After Solipsism

DAVID SCHMIDTZ

Here are two mundane facts. First, we are social beings. Second, we live among people who decide for themselves and among people whose decisions affect each other. These facts make morality what it is. They also make morality unlike what our current moral theories say it is. Or so I argue in this essay.

1. CHOOSING AN ACTION IS NOT LIKE CHOOSING AN OUTCOME

Because people decide for themselves, and because outcomes often are a function of what other people decide, we typically do not choose outcomes in the way that we choose actions. If we could choose outcomes in the way we choose actions, then consequentialist morality would reduce to a question of choosing whichever outcome maximizes utility. This premise, that choosing acts is the same thing as choosing outcomes, led to consequentialist morality being construed as it came to be construed in the twentieth century.

Some of the ways in which the premise is false are consequential. Consider that what you do affects people. You affect their payoffs, as we say in game theory. But we misunderstand the human condition if we suppose our ways of affecting people are all ways of affecting their payoffs. In particular, we affect how people *behave*, and sometimes how we affect their behavior will be consequential. That changes everything.

Scottish Enlightenment cared enough about humanity to study what has a history of working. Today's act-utilitarianism, by comparison, can seem remarkably inattentive to which of our ways of dealing with each other have 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.

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."[1] This principle requires something. But what? Singer specifies that, on his favored interpretation, the

---

[1] Singer 1972: 230.

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."[2]

Most critics consider it common sense that morality does not ask that much of us. 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 come to be interpreted as specifying what we should aim to sacrifice. We should maximize benefit to others, which in practice means … maximize cost to ourselves. How could anyone see this as straightforward?

Here is one factor. What Singer says is, after, all, what many of us were taught to believe. 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 forgivable.[3] I get that. Raised Catholic, I learned to do penance so that I'd be lovable, so that God would know I was serious. Interestingly, heretics who have the nerve to ask for an argument for the imperative to maximize personal cost find in Singer two thoughts. First, Singer says, suffering is bad; those who disagree need read no further. Second, we should act to minimize suffering whenever we can.[4] Singer offers the second thought as a "point" that is almost as uncontroversial as the idea that suffering is bad. Singer presents his "point" as an intuition, not an argument, but he defies readers to deny that it resonates. My honest reply is that, raised as I was, I get it. No question.

Here is another temptation us to equate what we want to accomplish with what we want to sacrifice. We are trained to brutally oversimplify consequentialist morality. In a parametric world, we can imagine a simple translation of inputs into output. What you accomplish can be a simple linear function of what you sacrifice. Singer sees two available actions (give vs. don't give) and asks which has more utility. If giving has more utility than not giving, then give. *Keep* giving until stopping would have more utility.

Give regardless of whether you have already given. What you gave in the past means precisely nothing when it comes to deciding how much more to give.

---

[2] Singer 1972: 241.
[3] See Tosi and Warmke (2016).
[4] Singer 1972: 231.

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. Bear with me.

The distinction between parametric and strategic games is as follows. 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? Answer: in a parametric game, there are no partners, so the question never arises. In a strategic world, by contrast, although you may *imagine* you can simply decide to have a pawn at K4, the reality is that it only looks that way until the next player moves.

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 up to ten years in jail. You need to decide whether to betray Jane by testifying against her. The prosecutor makes you an offer: If Jane keeps silent, you get a ten-year sentence reduction if you testify, or a nine-year reduction if you also keep silent. Alternatively, if Jane testifies against you, you get a one-year reduction for testifying, or zero reduction if you keep silent. Jane has received the same offer.

The essence of a Prisoner's Dilemma is that, collectively, all players are better off cooperating, while individually, each player is better off defecting. In the paradigmatic case, for each of you, 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 in this example) no matter what the other does.

	Jane	
	Betray (AKA *defect*)	Keep Silent (AKA *cooperate*)
Betray	1, 1	10, 0
You		
Keep Silent	0, 10	9, 9

Fig. 1. Prisoner's Dilemma[5]

Game theory predicts that individually rational players will defect in such cases, and thus fail to

---

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

realize their potential as cooperative social animals. In reality, we avoid this by devising ways to hold partners accountable for their choices. So, we contrive repeated games. In a repeated Prisoner's Dilemma, you still choose to cooperate or defect, but repeated play enables strategic play. You can play "tit-for-tat" (moving as Jane 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, yet *reciprocated* cooperation pays more than reciprocated defecting in the long run.[6]

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.[7] 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, it does not pay to work on Jane's *payoff* by unconditionally giving. What pays is working on Jane's *strategy*, by reciprocating. If you aim to do some good, you work toward an outcome: you act in ways that lead to people cooperating, not to people free-riding.[8]

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: having acknowledged the force of the Kantian "Ought implies can" dictum, Estlund goes on to urge us to keep in mind that, "*can't do* is requirement-blocking but *won't do* is not."[9] 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 unrealistic, but as Estlund notes, the supposition "that we shouldn't institute the Carens Market because people won't comply with it, doesn't refute the theory"

---

[6] See Axelrod 1984.

[7] 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).

[8] There are, of course, nonstrategic relationships—early child rearing, say—where allowing a free ride is the point. More generally, problem-solving is not always co-extensive with encouraging cooperation or discouraging free riding. There are times when you might see what a cooperative effort is adding up to and you might see that cooperation is an unnecessarily complicated means to an end, and then you might say, "Everyone, just relax. Let's not complicate this. I'm just going to write the check, fix the problem, and be done with it. We don't need to be working out the fine-grained details of our respective fair shares." I thank Meg xxx for discussion.

[9] Estlund 2011: 217.

that people should comply, or that the Carens Market is after all a moral ideal.[10]

So, if the undisputed fact that the Carens Market will not work is off-target as a refutation, could anything be *on* target? Here is one place where distinguishing between solipsistic theory and theory for political animals has bite. The strategic issue facing me as a political animal is not my faux-inability to command my own will but rather my perfectly real inability to command my partners' will.

To be clear, Estlund is right about parametric cases: if I am intensely averse to moving my pawn to K4, my aversion has nothing to do with whether K4 is the ideal move. So far so good. But in a social world, 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. This is not analogous to, and is in no way illuminated by, the faux-inability involved in my visceral reluctance to move to K4.

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. True imagination has everything to do with anticipating 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.

Crucial point: Imagining what *would* be ideal in a parametric world is no substitute for being able to see what *is* ideal in a strategic world.

4. WHICH NUMBERS COUNT?

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 acknowledge that this is not a story about them. The point is not for them to be the hero. Result-oriented humanitarians ask the fundamental question—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?

---

[10] Estlund 2011: 217.

The intuitive case for utilitarianism in the form of Singer (1972)[11] involves a hidden empirical premise: Other things equal, the *act* with the highest payoff (represented non-strategically, as in a one-shot payoff matrix) leads straight to the *outcome* with the highest actual utility. This premise 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 the action with the biggest number. Genuinely result-oriented consequentialist morality in a strategic world is about standing ready to walk away from the biggest number, because what counts is not numbers attaching to acts, but numbers attaching to outcomes. Outcomes are consequences not of particular acts but of patterns of cooperation.[12] (Check Fig.1. *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 her ideal response is cooperative, a result-oriented consequentialism holds you responsible for doing what you can to *encourage* her cooperative response. You cannot flat-footedly *choose* it.

In a social world, the problem is not that we lack the will to pull the lever that would end world famine. The problem is, we live in a world of levers, and whether to pull those levers is mostly someone else's call. We hope for. We aim. We work toward. 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 cell). You can work toward an outcome, but the way to do that is to play strategically and find a way to encourage Jane to mind the larger consequences of her corresponding choice of column.

---

[11] I once wrote:

> 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 exchange with other actors who respond in complex ways that are not easy to predict. These responses in turn present the original decision-maker with additional decisions. Strategic consequentialists thus tend to be sensitive to incentives that their choices create, to informational limitations and asymmetries they face, to opportunities for misrepresentation, and also to considerations having to do with motives and intentions, or anything relevant to predicting how others will behave.

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

5. DEMANDING TOO LITTLE

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."[13] But what if utopian justice is not a high standard? What if the utopian standard is so appallingly low that we could meet it without solving a problem, without making anyone better off, and without doing anything that needs doing? Aiming high, at a minimum, involves aiming to solve a problem. So, I want to stress that my worry about unconditional giving is not the sense in which unconditional giving demands too much, but the sense in which it demands *nothing*. Singer's Principle does not aim high. It is not even a response to what is after all a strategic problem. It is a way of imagining how great it would be not to have a strategic problem. Singer's 1972 principle doesn't demand that we solve the problem of learning how to move players toward mutual cooperation; it demands only that we contribute. For anyone who cares about consequences, 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 game-ending, it is not impressive for me to say, "I have high standards, but the world isn't ready for them. My partner could have done what would have translated my move to K4 into my vision of the ideal outcome, namely a victory for me, but human nature is too flawed for that." Imagine me saying that to chess players. They would infer that I am mocking high standards, because they live in that world and they know what taking the game seriously is like.

To articulate ideals relevant to a strategic world—to say what is worthy of aspiration in a strategic world—we need "smart" ideals. In contemporary engineering jargon, "smart" ideals learn from experience. To take an ideal seriously is to treat its content as provisional. Mature thinking about the ideal's content can reveal earlier thinking about the ideal to have been juvenile. Serious theorizing about ideals does not treat ideals as unresponsive to the logic of the problem. We learn as we go. Smart ideals evolve.

High standards are for judging an incentive structure in terms of how people respond to it. Moral theorizing is not a game you win by having the theory that demands the biggest sacrifice. The activity of moral theorizing is more demanding and needs to be more imaginative than that.

6. FAMINE-PROOF COMMUNITY

If human welfare depends on variables other than how you choose to act, those other variables could

---

[13] David Estlund, "Prime Justice," *Political Utopias*, in Vallier and Weber (2017) at 47.

be overwhelmingly more important. There is no consequentialist reason to assume that the natural subject matter of true consequentialism is your case-by-case decision-making. That would depend on circumstances. Consider this empirical possibility: your case-by-case decision-making might be uninteresting—inconsequential—from that sort of consequentialist perspective.[14]

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.

Consequentialism so construed takes an interest (as Singer to his credit 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.[15]

Number and percentage of world's undernourished:

Years	Number	
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 estimates the number of undernourished in 2015 at 795 million.[16]

Over that span, something was ramping up food production and, second, making producers more effective at getting food to consumers who need it. Society as a mutually advantageous cooperative venture was scaling up: advances in finance (micro-banks), communication (cell phones, the Internet, the "app"), transportation (global container shipping), etc...

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 decline chemotherapy, you

---

[14] To David Hume, a curious feature of justice is that it involves eschewing case-based reasoning in favor of being guided by principles. Thus, yielding the right of way might have zero utility at best in a given case while remaining useful or agreeable as a general practice, by virtue of enabling us to know what to expect from each other.

[15] Data gathered by the United Nations (see <http://www.fao.org/hunger/en/>).

[16] <http://www.fao.org/news/story/en/item/288229/icode/>.

want information!) Amartya Sen earned his Nobel Prize partly for his work on 20[th] century famines, showing that, amazingly, not one was caused by lack of food. Natural disasters can push a population over the edge, but they do not force a population to live on the edge in the first place. Famine is caused by eroding rights, not eroding soil. When local farmers lose the right to choose what to grow or where to sell it, they lose everything, and that is when people starve.

There are ways of structuring, refereeing, and playing the game that lead to war, 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. 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 anyone born into that low a caste) would be foreign. Yet famine-proof rules are not anarchic. While famine-proof rules don't presume to pick people's destinations, they *do* manage traffic.[17]

7. SOLIPSISM

What has a history of enabling people to work their way out of pits of famine? 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 "what 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.) Our task is not only, indeed not mainly, to decide what to expect from ourselves. Our task is to discover which of the evolving patterns of cooperation and coordination in which we are embedded are *observably* making our world a better place.

---

[17] 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 of utility 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 so long as we keep in mind that the point is to study which rules work better than others, the theory is rule-utilitarianism. Do the best rules allow exceptions? That is an empirical matter. Such matters typically cannot be settled by thought experiments.

However, this one can. 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 fact that what motorists need to know is that motorists around them are respecting simple, general, mutual expectations *not based on case by case reasoning*. See Schmidtz 1995: ch. 7.

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."[18] We are theorizing about the part of morality that is essentially social when we investigate what citizens around us actually expect from each other and which of those expectations actually are helping people to get out and stay out of pits of misery. On my view, those historically vindicated mutual expectations make up the social strand 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 clothes muddy, but this is insignificant, while the death of a child would presumably be a very bad thing."[19] This is the intuitive motivation for Singer's principle. Most people (including me) find it compelling. Some scholars try to qualify the intuition with further intuitions regarding separate personhood, personal projects, or agent-centered prerogatives. But for argument's sake, suppose we forget about ways of restricting consequentislism and straightforwardly ask, 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 that situation even once.[20]

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. By contrast, in a real life, if I literally pull a drowning baby out of a pond, I will get up the next morning to a life of my own.

So, I agree that being moral is about stepping up when emergencies like SHALLOW POND fall

---

[18] Foot 1978: 189.
[19] Singer 1972: 231.
[20] But see Stroud's (2013) call for an alternative to this seemingly "concessive strategy."

into our laps. At the same time, 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. Thus, we give ourselves budgets: a month to find a house, a thousand dollars for our Las Vegas weekend, another thousand for charitable giving. Limiting a given pursuit leaves room in our lives for other things. Budgets acknowledge that we have more than one goal and that we would not be better people if instead we were consumed by a single goal.

There is a reality here undreamt of by contemporary utilitarian moral theory. Namely, humanly rational choice is choice for essentially compartmentalized choosers. 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. Humanly moral choice is not oblivious to this central fact about what it takes to respond rationally to the human condition.

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. Even worse would be imagining that everyone has reason to join me in focusing on the single target that happens to grip me at this moment.

The truth in liberalism is *intra-* as well as inter-personal. There is no reason to believe we all should have the same destination, and likewise no reason to believe Singer should wake every morning with the same destination he had the day before. What makes it okay, even exemplary, for Singer 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—to find *several* workable ways—of making it good that we lived here, feeling no need to regiment all 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 to not take marching orders from anyone's theory, *not even our own*?

Further, social worlds are thick with arbitrary limits. Why does my community set limits at

30 miles/hour, or 18 years of age? Details seem arbitrary, yet we live better lives when we know what to expect from each other. We discover, inherit, and often *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

Some theorists find it mysterious that morality would incorporate any constraints beyond a requirement to maximize the good.[21] Notice, however: that is 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? Answer: *By making murder illegal*.

Consequential institutions constrain goal-directed actions of individuals. They impose constraints as a means to an end—namely the end of making it safe for people to trust each other.[22] Consequential institutions serve 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.[23]

---

[21] See Kagan 1989: 121–7. See also Scheffler 1982: 129. 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.

[22] A referee cautions me: readers will deem it analytic: consequentialist institutions aim at maximizing the good, not at making it safe for people to trust each other. Trusting the referee, I will risk belaboring the point: whether consequentialist plumbers aim to maximize the good cannot be stipulated, because whether adopting that aim has utility is an empirical question. A plumber may do more good by sticking to fixing the pipes. A surgeon may do more good by focusing on saving patients. Roles and institutions can have utility precisely by leaving the maximizing to someone else.

[23] Just to be clear, the idea of making room for the "suberogatory" is not my concern here. I am not aiming to make room for a category of the permissibly suboptimal. Instead, I am theorizing about "what works?" and not "what to do?" When people expect x from each other as a descriptive fact, and when we can see the further normative reality that they would be worse off if they did not expect x from each other, then we can see that x is not arbitrary. People have some reason to go along with x. Namely, expecting x works. —and not to treat x's normativity as needing to be underwritten by a philosophical theory. x is not suberogatory. Neither is it supererogatory. Neither is x optimal. Rather, it is the

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."[24] In 1972, taking Singer's conclusion seriously meant focusing on your input: on the variable you control.[25] 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 people's life expectancy more than it would ours, and we ship people our toxic waste, too, whenever it cuts their life expectancy less than it would cut ours. Could that be wrong? By what standard?

There is a crucial theoretical question here. Namely, what is the other standard—the other part of morality—that can 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 other part of morality?[26]

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 sacrificing patients 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 is the respecting, not the promoting, that enables us to trust each other enough to live

---

mutual expectation that is serving the common good in that particular time and place. I thank Nathan xxxx for helpful discussion.

[24] Singer 1972: 243. See generally Badhwar 2006.

[25] The essay is as clear as any philosophical essay ever written, seeming to leave no room to 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. Rather, it is desirable in itself.

[26] 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 maximands 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 progress in the direction of such reinvention, see Regan 1980, anything by Russell Hardin, such as Hardin 1991, and Woodward and Allman 2007 or Schmidtz 1995.

together, and work together in a way that is actually, not just 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.

To have utility, an institution must neither ignore individual rationality nor try to stamp it out; rather, it must function so that people's strategic responses to the institution (and people's subsequent responses to each other's responses) have the effect of promoting the common good.[27]

10. TAKING STOCK

1. There is a literature on whether Singer's Principle demands too much, and whether utilitarianism as a decision procedure leaves room for personal projects. If I had nothing to say beyond joining that fray, I would not bother. I do not say Singer made a small, obvious mistake. If Singer made a mistake, it is so huge that we need to step back to see it. We need to step back to see how different it would be to take a break from asking what to do and instead ask what works. The difference would be somewhat like Alvin Goldman asking what works when it comes to reliably forming true beliefs, and in the process re-inventing epistemology. Epistemology spent centuries trying to escape from inside your head, searching for tools to refute solipsism on solipsism's own terms, and prove you are not dreaming. There was an evidence-based alternative: study belief formation from outside. Ask which ways of acquiring and processing information are conducive to forming accurate beliefs.

2. Whatever we make of accusations that utilitarianism demands too much, I am struck that there is no literature on what Singer's Principle fails to demand. It fails to demand a response to the human condition. It fails to ask what has a record of fostering human flourishing. It is an empirical matter which patterns of rules and expectations are functional enough to command the respect implicit in deeming them moral.[28] But whatever social morality turns out to be, it does not go to heroic lengths to fool me into thinking that morality begins with me.[29]

---

[27] 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. But for further discussion, Schmidtz 1995 teases out various nuances of a proper definition.

[28] 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.

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

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, and ends famine, are patterns of cooperation and mutual expectation that observably, not hypothetically, are in place, facilitating cooperative ventures for mutual advantage. Respecting such conventions and expectations makes us fit for society.[30] So, even if social morality as depicted in Singer (1972) were in some theoretical way too demanding, toning down its demands would miss my point. I cannot tell how demanding social morality is by asking what I need to do to have a clean conscience. Instead, I learn how demanding my social morality is by observing and evaluating the traffic management scheme in which I live. *If* that scheme observably is making its world famine-proof, then that scheme (and the set of rules it encodes, other things equal) commands my respect.

 Observably, basic structures—the ones that have a history of working—largely trust people to be the separate agents they are, and 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.

 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, as a responsible moral agent, 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.

4. I close with a further speculation, in this case about deontology, the other main protagonist in our pantheon of moral theories about what to do. Deontologists regard 'What can be universalized?' as a foundational question. My thought: in a strategic world, the solipsistic interpretation of universalizability—imagining a choice between everyone cooperating and everyone declining to cooperate—is not universalizable. We cannot universalize being blind to the vast strategic difference between reciprocity and unconditional cooperation. A strategic deontology acknowledges that moral

---

[30] Singer nowadays asks people to give not to a point of marginal disutility but to give, say, 1% 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% 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% might be the most influential thing he could say. Without meaning to criticize, I observe only that moral theory as usually understood is an attempt to articulate truths about morality rather than to influence behavior. So, if the truth is that we should give 1%, where does that leave Singer's principle?

deliberation'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. They decide for themselves.

In a strategic world, imagining yourself unilaterally making *the* choice between everyone cooperating and everyone defecting is nothing like imagining yourself choosing for everyone in situations *relevantly like yours*. The relevant *essence* of your situation is that you are not choosing for everyone. So, 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 that difference had better not be invisible to a universalization test we hope to use to make serious decisions. Yet unconditional giving will misleadingly appear to be as universalizable as reciprocity if we interpret universalizability by imagining a single chooser representing a construct "all people together," as if the very essence of moral choice involved ignoring our separate agency.

My proposal here is to treat *strategic* deontology as an alternative to "act-deontology" and thus to envision choosing among strategies, not among 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 the moral truth in deontology: what is 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.[31]

Acknowledgments

---

[31] I doubt that this move solves all imaginable puzzles for the deontological approach. It may, however, 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 by asking whether we can will a world in which people stop taking each other's promises seriously.

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 & Limits (If Any) of Political Philosophy," *Phil & Pub Affairs* 39:207—37.

Estlund, David (2017). "Prime Justice," in Vallier and Weber. 35-55.

Foot, Philippa (1978). *Virtues and Vices and Other Essays* (Berkeley: U of California Press).

Hardin, Russell (1991). "Acting Together, Contributing Together," *Rationality & 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).

Regan, Donald H. (1980). *Utilitarianism and Cooperation* (New York: Oxford University Press).

Scheffler, Samuel (1982). *The Rejection of Consequentialism* (NY: Oxford University Press).

Schmidtz, David (1995). *Rational Choice & Moral Agency* (Princeton University Press). Revised Kindle edition (*Rational Choice & 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 & Brandon Warmke (2016). "Moral Grandstanding," *Phil & Public Affairs* 44:197-217.

Vallier, Kevin & Michael Weber, eds. (2017) *Political Utopias* (NY: Oxford Press).

Woodward, James and John Allman (2007). "Moral Intuition: Its Neural Substrates and Normative Significance," *Journal of Physiology—Paris* 101: 179-202.