking is a quasi-ordering—it is a transitive, reflexive ranking of histories. It might be incomplete. It is not *pervasively* incomplete, it is not massively incomplete, so there is some comparability. Including between the histories of different people. But there could well be pockets of incomparability.

At the end of the day, I think that all four cells in that table are admissible—I think that intra- and interpersonal levels and differences are comparable. I understand that formally—this is what I talk about in my chapter on extended preferences in the *Oxford Handbook of Well-Being and Public Policy* and also in my appendix to *Measuring Social Welfare*—as a quasi-ordering of the set of histories (which captures well-being levels) and a quasi-ordering of the set of history pairs (which captures well-being differences), but both those quasi-orderings can be incomplete. ⁴⁷ There can be pockets of incomparability.

Now, at the level of the social welfare function, what are you going to do? The simple approach, the approach that I use in *Measuring Social Welfare*, the approach generally used in the literature, is to assume that there is a single well-being measure. For my purposes, a single well-being measure that captures levels and differences. *That* presupposes completeness—if there is incompleteness in the comparisons of either levels or differences, the levels ranking and differences ranking cannot be captured by a single well-being measure. The simple approach is for tractability. There are ways to generalize the social welfare function beyond this operational assumption of completeness, which I discuss in my 2012 book. For example, in that book, what I do is to think of the well-being comparisons being represented by a *set* of well-being measures. Well-being comparisons of levels and differences correspond to a super-valuation over that set and the social welfare function is applied to all the well-

⁴⁷ See Adler (2016b; 2019, 259-297).

being measures in that set. This is going to be a way to allow for incompleteness within the social welfare function framework. But, again, it is simpler, more tractable, to just assume a single well-being measure.

The general point, if we are understanding this correctly, is that you are sympathetic to having a more fine-grained distinction between incomparability at one end—intrapersonal incomparability—and full ratio scale of comparability at the other.

Correct. Well, there are two issues. One issue is, is well-being measurable on an ordinal scale, an interval scale, or a ratio scale? The question there is, can we make comparisons of levels, comparisons of differences, and/or comparisons of ratios? I actually think that we can make comparisons of all three. Then, in making level, difference, and/or ratio comparisons, are those complete or not? There I would say 'no'. Plausibly, there are going to be pockets of incomparability in making level comparisons, difference comparisons, or ratio comparisons. We are going to be somewhere in the middle between full completeness and global incomparability. That is going to be the best theory of well-being and, again, just in terms of making the social welfare function approach implementable, we are going to assume, or impose, completeness by having a single well-being measure. But that is a simplification.

A related, but also distinct question, on the metaphysics of value: some of these approaches require strong scaling assumptions, which you also spoke about earlier. Minimally you normalize a lower bound at zero. How do you interpret that? Is that the worst possible life that you can lead, or is that just neutral between positive features and negative features?

It depends. Utilitarianism in the fixed population context only needs a well-being measure which is unique up to a positive affine transformation. It does not really matter what the zero is. That becomes different in the variable population context. In a variable population context, it depends on what the variable population view is: there is total utilitarianism, there is also critical level utilitarianism. If you are a total utilitarian, who says adding someone to the population is good as long as that person is better off than in the neutral life, it is natural to assign zero to the neutral life. If you are a critical level utilitarian, who says—this is Broome's view—there is some critical level that is better than neutrality, such that adding someone to the population is only good if the person is above the critical

level—this is designed to avoid the repugnant conclusion—then it might be natural to assign zero to the critical level.⁴⁸ You can do either one as long as you are aware of what you are doing.

For prioritarianism, it is more complicated. No prioritarian social welfare function, whatever the concave transformation function might be, is going to be invariant to a positive affine transformation of the well-being measure. With the Atkinson prioritarian social welfare function, even in the fixed population context, the interval scale of well-being is not enough. We have to fix a zero point. That zero point is the Atkinson zero, in the sense that it has got to be no better than, and probably worse than, any of the lives under consideration. The problem with Atkinson is that it does not work with negative well-being numbers. We are going to have to pick zero so that none of the lives under consideration are assigned a negative well-being number.

IV. THE ETHICAL DELIBERATOR AND THE ROLE OF THE PHILOSOPHER

Changing track a little bit. There are a number of value judgments central to the social welfare function approach, as you discuss in Measuring Social Welfare: "Choosing a correct theory of well-being, constructing a measure w(.), adopting a rule E, specifying an uncertainty module" (79). How are these questions to be answered? Your own answer to this consists in presenting reflective equilibrium as a moral epistemology and you write: "It is for each ethical deliberator to locate her own point of reflective equilibrium—to determine which SWF she favors, on balance and after reflection" (115), that is, which "SWF most fully embodies her ethical commitments" (96). Can you elaborate on how this is supposed to work in practice?

In practice, at the level of each of us as individual thinkers trying to decide what we think morally, the situation is just what I say. You have to first of all deliberate about whether you are inclined towards consequentialism or some kind of non-consequentialist view, say a deontological view; within consequentialism, towards welfarism as opposed to non-welfarist consequentialism; and then within welfarism, you need to think about what welfarist world-ordering you favor—that is going to be a matter of specific thought experiments, we can have thought experiments which are

_

⁴⁸ See Broome (2004); see also Blackorby, Bossert, and Donaldson (2005).

trouble for utilitarianism—but one of the nice things about the social welfare function literature is that it also shows the role of axioms in thinking about different world orderings. And again, E is a world ordering, more precisely it is a way to express the ranking of worlds in light of well-being. In choosing between the utilitarian E rule, the prioritarian E rules, egalitarian, leximin, and so forth, thinking about how those relate to different axioms is very helpful. And, indeed, what I talk about in chapters three and four in *Measuring Social Welfare*—here I am just synthesizing the existing literature—is that if we use the five axioms of Pareto, anonymity, Pigou-Dalton, separability, and continuity, we can really narrow down the space of different social welfare functions and crystalize our thinking. Most people in this space find Pareto and anonymity pretty compelling, I also find Pigou-Dalton quite compelling. Pigou-Dalton is, of course, the axiom that separates between utilitarianism and prioritarianism. So, part of the way that I come to my own, prioritarian, reflective equilibrium is by thinking about and endorsing Pigou-Dalton.

The follow-up question to that is how this idea of the ethical deliberator will play out in practice, who this ethical deliberator is supposed to be when we apply the social welfare function framework?

That is more complicated. The complication is that, if 'in practice' means in government, then we have to think about the relation between law and morality. Because government is constituted by different officials with different official roles. And each official is given legal powers and restrictions, and what an official does is constrained by those. An official cannot simply do what she thinks is the ethically best thing—she has to act within a role. For example, in the United States, we have civil servants in the federal government who are not elected. And I think a lot of people would think it would be legally problematic—not ethically problematic, legally problematic—for an unelected civil servant to decide, 'I am just going to operationalize whichever social welfare function I like as an ethical matter'. Rather, that seems to be the kind of judgement which legally should be made only by a legislator or, in the U.S., by the elected President. As an individual, or as a citizen, I am not legally constrained; I can act on my own social welfare function in voting or donating my money. The additional complication comes in to the extent that we have officials with a legal role and operating under legal constraint. There again we have to think about the connection between law and morality, which is a complicated topic. But here I would say at least that, in a democratic legal

system, there are going to be some decision makers with sufficient democratic authorization to be able to choose between different social welfare functions. At a minimum, that would be legislators; a legislator can pick a social welfare function to help him decide how to vote on proposed legislation. And there might be other officials—for example, in the U.S., presidents—operating within their zone of legal discretion.

The upshot would basically be that, whether we should have a prioritarian or utilitarian social welfare function is ultimately something that should be democratically informed as you see things?

Well, yes and no. I am a citizen. Let us imagine that everyone else in the country is utilitarian and I am prioritarian. I should not necessarily change my view—I might think everybody else is misguided. I might change my view: the fact everyone else is utilitarian might lead me to reflect, but if after reflection I still believe in prioritarianism, I have no democratic obligation to change my view simply because everyone else is utilitarian. Again, we are welfarists. I hate to say this, but it is the case: welfarists believe that democracy has no bedrock moral role. That itself is complicated because you might say that democratic values are somehow related to well-being. But at least to a first approximation, democratic legal structures have no bedrock ethical role for welfarists. It is rather that democratic legal structures are a good way in practice to advance well-being look how dictatorships are doing as compared to democracies. So, the point is that we have good welfarist reason to promote democratic legal structures, and part of democratic legal structures is that people's legal roles are limited and that officials without democratic authorization cannot legally make certain decisions that officials with democratic authorization can. But again, it seems to me that—and this is a very complicated topic—all of these democratic constraints are a matter of law, not morality, and that the welfarist—whether she is a utilitarian or prioritarian sees democratic legal structures simply as a useful instrument to promoting well-being.

Relatedly, how do you see your own role and the role of philosophers in the social welfare function approach? Is it to make the strongest possible case for particular views on these substantive (ethical) issues (that are involved in choosing these different functions, a well-being measure, and uncertainty modules)? Or is it to provide a 'taxonomy' of alternative specifications without really taking any side? We ask because

the book can be read as laying out the framework (assumptions and choices) and providing reasons for specific social welfare functions without really taking a side. Nevertheless, your own sympathy for the prioritarian class of social welfare functions does come through as well. The growth of knowledge over the last century, let us say, or 150 years whatever it is—about both economics and, within philosophy, about consequentialism, has been unbelievable. We have a much, much better formal and substantive understanding of these subjects—I mean, think about where philosophy was, say, before Sidgwick started writing, before the 19-century utilitarians started writing; let alone the development of economics with Samuelson, Arrow, Sen and all that, I mean the explosion of knowledge—genuine knowledge—has been fantastic. 49 So, I think philosophers can do their bit as part of that and that can be different things. It really depends on what you are good at and what you want to do. One thing is to present the pros and cons of different approaches and develop our understanding of possible axioms, the implications of those axioms, ways to characterize different views, and so forth. It is also valuable to advocate for a particular view. For example, think about Larry Temkin; his book *Inequality* is a great book.⁵⁰ He has devoted his academic efforts to arguing for the value of equality. That has made a huge contribution. Contrast John Broome. Broome is a utilitarian but it is not as if Broome has argued in a full-throated way for utilitarianism in the same way that Temkin has argued for equality. Rather, Broome has been exploring the axiomatic and substantive arguments for and against utilitarianism. I think Broome has made an incredible contribution. So, it is really what you can do. My advice would be to philosophers, work on problems that engage you and make the contribution where you can make it.

They also feed into each other, right. We also posed the question in a rather stark way. Clarifying axiomatic structure can through reflective equilibrium clarify for you what your position is, and the other way as well, Temkin is an instructive example here—the advocacy can also clarify what is at stake.

Consider Parfit. In the fourth part of *Reasons and Persons*, he identified some very important puzzles for population ethics—the non-identity problem, the repugnant conclusion—and a key contribution was to say, 'I cannot identify a theory that avoids all these unpleasant consequences; I

_

⁴⁹ See Sidgwick's *The Methods of Ethics* (1874).

⁵⁰ See Temkin (1993).

would like to come up with a theory *X* that does that, but I haven't yet done so'. That intervention has been incredibly fruitful—he launched an entire generation of people looking for theory *X*.

On the other hand, if you think about part three of *Reasons and Persons*, where Parfit adopts a particular view of personal identity—kind of a psychological continuity view—there he is less skeptical, the idea seems to be, 'this is what personal identity consists in'. That has also been very valuable in the literature on personal identity—people react to him. That is less skeptical, or less open, but it has also been very useful.

This ties in very well with the last question of this section. It will be clear to everyone who reads Measuring Social Welfare, and your other work as well, that John Harsanyi and Derek Parfit have had a significant influence on your thinking. Can you tell us about how you came across their work and what the nature of their influence has been on your thinking?

Even though I did do a fair bit of philosophy in college, I don't think I read either of them during college. I did grad work in history—didn't read them then. I was at Oxford, but I never met Parfit when I was there as a grad student in history. It's when I started law teaching, teaching seminars on the ethical foundations of cost-benefit analysis, that I started to read both Harsanyi and Parfit.

Parfit's *Reasons and Persons* is widely seen as perhaps *the* greatest work of moral philosophy over the last century—certainly the greatest work of moral philosophy in the utilitarian or welfarist tradition over the last century—so it has to be read and grappled with. As for Harsanyi: with reference to cost-benefit analysis, I was very interested in this problem of interpersonal comparisons. As mentioned, the view that I was developing with Eric Posner was that we should think about cost-benefit analysis as a rough proxy for overall well-being. I was also inclined towards a preference-based view of well-being. So then the puzzle was, 'how do we adopt a preference-based view of well-being and allow for interpersonal comparisons?', and of course Harsanyi proposes a solution to that. That's really it. I came to Harsanyi's work on extended preferences as a result of my interest in both preference-based views of well-being and interpersonal comparisons.

V. LOOKING AHEAD

What's next for you in terms of your research and the social welfare function approach? What future developments are your looking for in terms of the social welfare function approach in both theory and application? What research questions are you looking to tackle in the future? On the more practical side, we just finished this *Prioritarianism in Practice* book, which involved a whole team of scholars—with different chapter authors looking at the application of prioritarianism as contrasted with utilitarianism and other standard approaches (cost-benefit analysis, cost-effectiveness analysis), in different policy domains. The policy domains included climate policy, optimal taxation, risk regulation, health policy, COVID-19, education policy.

I am spearheading with Ole Norheim a follow-on effort, which we are calling PiP Health, continuing the work on the application of prioritarianism in health policy. QALYs and QALY-based cost-effectiveness have had a big role in health policy assessment in various countries. But for these frameworks, the measure of individual attainment is health. The QALY approach uses a measure of health, which combines longevity and health states, and then the impact on income is considered separately. That is cost-effectiveness analysis. What PiP Health is trying to do is thinking about health policies using a measure of lifetime well-being that integrates longevity, health, *and* income. And then using that as the input into a prioritarian social welfare function. The chapter on health policy in the *Prioritarianism in Practice* book does this already, but we are hoping for follow-on work and other related work on prioritarianism in health.⁵¹ That's the applied project.

The other thing right now: it may be that the United States is going to move to, finally, incorporating distributional weights into cost-benefit analysis—distributional weights are a way to approximate a utilitarian or even a prioritarian social welfare function.⁵² So, I may be involved in that effort, or writing about that effort.⁵³

More theoretically, I think there are a whole bunch of outstanding questions still for prioritarianism. For example, prioritarianism and non-

-

⁵¹ See Cookson, Norheim, and Skarda (2022).

⁵² See Adler (2016a).

⁵³ Subsequent to this interview, the Biden Administration adopted a new guidance document for cost-benefit analysis. See Office of Management and Budget, Circular No. A-4 (November 9, 2023), at https://www.whitehouse.gov/wp-content/uploads/2023/11/CircularA-4.pdf. This guidance allows for distributional weighting.

human animals. How we should take account of the well-being of nonhuman animals if we are prioritarians is a really important topic. There has been a broader interest in social welfare functions and animals, so that is something that I hope to write about.

There are other pieces of the puzzle that I think, philosophically, still need to be worked on. For example, the whole issue of separability. Separability is this technical axiom that distinguishes between, as I see it, prioritarianism and a kind of moderate egalitarianism. A moderate egalitarian social welfare function respects Pareto, anonymity, and Pigou-Dalton, but fails separability, while prioritarianism respects Pareto, anonymity, and Pigou-Dalton, and *satisfies* separability. I think more philosophical attention needs to be given to separability, and there are lots of complicated issues about separability not just under certainty in the fixed population case, but with variable population and under uncertainty. I'm hoping to work on these issues.

A third thing I'm working on now, in terms of philosophical fundamentals, is luck prioritarianism: the question of how to incorporate considerations of desert and responsibility into prioritarianism. The thought behind luck egalitarianism and this whole economic literature on equality of opportunity is that, at the end of the day welfarism is insufficient because we should care about the extent to which people are responsible for being badly off as opposed to being badly off through no fault of their own. How to make sense of that thought under the rubric of prioritarianism is an open question. I took a stab at it in an article in *Utilitas*, in 2018— "Prioritarianism: Room for Desert?"—but much more works needs to be done.⁵⁴ A final piece of the puzzle is the debate between luck egalitarians and relational egalitarians. Relational egalitarians—I am thinking here of the work of Elizabeth Anderson, for example, or more recently of Juliana Bidadanure—are non-consequentialists who care a lot about social relations, or social status, as opposed to the distribution of income or wellbeing.55 I think one question for, let's say, prioritarian welfarists is how we take account of relational egalitarian insights. The key relational egalitarian insight, as I see it, is that, to put it in welfarist terms, social status matters a lot for well-being. That the relations of social hierarchy, oppression, subordination—which of course we continue to see all around us have a huge impact on individuals' well-being. The question is, what exactly does that mean as a prioritarian—how do we take account of these

-

⁵⁴ See Adler (2018a).

⁵⁵ See, for example, Anderson (1999) and Bidadanure (2016).

problematic types of social relations within prioritarianism? Do we somehow funnel them into our well-being measure? How is that going to work? It has been avoided in practice, because typically when prioritarianism is implemented the inputs to well-being are things like income, longevity, health, and so forth, and not the quality of someone's social relations or social status. There are a range of issues, both theoretical and practical, that need to be engaged there.

And then the other thing I am doing, I am finishing up this book on *Risk, Death, and Well-Being: The Ethical Foundations of Fatality Risk Regulation*, looking at fatality risk regulation from a prioritarian as well as a utilitarian perspective. All of this will keep me busy over the next couple of years.

Your academic work is, as we have explored, highly interdisciplinary in nature. How do you look at the future of interdisciplinary work in law, economics, and philosophy? What developments are you seeing and what developments are you hoping for?

I think there is a lot of good work occurring at the intersection of economics and philosophy, certainly. In the normal way, generated either by problems that are coming up in the literature or externally. I can think of a number of examples. One is all of this work on variable population. The issue of variable population, or population ethics, has always been a fruitful source of interchange between economists—for example, Blackorby, Bossert, and Donaldson—and philosophers—obviously going back to Parfit or Gustaf Arrhenius—and that continues to happen. We have a new generation working on these problems, people like Dean Spears, Mark Budolfson, and others from the economics side and philosophers as well. On that is one exciting area. There are a lot of people working both on different possibilities and on advancing challenging new ideas; for example, an increased willingness perhaps to accept the repugnant conclusion. In any event, population ethics is one area where I see a lot of new work at the intersection of economics and philosophy.

The other thing right now, growing out of effective altruism, is the whole interest in long-termism, fanaticism. The idea that we should devote a lot of efforts to increasing the chance that humanity continues for a very long time, and should be willing to impose present costs for the

⁵⁶ See Blackorby, Bossert, and Donaldson (1995, 1997, 2005), Parfit (1984), and Arrhenius (2000).

⁵⁷ See, for example, Spears and Budolfson (2021).

sake of only a small chance of increasing long-term survival, has generated this literature on fanaticism—coming out of, for example, the Oxford Global Priorities Institute and other such think-tanks, and including publications in *Ethics*. This is also proving to be a very fruitful area for work at the intersection of economics and philosophy.

A third example is the work now among contractualists about risk regulation. How contractualists growing out of Scanlon's work on contractualism should think about permissible risk imposition.⁵⁸ It is philosophical but it is also at times fairly formal, and I view this as another exciting area of research that integrates formal methods with philosophy.

Topics for collective attention emerge organically in the literature, which is the way scholarship works. And then we have external stimuli; for example, for a long time now thinking about climate policy has been an important external stimulus to economic and philosophical work.

The EJPE is an interdisciplinary journal, and our readers are scholars who either do work at the intersection of philosophy and economics, or are at least open to such an interdisciplinary approach. What advice would you give to graduate students aiming to pursue an interdisciplinary academic career, such as your own?

I feel that I have been extremely lucky. The U.S. legal academy is very open to interdisciplinary approaches; I was fortunate to get a job at a U.S. law school and thus to have the freedom to pursue my research interests, educate myself in different fields—again I came in without a PhD in either economics or philosophy and was able to pursue those interests. As a matter of practical advice: the first thing is to get a job. If you want to be able to pursue interdisciplinary scholarship over a lifetime, the first thing is to figure out, depending upon your country, what sorts of long-term positions are available. In the U.S. that is heterogeneous—there are positions at law schools, there are positions at public policy schools, there are positions at philosophy departments, there are positions at econ departments—so think about that, and think about your own skill set. Earning a PhD or another advanced degree (in my case, a law degree) and landing a job is a twofold thing: on the one hand it is about pursuing your interests, on the other it is an exercise in finding a job and then, later, getting tenure. So, again, the most pragmatic advice is to keep in mind employment prospects as you pursue your studies.

⁵⁸ See Scanlon (1998).

Another point is that work at the margins of economics, law, and philosophy can proceed in lots of ways. It can be quite formal; it can be quite philosophical and conceptual; it can be focused on legal implementation, which is very complicated. So, figure out which of those you enjoy most and which of those you are most able to contribute to. If you don't have a lot of interest in, or aptitude for, math then don't try to do really formal work—do philosophical work and try to collaborate with formal people. And vice versa: if your value-added is the math side, that's great, try to do the philosophy too, but also be prepared to collaborate.

In short, it's a matter of figuring out where the jobs are and what it takes to get a job, figuring out what interests you, and figuring out what your skillset is.

Now, if you are fortunate enough, as I have been, to have a tenured job where you can pursue a bunch of research questions, at that point, the issue is not landing a job, but it is still figuring out where you can have an impact, which also relates to what interests you and what your capabilities are.

As a final question, if you had to name three philosophical works that any graduate student in law, economics, and philosophy should read, which would those be?

I'd certainly say Parfit's Reasons and Persons. 59 As I already said, the greatest work in moral philosophy—certainly welfarist moral philosophy—in at least the last century. And Broome's Weighing Goods. 60 Broome has been incredibly important in this field. It is an attempt to synthesize economics and philosophy, by someone who was first a chaired professor in economics then a chaired professor in philosophy. So, Reasons and Persons and Weighing Goods.

Now picking the third is a little tricky. I think there are a number of possibilities. Sen's Collective Choice and Social Welfare is incredibly important, it is more formal than philosophical. ⁶¹ I think that everyone needs to have read Rawls' A Theory of Justice. 62 Two other books that I think are really important. The best single book on the philosophy of well-being is Sumner's Welfare, Happiness, and Ethics—it's a great book. 63 Harsanyi did incredibly important stuff, but he didn't write a book-length work that I'd

⁵⁹ See Parfit (1984).

⁶⁰ See Broome (1991).

⁶¹ See Sen (1970a).

⁶² See Rawls (1971).

⁶³ See Sumner (1996).

recommend. I would consider Temkin's book *Inequality*, a really important defense of the value of equality but, more generally, just incredibly rich philosophically.⁶⁴ I realize this is two plus a bunch as opposed to three...

Professor Adler, thank you for sharing your time and ideas with us. Thanks a lot, I appreciate it.

REFERENCES

- Adler, Matthew D. 2012. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. New York, NY: Oxford University Press.
- Adler, Matthew D. 2015. "The Ethical Value of Risk Reduction: Utilitarianism, Prioritarianism, and Cost-Benefit Analysis." In *Ethics & Risk Management*, edited by Lina Svedin, 9–29. Charlotte, NC: Information Age Publishing.
- Adler, Matthew D. 2016a. "Benefit-Cost Analysis and Distributional Weights: An Overview." *Review of Environmental Economics and Policy* 10 (2): 264–285.
- Adler, Matthew D. 2016b. "Extended Preferences." In *The Oxford Handbook of Well-Being and Public Policy*, edited by Matthew D. Adler and Marc Fleurbaey, 476–517. New York, NY: Oxford University Press.
- Adler, Matthew D. 2018a. "Prioritarianism: Room for Desert?" *Utilitas* 30 (2): 172-197.
- Adler, Matthew D. 2018b. "Utilitarianism, Prioritarianism, and Climate Change: A Brief Introduction." In *Climate Change and Its Impacts: Risks and Inequalities*, edited by Colleen Murphy, Paolo Gardoni, and Robert McKim, 69–89. Cham: Springer.
- Adler, Matthew D. 2019. *Measuring Social Welfare: An Introduction*. New York, NY: Oxford University Press.
- Adler, Matthew D. 2023. "Age and the Social Value of Risk Reduction: Three Perspectives." In *Ageing without Ageism? Conceptual Puzzles and Policy Proposals*, edited by Greg Bognar and Axel Gosseries, 53–70. New York, NY: Oxford University Press.
- Adler, Matthew D., David Anthoff, Valentina Bosetti, Greg Garner, Klaus Keller, and Nicolas Treich. 2017. "Priority for the Worse-Off and the Social Cost of Carbon." *Nature Climate Change* 7: 443–449.
- Adler, Matthew D., Richard Bradley, Maddalena Ferranna, Marc Fleurbaey, James Hammitt, and Alex Voorhoeve. 2020. "Assessing the Wellbeing Impacts of the Covid-19 Pandemic and Three Policy Types: Suppression, Control, and Uncontrolled Spread." T20 Policy Brief; Task Force 4: Social Cohesion and the State. Available online: https://www.global-solutions-initiative.org/policy_brief/assessing-the-wellbeing-impacts-of-the-covid-19-pandemic-and-three-policy-types-suppression-control-and-uncontrolled-spread/.
- Adler, Matthew D., Richard Bradley, Maddalena Ferranna, Marc Fleurbaey, James Hammitt, Rémi Turquier, and Alex Voorhoeve. 2023. "How to Balance Lives and Livelihoods in a Pandemic." In *Pandemic Ethics: From COVID-19 to Disease X*, edited by Dominic Wilkinson and Julian Savulescu, 189–209. New York, NY: Oxford University Press.

_

⁶⁴ See Temkin (1993).

- Adler, Matthew D., Maddalena Ferranna, James K. Hammitt, and Nicolas Treich. 2021. "Fair Innings? The Utilitarian and Prioritarian Value of Risk Reduction over a Whole Lifetime." *Journal of Health Economics* 75: 102412.
- Adler, Matthew D, and Marc Fleurbaey, eds. 2016. *The Oxford Handbook of Well-Being and Public Policy.* New York, NY: Oxford University Press.
- Adler, Matthew D, and Marc Fleurbaey. 2018a. "In Pursuit of Social Progress." *Economics & Philosophy* 34 (3): 443–449.
- Adler, Matthew D, and Marc Fleurbaey. 2018b. "Response to Our Commentators on the Report of the International Panel on Social Progress 2018." *Economics & Philosophy* 34 (3): 447–482.
- Adler, Matthew D., James K. Hammitt, and Nicolas Treich. 2014. "The Social Value of Mortality Risk Reduction: VSL versus the Social Welfare Function Approach." *Journal of Health Economics* 35: 82–93.
- Adler, Matthew D., and Kenneth E. Himma, eds. 2009. *The Rule of Recognition and the U.S. Constitution*. New York, NY: Oxford University Press.
- Adler, Matthew D, and Ole F. Norheim, eds. 2022. *Prioritarianism in Practice*. New York, NY: Cambridge University Press.
- Adler, Matthew D, and Eric A. Posner. 1999. "Rethinking Cost-Benefit Analysis." *Yale Law Journal* 109 (2): 165–248.
- Adler, Matthew D., and Eric A. Posner, eds. 2000a. *Cost-Benefit Analysis: Legal, Economic, and Philosophical Perspectives*. Chicago, IL: University of Chicago Press.
- Adler, Matthew D., and Eric A. Posner. 2000b. "Implementing Cost-Benefit Analysis when Preferences are Distorted." In *Cost-Benefit Analysis: Legal, Economic, and Philosophical Perspectives. Supplement to Volume 29, Journal of Legal Studies*, edited by Matthew D. Adler and Eric A. Posner, 1105–1147. Chicago, IL: University of Chicago Press.
- Adler, Matthew D, and Eric A. Posner. 2006. *New Foundations of Cost-Benefit Analysis*. Cambridge, MA: Harvard University Press.
- Adler, Matthew D., and Eric A. Posner. 2008. "Happiness Research and Cost-Benefit Analysis." In *Legal Implications of the New Research on Happiness. Supplement to Volume 37, Journal of Legal Studies*, edited by Eric A. Posner and Cass R. Sunstein, 253–292. Chicago, IL: University of Chicago Press.
- Adler, Matthew D., and Chris W. Sanchirico. 2006. "Inequality and Uncertainty: Theory and Legal Applications." *University of Pennsylvania Law Review* 155 (2): 279–377.
- Adler, Matthew D., and Nicolas Treich. 2015. "Prioritarianism and Climate Change." *Environmental and Resource Economics* 62 (2): 279–308.
- Anderson, Elizabeth S. 1999. "What is the Point of Equality?" Ethics 109 (2): 287–337.
- Anscombe, G. Elizabeth M. 1957. Intention. Oxford: Basil Blackwell.
- Arneson, Richard J. 1989. "Equality and Equal Opportunity for Welfare." *Philosophical Studies* 56 (1): 77-93.
- Arneson, Richard J. 2000. "Welfare Should Be the Currency of Justice." *Canadian Journal of Philosophy* 30 (4): 497–524.
- Arrhenius, Gustaf. 2000. *Future Generations: A Challenge for Moral Theory*. PhD Dissertation, Uppsala University.
- Arrow, Kenneth J. 1951. *Social Choice and Individual Values*. New York, NY: John Wiley and Sons.

- Bidadanure, Juliana. 2016. "Making Sense of Age-Group Justice: A Time for Relational Equality?" *Philosophy, Politics & Economics* 15 (3): 234–260.
- Blackorby, Charles, Walter Bossert, and David Donaldson. 1995. "Intertemporal Population Ethics: Critical-Level Utilitarian Principles." *Econometrica* 63 (6): 1303–1320.
- Blackorby, Charles, Walter Bossert, and David Donaldson. 1997. "Critical-Level Utilitarianism and the Population-Ethics Dilemma." *Economics & Philosophy* 13 (2): 197–230.
- Blackorby, Charles, Walter Bossert, and David Donaldson. 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics.* Cambridge: Cambridge University Press.
- Bossert, Walter, and John A. Weymark. 2004. "Utility in Social Choice." In *Handbook of Utility Theory, vol. 2 (Extensions*), edited by Salvador Barberà, Peter J. Hammond, and Christian Seidl, 1099–1177. Boston: Kluwer Academic.
- Botzen, W. J. Wouter, and Jeroen C. J. M. van den Bergh. 2014. "Specifications of Social-Welfare in Economic Studies of Climate Policy: Overview of Criteria and Related-Policy Insights." *Environmental and Resource Economics* 58 (1): 1–33.
- Brandt, Richard. 1979. A Theory of the Good and the Right. Oxford: Clarendon Press.
- Broome, John. 1991. Weighing Goods: Equality, Uncertainty and Time. Oxford: Basil Blackwell.
- Broome, John. 2004. Weighing Lives. Oxford: Clarendon Press.
- Cohen, Gerald A. 1989. "On the Currency of Egalitarian Justice." Ethics 99 (4): 906-944.
- Cookson, Richard, Ole F. Norheim, and Ieva Skarda. 2022. "Prioritarianism and Health Policy." In *Prioritarianism in Practice*, edited by Matthew D. Adler and Ole F. Norheim, 260–316. New York, NY: Cambridge University Press.
- Dworkin, Richard. 1978. *Taking Rights Seriously*. Cambridge, MA: Harvard University Press.
- Dworkin, Ronald. 1981a. "What is Equality? Part 1: Equality of Welfare." *Philosophy & Public Affairs* 10 (3): 185–246.
- Dworkin, Ronald. 1981b. "What is Equality? Part 2: Equality of Resources." *Philosophy & Public Affairs* 10 (4): 283–345.
- Ferranna, Maddalena, James K. Hammitt, and Matthew D. Adler. 2023. "Age and the Value of Life." In *The Routledge Handbook of the Economics of Ageing*, edited by David E. Bloom, Alfonso Sousa-Poza, and Uwe Sunde, 566–577. New York, NY: Routledge.
- Goodin, Robert E. 1986. "Laundering Preferences." In *Foundations of Social Choice Theory*, edited by Jon Elster and Aanund Hylland, 75–102. Cambridge: Cambridge University Press.
- Harsanyi, John C. 1977. "Morality and the Theory of Rational Behaviour." *Social Research* 44 (4): 623–656.
- Hausman, Daniel M. 2012. *Preference, Value, Choice, and Welfare*. New York, NY: Cambridge University Press.
- Hicks, John R. 1939. "The Foundations of Welfare Economics." *Economic Journal* 49 (196): 696–712.
- Kaldor, Nicholas. 1939. "Welfare Propositions of Economics and Interpersonal Comparisons of Utility." *Economic Journal* 49 (195): 549–552.
- Kaplow, Louis. 1996. "The Optimal Supply of Public Goods and the Distortionary Cost of Taxation." *National Tax Journal* 49 (4): 513–533.

Kaplow, Louis. 2004. "On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy." *Journal of Economic Perspectives* 18 (4): 159–175.

Kaplow, Louis. 2008. *The Theory of Taxation and Public Economics*. Princeton, NJ: Princeton University Press.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimal Income Taxation." *Review of Economic Studies* 38 (2): 175-208.

Parfit, Derek. 1984. Reasons and Persons. Oxford: Oxford University Press.

Posner, Eric A. 2021. *How Antitrust Failed Workers*. New York, NY: Oxford University Press.

Railton, Peter. 1986. "Facts and Values." Philosophical Topics 14 (2): 5-31.

Ramsey, Frank P. 1927. "A Contribution to the Theory of Taxation." *The Economic Journal*, 37 (145): 47–61.

Rawls, John. 1971. A Theory of Justice. Cambridge, MA: Harvard University Press.

Raz, Joseph. 1986. The Morality of Freedom. Oxford: Clarendon Press.

Roemer, John E. 1996. *Theories of Distributive Justice*. Cambridge, MA: Harvard University Press.

Roemer, John E. 2020. "What Egalitarianism Requires: An Interview with John E. Roemer." *Erasmus Journal for Philosophy and Economics* 13 (2): 127–176.

Royal Swedish Academy of Sciences. n.d. "The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 1996." Accessed January 17, 2024. https://www.nobelprize.org/prizes/economic-sciences/1996/summary/.

Scanlon, Thomas M. 1998. *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press.

Scitovsky, Tibor. 1941. "A Note on Welfare Propositions in Economics." *The Review of Economic Studies* 9 (1): 77–88.

Sen, Amartya. 1970a. Collective Choice and Social Welfare. San Francisco: Holden Day.

Sen, Amartya. 1970b. "Interpersonal Aggregation and Partial Comparability." *Econometrica* 38 (3): 393–409.

Sen, Amartya. 1979. "The Welfare Basis of Real Income Comparisons: A Survey." *Journal of Economic Literature* 17 (1): 1–45.

Sen, Amartya. 1980. "Equality of What?" In *Tanner Lectures on Human Values, Volume* 1, edited by Sterling M. McMurrin, 195–220. Cambridge: Cambridge University Press.

Sen, Amartya. 1982. "Rights and Agency." Philosophy & Public Affairs 11 (1): 3-39.

Sidgwick, Henry. 1874. The Methods of Ethics. London: Macmillan.

Spears, Dean, and Mark B. Budolfson. 2021. "Repugnant Conclusions." *Social Choice & Welfare* 57 (3): 567–588.

Sumner, Wayne. 1996. Welfare, Happiness, and Ethics. Oxford: Clarendon Press.

Temkin, Larry S. 1993. *Inequality*. New York, NY: Oxford University Press.

Temkin, Larry S. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. New York, NY: Oxford University Press.