View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by The Australian National University

Published in *Fair Priority Setting and Rationing at the Bedside*, ed. Marion Danis, Samia Hurst, Anne Marie Slowther, Reidun Forde and Leonard M. Fleck (Oxford: Oxford University Press, 2015), pp. 224-37.

Chapter 13

The Collective Action Problem*

ROBERT E. GOODIN

A classic collective action problem has the following structure. Each actor would be better off if everyone were to perform a certain action. But each actor would be even better off than that if everyone *except her* were to perform that action. Each one of them is thus tempted to let the others perform the action, while not doing so oneself. Yet each of the others, being identically situated, does the same. So no one ends up doing it at all.

That is the "tragedy of the commons" (Hardin, 1968). The tragedy lies in the fact that an outcome that would have been better for all concerned, if only they could have organized to act collectively in pursuit of it; but that outcome is virtually impossible to obtain through uncoordinated private action (Olson 1965; Hardin, 1982; Ostrom, 1990).

For a familiar example, consider the case of fisheries. All fisherfolk would be better off if all of them restricted their catch to sustainable levels. Each of them, however, has a private motive to catch more than that. Yet if

* I am grateful to Marion Danis, Alan Wertheimer, and Dave Wendler for discussion of these issues and to Len Fleck and Annette Rid for comments on a previous draft of this chapter.

all of them do so, overfishing will lead to the exhaustion of the fishery. None of them want that to happen: quite the opposite. But no one of them (and no small group of them) can, by desisting, prevent that from happening. So each acts on his private motive, and overfishing ensues.

Real world collective problems are vastly more complicated than those examples might imply. Even as regards fisheries, restrictions might rightly concern not merely how many fish you can catch but what kinds, what sizes, and where. Likewise, restrictions on the use of a common pasture might have to stipulate a trade-off schedule among various different animals that might be put out to graze there, or specify tradeoffs between grazing and cultivating the common lands. In medical applications, the schedule for restricting allocation of resources according to conditions and prognoses will be even more complicated, by many orders of magnitude, than that. Without meaning in the least to downplay the difficulty of coming up with a remotely defensible schedule, I merely observe that the structure of the problem is the same across all these cases.

In connection with bedside rationing, the common pool resource that is most commonly discussed is health care funding, be it in the form of a fixed national healthcare budget or some smaller pot of funds in the case of more decentralized systems. (Asch and Ubel, 1997; Ubel and Goold, 1997). As such the issue is labeled 'cost containment.' Of course, just how much importance should be assigned to containing costs of health care is contentious. Some would insist that "economic... considerations are not germane to ethical

medical practice" at all (quoted in Aaron and Schwartz, 1984, p. 127). Indeed, Wyller has argued this point in chapter 15. Even those not disposed toward that extreme view must nonetheless acknowledge that pressure on the health budget could be reduced in at least two ways: one would be by reducing demand, through rationing; another would by increasing supply, through increased healthcare funding or making more effective use of available resources.

Cost containment looks like a classic common pool resource only if we regard the health budget as strictly fixed. That, many would say, is not something we should easily concede. Perhaps in the long term we should not. Perhaps we should instead campaign for increased health care funding, or even an open-ended commitment to meeting health care needs regardless of cost. Perhaps we should campaign for improved effectiveness of the medical use of resources. Those are important long-term projects that may eventually bear fruit.[1]

In the short term, however, matters are typically otherwise. The supply of health care resources available at any given time is ordinarily strictly fixed and limited. Where that is so, the supply of those fixed resources would indeed be a common pool resource, and bedside rationing to maximize the health benefits derived from it would pose a collective action problem of the sort just described.

[1] So too is improving health worldwide. But for purposes of this chapter I set aside issues of global justice and concentrate on issues of health care within one country.

The bulk of this chapter will indeed discuss bedside rationing primarily in terms of cost containment. But it bears emphasizing that it is not purely a matter of cost containment alone. Collective action problems of a strictly analogous sort arise in other respects as well. Let me offer just two other health-related examples, to suggest the generality of the issues involved.

Consider, for one other example, the role of overprescribing in leading to the emergence of antibiotic resistant bacteria. All patients have a common interest in that not happening. Assuming there is any chance at all that the drug might be even minimally efficacious, each patient has an interest in being prescribed the drug by his own physician. Each patient hopes that all other physicians exercise restraint, while his own does not. Yet if each gets what he wishes for in his own case, no restraint will be exercised by any, and the efficacy of the antibiotics will be eroded for all.

For another example, consider a vaccination that carries some slight risk of adverse consequences. Any given patient might prefer to avoid that risk by not receiving the vaccination, trusting that enough of the others around him will have been vaccinated to provide "herd immunity" (Anderson and May, 1985). But of course if everyone does that then there will be no herd immunity. If each physician respects the wishes of her own patient in this regard, then the common pool resource is extinguished.

Supposing the health budget to be strictly fixed, the same is true of bedside practices of prescribing treatments that are somewhat less efficacious but substantially less expensive, or refraining from the use of marginally

effective treatments altogether. Each patient hopes his physician will provide him with the very best treatment available, whatever the cost, while hoping that physicians prescribe lower-cost treatments (or no treatments) to others to protect the health budget that is paying for his treatment. But again, if each gets what he wants in his own case, then the overall health budget is seriously eroded.

Collective action problems lie at their heart of all of those cases. Much though each wishes she could be the sole exception, each would be better off if everyone's physician respected the general rule of treating everyone (herself included) as she ideally wishes her physician to treat everyone except her. The trick lies in how to secure that outcome, in the absence of some external enforcement.

The problem is that, if left to her own devices, each physician is naturally tempted – some would say (wrongly, I shall go on to argue) "professionally obliged" (Levinsky, 1984) – to promote the best interests of her own patient. Doing that in the context of a collective action problem would, however, mean deviating, in each and every case, from the general rule that would be best for all.

I. Particular Interests vs the Common Good

Collective action problems arise from the constellation of preferences of each of the actors involved and from the way in which those interact. Specifically, collective action problems arise from agents choosing to act on the basis of some "particular interest" rather than on the basis of the "common good" of all (Runciman and Sen, 1965).

In the case of the fisherfolk, that "particular interest" took the form of a self-regarding private motive: increasing their own catch and hence their own profits. In the case of the physicians as discussed above, that "particular interest" is presumed to take the form of an other-regarding motive: benefiting one of their own patients.[2] What is crucial in generating a collective action problem is not whether the agent acts from self- or other-regarding motives. What is crucial is, instead, that the agent acts with a view to doing what is best not for all those who are affected by the action, but rather for some subset of them.

That fact suggests one simple solution to collective action problems. If all agents act purely with a view to doing what is best for all agents, no collective action problem arises. With anything short of "all" and "purely" collective action problems can still arise, however.

For example, a collective action problem would persist, to some greater or lesser extent, even if physicians partially internalized the "common good" alongside the "particular interests" of their own patients. The more heavily they weigh the former as compared to the latter in their decision-making, the

---

[2] In a fee-for-service system there might also be self-regarding motives for physicians. But I set that possibility aside for purposes of my discussion here.

more often the former will trump the latter for them in deciding what to do, and the less severe the collective action problem will be in consequence. But collective action problems would be precluded altogether only if physicians give absolute priority to the "common good" over the "particular interests" of their own patients, in cases of conflict.

Ignoring altogether the "particular interests" of one's own patients – even if only where those actually conflict with the "common good" of all patients taken together – simply does not come naturally to medical practitioners. It clashes with the conceptualization of the first duty of physicians as being to their patients (Levinsky, 1984; Snyder, 2012, p. 86 and passim).

II. The Permissibility of Particularism

We might query whether that is the correct way to conceptualize the duty of the physician. Maybe physicians have no duty to (or maybe even a moral duty not to) help patients pursue every one of the "particular interests" that they might happen to have. That might include even some health-related interests.

I shall analyze that issue in two steps. First I shall address the question of whether it is morally permissible for any given patient to seek special treatment for her own health care needs, over and above that devoted to

everyone's similar health care needs. I shall then address the question of whether, even if it is morally permissible for any given patient to seek such special treatment, it is morally obligatory for physicians to give any patient such special treatment.

A. From the Patient's Perspective

Remember, patients are moral agents too. They are not merely bearers of interests, desires, and impulses; they are also the bearers of moral duties. The charge to a physician to attend assiduously to the particular interests of her patients is a moral charge. And one can be morally charged to assist others only in pursuing interests that it is morally permissible for those other agents to pursue.

Thus, in asking what a physician ought to do in furtherance of her patients' interests, we have to ask what sorts of interest the patient might himself take in his own health. What will be of most concern to us in relation to collective action problems is the sort of interest a patient might take in his "relative health status," that is, in his own health status and resources devoted to his own health care in comparison to that of others. What sorts of such interests are morally permissible for the patient to take?

Here are three sorts of attitudes one might take to the health of others, compared to one's own: (1) indifference; (2) comparative advantage; (3) no comparative disadvantage.

Take first the case of someone who is indifferent to the health of others. He wants purely to maximize his own health, and he does not care one iota about anyone else's health. If all patients are of this sort, and each physician internalizes the interests of their own patients thus understood, that gives rise to a collective action problem of the sort discussed above. Such patients (and physicians acting pursuant to their interests thus specified) want to maximize their own absolute well-being, without regard to the impact of doing so on others. If others exercise restraint, such a patient (and a physician acting on his behalf) would have no hesitation in taking advantage of that fact to further his own interests.[3]

Consider next the case of someone who is not indifferent to the health status of others but rather seeks comparative advantage over others in that respect.[4] He might seek to be healthier than others as an end itself. Or he might seek it as a means to some other ends. Being more malnourished than your competitors for the same job puts you at a comparative disadvantage in the labor market (Dasgupta and Ray, 1986/7). So too does being less healthy than your competitors. Insofar as health is a means to other ends that one harbors, one might seek comparative advantage over others in the realm of health in order to gain more of those other end-use goods (Sen, 1983).

---

[3] In the terms of game theory, these represent Prisoner's Dilemma style preferences (Luce and Raiffa, 1957, pp. 95-102).

[4] In terms of game theory, these represent Status Good preferences (Shubik 1971).

Finally, consider the case of someone who is concerned merely not to be put at a comparative disadvantage from exercising restraint that others do not reciprocate.[5] Someone thus motivated simply does not want to be "played for a sucker." Experimental economists have shown that the impulse toward "strong reciprocity" is ubiquitous, not only among people in Europe and North America, but also across a wide range of cultures (Bowles and Gintis, 2002).

Among those preference structures, which count as morally permissible? It would seem clearly impermissible to seek comparative advantage over others in the realm of health – certainly insofar as that is achieved by actually worsening the health status of others. Giving others some wasting disease or preventing them from being cured of one, just so one can get a job ahead of them, is morally unacceptable. Whatever interest a patient might have in doing that, it is not an interest that her physician can be morally obliged to help her pursue.

Displaying utter indifference to the health of others might be almost as bad. Giving complete weight to one's own health interests, and absolutely none to those of others, is morally obnoxious as a failing of human sympathy, most would probably agree. But giving somewhat more weight to one's own health than that of others might strike us as morally permissible. And remember, the collective action problem in view arises, to some greater or

[5] In terms of game theory, these represent Assurance Game preferences (Sen 1967).

lesser extent, whenever patients (and physicians acting on behalf of them) give any extra weight whatsoever to their own particular interests.

Consider finally the third stance: merely not wanting to be played for a sucker. That may not be the most attractive stance that is morally conceivable. Morally, surely it is good to do good for others, whether or not they do good for you. But while that is certainly morally good, not all actions that are morally good are morally obligatory for you always to perform. Some are supererogatory, above and beyond the strict call of duty, the stuff of saints and heroes. In other cases, the duty to perform a good action takes the form of an imperfect duty, one (like the duty of charity) that you should perform on some occasions but you need not perform on all occasions.

It is unclear precisely what description best fits the duty people (or physicians acting on their behalf) have to refrain from using health resources that would yield more health benefits if devoted to others. But as I have already said, there seems to be no strict duty to do so. It seems morally permissible for people to weigh their own interests more heavily than those of other people, at least to some extent, at least sometimes or in certain sorts of circumstances. And if it is permissible for them to do so unconditionally, it is unclear why it should not be permissible for them to do so conditionally – specifically, conditional on reciprocity, if and only if others do likewise.

B. From the Physician's Perspective

The upshot of the previous discussion is that it is morally permissible – at least sometimes, to some extent – for a patient to display a preference for his health needs to be met instead of those of other people with similar health needs. But it does not automatically follow from that fact, necessarily, that it would be morally obligatory (or maybe even morally permissible) for that patient's physician to act on those preferences.[6]

I have spoken above in terms of the physician being the agent of her patients, acting on their behalf. From the principal's point of view, the whole point of hiring an agent is to get someone to do what you want them to do for you. It is in the nature of the relationship, however, that the agent always acts with some latitude in pursuit of the principal's objectives, and she may well end up doing other than exactly as the principal would wish. It is not simply that that might happen: it may, in the sense that it is proper that it should. A large part of the point of hiring an agent is that the agent should have discretion to exercise her independent judgment in ways to pursue the principal's objectives better than the principal would have done on his own.

Furthermore, there are some things that you morally may not instruct an agent to do for you. Obviously, this includes things that you are morally prohibited from doing altogether: morally, it is wrong to hire an agent to commit a murder for you for the same reason, morally, it is wrong to commit the murder yourself. But there may also be things that would be morally

[6] There are various other situations, too, in which a physician might not be obliged to act on a patient's preferences, as when a patient requests an operation whose risks are excessive or amputation of a healthy limb in cases of body integrity disorder.

permissible for you to do yourself that it would not be permissible to instruct an agent to do for you. Morally, it may be permissible for you not to jump in and save every drowning child you see at a crowded beach yourself. But it would be morally impermissible for you, when hiring a lifeguard for that beach, to instruct her to do the same.[7]

The last two thoughts, taken together, suggest some considerable scope for conceptualizing the duties of the physician-cum-agent as being somewhat detached from preferences and particular interests of her patient. The patient may have an understandable preference for his interests to be served ahead of others'. But that does not automatically translate into a duty on his physician slavishly to do so, particularly not if the patient's own larger interests would be better served by the physician's doing otherwise in certain respects.

The latter is likely to be the case, in turn, because of the nature of the collective action problem in view. Everyone would be better off, if everyone (or, rather, everyone's physician acting on his behalf) exercised restraint. Everyone is symmetrically situated; no one can have any reasonable, realistic expectations of being treated differently than anyone else. If it must be the case either that all exercise restraint or that none do, each patient is clearly better off in the former case than the latter.[8]

---

[7] Permissions in general are not transferable. When I give you permission to enter my house, that does not give you any right to give permission to others to enter.

[8] If it is the case that some need not exercise restraint so long as others do, the strategic structure of the game changes: it is then not a pure collective action (Prisoner's Dilemma) game, but rather an instance of a Chicken game nesting within a Prisoner's Dilemma (Taylor and Ward, 1982).

III. Professional Norms Encompass All Particular Interests

As I said at the outset, collective action problems arise from people pursuing particular interests rather than the common good of all. One very standard solution to the problem is to create a collective group agent responsible for serving the interests of all and to charge that collective group agent with the task of making decisions for the group as a whole. Representing as it does the group as a whole, that collective group agent has no particular interests apart from the common good of all. The collective action problem is thus straightforwardly resolved.

That is the way Nordic countries traditionally avoided the counterproductive consequences of dispersed wage bargaining. Inflation arises from (among other things) wage competition. Trade unions strive to obtain competitive advantage for their own members in seeking higher wage rates than workers in other sectors of the economy. But if all independently pursued their particular interests in that respect, all workers would be worse off; the ensuing inflation would more than erode the purchasing power of all nominal wage gains. In Scandinavia, the solution historically took the form of a single umbrella organization (such as Sweden's LO, the Landsorganisationen) being charged with the task of negotiating wages for workers across the country as a whole. Encompassing the interests of all

workers as a whole, the LO avoided the problem of counterproductive pursuit of sectoral advantage (Olson, 1982, pp. 89-92; cf. Hernes, 1991). The common good of all was secured in this way by entrusting its pursuit to some organization whose sole interest was in the good of all.

Professions can be like that. Even if particular physicians feel duty bound to internalize the interests of their own particular patients, the profession (and associations representing the profession as a whole) can internalize the interests of all the patients of all the physicians.

Where what is good for each will be undermined by independently pursuing what is good for each, a better outcome for each can be achieved by collectively pursuing the good of all. That is the general thought.

As applied to the case of health care rationing, the thought is that every patient will be better off than he would otherwise have been if profession-wide norms of good practice in the effective use of medical resources are imposed. Such a system will be an improvement over the practice of entrusting rationing decisions to the discretion of particular physicians at the bedside of particular patients one-by-one. In bowing to those norms, the physician is doing precisely what a good agent ought to do – pursuing her principal's (i.e., patient's) interests in ways different from but better than they would be if the principal- cum-patient's own particular preference in the matter were slavishly respected.

Far from breaching their professional duties of care toward their patients, then, physicians would actually be best discharging them by

developing profession-wide standards of health-care rationing that all those physicians then all implement. Bedside rationing systematically governed by such professional norms would be not only for the good of all. It would also be for the good of each, seen from a broader perspective. Each trying to get more than that would lead, ultimately, to each getting less than that.[9]

IV. The Physician's "First Duty" – Not "Only Duty"

To say that "the physician's first and primary duty is to the patient" (Snyder 2012, p. 86) is not to say that that is the physician's only duty. Nor is it to say that that "first and primary duty" cannot itself be overridden by other weighty considerations, from time to time.

Wendler (2010) offers various examples of exceptions to the rule always to do what is in a patient's best interest. That occurs, for example, when an attending physician allows her intern to insert the central line rather than doing it herself, even though that poses somewhat more risk to the patient, on the grounds that learning-by-doing is the only way for new physicians to be properly trained.

More generally, the rule always to do what is in your patient's best interest is violated whenever you have multiple patients whose needs make competing demands on your time, attention, and resources. This happens

---

[9] That is, over a substantial run of relevantly similar cases. Of course, the Keynesian aphorism, "In the long run, we're all dead," can apply with special force in some medical cases.

most clearly in emergency medicine, whenever one patient's treatment is postponed to allow treatment of another patient with more urgent needs. In order to do what is in the latter patient's best interest, the emergency physician has to do something that is not in the former patient's best interest (postpone treatment).

Indeed, the rule always to do what is in your patient's best interest, when you have multiple patients with competing claims, involves a notorious logical fallacy familiar from Bentham's initial sloppy statement of the fundamental rule of utilitarianism as "the greatest good to the greatest number" of people. The rule, as stated, demands the maximization of two functions at once – such a double maximand is a nonsense. What does good for "the greatest number of people" might be one thing, what does the "greatest good" aggregating the good done to all people might be something else altogether. The rule always to do what is in your patient's best interest, when you have multiple patients whose interests are not necessarily perfectly aligned, involves an identical nonsense. It involves a maximization exercise that is logically impossible to perform.

The duty always to do what is in your patient's best interest is thus the physician's first but not only duty. Among other things, the physician has identical duties to her other patients. She has one duty qua "this patient's physician"; she has another duty qua "that patient's physician"; and insofar as their medical needs make conflicting claims on her, discharging those twin

duties as best she is able will typically preclude the physician from doing literally what is in the very best interests of either of the two patients.

Let us now generalize one step further, from: (1) "the duty of the physician, qua 'this patient's physician,' to this particular patient"; to (2) "the duty of the physician, qua 'these patients' physician,' to all of her patients"; to (3) "the duty of the physician, qua physician (to all patients)." The sorts of tradeoffs involved in (2) (between the competing claims of all of the physician's own patients) provide a model for the sorts of tradeoffs that are involved in (3) (between the competing claims of all patients of physicians in general) (Tavaglione and Hurst, 2012: 11). Those latter tradeoffs lie at the heart of the duty of "stewardship of resources" enunciated in the Ethics Manual of the American College of Physicians.[10]

What is undeniably true is that the sorts of tradeoffs involved in both (2) and (3) are best governed according to rules set down somewhere other than at the bedside (Snyder 2012, p. 90). At the bedside, the physician's "primary role" is indeed "as a [particular] patient's trusted advocate." But in determining what she may or may not legitimately do to best further the interests of that particular patient, the bedside physician must reflect on a wider set of laws, rules, and professional norms crafted with a view on how best to further the interests of all patients. The rule that the physician should not lie to insurance companies to maximize her patient's payout is one such

[10] "Physicians have a responsibility to practice effective and efficient health care and to use health care resources responsibly. Parsimonious care that utilizes the most efficient means to effectively diagnose a condition and treat a patient respects the need to use resources wisely and to help ensure that resources are equitably available" (Snyder 2012, p. 87, see similarly p. 90).

rule, from a very different realm (Snyder, 2012, p. 89).[11] Another is the rule that the physician should "use all health-related resources in a technically appropriate and efficient manner" and "plan work-ups carefully and avoid unnecessary testing, medications, surgery and consultations" (Snyder, 2012, p. 90). Rules like that are best made by reflecting upon the duties of the physician to all of her patients (in 2) or upon the duties of physicians as a whole to all of their patients (in 3), and then imposed as constraints on the decision-making of any particular physician at any particular patient's bedside.

Seen as a solution to a collective action problem, such profession-wide setting of rules and norms to govern bedside practice will have the effect of doing the best that realistically can be done for patients themselves. If every patient's physician tried to do better for her own patient than that, the upshot would be that everyone's patients (including, over the long term, that patient herself) would be worse off. Physicians binding themselves to a system of restraint in that way is what truly serves the best interests of their patients, individually as well as collectively, over the long haul.

V. Other Solutions

[11] At least not unless the physician is operating in an "unjust restrictive environment" (Tavaglione and Hurst 2012).

I have recommended a system of professional norms that could systematically guide bedside rationing, as my preferred mechanism for managing those tradeoffs collectively. That is not the only conceivable way to solve the collective action problem surrounding health care resources, of course. Considering the limits of the alternatives, however, I think it the best.

One class of alternatives would involve the direct regulation of medical practice by state authorities. One solution, for example, would be simply to enact legislation that restricts undesirable practices and/or imposes preferred medical practices. We might be able to do that in rare cases. We might legislate to compel all school children to be vaccinated against certain specified diseases, for example. Or in centralized healthcare systems such as the British National Health Service, authorities can simply "remove some services completely... off the list of services that they make available to patients in their regions" (Rogers and Braunack-Mayer, 2004, p. 84). But as a general strategy, that is politically unrealistic in most places. It is probably medically ill advised as well. Politicians writing general rules cannot hope to capture the nuance of particular cases; the exercise of some physician discretion, operating under looser forms of professional norms, would lead to better health outcomes (Goodin, 1982, ch. 5; Braithwaite et al., 2007).

A weaker version of that might be to require physicians to get specific permission from some central authority for prescribing particularly expensive treatment, requiring the physician to justify that treatment in the particular case. The Australian Pharmaceutical Benefits Scheme operates that way with

respect to a restricted list of particularly expensive medicines. Or maybe it might suffice just to impose a requirement for physicians to file a report to some central authority when opting for hyper-expensive treatments. Simply "knowing they're being watched" might encourage physicians to self-censor their overuse of expensive products and procedures. Such solutions might go some way toward ameliorating collective action problems, but they cannot be expected to eliminate them altogether.

A second class of alternatives works on the supply side to avoid overuse of especially expensive medical products or procedures. For capital-intensive branches of medicine, for example, you can simply not buy the machine. In that way you can guarantee that no one overuses a machine that is not available (Aaron and Schwartz, 1984, p. 128). But of course that guarantee comes at the cost of no one being able to use a machine that is unavailable, even if such a machine has been developed and marketed.

A third class of alternatives works on the demand side. Among a small and stable group of actors, collective action problems can often be reliably overcome through systems of reciprocal forbearance, enforced by breaches being punished by tit-for-tat retaliation in subsequent rounds of the game (Axelrod, 1984; Ostrom, 1990). Sometimes that might work in a health-care context. In a small community with only a handful of physicians, access to local hospital might be successfully controlled in that way. But as a general strategy for a very large municipal (much less national) community of patients and physicians, that solution is simply not viable.

A fourth class of alternatives involves partial implementation of the scheme described above for solving collective action problems through the use of all-encompassing organizations. If you want an organization that is literally all encompassing – that is responsible for looking after the interests of literally all patients – then you will need to be thinking in terms of something like the profession as a whole.[12] But partial versions of that strategy are available, and have sometimes been implemented with limited success.

An example of that is the GP Fundholder scheme introduced into the British National Health Service under the Tories in 1991 and its successor (the Primary Care Trusts) under the subsequent Labour Government. That scheme had the effect of making each medical practice internalize the interests of its portfolio of patients as a whole. Resources devoted to one of a practice's patients were resources that would not then be available to other patients of that same practice, toward whom physicians associated with that practice had just the same professional obligations.

Such schemes encourage physicians to think in terms of the common good of patients in their practice. But they still leave physicians pursuing the particular interests of the patients in their own practice, at the expense of patients in other practices. That is a flaw that is endemic to any partially encompassing solution to a collective action problem.[13]

---

[12] Or the state. The reason the state can resolve collective action problems is that it has coercive powers with which to enforce its edicts. The reason it does, when it does, is that it internalizes the interests of all those under its jurisdiction.

[13] That is not the only flaw with such schemes, of course. Worse is the risk that practices will "cherry-pick" good patients who are cheap and easy to care for, while patients with costly or complex conditions prove unable to find care.

VI. Conclusion

Physicians confront collective action problems all the time. Each vigorously pursuing the interests of each of their patients often leads to outcomes that are worse for all patients. Each physician's patients would typically be better off if all physicians exercised restraint in the use of healthcare resources. That might be accomplished, after a fashion, in any of many ways. But the best is almost certainly one that leaves room for the exercise of clinical judgment, whilst subjecting that to some collective discipline. The medical profession as a whole ought to develop norms to govern the appropriate use of scarce healthcare resources. Physicians engaged in bedside rationing can be guided by such norms, confident in the knowledge that in abiding by them they are indeed doing their best by their patients overall.[14]

---

[14] Some may be tempted to respond "there is no such thing as 'the medical profession as a whole, merely all sorts of competing sub-specialties trying to maximize resources available to them.'" In response to them, I would simply repeat the last two sentences of the text. Competition among sub-specialties for scarce resources is simply another collective action problem, in which all would benefit if all exercised restraint; and that is best accomplished via coordination orchestrated by the most all-encompassing group ("the medical profession as a whole"). Where that group is not yet collectively organized the first task, of morality and extended prudence alike, is to see to it that it gets effectively organized.

References

Aaron, Henry J. and William B. Schwartz. 1984. The Painful Prescription: Rationing Hospital Care. Washington, D.C.: Brookings Institution.

Anderson, Roy M. and Robert M. May. 1985. Vaccination and herd immunity to infectious diseases. Nature, 318 (#6044): 323-9.

Asch, David A. and Peter A. Ubel. Rationing by any other name. New England Journal of Medicine, 336 (#23: 5 June): 1668-71.

Axelrod, Robert. 1984. The Evolution of Cooperation. New York: Basic Books.

Bowles, Samuel and Herbert Gintis. 2002. Homo reciprocans. Nature 15 (#6868: Jan. 10): 125-8.

Braithwaite, John, Toni Makkai and Valerie Braithwaite. 2007. Regulatory Ritualism and the New Pyramid: Governance of Aged Care Across Three Decades in Three Nations. Cheltenham: Elgar.

Dasgupta, Partha and Debraj Ray. 1986/7. Inequality as a determinant of malnutrition and unemployment. Economic Journal 96: 1011-34 and 97: 177-88.

Goodin, Robert E. 1982. Political Theory and Public Policy. Chicago: University of Chicago Press.

Hardin, Garrett. 1968. The tragedy of the commons. Science, 162: 1243-48.

Hardin, Russell. 1982. Collective Action. Baltimore, Md.: Johns Hopkins University Press.

Hernes, Gudmund. 1991. The dilemmas of social democracies: the case of Norway and Sweden. Acta Sociologica, 34 (#4: Dec): 239-60.

Levinsky, Norman G. 1984. The doctor's master. New England Journal of Medicine, 311 (Dec 13): 1573-5.

Luce, R. Duncan and Howard Raiffa. 1957. Games and Decisions. New York: Wiley.

Mechanic, David. 1992. Professional judgment and the rationing of medical care. University of Pennsylvania Law Review, 140 (#5: May): 1713-54.

Olson, Mancur Jr. 1965. The Logic of Collective Action. Cambridge, MA: Harvard University Press.

Olson, Mancur Jr. 1982. The Rise and Decline of Nations. New Haven, CT: Yale University Press.

Ostrom, Elinor. 1990. Governing the Commons: The Evolution of Institutions for Collective Action. New York: Cambridge University Press.

Rogers, Wendy A. and Annette J. Braunack-Mayer. 2004. Practical Ethics for General Practice. Oxford: Oxford University Press.

Runciman, W. G. and A. K. Sen. 1965. Games, justice and the general will. Mind 74: 554-62.

Sen, Amartya. 1967. Isolation, assurance and the social rate of discount. Quarterly Journal of Economics 81: 112-24.

Sen, Amartya. 1983. Poor, relatively speaking. Oxford Economic Papers 35: 153-69.

Shubik, Martin. 1971. Games of status. Behavioral Science 16: 117-29.

Snyder, Lois, for the Ethics, Professionalism and Human Rights Committee of the American College of Physicians. 2012. American College of Physicians Ethics Manual: sixth edition. Annals of Internal Medicine, 156 (Jan 3, pt 2): 73-104.

Tavaglione, Nicolas and Samia A. Hurst. 2012. Why physicians ought to lie for their patients. American Journal of Bioethics, 12 (3): 4-12.

Taylor, Michael and Hugh Ward. 1982. Chickens, whales and lumpy public goods: alternative models of public-goods provision. Political Studies 30: 350-70.

Ubel, Peter A. and Susan Goold. 1997. Recognizing bedside rationing: clear cases and tough calls. Annals of Internal Medicine, 126 (#1: Jan. 1): 71-80.

Wendler, David. 2010. Are physicians obligated always to act in the patient's best interests? Journal of Medical Ethics, 36 (2: Feb): 66-70.