

## The Self-Enforcing Lottery

Antti Kauppinen

University of Amsterdam

Draft, August 25, 2009

Suppose things are very bad, so bad that some of us need to die: there simply are not enough resources for everyone to live. Being the sort of beings we are, moral but imperfectly so, we have tried to keep everyone on board, but have finally been forced to conclude that it is the way to all-around starvation. Our options are few. Either we will all die trying to keep everyone alive, or we somehow select some who will make the ultimate sacrifice for others. Clearly the latter route is the lesser of two evils. But how to select? We could let things run their course and allow the weak to perish in the war of all against all, which is sure to come when all that is left are the seeds for next year's meagre crop. This would not only be morally intolerable, but also irrational, for the best Hobbesian reasons – we all need to sleep sometimes, and no one can be sure of survival.

So instead, we should, both morally and rationally, hold a lottery that gives everyone an equal chance to survive. The elements of any lottery are the lots, the objects by means of which chances are distributed, a randomizing device that designates some lots as winning either before or after the draw, and the designed outcome, the distribution of prizes according to distribution of the lots designated as winning or losing. The actual outcome of the lottery may diverge from the designed outcome if there is cheating or mistakes involved at some stage of the process; this will be important. In ordinary lotteries, the lots are tickets, and the designed outcome is that those who have a ticket designated as winning receive some money or goods. In the

survival case, however, if you draw a winning lot, your prize is living, together with the other winners, with enough resources to go on. If you lose, you are to report at the next departing ice floe. Provided that everyone has an equal chance, this solution treats everyone with equal respect. Importantly, it is also in everyone's enlightened self-interest, given that it is the best deal we can expect other rational agents to agree to, and certainly far better than getting nasty and brutish.

Many a moralist agrees with this picture.<sup>1</sup> And there certainly is much to be said for a lottery in such circumstances. But there is also a serious challenge to it. For what happens after people have drawn their tickets? Those who find themselves with a losing ticket have nothing to lose, and nothing to gain by abiding by the rules. The consequences are predictable: it is a war of all against all after all, the very state of affairs that the lottery was designed to avoid. Worse, this is not just empirical speculation. It is the rational thing to do if you are one of the losers – even if war gives you only a small chance of survival, your expected utility will be higher than it would if you comply.<sup>2</sup> It may not be the moral thing to do, depending on what the correct relationship between morality and rationality is, but part of the argument for the lottery in this case is that it is the solution on which rationality and morality converge. Further, many philosophers who are entirely unpersuaded by Hobbesian considerations accept something like Owen Flanagan's Principle of Minimal Psychological Realism, according to which the demands of a moral ideal or theory should be such that it is psychologically possible for creatures like us to comply with them (Flanagan 1991: 32). Requiring a massive part of the population to quietly agree to self-sacrifice after drawing a losing lot, while still having a fighting chance of survival, surely violates this

---

<sup>1</sup> Survival lotteries in general are endorsed by Harris 1975 and, for example, Timmermann 2004. Lifeboat lotteries are endorsed by, for example, Broome 1984 and Rakowski 1993. My cases are inspired in particular by Mulgan (forthcoming).

<sup>2</sup> Even Hobbes would agree with this. When one's own life is the cost, no covenant can oblige.

constraint. I call this the *Compliance Problem* for life-and-death lotteries without external authority.

There is, I believe, only one possible solution to the Compliance Problem. Let me first illustrate it with an example. (I emphasize that this is *only* an example of what a concrete solution would look like, so I will avoid messy detail as far as I can.) Suppose, again, that we are living in a post-apocalyptic world in which all central authority has collapsed and crops have failed, so that some of us have to go for anyone to survive. It is not wildly unrealistic to assume that there is still one abundant resource: guns, let us say Kalashnikov-style assault rifles, and ammo. A Kalashnikov is in many ways an ideal personal weapon: extremely reliable, reasonably accurate with a nice rate of fire, and easy to use. As conflicts in the developing world have amply demonstrated, just about any man, woman, or child can employ it to deadly effect. This makes it the ideal *lot* for a lifeboat-type survival lottery. Here is one reasonably realistic way to organize it. Say that there are a hundred of us on an island that supports fifty over the next winter, as our best scientists have determined. What we do is load fifty Kalashnikovs with live ammunition and another fifty with empty clips, and randomly distribute them to a hundred huts or under a hundred trees, or whatever is the most convenient arrangement that gives everyone access to exactly one random gun.<sup>3</sup> (It is in everyone's interest to come up with some ingenious way to do this fairly.) At the agreed time, everyone goes to a hut and picks up a gun, no doubt quickly checking whether it has ammunition. If you get bullets, you are a winner. If you get an empty clip, too bad for you: you are to depart and perish. Your heart may rebel against your bad luck, but even without external authority, that is all you can do, for any one of the winners is, *by the very fact that she has drawn a winning Kalashnikov*, able to mow you down if you try

---

<sup>3</sup> For an additional dose of psychological realism, we might make the units of lottery families rather than individuals, but this and related complications are not pertinent to the main point.

anything. If there is to be a war of all against all, its outcome will match the designed outcome of the fair lottery.

To be sure, there is a wrinkle here due to the fact that a loser might pretend to be a winner to get close enough to pounce. This requires additional arrangements to make sure that the identity of winners and losers is known to all. Perhaps winners should fire a shot in the air when everyone is around. If your neighbour's gun does not fire, you know what to do. Alternatively, the lots could be clips instead of guns (which could themselves be pre-distributed without ammunition to everyone), the losing ones modified so that they will not slot in at all. (It happens to be clearly visible whether a Kalashnikov clip is in place or not.) This would have the result that everyone would wear their status in plain sight. Options are endless, and it is in everyone's interest when arranging the lottery to ensure the signalling issue is taken care of, since otherwise the lottery might as well not be held and the war of all against all would loom again.

The general solution, thus, is to *make the lots themselves the means of enforcement for the designed outcome*. This way, there is a minimal gap between conducting the lottery and bringing about the designed outcome: once the lottery has been held, its designed outcome is mutually known and automatically enforceable by the winners, so that no external authority is needed. To be sure, the gap could in principle be eliminated altogether, so that conducting the lottery would automatically *bring about* the designed outcome. For example, available technology permitting, everybody could be hooked up to wires at the same time and random fatal electric shocks administered to the right amount of people.<sup>4</sup> With my choice of weapon, the equivalent would be circular firing squads with randomly assigned live bullets. This, however, would be to dissolve rather than solve the Compliance Problem, since

---

<sup>4</sup> This was suggested to me by [reference omitted for blind review].

compliance would no longer be an issue, precisely because there is no gap within which the losers could choose either to comply or not. The problem with this alternative not rational as much as moral: it necessarily involves killing innocent people, while in the Kalashnikov-type lottery, nobody actually needs to be *killed*. Rather, the losers remove themselves from the population one way or another, as they have agreed to do *ex ante*.<sup>5</sup> The guns eliminate the *incentive* to deviate from the agreement, not the losers themselves. Thus, leaving a minimal gap leaves maximal room for free choice, and thus amounts to maximal respect for individuals as rational agents under the circumstances.

As before, it is *ex ante* rational for everyone to agree to such an arrangement in the specified circumstances, but with this method of lottery, it is also *ex post* rational for everyone to abide by the results.<sup>6</sup> Granted, in the real world Kalashnikovs or other weapons may not be found in the right amount, everyone may not be able to use them well enough, somebody's gun may jam, and so on, but in spite of their practical importance, these are relatively trifling issues, philosophically speaking. *In principle*, the solution is clear, and it seems clear that a self-enforcing lottery is the only possible solution to the Compliance Problem for survival lotteries with no external enforcement. Or is it?

---

<sup>5</sup> In most circumstances, removing themselves from the community will leave the losers with a small chance of survival. This is one reason why death by bullet need not be the absolutely certain outcome of deviating from the agreement: it is enough to ensure rational compliance that death by Kalashnikov is more likely than death by the beasts of the jungle.

<sup>6</sup> Indeed, it is more rational for people to agree to the lottery *ex ante* when they can be confident that its designed outcome will be the actual outcome.

**References**

Broome, John 1984. Selecting people randomly. *Ethics* 95, 38–55.

Flanagan, Owen 1991. *Varieties of Moral Personality: Ethics and Psychological*

*Realism*. Cambridge, Mass.: Harvard University Press.

Harris, John 1975. The survival lottery. *Philosophy* 50, 81–87.

Mulgan, Tim (forthcoming). *Ethics for a Broken World*.

Rakowski, Eric 1993. Taking and saving lives. *Columbia Law Review* 93, 1063–1156.

Timmermann, Jens 2004. The individualist lottery: how people count, but not their numbers. *Analysis* 64, 106–112.