

7

After Solipsism

DAVID SCHMIDTZ

1. CONSEQUENCES BY THE NUMBERS

As Adam Smith saw it, a “man of system”

seems to imagine that he can arrange the different members of a great society with as much ease as the hand arranges the different pieces upon a chess-board. He does not consider that the pieces upon the chess-board have no other principle of motion besides that which the hand impresses upon them; but that, in the great chess-board of human society, every single piece has a principle of motion of its own, altogether different from that which the legislature might chuse to impress upon it.¹

To Smith and other Scottish Enlightenment moralists, especially Hume, the human condition was defined by a handful of key features. We exhibit limited benevolence. We live in a world of manageable scarcity. For better and for worse, our possessions are subject to easy transfer.

Implicit in these observations is a prior and even more essential feature: namely, we are social beings. We live among people who decide for themselves and whose decisions affect each other. What you do affects people. What you do affects people’s payoffs, of course, but my theme in this chapter is a caution about how we misunderstand the human condition if we suppose our ways of affecting people can all be seen as ways of affecting their payoff.

Consider how plausible it would be, if the payoff—the utility—were all that mattered, to think consequentialist morality is about maximizing utility. My focus in this chapter is on how much less plausible that reduction is when we note that affecting people’s payoffs is only one among several ways of affecting people. What if the lasting consequence

¹ Smith 1984: 234.

of moving my pawn to K4 rather than K3 is not that my pawn sits at K4, but that other players play differently?

Scottish Enlightenment theorists focused on the nature and causes of the wealth of nations, or more generally, on what makes the world a better place. Today's act-utilitarianism, by comparison, sometimes seems remarkably inattentive to what has any robust history of good consequences. One prominent strand of today's utilitarianism is useless not because it is obsessed with consequences, but because it largely ignores them.

Yet, my target is not utilitarianism *per se*. I aim simply to ask: what would it be like to theorize about the nature and causes of good outcomes? How would theorizing about outcomes be unlike theorizing about acts? Parametric games (like *solitaire*) involve one decision-maker, one player. Strategic games (like *poker*) involve several decision-makers. The distinction matters here insofar as, in a *parametric* world, outcomes are straightforward consequences of the acts we choose. You decide on an outcome—whether to have a pawn at K4—and that is the end of it. But wait! Did moving your pawn to K4 give your partner reason to move her knight to K4? In a parametric game, there are no partners, so the question never arises. In a strategic world, by contrast, you may *imagine* you can simply decide to have a pawn at K4, but it only looks that way until the next player moves.

This is a call for theorizing about the morality of a strategic world.

2. SACRIFICE

In Peter Singer's words, "If it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it."² This principle seemingly requires something. But what? Singer specifies that, on his favored interpretation, the principle requires "reducing ourselves to the level of marginal disutility," which means, "the level at which, by giving more, I would cause as much suffering to myself or my dependents as I would relieve by my gift. This would mean, of course, that one would reduce oneself to very near the material circumstances of a Bengali refugee."³

² Singer 1972: 230.

³ Singer 1972: 241.

Most readers consider it common sense that morality does not oblige people to sacrifice so much. Morality is demanding, but not maniacal. *This is not my concern.* Instead, what I find odd is that a consequentialist principle, ostensibly specifying what we should aim to accomplish, would be interpreted as specifying what we should aim to sacrifice. We should maximize benefit, which in practice means, maximize cost. What would lead anyone to see this as straightforward?

One factor—not my focus in this chapter—is that this is, after all, what many of us were taught. We were taught that we are sinners by nature. Ethics was a project of proving to God that we have done all we can to have clean hands. The point of ethics was to sacrifice enough to be above reproach.⁴ I learned to do penance, so that I would be lovable, so that God would know I was serious. I was raised to see something like giving to a point of marginal disutility as literally unquestionable. Interestingly, heretics who have the nerve to ask for an argument find in Singer only this: “I will henceforth take this assumption as accepted. Those who disagree need read no further.”⁵

A second factor, which *is* my focus, is that in a parametric world, there can be a simple translation of inputs into output: what you accomplish can be a simple function of what you sacrifice. Singer concludes his essay by saying, “What is the point of relating philosophy to public (and personal) affairs if we do not take our conclusions seriously? In this instance, taking our conclusion seriously means acting upon it.”⁶ In 1972, taking Singer’s conclusion seriously meant focusing on your input: on the variable you control.⁷ His principle sees actions available (give versus don’t give) and asks which has more utility. If giving has

⁴ See Tosi and Warmke (forthcoming).

⁵ Singer 1972: 231.

⁶ Singer 1972: 243. See generally Badhwar 2006.

⁷ The essay is as clear as any philosophical essay ever written, seeming to leave no room for doubt that these direct quotations accurately reflect the essay’s thesis. But then Singer adds that even if we replace his favored interpretation with something more moderate, “it should be clear that we would have to give away enough to ensure that the consumer society, dependent as it is on people spending on trivia rather than giving to famine relief, would slow down and perhaps disappear entirely. There are several reasons why this would be desirable in itself” (Singer 1972: 241). He does not say consumer society disappearing entirely would be good for Bangladesh and presumably understands that it would not be. But (for reasons he finds too obvious to mention) it nevertheless is desirable in itself. I suppose the most charitable reading of this remark is that it must not be as clear as it looks.

more utility than not giving, then give. *Keep* giving until stopping has more utility than further giving.

Give regardless of whether you have already given. What you have already given means precisely nothing when it comes to justifying what you do next.

3. STRATEGY

What else matters, besides how much you are in a position to give? What matters *more*? My answer requires briefly touching base with game theory. The Prisoner's Dilemma (Fig. 1) models a key insight into the human condition. Here is the example from which the name derives: You and your partner Jane face criminal charges, and ten years in jail. You need to decide whether to betray Jane by testifying against her, or to keep silent. The district attorney makes you an offer: Supposing Jane keeps silent, you get a ten-year sentence reduction if you testify, or a nine-year reduction if you also keep silent. Alternatively, supposing Jane testifies against you, you get a one-year reduction for testifying, or zero reduction if you keep silent. Jane has been offered the same deal. For each of you, then, keeping silent optimizes your *collective* sentence reduction, while testifying optimizes your *individual* sentence reduction. Testifying is a dominant strategy: each of you *individually* is better off betraying the other (one year better off here) no matter what the other does.

Game theory predicts that in a Prisoner's Dilemma, individually rational players will defect, and thus fail to make the most of their potential as cooperative social animals. In reality, we often avoid this outcome by devising ways to make partners answer for their choices. For example, we contrive to play repeated games. In a repeated Prisoner's Dilemma, you still choose to cooperate or defect, but repeated play allows for strategic play. You can play "tit-for-tat" (moving as Jane

		Jane	
		Betray (AKA <i>defect</i>)	Keep Silent (AKA <i>cooperate</i>)
You	Betray	1, 1	10, 0
	Keep Silent	0, 10	9, 9

Fig. 1. Prisoner's Dilemma

Note: The payoffs are ordered pairs (Yours left, Jane's right). Numbers are years of sentence reduction.

moved in the prior round, responding to cooperation with cooperation and to defection with defection), in that way reciprocate, and thereby make it pay for Jane to cooperate. Defecting pays more than cooperating in a one-shot game, but *reciprocated* cooperation pays more than reciprocated *defecting* in the long run.⁸ Reciprocators see fellow players as players, and treat strategic environments as if they were strategic.

This bit of game theory can help us to distinguish real from spurious ideals. Arguably, there is an ideal strategy in a Prisoner's Dilemma: reciprocity. Why? Because (1) it matters whether your partner cooperates, yet (2) you do not choose whether your partner cooperates. However, (3) you can make it pay for your partner to cooperate. This logic makes reciprocity an ideal strategy in a repeated Prisoner's Dilemma.⁹ By contrast, no such logic favors unconditional giving. Unconditional giving is a spurious ideal because instead of making it pay to cooperate, unconditional giving makes it pay to free-ride. In strategic settings, working on Jane's *payoff* by giving unconditionally has limited (often negative) value. What *pays* is working on Jane's *strategy*, by reciprocating. If you aim to do some good, you work to induce cooperation, not free-riding.¹⁰

In a discussion of weakness of will, David Estlund says being unable to will my own cooperation has no bearing on whether I ought to cooperate: "*can't do* is requirement-blocking but *won't do* is not."¹¹ This is fine in the solipsistic case. But then Estlund extends his point to strategic contexts. In a Carens Market, everyone is taxed in such a way that everyone ends up with equal disposable income after taxes, and yet, despite this, everyone works hard to maximize gross income. It sounds unpromising, to put it mildly, but as Estlund rightly notes, the supposition "that we shouldn't institute the Carens Market because people won't comply with it, doesn't refute the theory" that people should comply.¹² So, if that is off-target as a refutation, could anything be *on* target? Here is one place where distinguishing between solipsistic theory

⁸ See Axelrod 1984.

⁹ There are, of course, harder problems than the Prisoner's Dilemma. In a tragic commons, we face an influx of new *players*, making it much harder to teach partners to cooperate (Schmidtz 2008: chs. 11 and 12).

¹⁰ I allow, as would Axelrod, that there are, after all, nonstrategic situations—early child rearing, to give one example—where providing a free ride is the relationship's point.

¹¹ Estlund 2011: 217.

¹² Estlund 2011: 217.

and theory for political animals has some bite. The strategic issue for me (let's say) as a political animal is not my faux-inability to command my own will but rather this perfectly real fact: commanding my partners' will is nowhere to be found in my set of strategic options. Commanding my partners' wills is a garden-variety example of exactly the kind of "can't" on my part that absolutely does rebut a presumption of "ought."

To be clear, Estlund is right about parametric cases: if moving my pawn to K4 would be ideal, but I am intensely averse to so moving, my aversion has nothing to do with whether K4 is the ideal move. But a fundamental contrast: although *my* reluctance to move to K4 has *nothing* to do with whether moving to K4 is ideal, my inability to choose my partner's response has *everything* to do with whether moving to K4 is ideal. That I do not choose for everyone is *the* political fact of life. It cannot be assimilated to the faux-inability involved in my reluctance to move my pawn.

I may imagine how ideal moving my pawn to K4 would be, but to chess players such so-called imagination is the classic failure of imagination. Real imagination does not overlook how other players will respond. In a strategic world, solipsistic methods of identifying ideals are grossly unimaginative. It takes imagination to be a realist. A player who anticipates what can go wrong is the one whose imagination chess players admire. Imagining what *would* be ideal in a parametric world is no substitute for being able to imagine what *is* ideal in a strategic world.

4. WHICH NUMBERS COUNT?

We live among agents who decide for themselves. Some of our moral reasons stem from the fact that ours is a strategic world: people *respond* to us as if we were agents. They anticipate how we respond to circumstances; they treat our anticipated response as part of *their* circumstances, and react accordingly. Wanting to make a difference in our world, true humanitarians do their homework. Seeing millions on the edge of famine, result-oriented humanitarians see that this is not a story about them. The point of the story is not for them to end up being the hero. Accordingly, they ask the question at the heart of how some societies have made famine a thing of the past: namely, what enables farmers

to develop and successfully act on an ambition to feed customers by the millions?

The intuitive case for utilitarianism in the form of Singer (1972), which I (along with Woodward and Allman 2007)¹³ call parametric utilitarianism, involves an empirical premise: Other things equal, the action with the highest parametrically-represented utility leads straight to the outcome with the highest actual utility. This empirical premise not only is not a priori; in strategic situations it is not even true. Consequentialist morality in a strategic setting is not a matter of a solitary agent picking whichever action has the biggest number. It is about being ready to walk away from the biggest number, because what counts is not numbers attaching to acts, but numbers attaching to outcomes, where outcomes are consequences not of particular acts but of patterns of cooperation.¹⁴ (That is, outcomes are particular cells in the matrix, but the only act available to you is a choice of row.) Where Jane decides for herself how to respond, and where an ideal response is cooperative, consequentialism holds you responsible for inducing cooperative responses, not for per se choosing them.

In a social world, the problem is not that we are too weak-willed to pull a lever that would end world famine. The problem is, we live in a

¹³ I once said:

Utilitarians sometimes model morality as the sole player in a parametric game. Utilitarian morality, so represented, maximizes utility by treating human agents as if they were otherwise inert pawns to be moved at will by the game's one true player, and thus as entities that could straightforwardly be *directed* to act in a utility-maximizing way. This approach makes sense as yielding prescriptions for ethical play in parametric worlds. In our world, however, human agents are players, each with their own ends, each making their own decisions, each somewhat responsive to how others are playing. Whatever an institution's purpose, it will not serve its purpose simply by directing human agents to serve its purpose. (Schmidtz 1995: 167).

Woodward and Allman (2007: 185) independently draw the same distinction:

Strategic consequentialists recognize that when they make moral decisions they are typically embedded in an ongoing interaction with other actors who will respond in complex ways that are not easy to predict, depending on the decision-maker's choices, and that these responses will in turn present the original decision-maker with additional decisions and so on—in other words, that they are involved in a complex repeated game of some kind. Strategic consequentialists thus tend to be sensitive to incentives that their choices create, to the informational limitations and asymmetries they face, and to the opportunities for misrepresentation these create, and also to considerations having to do with motives and intentions, since these are highly relevant to predicting how others will behave.

¹⁴ Interestingly, in the case of an omission, we don't necessarily equate the outcome with what the omission caused.

world of levers, and whether to pull those levers is mostly someone else's call. We hope for/aim at/work toward outcomes. But we do not choose outcomes. If you care about consequences, you make sure you understand the difference between choosing an act (a row) and choosing an outcome (a particular cell). You can work toward an outcome, but only if you play strategically and find a way to make it pay for Jane to mind the larger consequences of her choice of column. To say ideally we would not need to take this into account would be like saying ideally we would not need to drive defensively. It is a remark about a world whose problems are not like ours.

5. DEMANDING TOO LITTLE

Laura Valentini says the big objection to Singer's principle "is not that it is impossible to act on this principle, but that acting on it would make one's life much less rewarding."¹⁵ I agree that this is a typical response, suggesting that unconditional giving "is not feasible. It is therefore naïve, and ineffective, to hold existing societies to account on the basis of such demanding moral standards."¹⁶

My own worry is not the sense in which unconditional giving demands too much, but the sense in which it demands *nothing*. David Estlund supposes, "prime justice might be utopian, in the sense that the standards are so high that there is strong reason to believe they will never be met."¹⁷ But what if utopian justice is not a high standard? What if the utopian standard is so low that we could meet it without solving a problem, making anyone better off, or doing anything that needs doing? Singer (1972) doesn't demand that we move players toward mutual cooperation; it demands only that we contribute. That isn't good enough.

We need higher standards for what we call a high standard. If I move to K4 and my partner's response is devastating, it is not impressive for me to say, "I have high standards, but evidently the world isn't ready for them. My partner could have done what would have made K4 a winning move, but human nature is too flawed for that."

¹⁵ Valentini, forthcoming in Vallier and Weber.

¹⁶ Valentini (2012: 659) is here making a point about unrealizable ideals in general.

¹⁷ David Estlund, forthcoming in Vallier and Weber.

Further, if we sought to articulate ideals relevant to a strategic world—to say what is worthy of aspiration in a strategic world—we ought to have ideals that can learn from experience. To take an ideal seriously is to treat its content as provisional. Maturing thinking about the content of the ideal can reveal earlier thinking to have been juvenile. We need not treat ideals as untouchable. Real ideals evolve.

Putting the social and therefore strategic nature of the human condition front and center, we provisionally could start by acknowledging that a society's basic structure (formal institutions, culture, and informal norms that structure what people learn for expect from each other) just is an invitation to play in a certain way. A basic structure is an incentive structure. We then ask what high standards for incentive structures would be like. To pick an incentive structure is to pick the compliance problem and the consequent pattern of behavior that predictably goes with it. We can treat a basic structure and its characteristic compliance problem as conceptually separable, but we misunderstand the nature of a basic structure when we talk as if they can be picked separately—as if we can set aside as a distracting detail the price of picking structure x . There is only one thing to pick: to pick the structure is to pick the problem. To pick a bad problem is to pick a bad structure.

High standards are for judging an incentive structure in terms of how people respond to it. The ideal structure is the one that gets the ideal response, but ideal theorists sometimes talk as if the ideal structure were the one that makes the ideal demand in abstraction from how people respond to it. But moral theorizing is not a game you win by having the most demanding theory.

6. FAMINE-PROOF COMMUNITY

If human welfare depends on variables other than how you choose to act, those other variables could be overwhelmingly more important. There is no consequentialist reason to assume the natural subject matter of true consequentialism is your case-by-case decision-making. It would depend on circumstances. Here is an empirical possibility: your case-by-case decision-making might be uninteresting—inconsequential—from that sort of consequentialist perspective.

Consequentialism could start by investigating which variables have a history of mattering. Why are fewer people starving today than in 1972?

Which ways of organizing communities have a history of making famine a thing of the past? History is a complex, poorly controlled experiment, but its lessons are clear enough when it comes to detecting which communities secure reliable access to food, even in the face of periodic shocks that otherwise have lethal consequences.

That sort of consequentialist takes an interest (as Singer currently does) in the history of improving global trends. In the arena of world hunger, we are trending in a good direction. But happy trends come at the end of stories about what induces, expands, and sustains patterns of cooperation in a social world. Singer rightly acknowledges (in conversation in 2013) that the percentage of people starving, even the absolute number of people starving, has fallen since 1972.¹⁸ Something in the world is, first, ramping up food production and, second, making producers more consistently effective at getting food to consumers who need it. Something is expanding the scope of society as a mutually advantageous cooperative venture: advances in finance (micro-banks), in communication (cell phones, the Internet), in transportation (global container shipping), and so on.

I sometimes say, moral institutions are the ones you want your children to grow up with. As an empirical observation, the kind of research we do when we care is empirical research. (If you are helping a son or daughter choose a car or college, or deciding whether to forgo chemotherapy, you want information.) Amartya Sen won his Nobel Prize in Economics partly for his empirical work on the history of twentieth-century famines. He reports that famines never occurred in democracies, at least not in market democracies,¹⁹ and never were caused by lack of food. Natural disasters could push a population over the edge, but natural disasters are not what force a population to live on the edge in the first place. Famine is caused not by eroding soil but by eroding rights.

¹⁸ Data gathered by the United Nations on the number and percentage of the world's undernourished are as follows (see <http://www.fao.org/hunger/en/>): 1990–1992, 1015 million (19%); 2000–2002, 930 million (15%); 2006–2008, 918 million (14%); 2009–2011, 841 million (12%); 2012–2014, 805 million (11%). The UN estimated the number of undernourished in 2015 at 795 million. (<http://www.fao.org/news/story/en/item/288229/icode/>).

¹⁹ Sen classifies countries such as Nigeria and Zimbabwe, where famine did occur, as nominal democracies.

There are ways of structuring, refereeing, and playing the game that lead to war and famine and corruption on a genocidal scale, while other ways lead to peace and prosperity. What Sen learned was that the rules of famine-proof countries don't stop farmers from producing food and shipping it to places where they can get a good price for it. Famine-proof rules acknowledge that farmers have for generations been gathering and updating information regarding how to produce, store, transport, and sell particular crops in particular places. No one is more interested or more capable than farmers are when it comes to getting things done. Famine-proof rules don't take decisions out of their hands. Famine-proof rules don't route decisions through offices of distant Brahmins: people who may never have met a farmer, and for whom the thought of caring about a farmer (or about anyone born into that low a caste) would be foreign. Yet famine-proof rules are not anarchic. While they don't presume to pick people's destinations, they *do* manage traffic.²⁰

7. SOLIPSISM

We do not need to know whether moral institutions work necessarily, work perfectly, or are legally guaranteed to work. We don't need to know what *would* work under imaginary conditions if only we had no need to confront the strategic reality of life among agents who decide for themselves. We do need to know this: what has a history of enabling people to work their way out of pits of famine?

David Hume aimed "to introduce the Experimental Method of Reasoning into Moral Subjects."²¹ Epistemology analogously spent centuries trying to get out from inside your head, searching for tools to refute solipsism on its own terms and prove you are not dreaming. There

²⁰ Note that the principle of utility has been applied to several subject matters: which acts to choose, which rules to respect, etc. An underappreciated fact: the principle competently applied to one subject matter never "collapses" into the principle as applied to some other. So long as the principle is applied to rules, not acts, and the point is to study which rules have better consequences than others, the theory is rule-utilitarianism. Do the best rules allow exceptions? That is an empirical rather than conceptual matter. (But imagine deciding case by case, at each intersection, which assignment of respective colors to "go" and "stop" lights are optimal for *this* intersection. The principle of utility itself would say we are asking the wrong question. The principle itself is sensitive to the need of motorists for simple, general, mutual expectations *not based on case by case reasoning*.) See Schmitz 1995: ch. 7.

²¹ This is the subtitle of Hume's *Treatise of Human Nature*.

was an evidence-based alternative: study belief formation from outside. Ask which ways of acquiring and processing information are conducive to forming accurate beliefs. Like Hume, Adam Smith inquired into the nature and causes of the wealth of nations, studying what makes some societies prosperous and, in effect, famine-proof. Smith set aside the egocentric question of “how much does morality ask of me?” and went straight to the question of what has historically been conducive to prosperity in the world as it contingently is. (Singer has been studying this too, to his credit.)

Viewed through the lens of morality as a social phenomenon, our task is not only, indeed not mainly, to decide what to expect from ourselves. Our task is to discover what our world (our laws, our neighbors, and so on) expects from us, and critically evaluate those expectations—asking which of the evolving patterns of cooperation and coordination in which we are embedded are observably making our world a better place. A good person, a good political animal, is among other things a good citizen. Good citizens take their society’s basic structure seriously. On one hand, this is the opposite of being uncritical; on the other hand, it’s also the opposite of trying to settle what morality demands by deriving a priori limits from the analysis of the terms.

Philippa Foot once said, “When anthropologists or sociologists look at contemporary moral philosophy they must be struck by a fact about it which is indeed remarkable: that morality is not treated as an essentially social phenomenon.”²² We are theorizing about the part of morality that is an essentially social phenomenon when we investigate what the citizens around us actually expect from each other and which of those expectations actually are helping people to get out of, and stay out of, pits of misery. On my view, those historically vindicated mutual expectations make up the social category of genuine moral obligation.

8. BEYOND SHALLOW POND

One of philosophy’s most famous intuition pumps is Singer’s SHALLOW POND: “If I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. This will mean getting my

²² Foot 1978: 189.

clothes muddy, but this is insignificant, while the death of a child would presumably be a very bad thing.”²³ This is the intuitive motivation for Singer’s principle. Most people find it compelling. Some scholars try to hedge the intuition with further intuitions regarding separate personhood, personal projects, or agent-centered prerogatives. But for argument’s sake, what can an unrelentingly pure consequentialism say about SHALLOW POND?

It should say this: SHALLOW POND is a parametric situation. There is one player. The game is not repeated. Cooperation is not at issue. Reciprocated cooperation pays better than reciprocated withholding in the long run, but there is no long run in SHALLOW POND. There is no one who needs to be taught to reciprocate. Precisely because strategy is moot, what you need to do in SHALLOW POND is obvious. Wade in. Save that baby. Then get on with your life. You most likely will never be in that situation again, and hardly any of us have been in such a situation even once.²⁴

But note: SHALLOW POND is not world famine. There is no “end of story” when it comes to famine. The story of hunger will never be a story that ends with you wading in, saving the day, then getting on with your life. In a real human life, if I literally pulled a baby out of a pond, I would get up the next morning to a life of my own.

If Singer were to say morality is something beyond that, something obscured by SHALLOW POND, he would be right. Yes, being moral is about stepping up when emergencies like SHALLOW POND fall into our laps, but another part of the human condition is a moral responsibility beyond SHALLOW POND: namely, the challenge of embracing a cause. There is nothing arbitrary about the fact that we cannot function except within a framework of goals and constraints. At our best, we are undistracted. At our best, we focus on one goal at a time, and pursue it within constraints.

Crucially, the world itself is not constraining enough to give us a tractable framework for humanly rational choice. We impose constraints from inside so as to have problems we can handle. So, we give ourselves budgets: a month to find a house, a thousand dollars for our Las Vegas weekend. Limiting a given pursuit leaves room in our lives for other things, acknowledging that we have more than one goal and that we would not be better people if instead we were monomaniacal. Humanly

²³ Singer 1972: 231.

²⁴ But see Stroud’s (2013) call for an alternative to this seemingly “concessive strategy.”

rational choice is choice for essentially compartmentalized choosers, as is humanly moral choice. We stipulate constraints that help us fabricate the compartmentalized structure of separate pursuits that add up to a recognizably human life. Such constraints limit our pursuits even while helping us to be more or less undistracted within them.

It would be grossly counter-productive to think each day needs to focus on the same thing as the day before. The optimal number of projects for human beings is not necessarily one; neither is it typically one. If one felt compelled to work on the same project every day, one would be a model of neither rationality nor morality. Even worse—more oblivious to the descriptive fact of our separate personhood—would be imagining everyone has compelling reason to join us in focusing on the single target that happens to grip us on a given day.²⁵

Social worlds likewise are thick with arbitrary limits. Why would my community set limits at thirty miles per hour or eighteen years of age? The details can be arbitrary. Yet, we live better lives together when we know what to expect from each other. We discover, inherit, and as needed fabricate a framework of limited expectations (of each other and of ourselves) so we can *afford* to be social beings. Between nothing and too much is a point where we are responsible for choosing our own way (or ways—Singer is allowed more than one, as are we all) of making sure our world is better off with us than without us.

9. STARTING OVER

Where is the theoretical framework that makes a natural place for such limits? What would a habitable consequentialist theory look like? My answer in 1995: We could see the part of morality that is essentially social—the part that does not start with me—as largely a duty to respect roles assigned to us by institutional arrangements that work. I won't try to reconstruct the details.²⁶ But I did say this: some utilitarians find it

²⁵ Speculation: whatever truth there is in liberalism is an *intra-* as well as inter-personal truth. There is no reason to believe we all should have the same destination, and likewise no reason to believe Singer should wake up every morning with the same destination he had the day before.

²⁶ The following summarizes Chapter 7 of *Rational Choice and Moral Agency*. It now seems like archeology to me, but it locates the roots of my current thinking about the purpose of moral theory.

mysterious that morality would incorporate any constraints beyond a requirement to maximize the good.²⁷ Notice, however: that's only a mystery from the inside. From the outside, there is no hint of mystery regarding why moral institutions constrain individual action. If the good is to be realized, then institutions—legal, political, economic, and cultural institutions—must put persons in such a position that their pursuit of the good in a predictably partial manner is conducive to the good's production in general.

For example, if you ask why we need a law against murder, as opposed to a law requiring agents to minimize the number of murders, there is a simple answer. Legal institutions have their own unique way of minimizing the number of murders. How does our legal system play its special role in minimizing the number of murders? *By making murder illegal.*

Consequentialist institutions constrain the goal-directed actions of individuals as a means to an end—namely the end of making it safe for people to trust each other.²⁸ Consequentialist institutions exist in a strategic world. Hospitals, for example, serve their purpose in part by being safe. Hospitals save lives not by standing ready to sacrifice one patient to save five but by enabling people to see hospitals as places where patients can count on being treated as having rights.

Suppose we are serious about the requirement that “If it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it.” Suppose we take this principle *very* seriously: sending goods to where they do more good and sending *bads* to where they do less harm. We ship our food to wherever it extends their life expectancy more than it would ours, and we ship them our toxic waste, too, whenever it cuts their life expectancy less than it would ours. Could that be wrong? On what grounds? What is the other part of morality—the part that can

²⁷ See Kagan 1989: 121–7. See also Scheffler 1982: 129.

²⁸ A referee cautions me: readers will deem it analytic that consequentialist institutions aim at maximizing the good, not at making it safe for people to trust each other. Trusting the referee, at the risk of belaboring the point: whether a consequentialist plumber aims to maximize the good cannot be stipulated, because whether adopting that aim maximizes the good is an empirical question. A plumber may do more good by sticking to fixing the pipes. A surgeon may do more good if patients can trust her to focus on how to save them, not when to sacrifice them. Some roles and institutions have utility precisely by leaving the maximizing to someone else.

trump the imperative to “prevent something bad from happening whenever we can do so without sacrificing anything of comparable moral importance?” What is the nature and scope of that part of morality?²⁹

It is a historical fact that our only successful experiments in building functional societies involve treating rights as robust enough that people can count on their rights precluding case-by-case utilitarian calculation. Surgeons lack the right to sacrifice the optimal number of patients, for example, because that would collapse the ceiling of our global potential by turning surgeons into people whom we would not trust. Institutions that work enable us to count on our hopes and dreams being respected rather than collectively promoted; it turns out that the respect, not the promotion, is what empirically enables us to live together and trust each other, and work in the cooperative way that is actually, not merely theoretically, mutually advantageous.

A morality that serves purpose *x* as it works through institutions is one that induces the game’s genuine players—human agents—to act in ways that serve purpose *x*. In particular, if moral institutions serve the common good, they do so by inducing human agents to act in ways that serve the common good. The only institutions that have ever done this are institutions that put people in a position to pursue their respective goods in peaceful and constructive ways.

I argued in 1995 that an institution serving the common good is a matter of enabling and encouraging agents to choose destinations of their own that are apt to make the world a better place.³⁰ To serve the common good, the institution must neither ignore the nature and limits of individual rationality nor try to stamp it out; rather, it must function so that people’s strategic responses to the institution (and people’s

²⁹ Perhaps the most common response a reader will find to critiques of utilitarianism is that a utilitarian theorist can easily handle any possible criticism simply by changing the theory (adding extra maximals as needed, say). The puzzled critic is left wondering how insisting that a theory *can* change counts as responding to the charge that it *needs* to. To say the least, I have no problem with reinventing consequentialist theory so as to make it self-consciously and principally a response to the strategic reality of the human condition. But saying consequentialism could be thus reinvented is no substitute for actually reinventing it. For real progress in the direction of such reinvention, see Regan 1980, anything by Russell Hardin, such as Hardin 1991, and, as per note 14, Woodward and Allman 2007 or Schmidtz 1995.

³⁰ That is, better from just about anyone’s perspective, without making it worse from anyone’s perspective, or at least without making it worse by virtue of being exploitative—the details are tricky. See Schmidtz 1995.

subsequent responses to each other's responses) have the effect of promoting the common good.³¹

10. TAKING STOCK

1. There is a literature on whether Singer's Principle is too demanding, and on whether utilitarianism as a decision procedure leaves room for personal projects.³² If I had nothing to say beyond joining that fray, I would not bother. I certainly would not argue that Singer made any simple, small, obvious mistake. If Singer made a mistake, it is so large that we need to step back to see it. But therein lies what interests me. Thus, I do not say Singer's argument fails on its own terms.

Instead, I ask: what other terms are there? Are there terms that do not obscure what makes it okay, even exemplary, to focus on famine one day, on factory farms the next, on how laboratory animals are treated a day later, and on tending to an ailing mother the day after that? What makes it right to find our own way of making the world a better place, to indeed find several ways, feeling no need to regiment all of our days under the banner of a single project? Where is the theory that draws the line in such a way that not every life (including Singer's!) falls on the wrong side of the line? What makes it exemplary not to take our marching orders from anyone's theory, *not even our own*?

2. Whatever we make of accusations that utilitarianism demands too much, I am struck that there is no literature on what the theory incarnated in Singer's Principle fails to demand. It fails to demand a response to the human condition. It fails to track what has a history of being a foundation of thriving community. It is an empirical matter which social phenomena—which patterns of rules and expectations—are functional enough to command the respect implicit in deeming them

³¹ I thank a referee for cautioning me that readers will think this depends on how we define the common good. Not so. What I say here will be true on any definition of the common good I have seen, so long as we stick to the same definition throughout. Schmidtz 1995, however, teases out various nuances of a proper definition.

³² See Williams 1985 and Nagel 1986.

moral.³³ It is an empirical question how demanding they are. But whatever social morality turns out to be, it does not go to heroic lengths to fool me into thinking that morality starts with me.³⁴

3. A consequentialist might care enough about famine to theorize about something other than acts. A consequentialist can ask: why are some societies, but not others, famine-proof? The legacy of the social science launched by Hume and Smith boils down to the idea that what has massively good consequences, ends famine, and consequently is morally binding, are patterns of cooperation and mutual expectation that actually—observably, not hypothetically—are in place, facilitating cooperative ventures for mutual advantage. Respecting such conventions and expectations makes us fit for society.³⁵ So, even if social morality as depicted in Singer (1972) were in some way too demanding, toning down its demands would miss my point. I can't tell how demanding social morality is by asking what I need to do in order to have a clean conscience. I learn how demanding my social morality is by evaluating the traffic management scheme in which I live. *If* that scheme is observably making its world famine-proof, then it commands our respect.

Observably, basic structures that have a history of working largely trust people to be the separate agents they are, trusting them to mind the businesses that add up to a famine-proof society. As a general observation, functional structures tread lightly when it comes to dictating destinations.

³³ If we said whether *persons* command respect is a matter of how functional they are, that would be illegitimate. Suffice it to say, I do not assume institutions are persons. Schmidtz 1995 goes into detail.

³⁴ For admirable reassessments of morality's cosmopolitan demands, see Miller 2010 and Moellendorf 2002.

³⁵ Singer nowadays asks people to give not to a point of marginal disutility but to give perhaps 1 percent of their income, and build on that if, as Singer plausibly predicts, giving that much turns out to enrich the giver's life. From a utilitarian perspective, 1 percent may have nothing to do with the truth about how much we ought to give, but if Singer's job is to maximize how much he can get us to give, 1 percent might be the most influential thing he could say. Without meaning to be critical, I observe only that moral theory as usually understood is an attempt to articulate truths about morality rather than to influence behavior. From a listener's perspective, our having reason to believe what a speaker is saying typically presupposes that the speaker aims to tell the truth.

Yet, morality is not one-dimensional. The personal strand of morality, only alluded to here, demands more than the social strand, starting with its demand that we choose a destination and throw our lives at it (each of us deciding for ourselves what that means). The personal strand does *not* come strictly from outside. It is social only insofar as it demands that I take seriously my social nature. It concerns what I need to do to be treating myself with respect while operating within the observably demanding yet not suffocating confines of functional social morality. I've spent my career on that topic: locating the demands of morality's personal strand in relation to the demands of morality's essentially social strand. I'll say no more about it here.

I do not intend to colonize ethics and turn it into a social science. Neither do I feel any nostalgia for the Scottish Enlightenment. Still, we recently became and will remain badly over-specialized. As a result, following our premises to their logical conclusions will continue to seem bizarrely misleading to people who aim to get things done, if we go on treating the question of how to live as if we can set aside, as if it were the province of a different field, the fact that when we ask how to live, we are asking how to live as political animals.

4. I close with a further speculation, in this case about deontology, the other main protagonist in our pantheon of introductory theories. Can deontology adjust to the demands of a strategic world?

Think of it this way. In a Prisoner's Dilemma, we could will that everyone reciprocate. Could we also will that everyone cooperate unconditionally? The difference between the two strategies is momentous, so it had better not be invisible to a universalization test we hope to use to make moral decisions. Yet unconditional giving will misleadingly appear as universalizable as reciprocity if we interpret universalizability by imagining a single chooser representing the construct "all people together," as if the very essence of moral choice involved ignoring our separate agency.

In a strategic world, interpreting universalizability in solipsistic terms—imagining oneself choosing between everyone cooperating and everyone declining to cooperate—is not universalizable. We cannot universalize a test that is blind to the vast strategic difference between reciprocity and unconditional cooperation. A strategic deontology

acknowledges that we cannot universalize ignoring the fact that the exercise's point is to identify maxims fit for members of a kingdom of players. You choose how to live among ends in themselves, but ends in themselves are agents—beings who decide for themselves. Therefore, in a strategic world, imagining yourself choosing for everyone is nothing like imagining yourself choosing for people *in a situation relevantly like yours*. The very essence of your situation is precisely this: you do not choose for everyone.

So, my proposal: treat strategic deontology as an alternative to “act-deontology” and envision choosing among strategies, not actions. Do not see yourself as choosing among action-maxims “I should cooperate” versus “I should free-ride.” Instead, characterize alternative strategy-maxims as “I should encourage partners to cooperate” versus “I should encourage partners to free-ride.” Now you see that what is properly universalizable is acting so as to teach your partners to grasp their place in a kingdom of ends and thereby mature in the direction of moral worth. Teach them to cooperate.³⁶

Acknowledgments

This project is supported by a grant from the John Templeton Foundation. Opinions expressed here are those of the author and do not necessarily reflect views of the Templeton Foundation. For helpful discussion, I thank Jacob Barrett, Justin D’Arms, Liu Dongdong, Lisa Downing, David Estlund, Adam Gjesdal, Don Hubin, Brian Kogelmann, Win-chiat Lee, Keith Lehrer, Terry McConnell, Doug Portmore, Greg Robson, Abe Roth, Gary Rosenkrantz, Paul Ruffner, Bas Van der Vossen, Brandon Warmke, Michael Zimmerman, and two reviewers. I thank Shane Courtland for hosting me at the Center for Ethics and Public Policy at UMinn-Duluth, and Allan Silverman for hosting me at

³⁶ I doubt that this move solves all the puzzles that could be raised for the deontological approach. It does, however, seem to go some way toward getting us past certain “indeterminacy of description” problems in articulating the proper form of maxims as the subject matter of the universalizability test. Consider also the extent to which Kant himself was alive to a need for strategic deontology, insofar as we evaluate the universalizability of promise-breaking not merely by asking whether we can will everyone to stop taking their own promises seriously but also whether we can will a world in which people stop taking each other’s promises seriously.

the Mershon Center for International Security Studies at Ohio State. Finally, I thank Mark Timmons for his encouragement and counsel since I first conceived this project.

REFERENCES

- Axelrod, Robert (1984). *The Evolution of Cooperation* (New York: Basic Books).
- Badhwar, Neera (2006). "International Aid: When Giving Becomes a Vice," *Social Philosophy & Policy* 23: 69–101.
- Estlund, David (2011). "Human Nature and Limits (If Any) of Political Philosophy," *Philosophy & Public Affairs* 39: 207–37.
- Estlund, David (Forthcoming). "Prime Justice," in *Political Utopia: Contemporary Debates*, ed. Kevin Vallier and Michael Weber (New York: Oxford University Press).
- Foot, Philippa (1978). *Virtues and Vices and Other Essays* (Berkeley: University of California Press).
- Hardin, Russell (1991). "Acting Together, Contributing Together," *Rationality and Society* 3: 365–80.
- Kagan, Shelly (1989). *The Limits of Morality* (New York: Oxford University Press).
- Miller, Richard W. (2010). *Globalizing Justice* (New York: Oxford University Press).
- Moellendorf, Darrel (2002). *Cosmopolitan Justice* (Boulder, CO: Westview Press).
- Nagel, Thomas (1986). *The View From Nowhere* (Oxford: Clarendon Press).
- Regan, Donald H. (1980). *Utilitarianism and Cooperation* (New York: Oxford University Press).
- Scheffler, Samuel (1982). *The Rejection of Consequentialism* (New York: Oxford University Press).
- Schmidtz, David (1995). *Rational Choice & Moral Agency* (Princeton: Princeton University Press). Revised Kindle edition (*Rational Choice and Moral Agency*: 2015) available from Amazon.
- Schmidtz, David (2008). *Person, Polis, Planet* (New York: Oxford University Press).
- Singer, Peter (1972). "Famine, Affluence, and Morality," *Philosophy and Public Affairs* 1: 229–43.
- Smith, Adam (1984 [1790]). *Theory of Moral Sentiments* (Indianapolis: Liberty Fund).
- Stroud, Sarah (2013). "They Can't Take That Away From Me: Restricting the Reach of Morality's Demands," *Oxford Studies in Normative Philosophy* 3: 203–34.
- Tosi, Justin and Brandon Warmke (Forthcoming). "Moral Grandstanding," *American Philosophical Quarterly*.
- Valentini, Laura (2012). "Ideal vs. Non-ideal Theory: A Conceptual Map," *Philosophy Compass* 7: 654–64.
- Valentini, Laura (Forthcoming). "The Messy Utopophobia versus Factophobia Controversy: A Systematization and Assessment," in *Political Utopia: Contemporary Debates*, ed. Kevin Vallier and Michael Weber (New York: Oxford University Press).
- Vallier, Kevin and Michael Weber (eds.) (Forthcoming) *Political Utopia: Contemporary Debates* (New York: Oxford University Press).
- Williams, Bernard (1985). *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press).
- Woodward, James and John Allman (2007). "Moral Intuition: Its Neural Substrates and Normative Significance," *Journal of Physiology—Paris* 101: 179–202.