

First Come, First Served?*

Tyler M. John and Joseph Millum

Waiting time is widely used in health and social policy to make resource allocation decisions, yet no general account of the moral significance of waiting time exists. We provide such an account. We argue that waiting time is not intrinsically morally significant, but its use is justified across a range of pretheoretically compelling scenarios. First, there is a duty of fairness prohibiting line cutting where a sufficiently just queue exists. Second, where candidates are in relevantly similar circumstances, allocating by waiting time is efficient, maximizes distribution equality relative to other Pareto efficient distributions, and approaches the fairness of an equiprobable lottery.

I. INTRODUCTION

Waiting time is widely used in health and social policy as a means for deciding to whom scarce resources are allocated. It is prominently featured in many organ allocation schemes, where time on a waiting list is one criterion determining who among prospective recipients will get the next organ available.¹ Waiting time for operations is used as a rationing device in many public health care systems.² Government-subsidized housing is distributed on the basis of various principles across different US states, but waiting time is a constant factor. ICU beds and university courses are

* Thanks to David DeGrazia, Michael Garnett, Douglas MacKay, Holly Smith, Alec Walen, and David Wasserman; our audiences at the 2016 Annual Meeting of the American Society for Bioethics and Humanities, the 2016 Mid-Atlantic Bioethics Fellows' Workshop at the University of Pennsylvania, and the 2016 New Scholarship in Bioethics Annual Symposium; two anonymous reviewers; and the associate editors at *Ethics* for their comments on this article. The views expressed are the authors' own. They do not represent the position or policy of the National Institutes of Health or any other part of the US government.

1. OPTN/UNOS Board of Directors, *Organ Procurement and Transplantation Network (OPTN) Policies*, June 1, 2017.

2. Stephen Martin and Peter C. Smith, "Rationing by Waiting Lists: An Empirical Investigation," *Journal of Public Economics* 71 (1999): 141–64.

Ethics 130 (January 2020): 179–207

© 2020 by The University of Chicago. All rights reserved. 0014-1704/2020/13002-0002\$10.00

generally filled on a first come, first served (FCFS) basis among eligible candidates for a spot, and others are placed on a waitlist. Retail stores allocate their wares and their checkout lines according to the same criteria. Meanwhile, many people seem to regard waiting time or one's place in a queue as morally relevant to one's claim to a resource. "Queue jumping" or "cutting in line" elicits moral disapproval, while the assertion "I was here first" is commonly advanced as a putative justification for having a greater claim to a good or getting the good first.

Despite the widespread use of waiting time in allocation schemes, little has been said about when and why waiting time is morally relevant when allocating scarce resources.³ The few discussions of the ethics of using queues or waiting time focus on problems of application—for example, arguing that, in practice, waiting time is likely to be associated with societal privilege and therefore ought to be excluded from consideration.⁴ But such pessimism about the use of waiting time as a basis for resource allocation is not warranted. Focusing on just these issues of application misses important prior normative questions about the justification of allocation schemes that incorporate waiting time and cannot therefore give us summary moral guidance on the incorporation of waiting time into allocation schemes. Moreover, as we argue, FCFS is in fact a just way to allocate resources in many of the cases where it seems pretheoretically compelling, and waiting time has unique normative properties which frequently justify its incorporation into resource allocation schemes.⁵

In this article, we analyze the morally significant features of waiting time. We argue that the fact that someone has waited longer in a queue for a scarce good is not intrinsically morally significant. The first person in the queue does not in virtue of that fact have a right to the good, and the length of time they have waited does not directly factor into the moral importance of giving it to them. However, waiting time can and sometimes should play a role in justifying allocation decisions. First, while there is no general right to FCFS, there is a general right not to be set back by line cutting when a sufficiently just allocation system that uses waiting

3. One notable exception is Ronen Perry and Tal Z. Zarsky, "Queues in Law," *Iowa Law Review* 99 (2013): 1595–1658.

4. James F. Childress, "Putting Patients First in Organ Allocation: An Ethical Analysis of the US Debate," *Cambridge Quarterly of Healthcare Ethics* 10 (2001): 365–76; Norman Daniels, "Fair Process in Patient Selection for Antiretroviral Treatment in WHO's Goal of 3 by 5," *Lancet* 366 (2005): 169–71; Govert Den Hartogh, "Trading with the Waiting-List: The Justice of Living Donor List Exchange," *Bioethics* 24 (2010): 190–98; Govind Persad, Alan Wertheimer, and Ezekiel J. Emanuel, "Principles for Allocation of Scarce Medical Interventions," *Lancet* 373 (2009): 423–31.

5. FCFS and allocation (purely) by waiting time are perfectly coextensive where queue participants remain in line with no breaks to waiting. We here assume that this is how all queues operate.

time already exists. Line cutting violates a duty of fairness to play by the system's rules. Second, where an allocation scheme is being created or amended, waiting time has several morally attractive features that can justify its incorporation. Where candidates are in relevantly similar circumstances, allocating by waiting time is relatively efficient, maximizes distribution equality relative to other Pareto efficient distributions, and treats candidate recipients fairly (both in the sense of giving candidates equal chances and in the sense of expressing their equal status). Our analysis implies that waiting time is only morally relevant owing to these contingent features, which must in turn be weighed against other morally relevant considerations.

If allocating using waiting time is only indirectly valuable in this way, then it might be thought that its advantages could all be obtained using an alternative system that did not have the purported disadvantages of queuing systems. In particular, some have claimed that formal lotteries are a fairer way to select among otherwise equal beneficiaries. We argue that formal lotteries are no fairer than allocation based on waiting time when it is equiprobable (epistemically or objectively) how a prospective queue will be ordered. Moreover, in practice, lotteries share many of the disadvantages of queues; which is preferable will depend on contingent features of the allocation scenario. Even where problems of application cannot completely be resolved, we conclude, schemes which incorporate waiting time are frequently the morally best allocation schemes available.

II. THE CLAIM-RIGHT INTUITION

We start with the question whether there is some intrinsic moral significance to one's place in a queue or how long one has waited. If waiting time has intrinsic moral significance, then the length of time one has waited in a queue for a resource *ipso facto* gives one a stronger claim to that resource. For this to be the case, it must sometimes be true that when every other morally relevant consideration is held fixed, the strength of a candidate's claim to a resource increases because she has waited longer.⁶ Considering certain cases motivates the idea that waiting time has intrinsic moral significance:

Utilitarian Shopper: You are in line at a grocery store with a full cart. The person behind you has only two items. Your checkout will

6. Note that it is not sufficient to show that because waiting is typically unpleasant and unpleasant experience is bad, waiting time is often morally significant. This would not demonstrate that waiting time is intrinsically morally significant, but only that it very regularly correlates with something that is intrinsically morally significant: unpleasant experience.

take five minutes and her checkout will take one minute. If you check out before her, you must wait five minutes to complete your checkout, while she must wait six minutes. But if she checks out first, she must only wait one minute and you must wait six. The shopper behind you notes this inefficiency and corrects it by cutting in front of you.

How should we feel about the actions of this grocery shopper? If someone's position in a queue only matters morally insofar as it promotes agent-neutral goods, such as efficiency, then it seems that one ought, morally, to rearrange the queue. If the person behind you moves in front, things will go much better in terms of efficiency in that it will bring about a better aggregate result. It is true that the switch would increase inequality, but since one person must wait six minutes either way, from an agent-neutral perspective preferring the more equal outcome is equivalent to leveling down by imposing an additional four minutes of waiting time. Moreover, switching positions does not seem obviously unfair, since it would be done on the grounds of efficiency, and not on the basis of arbitrary or morally irrelevant reasons. However, for most of us, if a shopper behind us unilaterally corrected such an inefficiency by cutting in front, we would be outraged.

Intuitions about cases like this support the claim that people in a queue are wronged when others cut in front of them, even when doing so leads to better outcomes. They therefore constitute evidence that when someone gets to a resource whose possession is not settled before others, they have a right to use that resource before those who got to it later. It would therefore be pro tanto morally wrong for one of these others to take the resource before them. A fortiori, it would then also be true that waiting time has intrinsic moral significance, since the fact that one party had waited longer than others would itself strengthen their moral entitlement to a good.

Someone who would defend a general right to FCFS faces two challenges. First, a proponent of such a right must provide an explanation for the right that situates it within a plausible moral theory. Second, as we explain shortly, many of our intuitions about resource allocation appear inconsistent with any general right to allocation by FCFS. This apparent inconsistency would need to be resolved in favor of the right. We now argue that this dual argumentative burden is unlikely to be met. Instead, we conclude that there is no moral claim-right to FCFS and provide an alternative account of our intuitions about cases that preserves their normative force while explaining apparent inconsistencies across contexts.

We are not aware of any published arguments that defend a claim-right to FCFS. The closest analog would seem to be the Lockean notion

that someone who first reaches unowned property can thereby stake a claim to it.⁷ People's behavior in some social situations suggests that an analogous principle might underlie attitudes to queuing. It might be possible to stake a claim to first use of a resource by being there first, to stake a claim to use the resource second by getting to the resource second, and so on.

According to standard readings of Locke, property can be acquired when labor is mixed with unowned natural resources. Along with the requirement that the property be unowned when I mix my labor with it, Locke's account includes a proviso. One can only acquire a property right in some thing when there is "enough, and as good, left in common for others."⁸ This proviso explains how the acquisition of private property can be justified to other people without their explicit consent: provided they are not made worse off by my acquisition of the property, they have no grounds for complaint.⁹

Lockean accounts apply quite naturally to the sorts of goods with whose allocation we are concerned, since Locke was concerned with how we move from a situation in which natural resources are held in common or unowned to one in which individuals have private property in those resources.¹⁰ The Lockean proviso complicates the matter. Robert Nozick interprets the proviso—that one can only acquire a property right in some thing when there is "enough, and as good, left in common for others"—to mean that a person can only acquire a property right in a thing when doing so will not "worsen the situation" of others by depriving them of something they would otherwise possess.¹¹ I cannot acquire a property right in the Atlantic Ocean, since doing so would prevent others from enjoying the ocean, participating in international trade, and so forth. But, presumably, any time allocators distribute a valuable and absolutely scarce resource, whoever acquires that resource prevents someone else from acquiring it, and thereby worsens their situation. Thus, whenever we have a genuinely scarce resource, such that not everyone who could benefit from the resource will receive it, the distribution will not leave enough and as good for others. Even in cases where

7. Compare the discussion by Perry and Zarsky regarding whether FCFS tracks desert ("Queues in Law," 1614–20).

8. John Locke, *Second Treatise of Government: An Essay Concerning the True Original, Extent and End of Civil Government* (New York: Wiley, 2014), 18.

9. Robert Nozick, *Anarchy, State and Utopia* (New York: Basic, 1974), 175–82; Gopal Sreenivasan, *The Limits of Lockean Rights in Property* (Oxford: Oxford University Press, 1995), 34–50.

10. It would not seem to apply to goods that are already privately owned. Where ownership has already been acquired, those who own the resources will surely have some latitude to decide how the resources are distributed; they are not forced on pain of injustice to allocate their goods on an FCFS basis, when they give (or trade) goods away.

11. Nozick, *Anarchy, State and Utopia*, 182.

receipt of a benefit is merely delayed by being later in the queue—like the grocery store—if there is a cost to waiting, then giving priority to the first person in line seems to make those behind her worse off. It is hard to see how any plausible interpretation of the proviso can be squared with a claim-right in the case of allocating scarce resources.

A further challenge for proponents of a claim-right is that not all of our intuitions seem to be consistent with there being a general pro tanto claim-right to FCFS. In fact, intuitions about FCFS schemes appear highly sensitive to cultural and social context. For example, in most parts of the United Kingdom bus seats are allocated on the basis of FCFS. Whoever gets to a bus stop first is first in line for a seat on the bus, and any queue jumping is highly inappropriate. If you are in the United Kingdom and you board a bus ahead of a Brit who arrived at the bus stop before you, she will feel as though she has been treated unfairly. In many parts of the United States, this principle does not apply. The person who arrives first at a bus stop in the United States does not believe that she thereby has a right to board the bus first. Indeed, if she were to try to enforce that right against other Americans, she would be greeted with bafflement. Such examples suggest that our intuitions about queuing do not always track general ethical reasons—they seem better described as tracking cultural mores. This gives us reason either to be skeptical about the veracity of the intuitions or to believe that the ethical facts in question are contingent upon cultural facts.¹²

Another example may help to show that where no social norm has been established mere waiting does not itself ground an entitlement.¹³ Suppose that the US government announces that it will distribute a scarce supply of Tamiflu in the face of an impending epidemic and appoints a blue-ribbon commission to establish allocation criteria by the end of the week. At the week's start, frightened people begin queuing outside of the site where the drug will be dispensed. We doubt that these people's mere act of queuing forces the government's hand such that they must now allocate Tamiflu according to FCFS, or even that they must assign any positive moral weight to the time that these candidates have waited in their allocation scheme. As such, we have further reason to doubt that mere waiting grounds any moral entitlements.

Second, other cases suggest that our intuitions about the right to FCFS in specific cases are inconsistent with that right being ultimately grounded in a claim-right to priority based on arrival time. Consider

12. An anonymous reviewer notes—and we agree—that our divergent intuitions may be explained by status quo bias combined with different status quos across different contexts. The justification we give in the following section for a duty to respect queuing systems is not undermined by such an explanation—these intuitions will still track the ethical facts because they are explained by the same phenomena that justify those ethical facts.

13. Thanks to David Wasserman for this example.

again the case of the checkout line. You are in a grocery store, waiting in line with a cart full of groceries, and someone with only a few groceries unapologetically cuts in front of you. If this happened to you, you would probably be angry and feel as though you had been wronged. Now consider a similar case. You are in a grocery store, waiting in line with a cart full of groceries, and someone with only a few groceries unapologetically but indirectly cuts in front of you by choosing to go through the express checkout that you are ineligible to use. In such a case, you got to the checkout line first, but you were not served first. It does not appear that people feel wronged by such shoppers or by the store's management in cases like these.

The two grocery store cases are analytically equivalent in several important ways. Both involve the same resources being distributed, include the same conditions of scarcity, and violate FCFS for reasons of efficiency.¹⁴ They differ in that in the first case the line jumper is not following generally accepted social norms, while in the second case they are. Since the two cases are analytically equivalent with respect to their violation of FCFS and both yield different intuitive judgments, if we want to preserve our judgments about both cases, we should not accept a general claim-right to FCFS.

In the next section, we defend a different moral foundation for the requirement not to cut in line. Our explanation allows that there is no general moral claim-right to FCFS and that waiting time is not intrinsically morally significant, while preserving our intuitions about most cases.

III. WHY CUTTING IN LINE IS IMPERMISSIBLE

We believe that our intuitions about our various cultural practices, about allocation schemes that circumvent FCFS for the sake of efficiency, and about the insignificance of waiting time in itself are vindicated by the following independently plausible explanation. While there is no general right to FCFS, there is a general right not to be set back by line cutting when a sufficiently just allocation system that uses waiting time already exists. Line cutting violates a duty of fairness to play by the system's rules.¹⁵

14. It might appear that the cases are disanalogous in that in the first case the violation of FCFS requires you to wait a longer amount of time than you otherwise would have had to wait, while this is not true in the second case. This is illusory. The grocery store has made a decision about how to allocate its staff to checkouts and whether those checkouts will be standard or express. The person with a full cart completes her purchase more slowly when there are express checkouts than when there are not because she is not allowed to use those checkouts.

15. John Rawls, *A Theory of Justice*, rev. ed. (Cambridge, MA: Belknap Press of Harvard University Press, 1999), 96–98. See also Herbert L. A. Hart, "Are There Any Natural Rights?," *Philosophical Review* 64 (1955): 175–91; A. John Simmons, "The Principle of Fair Play," *Philosophy and Public Affairs* 8 (1979): 307–37.

John Rawls articulates the principle underlying this duty as follows: “When a number of persons engage in a mutually advantageous cooperative venture according to rules, and thus restrict their liberty in ways necessary to yield advantages for all, those who have submitted to these restrictions have a right to a similar acquiescence on the part of those who have benefited from their submission. We are not to gain from the cooperative labors of others without doing our fair share.”¹⁶ Consider how this might apply to a simple physical queue, such as the queue at a ticket machine. First, note that the social practice of queuing for tickets is an active scheme of cooperation. We follow A. John Simmons in thinking that Rawls intended for this to be interpreted broadly, such that “both a tenant organization’s program to improve conditions in an apartment building and an entire political community’s cooperative efforts to preserve social order” qualify as such schemes of cooperation.¹⁷ Second, the practice of queuing requires everyone in it (bar the very first person) to restrict their liberty by standing in line, rather than going directly to the machine to get served each time they want a ticket. Third, the practice is mutually advantageous in the sense that, over time, it allocates the time and effort of those queuing more efficiently and peacefully than if there were no system and participants simply mobbed the ticket dispenser. As a consequence, willing participants in the cooperative social practice face fewer harms and opportunity costs and receive more benefits than they would were there no practice of queuing.¹⁸ Fourth, as with other cooperative endeavors that give rise to the duty of fairness, FCFS schemes allow for the possibility of free riding. Suppose that the queue is in place and someone arrives at the ticketing location and skips to the front of the line, slipping in while the current user of the machine is just exiting. The line cutter clearly gains from the fact that the others are regulating their behavior according to the social norm: he does not now have to push through a mob of people trying to use the machine. Meanwhile, the interests of every other participant in the queue are set back, since each must now wait for a longer duration. Consequently, they can complain that he is taking advantage of them.¹⁹

This duty of fairness (or “fair play”) has primarily been discussed in relation to the question whether citizens have obligations to obey the laws of their countries, even if they did not give consent to live under those

16. Rawls, *Theory of Justice*, 96. See also Hart, “Are There Any Natural Rights?,” 185.

17. Simmons, “Principle of Fair Play,” 310.

18. We explain the benefits at length in Sec. IV.

19. While we follow Rawls and Simmons in describing the agent as benefitting in our explanation of the duty of fairness, we do not think that it is a necessary condition on violating the duty of fairness that an agent benefit. A line cutter who fails to successfully benefit from cutting in line has still wronged others in the queue. This is one way in which the duty of fairness differs from a duty not to exploit others.

laws. A challenge to those who would ground such political obligations in this duty is that people who benefit from the submission of others to a joint enterprise are not necessarily themselves participants in that enterprise and therefore may reasonably not view the enterprise as a cooperative endeavor in which they participate. Others have argued that the underlying principle of fairness cannot apply to cooperative enterprises which produce unrefusable benefits and therefore would not generate all of the political obligations to the state that we standardly think exist, such as to support public goods.²⁰ Neither of these concerns applies to allocation schemes that use waiting time. Anyone seeking to obtain the goods that are being allocated using waiting time will have reason to see herself as a participant in the relevant social scheme, and anyone can refuse the benefits of waiting in a queue for a resource by not joining the queue. Thus, even those skeptical of the principle of fairness as a ground for political obligations in general need not be skeptical of it as a ground for an obligation to follow queuing rules.

A related issue concerns the duties of allocators of scarce goods where queues are already in place. If waiting time is of no intrinsic moral importance, why are allocators not permitted to violate queuing rules by allocating goods to whoever will most benefit from them? Our view is that when allocators of a scarce good have set up an allocation scheme according to distribution rules, they thereby undertake an obligation to respect the rules of this scheme in virtue of the commitment that they have made to candidates to the good. We take no stance on whether the obligation is possessed for direct moral reasons (i.e., in virtue of the commitment communicated) or indirect moral reasons (i.e., due to the outcome inefficiency that results from creating and then violating expectations).²¹ In either case, there are strong reasons not to revoke benefits where commitments have been made to candidates. Consequently, as with other procedures that have been set up to allocate goods, those who choose to allocate using FCFS should respect the rules of their own queues.

In summary, while there is no general right to FCFS, there is a general right not to be set back by line cutting when an allocation system that uses waiting time exists. Line cutting violates a duty of fairness to play by the system's rules. If a queuing system is sufficiently just, those who participate in that system have such a duty. Moreover, where commitments have been made to candidates, allocators possess a strong reason to respect the rules of allocation in place. However, whether such allocators

20. For variants on these objections, see Nozick, *Anarchy, State and Utopia*, 93–95.

21. Cf. Kamm, who likewise claims that it is wrong to allocate a resource to a line jumper in virtue of the fact that that would break a commitment we have made to a person in line before them. F. M. Kamm, *Morality, Mortality* (Oxford: Oxford University Press, 1993), 1:294–95. See also the discussion of legitimate expectations in Rawls, *Theory of Justice*, 273–77.

ought to set up an FCFS scheme in the first place is a matter governed by considerations of distributive ethics. We now turn to this question.

IV. THE MORAL VIRTUES OF WAITING TIME

In the previous section, we argued that where a sufficiently just system that incorporates queuing exists, participants in the system have a duty to follow the system's queuing rules. In Section II we denied that waiting time is intrinsically important. Being the first person to get to a scarce resource is not ipso facto a moral reason for you to get that resource. This prompts the question whether and how waiting time should be taken into account when setting up an allocation scheme *de novo* or when revising an existing scheme. In this section, we analyze the virtues of allocation schemes that incorporate FCFS with an eye to providing guidance to when and how waiting time should be used.

The cases in which FCFS is invoked very frequently have a similar structure: either the participants in the queue are taken to be roughly equal candidates for the good—they have equal claims to it and will derive similar benefits from it—or the allocator of the good has access to little information about prospective candidates. Where there are known, or easily knowable, morally relevant differences between candidates, FCFS is rarely invoked or is supplemented with additional principles. For example, in triage and organ allocation scenarios candidates' prospects of benefitting are considered in addition to waiting time, in order to separate out classes of roughly equal candidates. When pure FCFS is invoked and there are relevant differences between candidates that could easily be identified, those setting up the scheme may be criticized, for example, on the grounds that it fails to prioritize those in greatest need. Unsurprisingly, we have intuitions to match these social facts: we tend to regard FCFS as a morally justifiable allocation scheme to the extent that candidates have roughly equal claims to the good and candidates are equally well positioned to compete for spots in the queue.

Analyzing the moral virtues of FCFS vindicates these intuitions. We now show that when candidates have equal prospects, pure FCFS schemes are frequently efficient, distribute resources optimally equally among efficient alternative distributions, and treat candidate recipients fairly insofar as it is epistemically equiprobable who will be positioned at each place on a waiting list. These are, we propose, the most morally significant features of queuing systems. This suggests that it will be appropriate to include waiting time as a consideration within an allocation scheme precisely when doing so will better promote efficiency, distribution equality, and fairness than alternative allocation schemes. These virtues will tend to correlate with the degree to which candidates have equal claims to the good and are equally well positioned to compete for spots in the queue. However,

because waiting time has no intrinsic significance, when two allocation schemes differ with respect to their treatment of waiting time but are equivalent in every other morally relevant respect (such as efficiency, equality, and fairness), there is no moral reason to prefer the scheme that prioritizes candidates who have waited for more time over those who have waited for less time.

Three things are important to foreground before our analysis. First, in the following discussion we primarily consider schemes that distribute resources either strictly on the basis of waiting time or not at all on the basis of waiting time. This is a simplification for clarity of exposition. For many allocation contexts, neither sole reliance on waiting time nor ignoring waiting time altogether would be optimal ways to organize an allocation scheme. Instead, a mixed model should be adopted that incorporates waiting time alongside other considerations. Consider the case of allocating scarce organs for transplantation. It is frequently possible to assess the likelihood that an individual will benefit from a specific organ. For example, when a kidney becomes available, candidate recipients can be assessed for how well they are matched to that kidney, which predicts facts such as how likely they are to reject the organ. Relative to the differences in prospects for benefit, the costs of gathering this information about candidate recipients are low. Thus, it might be appropriate to give priority to those who are (much) better matched and so expected to benefit more while still taking waiting time into account when adjudicating between candidate recipients who are (roughly) equally good matches.

Note that strict FCFS schemes, such as individual grocery store checkout lines, necessarily take waiting time into account ordinally. They factor waiting time into allocation decisions only on account of the order in which people are queued, that is, giving no consideration to the length of time candidate recipients have been waiting. Where mixed models are employed, however, there is the option of taking waiting time into account ordinally or cardinally. If we consider waiting time cardinally, then the morally relevant feature of waiting time is not someone's relative position in a queue, but the amount of time they have been in that queue. For example, under the US Organ Procurement and Transplantation Network (OPTN) policy on the allocation of kidneys, patients on the waiting list receive additional points proportional to the amount of time they have spent on the list. We do not analyze the issue of adjudicating between ordinal and cardinal uses of waiting time in mixed models. This, too, can be decided by the moral virtues of waiting time we specify, but there are too many possible mixed models to offer a general analysis in this article.

Second, we do not take up questions about how to conceptualize the goods allocators are distributing. Many resources available for distribution are divisible in multiple ways, such that we can ask how much of a good each successful candidate should receive. For example, we can ask

whether ICU beds ought to be distributed by day or by stay (defined as staying until one's prospect of benefit falls below some threshold).²² Such decisions can make a difference not only to the size of the benefit to be distributed and the expectations that allocators are required to satisfy but also to the time candidates spend in a queue. Such questions are beyond the scope of this article, but given the generality of our account, which applies to any goods to be distributed, the virtues of waiting time apply regardless of the way in which divisible goods are divided. Likewise, we do not explicitly consider the question whether someone who has received a good has more or less claim to receive a similar good (e.g., someone who received a liver transplant but whose new liver then failed). Whether an allocation scheme should provide someone with more of a good (or another good of the type) will depend on whether doing so successfully promotes the values of efficiency, equality, and fairness.

Third, we include little discussion of another putative virtue of allocation schemes that incorporate waiting time—that waiting time tracks desert. This is for two reasons. First, as Perry and Zarsky show, there are many problems with correlating desert and waiting time.²³ The amount of time someone has waited is often involuntary or in no way morally commendable. It is as likely to track societal privilege or the amount of free time someone has as desert. Moreover, even where FCFS tracks desert, it may not track desert's proportionality: if someone gains a little bit of desert for getting out of bed earlier, this is little reason to prefer them for a lifesaving resource, and other considerations will ordinarily trump. Second, many ethicists (including ourselves) doubt that desert should be considered in the context of scarce resources at all. The same is not true of efficiency, equality, and fairness. In summary, while almost everyone agrees that efficiency, equality, and fairness are important in the context of scarce resource allocation, and while there are structural regularities that correlate these virtues with FCFS in the cases in which we find FCFS most pretheoretically compelling, the same is not true of desert.

A. *Waiting Time and Efficiency*

In many cases allocating by waiting time is highly efficient with respect to the benefits generated by the resource that is allocated. This is because many other allocation schemes, such as prioritizing the recipients who will benefit most or conducting a lottery, require us to first identify all

22. For an argument that ICU beds should be distributed by stay, until the point where the patient's prospects of benefitting fall sufficiently far below other candidate recipients' prospects of benefitting, see Leonard M. Fleck and Timothy F. Murphy, "First Come, First Served in the Intensive Care Unit: Always?," *Cambridge Quarterly of Healthcare Ethics* 27 (2018): 52–61.

23. Perry and Zarsky, "Queues in Law," 1614–20.

the potential recipients so that our allocation scheme can take each into account.²⁴ But allocating on the basis of waiting time means that we can begin distributing a resource as soon as the first potential recipient is in line. In situations where either it is an urgent matter that resources are distributed quickly or it is difficult to predict in advance who the candidate recipients for a resource will be, we often do not have the time or the means to identify the potential recipients. For example, when allocating ICU beds, both considerations apply. Eligible patients must be treated quickly to prevent their rapid deterioration, and physicians usually do not know which other people will need an ICU bed or when. Similarly, FCFS's distributive efficiency is of great significance in cases where the value of the resources to be allocated decreases over time. For example, because food, some vaccines, and solid organs all gradually deteriorate, allocating these resources as quickly as possible maximizes the benefit they can provide. In such cases, it is eminently sensible to allocate on an FCFS basis, for doing so most efficiently helps those who are in need.

FCFS is also efficient where it facilitates social coordination, decreasing waste and overhead costs. People often automatically form queues when waiting for some resource—strangers at a buffet may very quickly form a line and settle on its direction, for example—and in such circumstances everyone seems to agree that FCFS is an appropriate allocation scheme. This coordination for mutual benefit relies on there being a small set of salient points at which to start a queue and sufficient background agreement on the practice of queuing.²⁵ However, these conditions are frequently met, allowing resources to be distributed to potential recipients quickly, without argument, and with a minimum of managerial oversight. We can contrast this with lotteries and cost-effectiveness analyses. Such schemes almost never arise spontaneously, and some overhead costs and delays are inevitable.

In some contexts, FCFS allocation schemes are also efficient in virtue of their predictability relative to alternative schemes. If I am in line to board a ride at the carnival, I can see how long the line in front of me is and make a judgment about how long I would have to wait for the ride.

24. An anonymous reviewer suggests the alternative of a system whereby a lottery is conducted if, but only if, there is actual competition. Such a system could be equally efficient in terms of speed of distribution. We would note three points about such an alternative. First, it is no more efficient than FCFS, and likely less efficient given overhead costs. Second, it relies on candidate recipients arriving in batches, such that the participants in each lottery are clearly differentiated from participants who arrive a little later and belong in the subsequent lottery. Third, this is effectively a mixed system: it uses FCFS but with a lottery to decide among candidates who are roughly equal with regard to waiting time or arrival time.

25. Cf. Thomas Schelling's discussion of tacit coordination and tacit bargaining in *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1980), 54–67.

There may even be a sign telling me how long I must wait to board when I am at a given point. Contrast this with a carnival ride whose seats are allocated by lottery among the people waiting. In this case, I do not know how long I must wait to board the ride. In many FCFS schemes—especially physical queues—candidates who know they are unlikely to be first in line can rationally infer that they are either unlikely to get the good under allocation or that they will have to pay a high opportunity cost by queuing to get this good. The predictability of these schemes allows candidates to make decisions about whether to take on the opportunity costs of entering a queue, saving people time and allowing them to get more of what they want. Such systems therefore have, *ceteris paribus*, greater efficiency than schemes such as lotteries, which involve the same costs of entry but whose benefits to each candidate are less predictable.

Perry and Zarsky note a further way in which queuing systems may be more efficient than alternative schemes. They argue that “the fact that someone has achieved a temporal advantage may indicate that his or her efforts were greater than those of the competitors.”²⁶ This exertion of effort may, in turn, serve as a proxy for the individual’s expectation of benefiting, for, in general, individuals will exert greater effort to gain goods that they will benefit from more. For example, the most zealous football fans are those who are most likely to camp out for days at the box office to guarantee good seats.

The features just described tend to make allocation schemes that incorporate the use of waiting time more efficient. However, as noted, this does not entail that the most efficient way to allocate a given scarce resource in any context is to do so solely on the basis of waiting time. For example, for some candidate pools it may be easy to gather further information in order to identify and prioritize groups that are likely to benefit a great deal from the resource over groups that will benefit only a little. FCFS schemes will therefore tend to be more efficient when candidates have similar prospects or when acquiring the information needed to distinguish their prospects is costly. Social coordination problems can be solved in multiple other ways given sufficient explicit directions (though this is rarely costless). The predictability of a queue depends to a great extent on the information that individuals can easily gather about it. And queuing time is not a good proxy for effort exerted if some individuals have an easier time reaching the queuing point than others. These considerations illustrate why mixed models are often preferable in practice and why alternatives to queues, like lotteries and cost-benefit analysis, will sometimes be more efficient.²⁷ Finally, physical queues can sometimes impose negative externalities, in the form of opportunity costs on participants

26. Perry and Zarsky, “Queues in Law,” 1624.

27. See Sec. VI for further discussion of queues versus lotteries.

or by occupying physical space.²⁸ However, these costs are typically small (since large opportunity costs disincentivize queuing) or avoidable (where virtual queues can be employed).

While what is optimally efficient will inevitably depend on context, well-designed queues are generally highly efficient when candidates are roughly equally well positioned to compete for spots in the queue and when candidates have roughly equal prospects or information about candidates is scarce and difficult to acquire.

B. Waiting Time and Equality

Allocation based on waiting time tends to optimize distribution equality.²⁹ In cases in which candidates are relevantly similar—in a way the following analysis will allow us to specify precisely—when a scarce resource is distributed to whoever has waited longest, the resultant distribution of benefits across potential recipients will minimize inequality in the distribution of the goods being allocated compared to other Pareto efficient distributions. For example, allocating by waiting time will produce much more equal outcomes than allocating by lottery. Under divergent conditions, alternative allocation schemes may distribute resources more equally. The closer we are to perfect similarity between candidates, the better FCFS allocation schemes will satisfy the moral desideratum of distribution equality.

To demonstrate these points, we analyze several simple examples and then show how our conclusions about the examples generalize. Consider an organ allocation scenario, S_0 . A neutral body is distributing scarce kidneys among a patient population. Currently, ten people require a kidney transplant, and there are no kidneys available. Each day, one additional person requires a kidney transplant, and one additional kidney becomes available. Each patient deteriorates at the same rate and will not die in the first ten days they require a transplant. Each day spent without a transplant adds additional suffering to the life of the patient with kidney disease. If resources are allocated according to waiting time, after the first ten people receive a kidney, each person must wait precisely ten days to receive a kidney. Given that the patients deteriorate at the same rate, this means that the benefit to each patient is precisely the same. In S_0 , allocation based on waiting time is perfectly egalitarian. No efficient, competing allocation scheme could do better than allocation based on waiting time in terms of distribution equality in S_0 . Indeed, allocation based on a lottery, for example, would perform far worse.

28. For example, if a rivalrous resource is allocated on the basis of FCFS and some effort must be expended to stake a claim, then this may give rise to wasteful races (see Perry and Zarsky, “Queues in Law,” 1628–30).

29. Here “distribution equality” is understood as equal distribution of benefits.

However, even though there is a wait involved, S_0 is not a case of organ scarcity, since each person who needs an organ eventually receives one. To see how allocation based on waiting time operates in conditions of resource scarcity, consider S_s . In S_s , as in S_0 , every person who needs an organ deteriorates at the same rate. And, as in S_0 , in S_s one new organ becomes available every day. However, in S_s , *two* additional patients require an organ each day. In S_s , so long as no one dies from the lack of an organ, allocation based on waiting time produces the matrix given in table 1. Allocation based on waiting time follows a simple pattern in S_s . Each person must wait the number of days equal to half of their numeric position in the queue, rounded down to the next whole number. It is now worth considering whether this allocation is optimally egalitarian. To make this determination, we can simply reallocate an organ and see how it alters the equality of the distribution of burdens and benefits. So, suppose that on day 3, instead of giving the organ to C, we gave that organ to E, and we did not give C an organ until day 5. This would yield the distribution given in table 2. Note that such a decision does not affect the distribution's mean or median waiting time. However, it does affect the distribution's variance.³⁰ The first distribution has a variance of 5.5. The second distribution has a variance of 9.5. We think that any plausible egalitarian theory will imply that one distribution is more equal than another if its variance is lower. Any view that did not have this implication would imply that a distribution can sometimes be made more equal by redistributing resources from the worse-off to the better-off, violating the plausible Pigou-Dalton principle.³¹ So, the first distribution is more equal than the second distribution. In fact, among allocations that minimize total waiting time (i.e., that are efficient), the first of these two allocations is optimally equal.³²

There are several distributions that are more equal than the first distribution in S_s . Such distributions involve holding on to organs and not distributing them to potential recipients immediately. Consider one such distribution, where allocators hold on to the organs they acquire on the first two days and distribute all organs on the third day (see table 3). Assume, for ease of analysis, that the organs themselves do not deteriorate over the course of a few days. This allocation distributes waiting time perfectly equally and so is more egalitarian than allocating by waiting time.

30. The variance of a data set provides a measure of the dispersal of the data points from their mean.

31. Hugh Dalton, "The Measurement of the Inequality of Incomes," *Economic Journal* 30 (1920): 348–61.

32. See J. F. C. Kingman's proof of the same in "The Effect of Queue Discipline on Waiting Time Variance," *Mathematical Proceedings of the Cambridge Philosophical Society* 58 (1962): 163–64.

TABLE 1
ALLOCATION BY WAITING TIME

| Patient | Waiting Begins | Receives Organ | Wait Time (days) |
|---------|----------------|----------------|------------------|
| A | Day 1 | Day 1 | 0 |
| B | Day 1 | Day 2 | 1 |
| C | Day 2 | Day 3 | 1 |
| D | Day 2 | Day 4 | 2 |
| E | Day 3 | Day 5 | 2 |
| F | Day 3 | Day 6 | 3 |

However, the only reason it is egalitarian is because we have leveled many people down. In this distribution, no one has been made better-off and everyone except for F has been made worse-off. Most philosophers believe that we should not level down merely for the sake of outcome equality and would therefore argue that we do not, all things considered, have moral reason to hold the organs and distribute them later.³³ We should instead, as a minimal condition, bring about a Pareto efficient distribution which has less variance than alternative distributions. From all of this, we can see that allocating on the basis of waiting time in S_S is preferable, on account of its efficiency and promotion of equality.³⁴

The models above suggest that allocating scarce resources on the basis of waiting time optimizes distribution equality when each person on the waiting list deteriorates at the same rate and would benefit equally from the resource at each duration waited, or, to generalize, when time spent

33. Nils Holtug, "Egalitarianism and the Levelling Down Objection," *Analysis* 58 (1998): 166–74; Larry Temkin, "Equality, Priority, and the Levelling Down Objection," in *The Ideal of Equality*, ed. Matthew Clayton and Andrew Williams (Basingstoke: Palgrave Macmillan, 2000), 126–61; Derek Parfit, *Equality or Priority?* (Lawrence: University Press of Kansas, 1995). Note that most discussions of the "Levelling Down Objection" focus on the moderate egalitarian view that sometimes leveling down a group's welfare makes that distribution in one way better, not on the logically stronger and more radical egalitarian view that leveling down the welfare of a group may sometimes make that distribution better, all things considered. This stronger view, which regards some Pareto improvements as morally bad, is ordinarily ignored in such discussions, but it is thought to be a worse implication than that of moderate egalitarianism. On the stronger view, see Marc Fleurbaey, "English or Esperanto: A Case for Levelling Down?," in *Arguing about Justice: Essays for Philippe van Parijs*, ed. Axel Gosseries and Yannick Vanderborght (Louvain-la-Neuve: Presses Universitaires de Louvain, 2011), 229–36. Fleurbaey claims that there is "no doubt that efficiency considerations preclude the judgment that a leveling down operation improves the situation" (233).

34. Even if people are liable to die on the waiting list, provided that their future prognoses are roughly equal at the time they get onto the waiting list, allocating according to waiting time will still produce the most equal distribution of benefits.

TABLE 2
ALLOCATION BY WAITING TIME WITH REALLOCATION

| Patient | Waiting Begins | Receives Organ | Wait Time (days) |
|---------|----------------|----------------|------------------|
| A | Day 1 | Day 1 | 0 |
| B | Day 1 | Day 2 | 1 |
| C | Day 2 | Day 5 | 3 |
| D | Day 2 | Day 4 | 2 |
| E | Day 3 | Day 3 | 0 |
| F | Day 3 | Day 6 | 3 |

waiting for a resource is (cardinally) an equal counterfactual harm for each person compared to receiving the resource.³⁵ Other efficient allocation schemes, such as lotteries, yield distributions that are less equal or, at best, exactly as equal. This has the following practical upshots: When an allocation scenario meets the conditions set out in this section, such that time spent on the waiting list is cardinally equally counterfactually bad for each person, distributing according to waiting time produces the most equal, efficient distribution of benefits. And while we have argued that many commonsense FCFS schemes approach these conditions in practice, different FCFS allocation scenarios can be closer to or farther from this idealization. The greater the difference in how bad waiting is for each person, and the more information we have about the differences in how bad waiting is for each person, the less equal (and efficient) is allocation on the basis of waiting time. The more similar prospective recipients are to one another, and the less information we have about their differences, the better allocation by waiting time does in terms of expected distribution equality.

To illustrate the range of allocation scenarios, consider the information that may be readily available about candidate recipients for kidney transplants. As mentioned above, the entity making kidney allocation decisions will receive data about patients—such as regarding their blood type and sensitivity to specific antigens—that allows rough estimates of how likely they are to reject a particular donated kidney. Certain candidate recipients are especially hard to find a suitable donor for. In the United States, those patients are given higher priority on the rare occasions when a good match is found, since other patients are likely to match many more kidneys.³⁶ Given how great the effect of matching is on the extent of the benefit from a donated kidney, this seems to be a

35. Time spent waiting for a resource is cardinally an equal counterfactual harm for each person compared to receiving the resource precisely when, for any duration of waiting time d , if two candidate recipients have waited for d , they would be worse-off by the same amount if they failed to receive the resource.

36. OPTN/UNOS, *OPTN Policies*.

TABLE 3

ALLOCATION WITH EQUAL WAITING TIME

| Patient | Waiting Begins | Receives Organ | Wait Time (days) |
|---------|----------------|----------------|------------------|
| A | Day 1 | Day 4 | 3 |
| B | Day 1 | Day 4 | 3 |
| C | Day 2 | Day 5 | 3 |
| D | Day 2 | Day 5 | 3 |
| E | Day 3 | Day 6 | 3 |
| F | Day 3 | Day 6 | 3 |

clear case where we might do better in terms of equality by taking more than waiting time into account. On the other hand, sometimes the information available about competing potential recipients might be trivial compared to the benefits of allocation by FCFS. For example, if the entity allocating a kidney knew only that half of the good matches have had a seasonal cold, we would likely judge that information irrelevant. Equalizing disparities in quality of life from colds will be of trivial importance compared to equalizing disparities in waiting time, which tracks time spent under fatigue, high blood pressure, malaise, and a general failure to thrive due to kidney disease, as well as the severity of these states due to the progression of the disease.

Finally, it is worth noting a second way that FCFS—like lotteries—tends toward egalitarianism that is not shared by some other allocation schemes. Many allocation schemes could be perfectly fair in one-off allocation scenarios but would lead to very inequalitarian results in iterated schemes. For example, on the account of fairness we defend in the proceeding section, an allocation scheme which selects the candidate whose name comes first in the alphabet may in principle be perfectly fair, since it is possible that allocators will have equal rational credence in each candidate having the earliest name. However, allocating repeatedly on the basis of who has the earliest name would systematically favor certain people (as those with last names late in the alphabet will attest). While there may in some contexts be background injustices which systematically advantage some people in FCFS schemes (which we discuss further in Sec. VI), there is no general correlation between individual people and particular positions in queues. As such, absent these background injustices, the use of FCFS in iterated schemes leads to greater equality than the use of many other allocation criteria.

C. *Waiting Time and Fairness*

We believe that allocation based on waiting time, among equally deserving candidates for a good, is fair. While there are many normatively

important concepts of fairness (including the notion previously invoked in “fair play”), we mean this in two particular senses: it offers all candidates equal chances without regard to morally irrelevant characteristics, and it expressively signals the equal moral standing of all candidate recipients of a good.

First, allocation based on waiting time tends to be fair in that it does not privilege certain people over others on the basis of morally irrelevant factors. In this regard it is preferable to many actual allocation schemes in which priority is given to those who have more resources or greater social status. Consider, for example, the discretion that most medical specialists in the United States have with regard to scheduling patients. If a dermatologist wants to give preference to her golfing buddy or her spouse’s business contact, then she can book him in the next day—before the first regularly scheduled appointment or by rescheduling someone less well connected. An FCFS system does not allow such favoritism. In the presence of economic inequality, FCFS allocation schemes are also fairer than schemes that allocate goods based on willingness to pay. Such schemes favor those who have more money at their disposal. We have a reason based on fairness to prefer FCFS allocation schemes over alternative schemes like these, which arbitrarily advantage some candidates for a scarce resource over others.

FCFS tends toward perfect fairness in cases in which it is *epistemically equiprobable* which potential recipient will receive the scarce resource that is being allocated. This means that FCFS is perfectly fair when the allocator does not know ahead of time who will receive the goods at stake, such that from the perspective of the allocator each prospective recipient has an equal expectation of benefit. For example, suppose we set up the refreshments table outside a meeting room so that the first people out of the room and into the line will get their tea first. We do not know how people will arrange themselves in the room, nor do we have reason to think that certain people or classes of people are more or less likely to be far from the exit or much slower to leave. The allocation of refreshments by FCFS is therefore fair. Cases of perfect epistemic equiprobability are likely to be rare. However, it is much more common that allocators can achieve rough epistemic equiprobability within strata of eligible recipients. So, for example, among the patients who have been triaged as less urgent in an ER, it may be pretty much random from the point of view of the triage nurse who had to wait the longest or shortest amount of time from the point when they decided to seek care.

The claim that allocation based on waiting time is fair in the sense we have described is controversial. Some philosophers argue that when selecting between equally deserving potential recipients, fairness requires statistical equiprobability, paradigmatically satisfied by formal lotteries. Moreover, a common objection to the use of waiting time in allocation

decisions is that it is unfair in practice. We address these concerns in the following two sections.

FCFS allocation schemes additionally play an important expressive function, signaling the equal moral and political standing of each candidate recipient.³⁷ If everyone who wants a good must get into the same line for that good and follow the same waiting procedure, then everyone's claim is treated equally. This contrasts with allocation schemes in which people of high social status can simply bypass the system that ordinary people must operate within. The difference is nicely illustrated by Stuart Corbridge's description of queues for government services in rural eastern India:

Nevertheless, many poorer people—and poor women especially—can be seen waiting patiently outside government buildings like the Block Development Office. . . . It is not unusual for them to be kept waiting for hours, and sometimes for days. During the course of this wait they will often see local political bosses (*netas*) storm into the office of the highest ranking local government official to demand an audience. . . . Waiting is something that poorer people do, and more than once we witnessed *adivasi* (tribal) women standing in the sun or rain for hours waiting their turn to see “*sarkar*” (government), sometimes refusing to go for lunch in case they “lost their place.”³⁸

Here the queue functions as an allocation scheme only among people of equal social standing. Those queuing are treated equally and so are accorded equal respect. On Corbridge's interpretation, those who ignore the queue thereby express a contrary view about the relative status of the people who want the good being allocated: “Members of the upper castes (mainly men, it should be said) simply do not recognize the claims to equality that standing in line would seem to imply. Just as importantly, nor do many government officers.”³⁹ This provides an additional reason based on justice to prefer FCFS allocation schemes to allocation schemes which signal the unequal status of candidate recipients.

37. David Wasserman argues that lotteries have a similar expressive function: “In order to express an equal commitment to each claimant's receiving the scarce good, those who control the allocation process must have no reason to believe that any claimant is more likely than any other to satisfy the payoff condition.” David Wasserman, “Let Them Eat Chances: Probability and Distributive Justice,” *Economics and Philosophy* 12 (1996): 29–49, 31.

38. Stuart Corbridge, “Waiting in Line, or the Moral and Material Geographies of Queue-Jumping,” in *Geographies and Moralities: International Perspectives on Development, Justice and Place*, ed. Roger Lee and David M. Smith (Malden, MA: Blackwell, 2009), 183–98, 184.

39. Corbridge, “Waiting in Line,” 186.

There is much to be said in favor of FCFS. Under the conditions where we ordinarily find FCFS most compelling, allocation by waiting time is efficient, produces equal outcomes, and treats potential recipients fairly. For these reasons, waiting time ought to be incorporated into allocation schemes when doing so improves the balance of efficiency, equality, and fairness, and FCFS is a morally credible default allocation criterion in the absence of known background inequality.

D. When Does Waiting Time Begin?

For simplicity of analysis, we have assumed that FCFS identifies a unique, determinate ordering of potential candidates by arrival time. But it is sometimes the case that FCFS does not identify a unique ordering of candidates, since who is first in the queue depends on when the queue begins and that is not a natural fact.⁴⁰ For example, in the context of ICU care, “the first person” could mean the first person to have become ill, the first person to have arrived at the hospital, or the first person to have been diagnosed as in need of ICU care. FCFS would have different distributive implications depending on how we took the queue to be ordered. Although each of these three variants of FCFS would have some of the delineated moral virtues of FCFS—to the extent that the queue has low administrative costs, waiting is cardinally equally counterfactually bad for each person, the system satisfies epistemic equiprobability, and so on—they would have these virtues to varying degrees. This means that the moral virtues of allocating according to FCFS depend on how we specify the criteria for determining waiting time.

A key advantage of the account that we have offered is that it provides selection criteria for making such determinations. If waiting time had intrinsic moral significance, as we considered in Section II, then there would need to be a privileged, unique way of determining how long people have waited for a good. We have no need to make such problematic determinations. Given that waiting time has no intrinsic moral significance, and given that the reasons that allocators could have in favor of starting an FCFS scheme *de novo* are based on the efficiency, fairness, and equality of such a scheme, the appropriate waiting criteria are decided by distributive ethics. Where there are multiple candidate starting points for an FCFS allocation of scarce goods, allocators ought to choose the starting point that they judge achieves the optimal balance of the moral virtues we have considered in this section. Where there is an apparent tie, the benefits of efficient coordination support allocators choosing one of the optimal starting points and consistently adhering to it.

40. Thanks to an anonymous reviewer for raising this point.

V. WAITING TIME AND THE LOTTERY REQUIREMENT

We have claimed that use of waiting time in allocation decisions may be justified on account of, inter alia, its fairness. However, many philosophers endorse what Tim Henning has called “the Lottery Requirement”: when deciding between equally deserving candidates, procedural fairness requires us to give each candidate “an equal chance . . . by holding a lottery.”⁴¹ Such lotteries are typically characterized by ostension, and standard examples of fair lotteries in the literature include flipping a coin, spinning a wheel, and drawing lots.⁴² Among these philosophers, some are explicit that only statistically equiprobable lotteries can satisfy requirements of fairness.⁴³ For example, Kornhauser and Sager argue that “an equiprobable lottery has the unique capacity to divide the good probabilistically, making possible the allocation of a ½ chance of receiving the good to each candidate.”⁴⁴ What such writers have in mind here is that “well-chosen fixed rules” (in John Broome’s phrasing), such as “choose the candidate whose birthday is closest to today” or “choose the candidate with the warmest shirt tone,” are unfair in principle.⁴⁵ Although such rules may provide candidates with equal epistemic chances from the perspective of allocators, the objective probability that a particular candidate has the closest birthday or the warmest shirt tone is either 1 or 0. Such

41. Tim Henning, “From Choice to Chance? Saving People, Fairness, and Lotteries,” *Philosophical Review* 124 (2015): 169–206, 169. On the Lottery Requirement, see John Broome, “Fairness,” *Proceedings of the Aristotelian Society* 91 (1990): 87–101; James F. Childress, “Who Shall Live When Not All Can Live?,” *Soundings* 53 (1970): 339–55; Norman Daniels, “Reasonable Disagreement about Identified vs. Statistical Victims,” *Hastings Center Report* 42 (2012): 35–45; Peter A. Diamond, “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment,” *Journal of Political Economy* 75 (1967): 765; Jon Elster, *Solomonic Judgements: Studies in the Limitation of Rationality* (Cambridge: Cambridge University Press, 1989); Barbara Goodwin, *Justice by Lottery* (Luton: Andrews UK, 2013); Kamm, *Morality, Mortality*; Lewis A. Kornhauser and Lawrence G. Sager, “Just Lotteries,” *Information (International Social Science Council)* 27 (1988): 483–516; Ben Saunders, “The Equality of Lotteries,” *Philosophy* 83 (2008): 359–72; George Sher, “What Makes a Lottery Fair?,” *Noûs* (1980): 203–16; Peter Stone, *The Luck of the Draw: The Role of Lotteries in Decision Making* (Oxford: Oxford University Press, 2011); Wasserman, “Let Them Eat Chances.”

42. These are the examples invoked by both Henning (“From Choice to Chance?”) and Kornhauser and Sager (“Just Lotteries”).

43. This view is defended by Broome (“Fairness”) and Kornhauser and Sager (“Just Lotteries”). Kornhauser and Sager distinguish equiprobability (nonepistemic equiprobability) from impersonality (epistemic equiprobability), arguing that “some moral theories, under some circumstances,” are such that “equiprobability is a necessary as well as a sufficient condition of a fair lottery” (“Just Lotteries,” 491). It is also discussed in Elster, *Solomonic Judgements*; Sher, “What Makes a Lottery Fair?”; Stone, *Luck of the Draw*; and Wasserman, “Let Them Eat Chances.”

44. Kornhauser and Sager, “Just Lotteries,” 491.

45. Broome, “Fairness,” 88.

a view motivates the concern that FCFS schemes may also necessarily involve some unfairness. We here argue that this is not so.

In what sense do randomization procedures like rolling a die or tossing a fair coin give candidate recipients equal chances? We think that such randomization procedures have two morally relevant features. The first is that they satisfy epistemic equiprobability. If it is epistemically equiprobable that either of two candidates will receive a good, then, from the perspective of the allocators of that good, it is equally likely that either will receive it. The notion of epistemic equiprobability that we employ is derived from Johann Frick's evidence-relative notion of epistemic probability: "To say that, from the perspective of an agent, some event e has an epistemic probability p of occurring is to say that, given the evidence available to the agent, her rational degree of credence in the proposition ' e will occur' would be p ."⁴⁶ If there are two equally deserving candidates for a good, then an unbiased coin can give them equal chances in the sense that the allocator is rational to believe of neither that they are more likely to get the good.

The second normatively important feature is that when an allocator uses these randomization procedures, the facts which determine who wins are not fixed characteristics of the potential recipients.⁴⁷ Nothing about Andrea makes her any more likely to win an unbiased coin toss than Bathsheba. To see why this matters, recall how allocating according to a rule like "choose the person whose name comes first in the alphabet" can lead to inegalitarian results in iterated schemes. Andrea will always win when a rule like this is used because it is a fixed feature of her name that it comes before Bathsheba in the alphabet.

As we noted in Section IV, well-chosen fixed rules like FCFS are able to achieve epistemic equiprobability. It is sometimes the case that when allocators of a good distribute it among potential recipients according to a fixed rule they have no undefeated evidence that any candidate is more likely to receive that good than any other candidate. Recall the refreshments table. If I distribute refreshments by letting the first people out of the conference room get tea first, so long as I have no evidence that certain people are less likely to be close to the exit or able to exit swiftly, and so my rational degree of credence in each person's getting to the table first is equal, it is for me epistemically equiprobable who among attendees will get tea first.

46. Johann Frick, "Contractualism and Social Risk," *Philosophy and Public Affairs* 43 (2015): 175–223, 182. We are skeptical of the possibility of nonextremal, nonepistemic chances as a result of arguments from David Lewis and Jonathan Schaffer. See David Lewis, "A Subjectivist's Guide to Objective Chance," *Philosophical Papers* 2 (1986): 83–132; and Jonathan Schaffer, "Deterministic Chance?," *British Journal for the Philosophy of Science* 58 (2007): 113–40.

47. Thanks to an anonymous reviewer for making this point.

Unlike many other fixed rules, FCFS can also avoid correlations between fixed characteristics of potential recipients and their place in a queue. To achieve this, it must be the case that there is no relationship between people's fixed characteristics and their access to information about desirable queues or their ability to act on this information, and the order of candidate recipients in a queue must be the same as the order in which they decided to join the queue. For example, suppose that a short-staffed hospital gives all of its inpatients buttons to press when they require some nonurgent attention from a nurse. Patients are seen in the order in which they press the button. On the plausible hypothesis that each can press their button with ease upon deciding to request assistance, none is disadvantaged in getting a place in the queue by fixed characteristics they possess.⁴⁸

A proponent of the Lottery Requirement might argue that these two features do not exhaust the morally relevant senses in which lotteries can give candidates equal chances in ways other allocation schemes do not. Lotteries might also instantiate statistical equiprobability in the sense that they might give candidates statistically, or objectively, equal chances to receive the good, such that n candidates each have a $1/n$ chance in some nonepistemic sense.⁴⁹

Here we think that defenders of the Lottery Requirement face a dilemma. If they accept that the Lottery Requirement is satisfied by actual lotteries, such as drawing lots or tossing a fair coin, then they must also accept that it is in principle satisfiable by FCFS schemes. This is because queuing behavior, like any actual lottery, is constituted by ordinary physical events subject to statistical generalizations. Inasmuch as there are true statements assigning nonepistemic fractional odds to future physical events such as lottery outcomes—on account of indeterminism in either fundamental or nonfundamental ontology—there are also true statements assigning nonepistemic fractional odds to other ordinary future physical events such as queuing arrangements.⁵⁰ The same physical

48. One might be concerned that even in this case there will be some slight difference between patients in the time it takes for them to move from deciding to get in the queue for assistance and successfully pressing a button. Suppose instead, then, that brain sensors detect when patients decide to request assistance. At the point of decision they are put into the queue and the delay between decision and allocation to a place in line is precisely the same for each patient. This more fanciful example shows that there is not a problem in principle with FCFS meeting the criterion of fairness.

49. Ian Hacking, "Duality," in *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction, and Statistical Inference* (Cambridge: Cambridge University Press, 1975), 11–17, 12–13.

50. For one account of nonepistemic, nonfundamental, innocently emergent chances, see Barry Loewer, "David Lewis's Humean Theory of Objective Chance," *Philosophy of Science* 71 (2004): 1115–25.

laws from which we can derive a 0.5 statistical likelihood of a coin landing heads when tossed generate fractional statistical likelihoods for the occurrence of future events such as weather states, neural action potentials, dropped objects shattering, and people gaining information about desirable queues, given some initial conditions. In some physically possible cases, under the initial conditions where an allocating body sets up an FCFS scheme, it will be true that each candidate has identical fractional odds of finally occupying each respective position in the queue, thus satisfying statistical equiprobability.⁵¹ If, on the other hand, defenders of the Lottery Requirement deny that it is satisfied by actual lotteries, they must embrace the absurd result that it is unfair to allocate an indivisible good among two equally deserving parties by tossing an unbiased coin.

We claim that the two virtues of actual lotteries that support the Lottery Requirement are epistemic equiprobability and their independence from fixed characteristics of people. No further requirement of statistical equiprobability can be met by actual lotteries without also, in principle, being met by FCFS. On any acceptable theory of equiprobability, FCFS is in principle capable of offering equal chances to equally deserving recipients.

VI. WAITING TIME IN PRACTICE: THE UNFAIRNESS OBJECTION

Even if the use of FCFS in an allocation scheme offers fair chances in principle, it might be unfair in practice. Indeed, this is the chief objection philosophers have raised against allocation based on waiting time. Norman Daniels claims that “‘first-come, first-served’ . . . favours patients who are better off and better educated.”⁵² Similarly, Govind Persad et al. argue that “first-come, first-served allows morally irrelevant qualities—such as wealth, power, and connections—to decide who receives scarce interventions, and is therefore practically flawed.”⁵³ These writers prefer allocation by lottery, which they claim is more resistant to corruption.

Daniels and Persad et al. are correct that FCFS is sometimes unfair, as it is sometimes true that who gets in line first correlates positively with societal privilege. For instance, people who are more mobile—say, on

51. This will be true in various physically possible cases, such as cases where indifference is satisfied (i.e., everyone is equidistant from the queuing point, everyone approaches the queue at the same rate, etc.) and probabilistic independence is satisfied (i.e., no two candidates tend to arrive together). Such cases, iterated, will result in identical frequencies for each possible queuing outcome. Thanks to Isaac Wilhelm for helping us clarify this point.

52. Daniels, “Reasonable Disagreement,” 169.

53. Persad, Wertheimer, and Emanuel, “Principles for Allocation,” 424. See also Den Hartogh, “Trading with the Waiting-List,” 196.

account of being healthy and nondisabled—may have an advantage in physical queues, where they may reach the queuing point more quickly than others and experience less of a burden from waiting for extended periods. The wealthy and well-connected also may be able to get onto waiting lists earlier than other people. For example, one criticism of the use of waiting lists for solid organ transplants in the United States is that some people who need transplants get onto the list more quickly than others because they have better access to medical care.⁵⁴ In such cases, certain individuals and groups are likely to benefit more than others. Equally deserving prospective recipients of the scarce resource do not have epistemically or statistically equiprobable chances of getting the resource.

However, these problems with allocation schemes that employ FCFS are neither universal nor inevitable. For example, in scarce organ allocation schemes, if everyone has access to good-quality health care, equal opportunities to make checkup appointments, and education about when such appointments are appropriate, a rule incorporating FCFS will not discriminate between candidates on the basis of morally irrelevant features. When these conditions are not satisfied, mitigating policies can be incorporated to make the system fairer. For example, patients with lower levels of health literacy can be given greater support. To take a more quotidian example, a physical queue may constitute a barrier for people with mobility impairments and thereby exacerbate the disadvantages they experience. But the scheme might allow line jumping for such people to decrease waiting time and increase their expectation of benefit, or virtual queues might be employed rather than physical queues to increase equality of access.

Moreover, many problems of unfairness that plague some FCFS allocation schemes also plague allocation by lottery. For even lotteries can be gamed by the rich and powerful. For example, very wealthy people may game the US organ allocation system by getting onto waiting lists in multiple locations.⁵⁵ In effect, they are gaming a natural lottery for scarce organs: in each location there is a chance that a matched organ will become available, and so being on multiple lists increases one's chances.

Lotteries and well-chosen fixed rules like FCFS can be impartial ways to distribute scarce resources, but all such schemes can nonetheless be unfair in practice. The appropriate system for allocating scarce resources will depend on the context. FCFS is not universally less fair than lotteries, in principle or in practice; neither are lotteries universally less

54. Joy Victory, "Need an Organ? It Helps to be Rich," *ABC News*, January 20, 2006.

55. Raymond C. Givens et al., "Outcomes of Multiple Listing for Adult Heart Transplantation in the United States: Analysis of OPTN Data from 2000 to 2013," *JACC: Heart Failure* 3 (2015): 933–41.

fair than FCFS. To best satisfy the aim of procedural fairness, allocators of a resource must pay careful attention to the ways in which a specific scheme may systematically benefit some groups more than others. Careful oversight and leveling policies can help to ensure that the socially advantaged do not inappropriately exploit their privilege to gain an edge.⁵⁶

In evaluating alternative allocation schemes, it is also important to bear in mind the different values that the scheme ought to take into account. Consider the ICU triage case discussed in Section IV. Due to the fluid nature of arrivals, among patients of roughly equal urgency it will be much more efficient to have them seen by a physician according to FCFS than to wait until a group of patients are gathered and then conduct a lottery. Or, for a more everyday case, consider a bakery where customers take a ticket and are served in order. This is likely to be much more efficient (and much more egalitarian) than having a lottery every twenty minutes. In many cases, we suspect that this increased efficiency will be preferable even if some more privileged people are able to obtain the benefits somewhat more quickly.

VII. CONCLUSIONS

Despite the widespread use of waiting time in allocation schemes, the normative properties of queues have received little attention from normative theorists. In this article, we have attempted to remedy that gap. We have argued that waiting time is not intrinsically morally relevant. Nonetheless, where waiting time has been incorporated into allocation schemes that are sufficiently just, participants in those schemes have a duty of fairness to follow the rules. The use of waiting time has moral virtues that are relevant to policy makers considering creating or revising allocation schemes. Waiting time is valuable on account of its promoting efficiency, equality, and fairness among roughly equal candidates for a good. FCFS may be the most efficient means of distributing a scarce resource, as in the case of allocating ICU beds. It may be optimally egalitarian, as when it is used in certain organ allocation schemes. It may be perfectly fair, as in the case of the refreshments table. However, waiting time is never more than pro tanto morally important, and it is sometimes morally permissible or required to use multiple allocation principles for a given resource, or to exclude waiting time entirely in favor of an alternative scheme. Typically, considerations additional to waiting time ought to be taken into account, either because other morally relevant factors need

56. For example, UNOS changed their listing procedures for heart transplant recipients in 1999 in order to reduce “gaming” of the system. This appears to have been successful. See Dennis P. Scanlon et al., “Does Competition for Transplantable Hearts Encourage ‘Gaming’ of the Waiting List?” *Health Affairs* 23 (2004): 191–98.

to be captured or because waiting time is an imperfect proxy for efficiency, fairness, and equality. Despite disparagement of FCFS allocation schemes in the literature on account of their presumed unfairness, FCFS is frequently a just scheme of resource allocation, and waiting time has unique normative properties which frequently justify its incorporation into allocation schemes.